DISS. ETH NO. 27846

Next Generation Active Magnetic Shielding for n2EDM and Axion-Like Particle Search

A thesis submitted to attain the degree of

DOCTOR OF SCIENCES of ETH ZÜRICH

(Dr. sc. ETH Zürich)

presented by

SOLANGE FRANCINE EMMENEGGER

Msc. Sc. ETH in Physics

born on

11.01.1995

citizen of

Schüpfheim (LU)

accepted on the recommendation of:

Prof. Klaus Kirch

Prof. Stéphanie Roccia

Prof. Hans-Arno Synal

2021

Summary

There are still many open questions to which physicists are hunting for answers all over the world. Two of the most interesting ones concerning the creation of the universe are: Why is there more matter than anti-matter? What is dark matter? The Standard Model of particle physics, which was established and formed by many physicists does not have an answer to those questions. Nevertheless, it is the most accurate theory known, and provides precise description of a lot of phenomena.

One of the criteria laid out by Sakharov which needs to be met in order to cause an excess of matter over antimatter is the violation of the CP-symmetry (short for charge and parity). One observable which is particularly sensitive to CP violation is the neutron electric dipole moment (nEDM). It is of great interest as the electro-weak sector of the Standard Model predicts an exceedingly small value, orders of magnitude smaller than the range of current state-of-the-art experiments. However, are the predictions coming from Beyond Standard Model theories for the nEDM within the reach of those experiments, and they can therefore exclude certain theories which try to explain alternative sources of CP violation.

The experiment holding the current world's best upper limit on the nEDM was hosted at the renowned Paul Scherrer Institute (PSI) in Villigen, Switzerland. The nEDM collaboration standing behind this experiment is already working on the construction of the follow-up experiment n2EDM, which seeks to improve the sensitivity to the nEDM by another order of magnitude, with possibilities for even further improvement. Both experiments use ultra-cold neutrons and stand in the long tradition of nEDM experiments since the fifties of the last century. Over the course of time it became evident that the control of magnetic fields in- and outside the experiments is crucial to their success. The working principle of the nEDM experiments and their motivation in connection with CP-violation is the subject of Part I.

Apart from the magnetically shielded room (MSR), which provides a quasi-static shielding factor of 10^5, another innovation of n2EDM is the next generation active magnetic shielding (AMS). The strength of the system is that it is able to attenuate the earth magnetic field and disturbances from nearby magnetic source to around 1 µT on the surface of the MSR, and this with remarkably little space available. Because it consists of eight coils, each producing a basis field of the Cartesian harmonics, it is able to react flexibly to a large range of magnetic fields.

Nevertheless, prior to the design phase, the expected magnetic fields in the experimental area were measured with a 10 m high mapping tower. The outcome of this mapping campaign provided precious input to the design which is based on the "Method of Simple Coil Design". Many constraints, such as interfaces with other parts of the apparatus, feasibility of the construction and thermal output had to be considered. The coil system, the construction of which was finished in March 2021, weighed more than 1.3 t and approximately 50 km of cables were installed.

Prior to the construction, the validation of the working principle was undertaken with a prototype system at the ETH laboratory. It consisted of similar eight coils but was built in a small-scale version. It could be validated that the design goal of 1 µT homogeneity can be met and as a preparation for the commissioning of the n2EDM AMS, the positioning of the magnetic feedback sensors was studied. The n2EDM AMS and the prototype are treated in Part II.

At the opening of this section we also considered the nature of dark matter. At the moment it is only known that dark matter makes up for 85% of the matter content of the universe. There are many suggestions for models and hypothesised particles which could be dark matter, among which

the axion is a promising candidate. A sufficiently light axion would behave as a classical coherently oscillating field and can couple to the gluon. When it interacts with the gluons of the neutrons in our experiment, the axion field would imprint an oscillating signature in the dataset. This signature was sought for in 2017 in the still blinded nEDM dataset and the first laboratory constraint of the axion-gluon coupling emerged from this analysis. In Part III the improvements on this limit by the reanalysis of the unblinded nEDM dataset are presented.

Zusammenfassung

Es gibt immer noch viele offene Fragen für welche Physiker weltweit nach Antworten suchen. Die zwei interessantesten Fragen, welche die Entstehung des Universums betreffen sind: Wieso gibt es mehr Materie als Antimaterie? Was ist dunkle Materie? Das Standard Modell der Teilchenphysik, welches über Generationen von Physikern entstanden ist, hat keine Antwort auf diese Fragen. Trotzdem ist es die genauste bekannte Theorie und liefert akkurate Beschreibungen von vielen Phänomenen.

Eines der von Sakharov aufgestellten Kriterien, das erfüllt sein muss, damit ein Überschuss von Materie über Antimaterie entstehen kann, ist die Verletzung der CP-Symmetrie (für charge - Ladung und parity - Parität). Eine Observable welche besonders sensitiv auf die CP-Verletzung ist, ist das elektrische Dipolmoment des Neutrons (nEDM). Es ist von grösstem Interesse, da das Standard Modell nur ein verschwindend kleines Dipolmoment voraussagt, welches mehrere Grössenordnungen kleiner ist als die Genauigkeit von heutigen modernen Experimenten. Jedoch sind Voraussagen von Theorien jenseits des Standardmodells der Teilchenphysik (BSM) in Reichweite von aktuellen Experimenten und können deshalb gewisse Theorien, welche versuchen alternative Quellen von CP-Verletzung zu erklären, ausschliessen.

Das Experiment welches das aktuell weltbeste obere Limit des nEDM hält, stand am renommierten Paul Scherrer Institut (PSI) in Villigen in der Schweiz. Die nEDM Kollaboration welche hinter diesem Experiment steht, arbeitet am Aufbau des Nachfolgeexperiments n2EDM, welches die Sensitivität nochmals um eine Grössenordnung verbessern will. Zusätzlich gibt es sogar Möglichkeiten für weitere Verbesserungen. Beide Experimente verwenden ultra-kalte Neutronen (UCN) und stehen in der langen Tradition von nEDM Experiment die seit den fünfziger Jahren des letzten Jahrhunderts gebaut wurden. Im Laufe der Zeit hat sich herauskristallisiert, dass die Kontrolle von den Magnetfeldern im und ums Experiment entscheidend ist für den Erfolg der Experimente. Die Funktionsweise des nEDM-Experiments und seine Motivierung in Verbindung mit der CP-Verletzung ist Thema in Teil I dieser Arbeit.

Neben der magnetisch abgeschirmten Kammer (MSR) welche einen quasi-statischen Abschirmungsfaktor von 10^5 bietet, ist die aktive magnetische Abschirmung (AMS) nächster Generation eine andere grosse Innovation. Die Stärke des Systems liegt in der Fähigkeit das Erdmagnetfeld und magnetische Störungen von Quellen in der Nähe auf etwa 1 µT auf der Oberfläche des MSR zu dämpfen und dies erst noch auf sehr engem Raum. Weil es aus acht Spulen besteht, welche je ein Basisfeld der kartesischen Harmonischen Polynome produzieren, ist das System fähig, flexibel auf eine grosse Bandbreite von magnetischen Feldern zu reagieren. Trotzdem wurde vor dem Start der Designphase das erwartete Feld in der Experimentierhalle mit einem 10 m hohen Mappingturm vermessen. Das Resultat dieser Mappingkampagne hat wertvollen Input für das Design, welches auf der "Methode für einfaches Spulendesign" basiert, geliefert. Viele Einschränkungen, wie z. B. die Schnittstellen zu anderen Teilen das Experiments, Realisierbarkeit und Wärmeentwicklung mussten berücksichtigt werden. Das Spulensystem, dessen Bau im März 2021 vollendet wurde, wiegt mehr als 1.3 t und etwa 50 km Kabel wurden installiert.

Vor der Installation des Spulensystems wurde die Validierung des Designprinzips an einem Prototypen im ETH Labor durchgeführt. Der Prototyp besteht aus ähnlichen acht Spulen wie das n2EDM AMS wurde jedoch in einer kleineren Version gebaut. Es konnte bestätigt werden, dass das Designziel einer 1 µT Feldhomogenität erreicht werden kann und als Vorbereitung auf die Inbetriebnahme des n2EDM AMS wurde die Positionierung der magnetischen Feedbacksensoren untersucht. Das n2EDM

AMS und der Prototyp werden im Teil II behandelt.

Zur Einleitung dieser Zusammenfassung wurde auch nach der Natur der dunklen Materie gefragt. Zurzeit weiss man nur, dass dunkle Materie mehr als 85% des Materieinhalts des Universums ausmacht. Es gibt viele Vorschläge für Modelle und hypothetische Teilchen welche dunkle Materie ausmachen können. Unter diesen ist auch das Axion ein vielversprechender Kandidat. Ein genügend leichtes Axion würde sich als klassisches oszillierendes Feld verhalten und könnte an das Gluon koppeln. Wenn es mit den Gluonen in den Neutronen in unserem Experiment interagiert, würde das Axionenfeld eine oszillierende Signatur im Datensatz hinterlassen. Nach dieser Signatur wurde 2017 im verblindeten nEDM Datensatz gesucht und das erste Laborlimit der Axionen-Gluonen Kopplung ging aus dieser Analyse hervor. In Teil III wird die Verbesserung dieses Limits durch eine Reanalyse des entblindeten nEDM Datensatzes präsentiert.

Contents

Part I

Neutron Electric Dipole Moment

Chapter 1

Motivation

The beauty of nature often manifests itself in symmetries. Looking at the underlying laws, the world is more thrilling when some of them are broken. The world as we know it for example, would not exist if there was not an excess of matter over antimatter. If matter and antimatter came together they quickly would annihilate and nothing than photons would be left. So, we are the remnants of a small imbalance after the Big Bang called the baryon asymmetry [1].

1.1 CP-violation

It is hard to answer why this asymmetry exists, but in 1967 Sakharov [2] established three criteria which need to be met such that an excess of matter over antimatter can evolve. First, the Baryon number symmetry, which counts the number of quarks and antiquarks, needs to be violated, because clearly for the universe to evolve from a universe with B=0 to a state B$\neq$ 0 this needs to hold.

Second, there must be C- and CP-violation. C and P together with T are the three fundamental discrete symmetries of Physics. The charge (C) transformation turns all particles into antiparticles. Parity (P) inverts space, where as time (T) reverses the time arrow. CP transformation is then the concatenation of the first two. Thinking about mechanisms producing an abundance of matter, the production rate for baryons must be larger than the inverse process. It can be proven that this is only the case if the C- and the CP-symmetry are broken. Third, the whole system must be departed from thermal equilibrium.

At first sight, Sakharov's criteria can be fulfilled easily, especially as breaking of the fundamental symmetries is well known in the standard model. Already in 1957, Wu [3] measured for the first time P violation in the weak sector, by detecting preferred directions of electron emission relative to the nuclear spin in the ^{60}Co beta decay. After establishing the theory of the weak interaction which is maximally P-violating, it also became evident that the C-symmetry is broken, because the left handed neutrinos would transform into left-handed anti-neutrinos, which have not been observed yet [4].

CP-violation is thus one of the ingredients in understanding the matter-antimatter asymmetry and was first detected in the neutral Kaon system [5]. The measured CP-violation is originating from the fact that the mass eigenstate of K_L^0 is not identical with the CP eigenstate K_2^0 but much rather a mixture with the K_1^0 state [4]. This kind of CP-violation which comes from the mixing of the eigenstates is called indirect.

Additionally, one can also search for direct CP-violation, which corresponds to looking for a difference in the decay amplitude of a meson into a final state and its CP conjugate counterpart. In the Kaon sector the direct CP-violation is around 10^3 times lower than the indirect one [6, 7].

Both kinds of CP-violations in the standard model arise from a complex phase in the CKM matrix which is only a valid parameter in the case of three quark families. Therefore Kobayashi and Maskawa could postulate the third quark family, even before it was detected. As the CKM matrix involves mixing of all three quark families direct and indirect CP-violation can not only be observed in kaonic systems but also in other meson anti-meson pairs such as B^0 and D^0. For both systems CP-violation

was found [8–13]. Consistent with the CKM model, the direct CP-violation found in the B-sector was larger than in the K-mesons.

Even though many CP-violating processes were detected in the weak sector of the Standard Model, it was soon found out that it is not enough to explain a substantial amount of the baryon asymmetry [14, 15]. Also, CP -violation arising in the strong sector via the QCD theta angle has a substantial problem, as it is limited to values below 10^{-10}, due to most recent measurements of electric dipole moments [16, 17]. This issue is known as the strong CP-problem and will be treated in more details in Chapter 9.

Because of the failure of the Standard Model to explain CP-violation to a level where it becomes relevant to baryogenesis, rise was given to many Beyond the Standard Model (BSM) theories [18–21]. Many of these theories are building upon extensions of the Minimal Supersymmetric Standard Model (MSSM) and have to be fine-tuned because of stringent constraints coming from electric dipole moment measurements of the electron, neutron and atoms. Therefore the measurement of those quantities is important in shaping and excluding theories which try to provide new sources of CP-violation.

1.2 Electric Dipole Moments

The Hamiltonian of a particle in an electric and magnetic field is given by

$$H = -\left(\mu\mathbf{B} + d\mathbf{E}\right) \cdot \frac{\mathbf{S}}{S}, \tag{1.1}$$

where μ is the magnetic moment of the particle and d the electric dipole moment. The magnetic moment of many particles is finite and aligned along the spin direction as this is the only available axis in such a system. If we now apply the fundamental transformation laws as listed in Table 1.1 we can see that under the parity transformation the Hamiltonian becomes

$$H_{\mathrm{P}} = -\left(\mu\mathbf{B} - d\mathbf{E}\right) \cdot \frac{\mathbf{S}}{S} \tag{1.2}$$

and under the T transformation it becomes as well

$$H_{\mathrm{T}} = +\left(-\mu\mathbf{B} + d\mathbf{E}\right) \cdot \frac{\mathbf{S}}{S}. \tag{1.3}$$

Because the two Hamiltonians H_{P} and H_{T} differ from the initial one by a minus sign in front of the electric field part it can be concluded that an electric dipole moments violates P and T and thus because of the CPT-theorem [22] also CP.

	C	P	T
E	$-\mathbf{E}$	$-\mathbf{E}$	**E**
B	$-\mathbf{B}$	**B**	$-\mathbf{B}$
S	$-\mathbf{S}$	**S**	$-\mathbf{S}$

Table 1.1: Transformation of some relevant vectors under the fundamental C, P and T operators.

Measurements can be conducted in a variety of particles, probing different electric dipole moments. In experiments with atoms the Schiff screening theorem has to be considered [23], which states that when applying an electric field to neutral atoms, the charged constituents will rearrange such that the external electric field is perfectly cancelled at the place of the nuclei. However is the underlying assumption of this theorem, that the nucleus is point-like and the electrons are non-relativistic. Thus, atomic net EDMs can arise from relativistic corrections and finite size effects [24, 25].

When employing atoms for an EDM search, paramagnetic and diamagnetic systems are distinguished. On the one hand paramagnetic atoms mainly are sensitive to the nuclear EDM. An example is ^{199}Hg which is currently giving the lowest bound on an electric dipole moment with a value [17] of

$$|d_{\mathrm{Hg}}| < 7.4 \cdot 10^{-30}\, e\,\mathrm{cm}. \tag{1.4}$$

On the other hand diamagnetic systems are susceptible to the electron EDM and can be measured either in traps (HfF [26]) or using cold molecular beams (ThO [27], YbF [28]). The experiment with ThO yields the current world best upper limit. The electric dipole moment of charged elementary particles such as muons, protons and deuterons can also be measured in storage rings for example as by-product of the muon g-2 experiment [29] or as planned by the JEDI experiment [30].

Even though the nuclear electric dipole moment of Mercury gives the most stringent constraint on an EDM value, it is important to measure the EDM of a variety of particles. Because even if an electric dipole moment would be measured, it only sets a scale on the CP violation but does not immediately tell of what nature the source is. Additionally, especially electric dipole moments of large nuclei are hard to tie to an underlying physics process as atomic calculations as well as nuclear many-body-interaction calculations are involved. Hence electric dipole searches as for example of the neutron, which will be treated in details in the next chapter, are of utter importance.

Chapter 2

Neutron Electric Dipole Moment Measurement

2.1 Measurement Principle

To measure the neutron electric dipole moment, we need to investigate the influence of an electric field on the neutron. If the nEDM is not equal zero, its quantum mechanical states will split in presence of an electric field, due to the linear Stark effect [31]. The Hamiltonian of this system can then be written as

$$H = -2 \cdot (\mu \mathbf{B} + d\mathbf{E}) \cdot \mathbf{S}, \tag{2.1}$$

where $\mathbf{B}$ and $\mathbf{E}$ are the magnetic and electric field respectively, μ the magnetic and d the electric dipole moment and $\mathbf{S}$ denotes the spin of the particle. We already used that the neutron is a spin-1/2 particle.

Theoretically, this measurement can be conducted with any sort of particle, however the neutron has two substantial advantages. On one hand the particle needs to be exposed to an electric field, every particle with non-zero charge will consequently experience acceleration and is therefore much harder to be contained. On the other hand, because the neutron is a spin-1/2 particle, the only other electro-magnetic-moment it can have is the magnetic dipole moment [31].

From the Hamiltonian in Equation 2.1, it can be derived that by comparing the measurement with the electric field directions upward and downward relative to the B-field direction, the difference will be directly proportional to a possible nEDM.

$$H_{\uparrow\uparrow} = -(\mu B + dE) \cdot 2 \tag{2.2}$$

$$H_{\uparrow\downarrow} = -(\mu B - dE) \cdot 2 \tag{2.3}$$

As the Hamiltonian corresponds to the energy of the system, it is replaced in the next step by $\hbar$ and the Larmor precession frequency.

$$\hbar(\omega_{\uparrow\downarrow} - \omega_{\uparrow\uparrow}) = \hbar\Delta\omega = 4dE \tag{2.4}$$

The arrows indicate the (anti-)parallelism between the magnetic and electric field, the value of the spin ($\hbar/2$) has already been inserted at this point. The above connection can also be seen more easily by looking at the quantum states level diagram shown in Figure 2.1.

From equation 2.4 we can see that the measurement of nEDM can be ultimately reduced to a frequency measurement [32]. In the presence of only a magnetic field, the neutron will perform a precession movement with the frequency equal to the Larmor precession frequency. When the nEDM is existent and an electric field is additionally applied, the precession frequency will be modified due to the nEDM. Depending on the sign of d_n, which is yet unknown, the neutron will precess faster in one relative electric and magnetic field direction and slower in the other one. The art of the nEDM measurement is then the extraction of this difference of precession frequency.

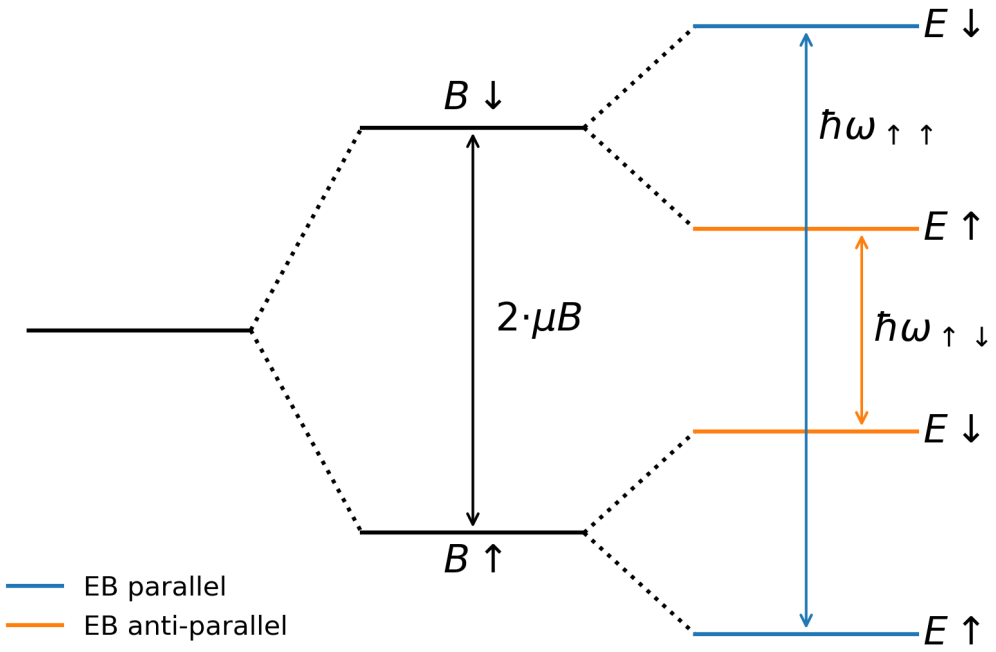

Figure 2.1: Quantum mechanical level splitting for the neutron in presence of magnetic and electric fields. The polarisation of the neutrons is pointing upward.

2.1.1 Ramsey's Method of Separated Oscillatory Fields

In many nEDM experiments Ramsey's method of separated oscillatory fields is employed to retrieve the precession frequency of the neutron. It will later be seen that the method can be used on both, beams of cold neutrons or stored ultra-cold neutrons. For the explanation of the method we first focus on the use of cold neutron beams, which was historically the earlier application.

The method was first described to be used in molecular beam resonance [33] but is here applied to a polarised neutron beam. The following explanation is visually underlined by Figure 2.2. First, the neutrons need to be polarised in one direction. Then, at the beginning of the polarised beam an RF magnetic field is applied to the particles. The RF frequency should match the Larmor precession frequency of the neutrons to have best efficiency. The applied RF field makes the neutrons tip over and the duration of the RF pulse can be adjusted such that the particles perform a 90° rotation, which is called a $\frac{\pi}{2}$-flip. After this $\frac{\pi}{2}$-flip a region of static homogeneous magnetic field follows, in which the neutrons undergo Larmor precession, in evolving combination of the spin up and down states. At the end of this homogeneous field region another $\frac{\pi}{2}$-flip is applied. If the Larmor precession frequency is exactly equal the RF frequency at the end of the free precession region the neutrons will be in phase with the RF pulse and all of them will in the end be pointing exactly downwards.

By detuning the RF frequency the neutrons will accumulate a phase shift during the Larmor precession and the $\frac{\pi}{2}$-flip will not make them point exactly downwards. The neutrons will be in a combined quantum state of spin-up and spin-down. By measuring the spin of the neutrons at the end of the apparatus the asymmetry between up- and down-states is measured. Close to the resonance, the more detuned the RF frequency is from the Larmor precession frequency, the less neutrons will be measured in the down state. With this procedure several frequencies can be scanned and resonance can be found. The first measurement of such a resonance curve with neutrons was conducted by Ramsey et al. [34] and can be seen in Figure 2.3.

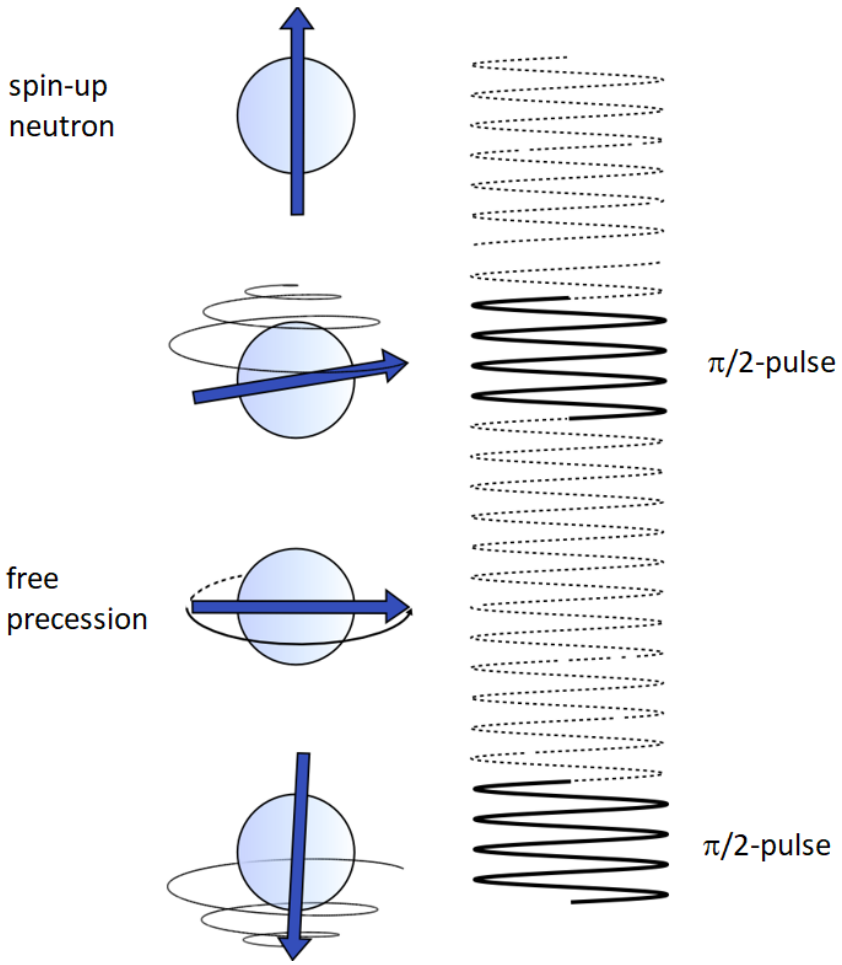

Figure 2.2: Visualisation of Ramsey's method of separated oscillatory fields. First, the neutrons are polarised such that their spin is pointing upwards. Then a magnetic $\frac{\pi}{2}$-pulse is applied, such that the neutrons are rotated into the vertical plane. Due to the magnetic holding field, the neutrons then start to precess. After a certain amount of time a second $\frac{\pi}{2}$-pulse is applied such that the neutrons are in the end pointing downwards. Courtesy of Philipp Schmidt-Wellenburg.

Besides the oscillation of the neutron counts coming from the fact that every 2π the detuned RF frequency is again in phase with the Larmor frequency, in Figure 2.3 it can be clearly seen that the counts are in total decreasing towards lower and higher frequencies. This is due to the fact, that the RF pulses are less efficient at rotating the spins of the neutrons if their frequency is not close to the Larmor precession frequency.

If an additional electric field is applied during the free precession phase and the neutron has an electric dipole moment the precession frequency will be faster or slower depending on the sign of the electric field and the sign of d_n. Scanning the frequency for both electric field polarities, the difference of the two detected central frequencies will then be directly proportional to a nEDM.

Because up to now no nEDM was measured, the sensitivity of such a experiment is important as it poses an upper limit on the nEDM value. In the case of Ramsey's method of separated oscillatory fields the statistical sensitivity is given by:

$$\sigma\left(d_\mathrm{n}\right) = \frac{\hbar}{2\alpha E T \sqrt{N_0}} \tag{2.5}$$

where α is the visibility, T the free precession time, N_0 the number of neutrons and E the strength of the electric field applied. The visibility is given by the normalised difference between neutrons

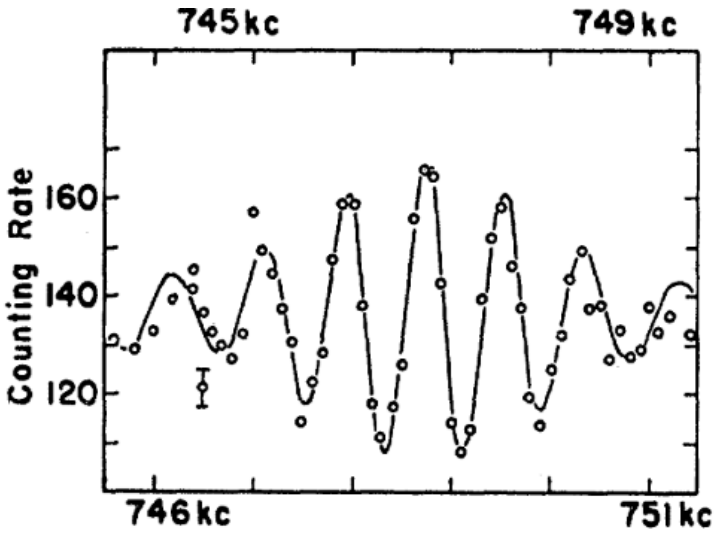

Figure 2.3: Ramsey resonance curve as reproduced from [34]. The x-axis shows the applied RF-frequency in kc which corresponds to kHz.

measured in the same polarisation state at the maximum and at the minimum of the central fringe. The better this separation or height of the fringe is the more sensitive the measurement.

On the basis of Eq. 2.5 we can formulate the analogy to the maxim of the Olympics "Citius, altius, fortius" ("faster, higher, stronger") for neutron electric dipole moment measurements: "Stronger electric fields (E), more polarisation (α), longer interaction time (T), more neutrons (N_0)"

2.2 Historical Efforts

Since the first measurement of the neutron electric dipole moment [34], with the magnetic resonance method in 1951 many efforts were undertaken to measure it with a better sensitivity. In the last 70 years, the upper limit on the nEDM could be decreased by more than six orders of magnitude. This journey is graphically documented in Figure 2.4.

The early experiments follow the same general scheme and consisted of an iron mirror polariser at the entrance of the beam, a region with homogeneous B-field, an iron transmission analyser and a neutron counter at the end [34]. In the region where the homogeneous magnetic field was applied, electrodes were installed, such that an electric field could be applied. Also at the end and at the beginning of this region, radio frequency coils were installed to perform the spin flip. Depending on the experiment the homogeneous region was between one and two meters long.

Until the seventies of the last century, the nEDM experiments were conducted with cold or thermal neutron beams. Cold neutrons have a velocity of less than 1000 m/s, the thermal neutrons which were used in certain experiments even had velocities of up to 2200 m/s. The velocity of neutrons poses two problems the beam experiments had to fight with [37]. The obvious problem was that the sensitivity of the measurement scaled with the inverse of the interaction time: the faster the neutrons are, the less sensitive the measurement is. This could have been compensated by increased length of the magnetic and electric interaction region, however increasing this region causes mechanical and stability problems.

The second challenge using cold/thermal neutrons arises from the so called $E \times v$-effect. The neutrons moving through the electric field E see an effective magnetic field $E \times v/c$ [38], coming from special relativity. If the electric field would be perfectly parallel to the magnetic field this effect can be completely eliminated. However, considering the length of the region in which the fields need to be parallel, this is nearly impossible, and leads to a significant systematic effect.

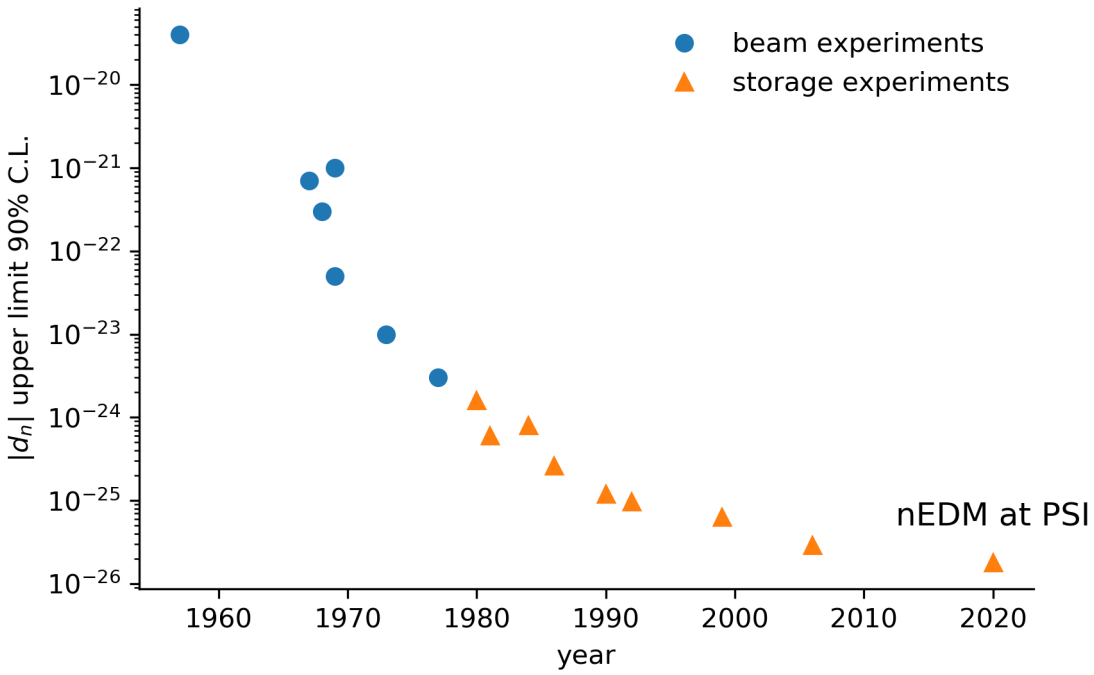

Figure 2.4: Historical overview of measured upper limits (90% C.L.) on d_n. Values are taken from [16, 35, 36].

At the end of the seventies, the above mentioned problems could be relaxed by the use of ultra cold neutrons. These are neutrons which have very low energies and can be contained, because they undergo total reflection at any incident angle. This significantly increased the observation time and eliminated the $E \times v$-effect. Thus, magnetic resonance experiments with stored ultra cold neutrons were the new way to go.

The layout changed fundamentally, and the beam lines were replaced with storage vessels which contained the neutrons. Still, before the neutrons entered the vessel they travelled through a polariser. The top and bottom of the storage volume was formed by electrodes, which provided the electric field. Around the vessel, a coil system was installed to produce the magnetic holding field, as well as the RF pulses for the application of the neutron spin flips. Those spin flips where now separated by time and no longer by distance as in the beam experiments. After the precession time which can last several dozens of seconds depending on the experiment the neutrons were emptied into a spin analyser and counted.

Substantial improvements on different parts of the apparatus led to the significant decrease of the upper limit over the last 40 years. First of all, there was the improved storage time of the neutrons. This was achieved by better coatings of the precession chamber, which led to less neutrons losses at the walls. Additionally, the magnetic field control was improved by the use of magnetic shields and magnetometry. Last but not least, the number of neutrons could be augmented by more powerful neutron source facilities.

2.3 Current World's Best Upper Limit

The nEDM experiment at PSI stands in the long tradition of nEDM measurements with stored ultra-cold neutrons and published the current world's best upper limit at the beginning of 2020 [16]. The

value is

$$d_{\mathrm{n}} = (0.0 \pm 1.1_{\mathrm{stat}} \pm 0.2_{\mathrm{sys}}) \times 10^{-26} e\,\mathrm{cm} \tag{2.6}$$

and translates into an upper limit (90% C.L.) of

$$|d_{\mathrm{n}}| < 1.8 \times 10^{-26} e\,\mathrm{cm}. \tag{2.7}$$

The measurements were taken in 2015 and 2016 with an upgraded version of the Sussex-RAL-ILL spectrometer which set the previous best upper limit [39]. By moving the apparatus from the neutron source at ILL to PSI one could profit from higher UCN densities [40–42].

Many parts of the apparatus have been replaced, while some of them as for example the passive magnetic shield has been continuously in use. The electrodes for example were replaced and a 11 kV/cm electric field could be applied instead of the 7 kV/cm.

Another major improvement was made with the increase of the precession time and the visibility. It is clear that a longer precession time directly impacts the visibility as the neutrons have more time to depolarise. The excellent analysing power of the spectrometer was $\alpha(T = 0) = 0.86$). Additionally, depending on what magnetic field gradients were applied the visibility was between 0.71 and 0.85 after 180 seconds, which is much higher than 0.58 when the experiment was run at ILL. These improvements could be achieved due to field optimisations.

The analysing power also profited from the 5 T superconducting magnet which provided excellent initial polarisation. At ILL a silicon foil with a 1 µm layer of iron was used as polariser, which only provided a 90% polarisation.

A system which was additionally installed at PSI is the array of 16 optically pumped Caesium magnetometers [43]. They were mounted on the outside of the top and bottom electrode and could therefore provide important information on the vertical magnetic field gradients. They were used to detect drifts of the vertical field gradients but also helped homogenising the magnetic field. nEDM also featured a Hg-comagnetometer, with which the neutron precession frequency can be corrected by the Hg-frequency and thus becomes insensitive to magnetic field changes. The ration of the neutron and the Hg-precession frequency is called R ratio and is an important piece in the analysis 10.2.

One of the drawbacks of moving the nEDM experiment to PSI was the much more unstable background magnetic field, due to neighbouring magnets. A solution was provided by the installation of a surrounding field compensation (SFC) [44], consisting of six Helmholtz-like coils which surrounded the experiment. They were used to stabilise the magnetic field around the volume of the experiment. This system is the predecessor of the elaborate active magnetic shielding which is going to be presented in Part II of this thesis.

2.4 Auxiliary searches

With the nEDM apparatus, not only could a more stringent upper limit be placed on the neutron electric dipole moment, but also other interesting physics can be extracted from the data. With the nEDM data collected a directional dependence of the neutron electric dipole moment can be searched for, which would imply Lorentz invariance violation. Such an analysis was performed with our apparatus in 2010 [45, 46].

Additionally, hypothetical short range spin dependant forces between the freely precessing neutrons and the nuclei in the precession chamber wall were examined and lead to a limit [47, 48].

In 2019 a limit on neutron mirror-neutron oscillations with the nEDM apparatus was set, where the disappearance of neutrons as a function of the applied magnetic field was studied [49].

In the third part of this thesis we are going to focus on interactions of the neutrons with hypothetical dark matter particles, called axion-like particles. A first laboratory constraint on these interactions was set from our dataset in 2017 [50].

2.5 Next Generation Measurement

A detailed description of the next generation experiment n2EDM is given in [51]. It is currently under construction at PSI. n2EDM is a totally new apparatus which profits from the knowledge gained from the running of the nEDM experiment.

The experiment will feature a double precession chamber as shown in Figure 2.5, which is the current state of the art. As the ground electrode will sit in the middle of the two chambers, it will allow for simultaneous measurement of both electric field states, which reduces some systematics. At the same time the electrodes are designed to run at an even higher electric field of 15 kV/cm. The diameter of the chamber will also be enlarged to 80 cm which allows for the storage of more neutrons at the same time.

Nevertheless, magnetic field control is still the most important task in the new experiment. A long-term stability of 30 fT over the free precession time of approximately 300 s is required within the precession chamber [51]. To achieve this ambitious goal an active magnetic shielding as discussed in Part II is indispensable.

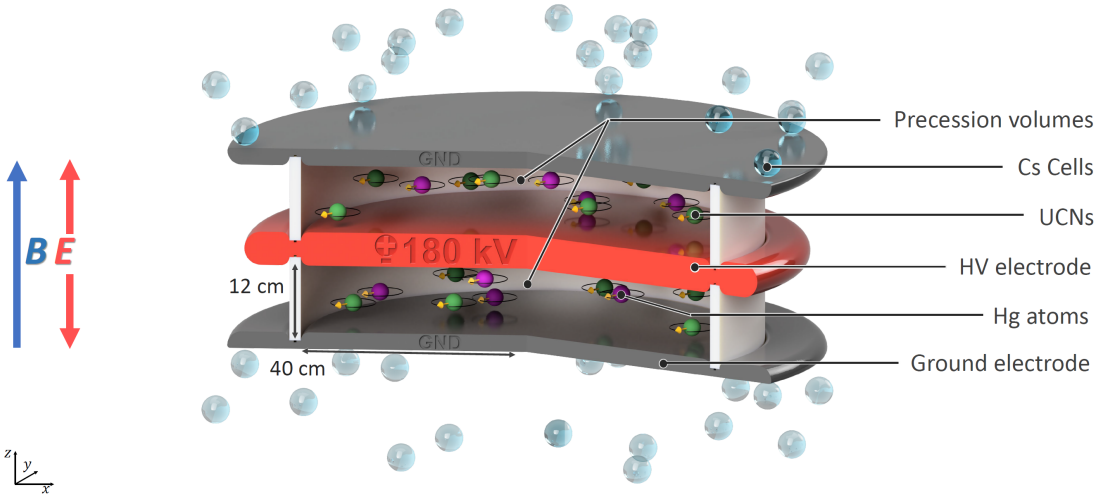

Figure 2.5: Depiction of the double precession chamber for n2EDM [51]. It can be seen, how the electric field is pointing in opposing directions in the upper and lower chamber.

Part II

Active Magnetic Shielding

Chapter 3

Introduction

The n2EDM experiment resembles its predecessor experiment nEDM regarding its design and includes similar further developed subsystems of it. In nEDM there was already an active shield against outside magnetic field changes installed, which was called surrounding field compensation system (SFC). By learning from the strength and weaknesses of the SFC in operation and design we created a next generation active magnetic shielding (AMS) for n2EDM. In the following sections the motivation for such a high performance system is presented and the major conceptual updates to its predecessor are highlighted.

3.1 The Need for an Active Magnetic Shielding

The n2EDM experiment is located at the Paul Scherrer Institute in Villigen, Switzerland, which is a dynamic research laboratory, where many different disciplines and scientific interests come together. The experimental site of n2EDM is on one side located directly next to the huge experimental hall, which accommodates the world's most intense proton accelerator HIPA and all the corresponding beam lines. On the other side n2EDM is facing COMET, the accelerator of the proton therapy facility at PSI and a strong magnet, named SULTAN, dedicated to tests of superconducting materials is also located nearby. It can easily be seen that n2EDM is situated in a magnetically vivid environment.

In Chapter 2, we discussed the importance of a magnetically stable environment for any experiment measuring the neutron electric dipole moment. In the predecessor experiment nEDM the daily ramping of the SULTAN magnet was visible in the data. Therefore, without an active magnetic shielding data taking would not have been possible during at least two hours every day, which would result in loss of approximately 10% of the data. Additionally, the unknown magnetisation state of the mu-metal could induce systematic effects. These facts strongly underline the importance of a performant active magnetic shielding.

The standard approach of magnetic field reduction in many applications, not only nEDM experiments, is a passive magnetic shield. Such a shield consists in its most basic version of a box made of mu-metal which houses the experiment or magnetically sensitive parts.

Mu-metals are alloys with a very high magnetic permeability often consisting of nickel and iron. The high magnetic permeability bends the surrounding magnetic field lines around the enclosed volume, guiding them rather through the mu-metal, than the enclosed space. This way, the inside of the box is protected from magnetic fields.

For n2EDM, a high performance shield consisting of six layers of mu-metal, with outer dimensions of $5\,\mathrm{m} \times 5\,\mathrm{m} \times 5\,\mathrm{m}$ was built. It is specified to provide a quasi-static shielding factor of 70'000 at $0.001\,\mathrm{Hz}$. Subsequently the passive shield of n2EDM, will be referred to as the magnetically shielded room or short MSR.

Despite the excellent shielding factor many experiments choose a two stage approach for their magnetic shielding. This has two main reasons which strongly relate to the particularities of mu-metals.

First, for mu-metals the shielding factor depends on the frequency of the magnetic field change. The shielding factor of mu-metals is impeccable at high frequencies, but drops dramatically off for lower frequencies. In Figure 3.1 the measured shielding factor of the MSR versus the frequency of the applied disturbance is plotted in all three directions.

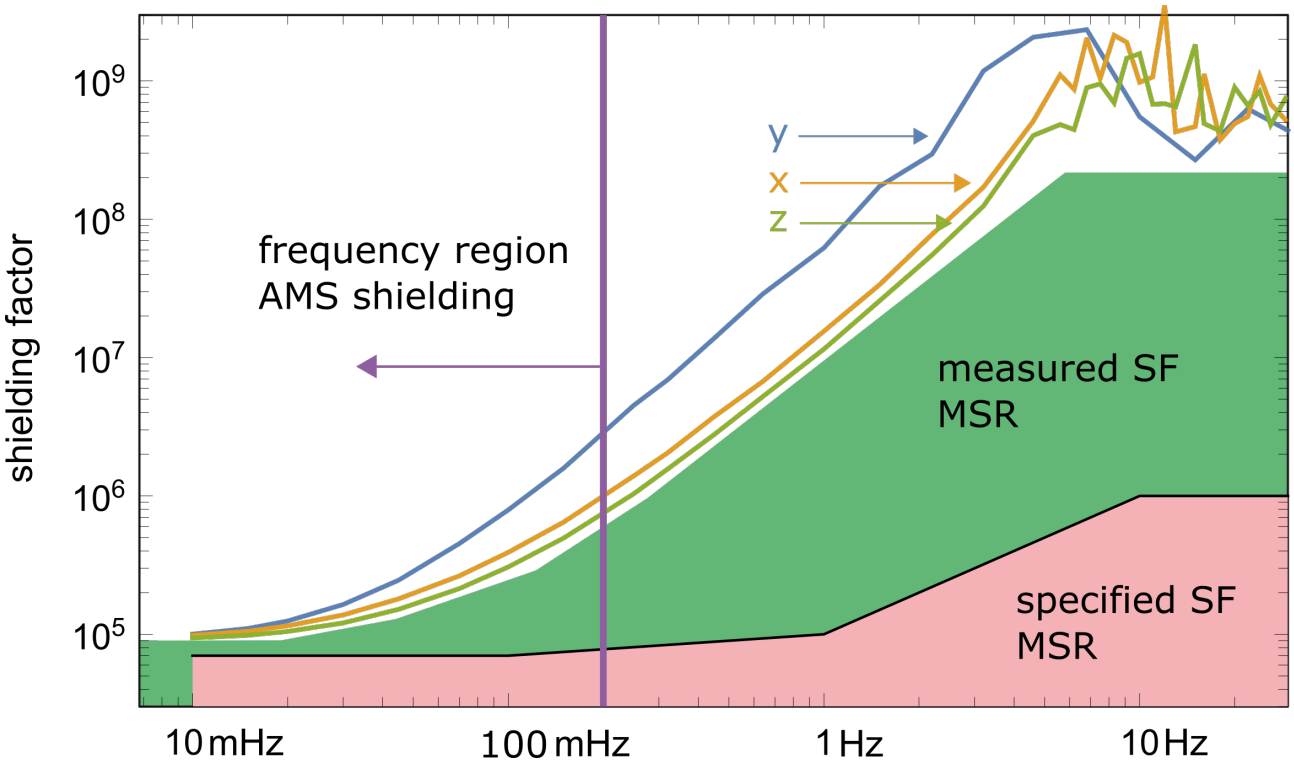

Figure 3.1: Shielding factor measurement in all three spatial directions for the MSR used in the n2EDM experiment. The green area encloses the minimal shielding that can be achieved in any direction. It can be seen that it outperforms the specified shielding factor (in red) at all frequencies. The region in which the AMS can provide additional shielding is indicated by the purple line. Courtesy of Georg Bison.

Over the range from 10 Hz down to 10 mHz the shielding factor of the MSR drops around three orders of magnitude. One particularly relevant timescale in the nEDM experiment is the frequency corresponding to the free precession time, which is $1/180\,\mathrm{s} \sim 5.6\,\mathrm{mHz}$. This frequency also falls in the low shielding regime of the MSR.

In Chapter 4 it will be shown, that the behaviour of the shielding factor of an active magnetic shielding is directly opposite. This means that an AMS provides good shielding at frequencies below 100 mHz which then decreases at frequencies above due to combined effects of the inductance of the coil system and DAQ timing. As it can be seen in Figure 3.1 the additional shielding from the AMS is primarily needed at low frequencies, therefore an AMS complements the performance of the MSR very well.

The second reason for a combined shielding approach is that mu-metal can be easily magnetised in strong fields, meaning that the shield becomes and acts as a magnetic source itself and thus can induce systematics. This mainly happens in slow changing or even static fields. Whenever the MSR is magnetised it needs to be degaussed in a lengthy process, on the cost of data taking. Therefore, it is crucial that the AMS provides a low field environment, such that a magnetisation of the shield can be suppressed as much as possible.

3.2 Working Principle

In the following, we are going to explore how an active magnetic shield in its simplest form works. It consists of a number of magnetic field sensors, coils and a data acquisition system which can readout the magnetic sensors and can output voltages. An example of such a system is depicted in Figure 3.2.

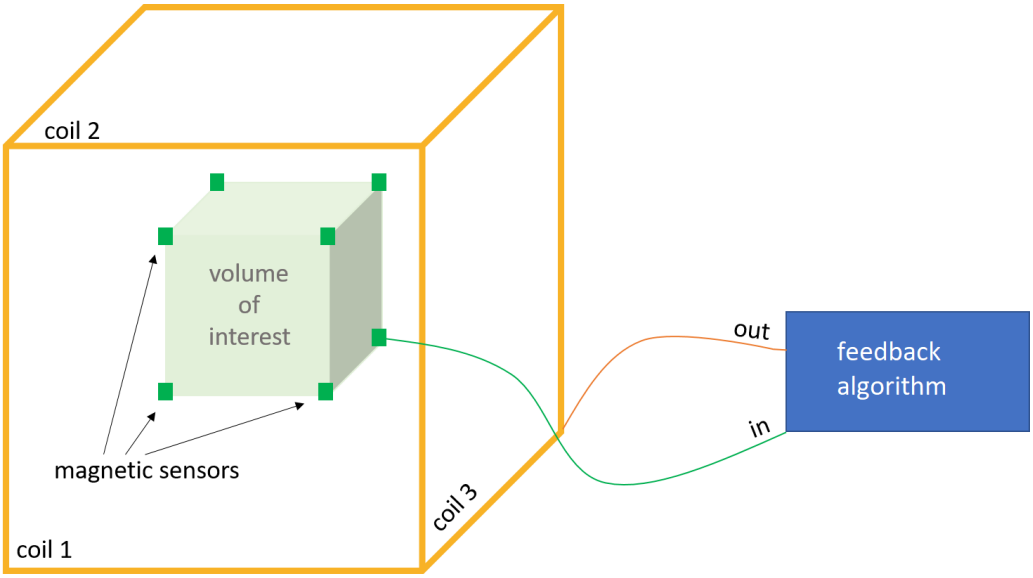

Figure 3.2: An active magnetic shielding consists of a set of magnetic sensors (green), which are enclosing the volume where the surrounding field should be attenuated. The volume of interest is surrounded by a set of coils (orange). The field values are then read out with a DAQ and a feedback algorithm calculates the currents which need to be output into the coils, such that the surrounding field is compensated best.

The goal of such a system is to reduce the magnetic fields in a volume of interest. In order to be able to do so, the fields that can occur there have to be known. For this purpose, magnetic field sensors are needed. On the positions where the sensors are located the field can be compensated with the highest accuracy. However, if the shapes of the produced magnetic fields are known, and if the feedback sensors are placed smartly, the field in the whole volume of interest can be compensated to a sufficient level without putting a sensor in every point.

In the example depicted in Figure 3.2 the sensors are placed on the corners of a cube in order to compensate the field in the enclosed volume. The magnetic field readings are then fed to the DAQ system where the information is processed. As the field often has to be compensated to zero, the system will output the currents which are needed to perfectly counteract the measured fields. We will later see that the relation between the current in each coil and the magnetic field at each sensor position can be measured beforehand. Because the underlying physics of magnetic field generation is linear (up to the point where an enclosed mu-metal cube is not appreciably magnetised), the algorithm can then calculate the correct output currents. The procedure described above can be utilised in a feedback loop, in which the magnetic fields will be measured continuously and the currents are updated accordingly.

Understanding this principle, we can now explore what happens in details when multiple magnetic sensors and coils are used. Ultimately the system is working according to the Biot-Savart law, which can be rewritten to state that a magnetic field created by a coil is linearly proportional to the current in the coil. Therefore all the measurements taken in the system can be summarised in matrix notation:

$$\mathbf{B} = \mathbf{MI} \tag{3.1}$$

Where B is a vector of all the magnetic fields measured by the magnetic sensors and I is the vector containing the currents in the coils. Consequently, M is a matrix which contains the physics described by the Biot-Savart law. M is called the correlation or response matrix and accommodates the relation between the current in a specific coil, and the magnetic field produced at the positions of the magnetic sensors. Equation 3.1 can be rewritten components-wise in the following way:

$$\begin{pmatrix} B_{1x} \\ B_{1y} \\ B_{1z} \\ \vdots \\ B_{nz} \end{pmatrix} = \begin{pmatrix} M_{11x} & M_{21x} & \dots & M_{m1x} \\ M_{11y} & M_{21y} & \dots & M_{m1y} \\ M_{11z} & M_{21z} & \dots & M_{m1z} \\ \vdots & \vdots & \ddots & \vdots \\ M_{1nz} & M_{2nz} & \dots & M_{mnz} \end{pmatrix} = \begin{pmatrix} I_1 \\ I_2 \\ I_3 \\ \vdots \\ I_m \end{pmatrix} \tag{3.2}$$

It can be seen that the correlation matrix has dimensions m x 3n where m is the number of coils i.e. currents and n is the number of magnetic sensors. Since often sensors measuring the fields in all three spatial directions are used (so called 3-axis magnetic sensors), one sensor provides B_x, B_y and B_z field values for each sensor position. Later it will be described in more detail how the information is stored in the correlation matrix.

3.3 Method of Simple Coil Design

In general the n2EDM experiment is much larger than its predecessor experiment, as already the precession chamber is upscaled in order to accommodate more neutrons. Additionally, the passive magnetic shield has now six layers instead of four. Nevertheless the size of the experimental area stayed the same, which causes the MSR to be much closer the active magnetic shielding than before.

The previously used Helmholtz-like coil structure is thus no longer an option, as the magnetic field homogeneity on the MSR will be insufficient. Therefore a new method of coil design was developed by M. Rawlik [52]. Its strength is, that it can provide good magnetic field homogeneity even if the distance between the coil system and the volume of interest is small. In the following chapter the basic principles of the method of simple coil design will be introduced briefly as it will later form the foundation of the n2EDM AMS design described in Chapter 6.

When it comes to implementing an active magnetic shielding system as described above there are several unknowns which need to be tackled. First it needs to be clear what fields the system must be able to compensate. It is possible that the surrounding field can change significantly over the course of an hour, days and even weeks. This includes not only different field strengths but also different shapes of magnetic fields. What kind of magnetic fields will be encountered in the n2EDM experiment will be thoroughly investigated in Chapter 5.

For now we assume that the magnetic fields we want to create and thus compensate are known. The better we can design coils which match the shape of the field, the better the compensation will be not only in the positions of the magnetic feedback sensors but all over the fiducial volume.

It is clear that for certain common magnetic fields it is known what coils should be used. For example homogenous fields are directly associated with Helmholtz coils. But even if such "ready-to-use" solutions are known, it does not mean that they show the best performance on a given volume of interest. So the question is how to design the coils producing the best compensation of a known magnetic field on a given volume and not only at the position of the sensors.

Providing a Grid Theoretically any imaginable shape of coil can be used. To simplify the problem, the number of possibilities can be constrained to a grid, consisting of many small rectangular coils which are called **tiles**. A tile is the smallest building block of the grid, the smaller this elementary building block, the more homogeneous the designed field can be.

Generally, these tiles can be of any shape, but with the perspective of a system that will be physically constructed, it is sensible to make use of rectangles. For the description of the algorithm

we are going to follow the illustrations in Figure 3.3. In the first part 3.3a we see an example system, where a grid (orange), a volume of interest (green dots) and field that needs to be created (blue) are given. As for the numerical algorithm to work the volume of interest needs to be discretised, it is effectively represented by points of interests (**POI's**) covering the fiducial volume densely. These are all the inputs needed in the method of simple coil design.

Finding Currents The question to ask is: What currents would flow in each of the tiles, if a goal field B_0 (blue) has to be compensated? This can easily be calculated by inverting Equation 3.1 and reducing it to a least squares problem.

$$\mathbf{I} = \mathbf{M}^{-1} \cdot \mathbf{B} \tag{3.3}$$

Here **B** represents the magnetic fields in x-, y- and z-direction at each point of interest for the field which needs to be created. The solution is then the vector I containing one current per coil to be applied, in order to achieve the target field. It has to be considered, that the solution not only depends on the target field and grid provided, but also on which and how many points of interest the least squares problem is solved on. In Figure 3.3b the current in each tile to achieve the target magnetic field is depicted.

Simplification From Figure 3.3b it can also be seen, that some of the currents are redundant, as every edge of a tile is shared with a neighbouring tile and some of those currents are counteracting. So in total there is much more current in the system than is actually necessary to create the target field. By treating the whole grid as a system of flowing currents, rather than restricting the currents to the individual tiles, the counteracting currents can be cancelled. This results in a grid, where every edge carries a current in a certain direction, thereby producing the target field. This third step is visualised in Figure 3.3c.

Finding Simple Loops Left with this net of currents, the algorithm needs now to find closed paths with the same current. These are called **simple loops** and are depicted in different colours in Figure 3.3d. The algorithm is heuristic and starts with the highest current for which a simple loop can be found, and removes the current from the current net. It then proceeds to look for a closed path with smaller currents amongst the currents left in the grid. This procedure is carried on until no simple loops with a current higher than a specified cut off current are found any more.

To a certain precision, this collection of simple loops operated in series produces the same magnetic field as the tiles with their individual currents. The precision is mainly depending on the cut off current, which gives the smallest current that can be driven.

This procedure can now be repeated for other target fields. A selection of simple loops which, when operated in series, produce a target field are called a **coil**. How a working system can be built from selection of simple loops is elaborated in Chapter 7.

3.4 Decomposition

In a rapidly changing magnetic environment such as a laboratory it can not be assumed that the shape of the background field will always stay the same. Therefore, one needs to design a system of coils which adapts to most of the external field configurations. Here, a good approach is to decompose the magnetic field into basis functions and then design a coil for each of these basis functions. A suitable decomposition for magnetic fields in a current- and magnetisation-free environment are the Cartesian harmonic polynomials. Those directly fulfil the Maxwell Ampere Equation $\nabla \times B = 0$ and the Maxwell Gauss equation $\nabla \cdot B = 0$. A general magnetic field can then be decomposed in the following way:

$$\mathbf{B}(\mathbf{r}) = \sum_{n=1}^{n_{max}} H_n \mathbf{P_n}(\mathbf{r}) \tag{3.4}$$

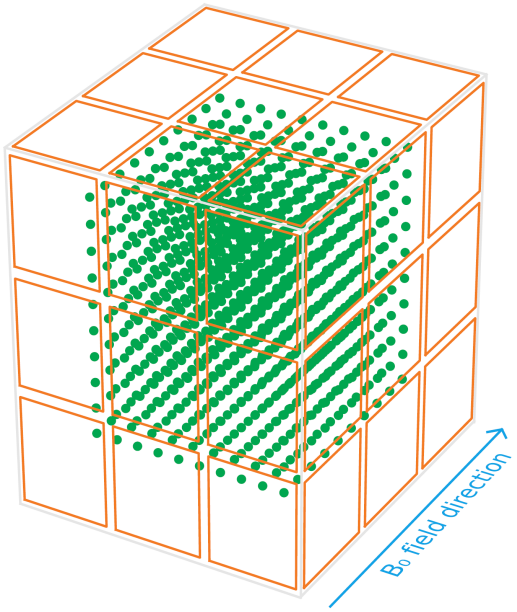

(a) Grid for a simple AMS system.

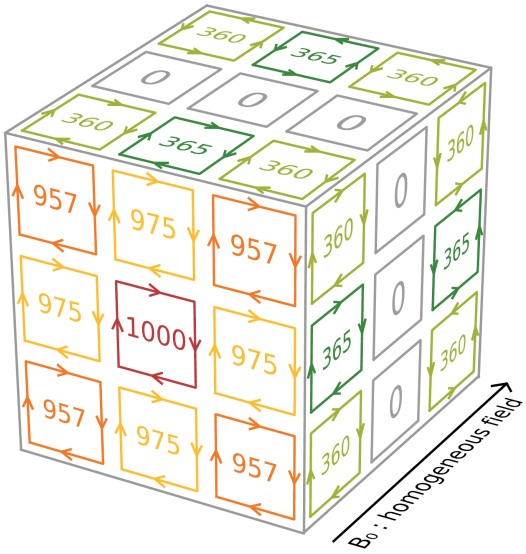

(b) Currents in the tiles.

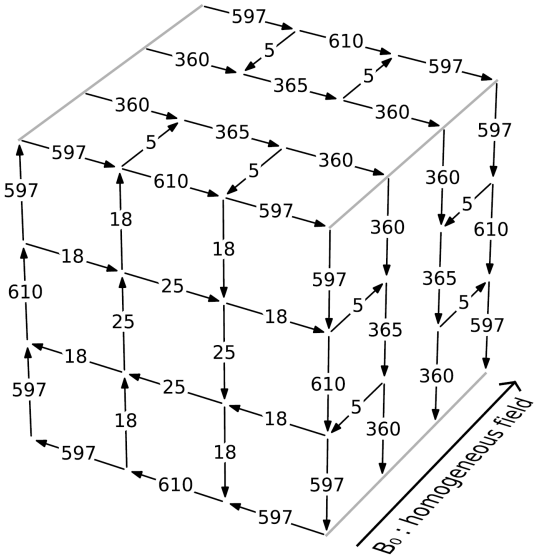

(c) Net of currents after cancelling counteracting currents.

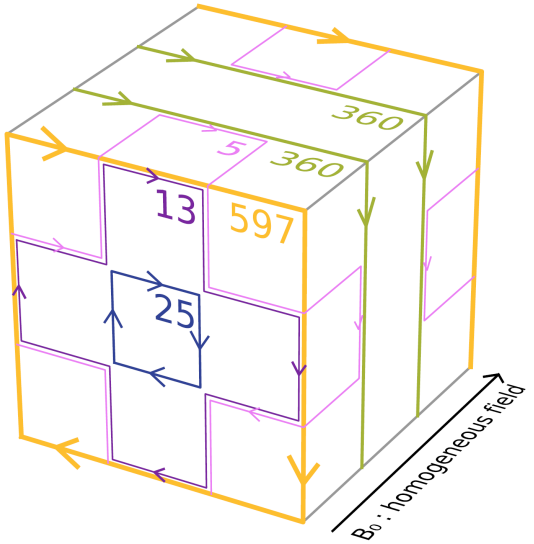

(d) Simple loops after simplification of the system.

Figure 3.3: Working principle of the method for simple coil design explained step by step. Courtesy of Michał Rawlik.

n	P_{x_n}	P_{y_n}	P_{z_n}
1	1	0	0
2	0	1	0
3	0	0	1
4	x	0	-z
5	y	x	0
6	0	y	-z
7	z	0	x
8	0	z	y
9	$x^2 - z^2$	0	$-2xz$
10	$2xy$	$x^2 - z^2$	$-2yz$
11	$y^2 - z^2$	$2xy$	$-2xz$
12	0	$y^2 - z^2$	$-2yz$
13	$2xz$	0	$x^2 - z^2$
14	yz	xz	xy
15	0	$2yz$	$y^2 - z^2$

Table 3.1: List of the Cartesian harmonic polynomials as used in the nEDM collaboration, up to 2nd order.

Where P_n are the harmonic polynomials and H_n are the corresponding harmonic coefficient. There are different possibilities how to rearrange the P_n. The polynomials used in the context of the nEDM and n2EDM experiment are shown in Table 3.1.

The advantage of this decomposition is, that the polynomials are orthogonal and form a complete basis. It was shown that in nEDM the non-orthogonality caused stability problems when operating the system [53].

Chapter 4

Prototype

In the laboratory at ETH Hönggerberg our group has a lot of possibilities to test and develop parts needed for the n2EDM experiment. An important work package is the investigation and validation of the working principle of the proposed active magnetic shielding. For this purpose a first prototype AMS was built during the doctoral thesis of my predecessor Michał Rawlik. In this thesis the shielding ability of system consisting of three homogenous coils on a cubical grid structure was demonstrated [54]. In order to ensure the smooth operation of the n2EDM AMS it is important to investigate the functioning of the prototype more deeply. In a first step the prototype was upgraded to five gradient coils and was redesigned to allow the insertion of a mu-meal cube inside to study the interplay between mu-metal and the active magnetic shielding. Additionally, the positioning of the magnetic feedback sensors was studied to ensure a high stability and high shielding factor of the system. In the following these continued investigations on the active magnetic shielding prototype are presented in detail.

4.1 Design

The prototype is a downscaled, symmetrical version of the n2EDM AMS. It encloses a volume of 1.3 m by 2.3 m by 1.3 m, divided into five, nine and five tiles in the corresponding directions x, y, and z as visible in Figure 4.1. On one side the cube is open and the holding structure is missing, such that later on a 1 m by 1 m by 1 m mu-metal cube can be inserted. The described geometry is used in the algorithm of simple coil design to design one homogeneous coil in every spatial direction and additionally five first-order gradient coils corresponding to the linear gradients in the Cartesian harmonic decomposition shown in Table 3.1. The following target fields have been chosen:

- 50 µT for the x-, y-, and z-coil
- 20 µT/m for the first order gradient coils

The strengths of the fields were chosen such that the occurring magnetic fields plus some external disturbances can be compensated easily. The homogeneous and first order gradient fields do not have the same strength, similar as in the n2EDM AMS. This has an influence on finding the best positions for the magnetic sensors, which will be shown later on.

The holding structure of the prototype is made out of thin laser cut aluminium plates onto which cable trays are mounted. The aluminium plates were reused from the previous design, however some bigger cable trays had to be mounted in order to accommodate the additional gradient coils. A picture of the system is shown in Figure 4.1.

The control of the prototype is realised by a Julia program running on a Linux machine providing commands to Beckhoff EtherCAT modules. The output of the EtherCAT modules is connected to amplifiers which output the currents to the coils. Additionally, eight three-axis magnetic sensors measure the magnetic fields which are read out by different EtherCAT modules, forwarding it to the algorithm which was described in detail in Section 3.2. Each coil is decomposed into three circuits in order to minimise the winding effort. Therefore 24 amplifiers (three for each coil) are needed in total.

Figure 4.1: Photograph of the AMS prototype in the ETH laboratory at Hönggerberg. The side facing the window front is open in order to enable to place a mu-metal cube inside. The origin is in the bottom left corner closest to the windows.

4.2 Static Performance

The evaluation of the performance of the prototype system can not be reduced to a single measurement, but is rather described by various measurements which give an overall picture. Nevertheless, in a first step it is important to know if the coils produce the required target field. By checking that the coils produce the target field, it is also made sure, that the envisaged homogeneity criteria are met. In the case of the prototype the coils were designed such that that the mu-metal cube later lies in a region with a $\pm 1\%$ homogeneity.

Compared to the situation which will be encountered at n2EDM, where the MSR is installed before the coils and can not be removed, the magnetic field produced in the laboratory setup can be measured. Therefore the prototype provides a good proof of concept.

Inside the aluminium holding structures which support the coils, a three-axis mapper is installed. It has a fluxgate magnetic sensor attached to it, which can explore most of the volume enclosed by the coils. To obtain a map of the magnetic field produced by the coil both a map with the coil turned on and off was taken. These two measurements were subsequently subtracted from each other in order to extract only the contribution of the fields created by the coils. In Figure 4.2a and 4.2b the measured magnetic fields in x-, y- and z-direction for the first gradient coil are compared to the simulation using the Biot-Savart law. In Figure 4.3, the target field was subtracted and it can be seen that the residuals stay below $1\,\mu\mathrm{T}$ in the whole plane for a target field of $20\,\mu\mathrm{T/m}$.

Most often planar maps were taken in the planes most relevant to the coil under investigation. All eight coils were mapped in the described way after their completion to spot eventual defects. Additionally, every circuit of each coil was mapped separately in order to compare it to the calculated

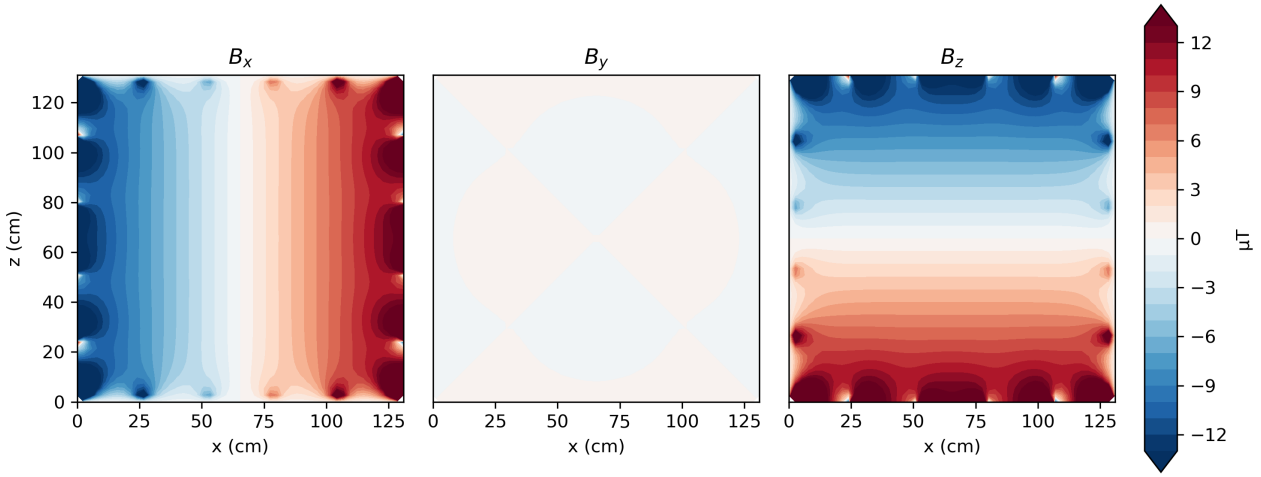

(a) Simulation of the magnetic field of the first linear gradient at the xz-plane at $y = 115\,\text{cm}$. The B_x, B_y and B_z components are plotted from left to right.

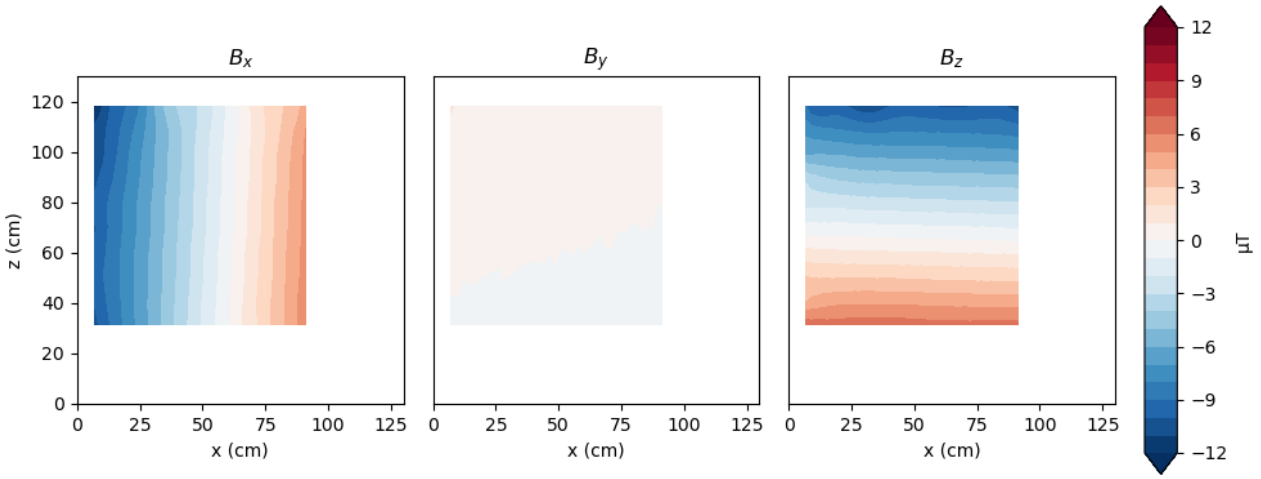

(b) Measurement of the magnetic field of the first linear gradient at the xz-plane at $y = 115\,\text{cm}$. The B_x, B_y and B_z components are plotted from left to right.

Figure 4.2: Comparison of the first linear gradient magnetic field simulation to the mapping.

fields and troubleshoot eventual disagreements.

In Table 4.1 the calculated field values for all the coils are listed and compared to the goal fields. It can be seen that generally, the measured fields are lower than the target fields. For the gradients this difference can measure up to 10%. The power supply can output twice as much current as needed to achieve the goal field, thus reaching the target fields is not a problem.

To infer the field value produced by the homogeneous coils, the measurement points were averaged over the area 20 cm inside the grid structure, to ensure that it is not dominated by the natural deviations close to the grid. To calculate the gradients, a line was fitted to the measurement points with the least squares method.

The fitted field values can then be subtracted from the mapping data in order to evaluate the deviation to the target field. In Figure 4.3, the magnetic field residuals are plotted for the first linear gradient. The deviations from the target field do not exceed $\pm 1\,\mu\text{T}$ in the volume where the mu-metal will be.

coil	measured	goal
x	50.0 µT	50 µT
y	49.6 µT	50 µT
z	49.2 µT	50 µT
1st grad	18.2/18.7 µT/m	20 µT/m
2nd grad	17.7/18.4 µT/m	20 µT/m
3rd grad	20.0/19.9 µT/m	20 µT/m
4th grad	17.7/19.1 µT/m	20 µT/m
5th grad	18.0/18.0 µT/m	20 µT/m

Table 4.1: The field value for each coil of the prototype calculated from the mapping data. 1st to 5th grad corresponds to the first order gradients labelled 4-8 in Table 3.1.

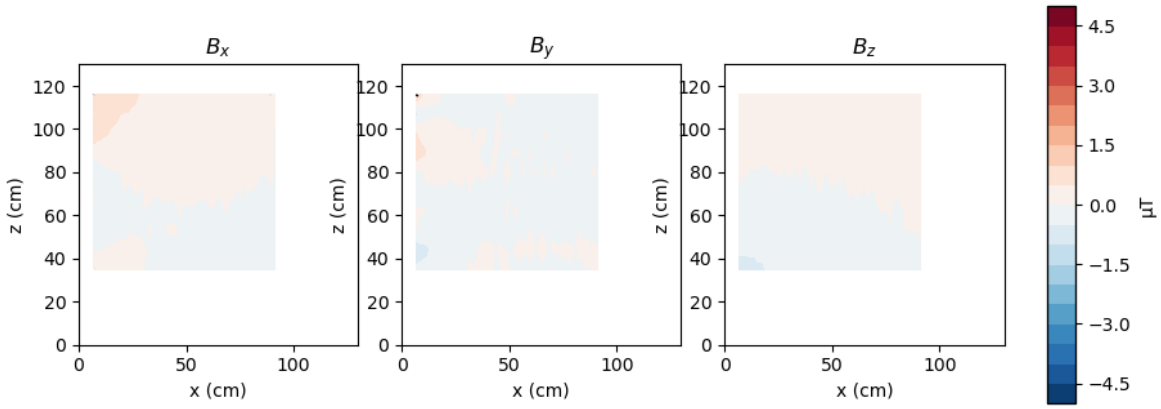

Figure 4.3: Measurement of the magnetic field of the 1st gradient being subtracted from the target field in the xz-plane at $y = 115\,\text{cm}$. In the whole plane, the field residuals are not exceeding $\pm 1\,\mu\text{T}$.

4.3 Sensor Positioning

As soon as dynamic changes need to be compensated, the dynamic properties of the AMS need to be investigated. In Section 3.2 it was discussed that the control algorithm of the AMS relies on the inversion of the correlation matrix M. In the following we are going to investigate the characteristics of the correlation matrix M.

4.3.1 Underlying Principles

For the overdetermined system described by the correlation matrix M, the Moore-Penrose pseudo-inverse needs to be calculated in order to solve for the currents in the Least-Squares problem given by $||MI - B|| \stackrel{!}{=} 0$. First, the matrix undergoes singular value decomposition (SVD)[55] and M can be written as:

$$M = UVW^T, \tag{4.1}$$

where U and W are two unitary matrices and V is a diagonal matrix. Then the Moore-Penrose pseudo-inverse M^{-1} can be calculated in the following way.

$$M^{-1} = WV^{-1}U^T \tag{4.2}$$

The spotlight thereby lies on the diagonal matrix V and its inverse V^{-1}. The diagonal entries of V are its singular values, which describe the the effect of combinations of coils on the magnetic sensors. The inverse V^{-1} is calculated by simply inverting the singular values v_i.

$$V_{ii}^{-1} = \frac{1}{v_i} \tag{4.3}$$

The condition number of M can then be calculated to be:

$$\text{cond}(M) = \frac{v_{max}}{v_{min}}, \tag{4.4}$$

where v_{max} and v_{min} are the maximal and minimal singular value of M. Problems emerge if the condition number is high, meaning that the system ill-conditioned. In this case, the inversion process generates small values, signifying small magnetic field changes, which cause a strong change in current and therefore a destabilised system. Consequently, this can even lead to oscillations.

This behaviour was observed in the surrounding field compensation of the nEDM experiment, where the condition number was 18.2. In this case the field was less stable with the surrounding field compensations than without [53, 54]. As the condition number is an intrinsic property of the system, it roots in its geometry. The solution to the problem is to design coils which produce intrinsically orthogonal magnetic fields, hence the choice of the Cartesian harmonic polynomials.

Even though the choice of coils for the prototype should not allow for high condition numbers, there is a problem regarding the value with which the Cartesian harmonic polynomials will be scaled. In reality this means that, the prototype with 50 µT homogeneous fields and 20 µT/m has an intrinsic factor of approximately 2 between the minimal and maximal singular value. Because building coils with a higher target field immediately implies a more complex system and thus a higher installation time it does not make sense to build 50 µT/m gradient coils for the sake of equal coupling. This is indeed unfortunate, thus will be mitigated by an optimised positioning of the magnetic feedback sensors.

4.3.2 Choice of Sensor Positions

For the feedback sensors, appropriate positions need to be chosen. This was done by using an algorithm to minimise the condition number, which was developed in a semester thesis [56], relying on a gradient descent method. Additionally, it was found that eight, according to this metric, is the optimal number of sensors.

A set of positions for the eight sensors is called a configuration. Several configurations were tested in the prototype. The positions of the three main configurations which were investigated more closely are depicted in Figure 4.4.

The sensor positions in configuration 1 were found by running the condition number optimisation algorithm on a volume with 20 to 40 cm distance to the mu-metal cube along the y-axis. After running some test shielding measurements, which are described in more detail in the next section, it has been discovered that for the compensation of slow fields it is crucial to minimise the offset of the residuals field at the sensor positions. Therefore, placing the sensors in a more homogeneous region of the compensation coil field is favourable. For the positions in configuration 2, the sensors were placed closer to the centre of the prototype, which automatically implies more homogenous fields. In a last attempt, the positions of configuration 3 were calculated by constraining the volume of possible sensor positions to a region where the fields of the coils differ less than 0.5 µT from the target fields and conducting the optimisation again.

In Table 4.2 the theoretical and measured condition numbers are compared for the three configurations of feedback sensor positions. Even though the measured condition numbers are generally lower, most importantly, the ordering of the two data sets is the same. In both data sets the condition number of configuration 1 is the lowest, as it was optimised over the largest volume, which leaves the most choices of positions. The constraint of the possible positions in the optimisation for configuration

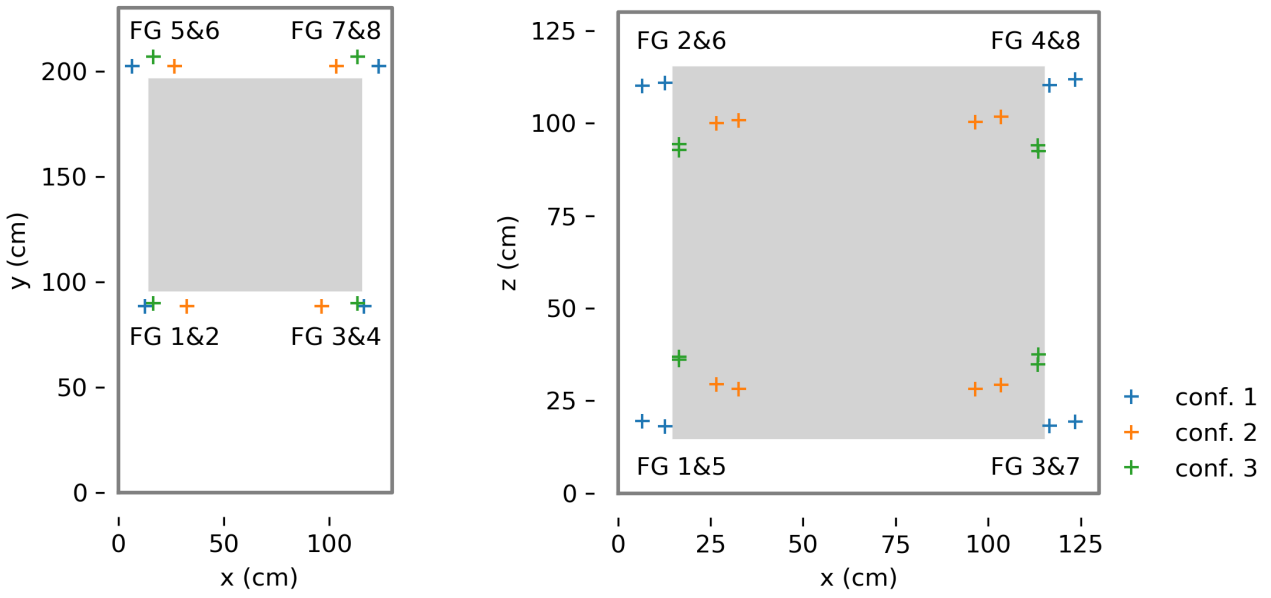

(a) xy-plane

(b) xz-plane

Figure 4.4: In dark grey the outline of the prototype coil system is shown. The cube in light grey shows the position of the mu-metal cube. The crosses mark the positions of the fluxgates in each configuration. In each of the four marked positions two sensors are mounted at different heights, the detailed coordinates of the positions belonging to configurations 1 and 2 are given in the Appendix of [57]. Coordinates of configuration 3 are listed in Table 4.3

	config. 1	config. 2	config. 3
calculated	4.69	6.72	5.41
measured	4.56	5.92	5.09

Table 4.2: Comparison of the theoretical and measured condition numbers for the three configurations of sensor positions.

3, leads to slightly higher condition numbers. Naturally the condition number of configuration 2 are the highest, as the positions were not subject to an optimisation.

The correlation matrix M which was measured for configuration 3, is shown in Figure 4.5. In the first three rows it can be seen how the homogeneous x-, y- and z-coils couple with approximately 50 µT each to the x-,y- and z-axis of each fluxgate magnetic sensor. For the gradient G1-G5 it is harder to spot a pattern, as there the coupling depends on the position of the respective sensor. Nevertheless it can be seen that the coupling is intrinsically smaller, around 20 µT.

Additionally the inspection of the feedback matrix is as well a tool which allows checking the correct functioning of the AMS prototype.

4.4 Shielding Measurement

The standard procedure to quantify the performance of so called magnetically shielded rooms (MSR) is to measure the shielding factor [58]. As the prototype is the active counterpart of these passive shields the shielding factor can also be used to describe the latter. Additionally, not only the shielding factor of the AMS prototype and the mu-metal cube will be measured but also the shielding factor of the combined system. Therefore it is ideal to use the same performance estimator in all measurements.

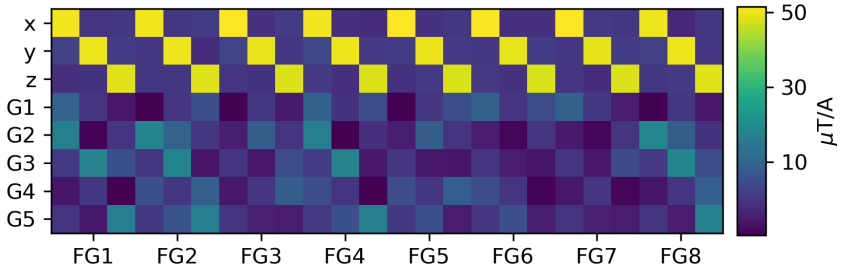

Figure 4.5: Correlation matrix M for configuration 3. The rows designate the coils of the AMS prototype, where the columns correspond to the readings in x-, y- and z-direction of the eight flux gate (FG) magnetic sensors as numbered in Table 4.3.

	x[cm]	y[cm]	z[cm]
FG1	16.5	90	36.1
FG2	16.5	90	92.8
FG3	113.3	90	37.6
FG4	113.5	90	92.4
FG5	16.5	207	36.9
FG6	16.5	207	94.3
FG7	113.4	207	34.8
FG8	113.4	207	94.1

Table 4.3: Coordinates of the optimal positions for eight magnetic feedback sensors in the prototype.

Essentially the shielding factor S describes the ratio between the size of a magnetic field perturbation caused by a modulated external field at a position in the middle of the passive/active shield with no shielding to the size of the perturbation at the same position with the shielding in place. The shielding factor S is therefore given by:

$$S = \frac{|B_\mathrm{o}|}{|B_\mathrm{i}|},$$

where B_o is called the outer field and corresponds to the amplitude of the disturbance without the shielding and B_i is called the inner field and corresponds to the value with the shielding. In the case of a performance measurement for large magnetically shielded rooms, the outer field at the position in the centre can only be calculated and not measured. In our prototype setup we have the advantage that the passive shield can be removed or the active shield can be turned off and therefore no calculations are needed.

We are following the measurement procedure given by [59] with a few modifications:

i Apply a field with sinusoidal modulations to the mu-metal cube

ii Detect the field for x-,y- and z-axis at one point

iii Fit amplitude and phase with the least square method

iv Compare calculated to measured values

This procedure has to be repeated several times for different frequencies of the applied external sinusoid field, as the shielding factor of active shields depends on the frequency of the disturbance. For the AMS of the n2EDM experiment, lower frequencies are of interest because in this range the MSR has a decreasing shielding capacity as already shown in Figure 3.1. The most important frequency range is around $1/180\text{s} = 5.6\,\text{mHz}$ as it corresponds to the inverse of the precession time. In our prototype measurement we therefore mainly investigated frequencies from 1 mHz to 1 Hz .

In the prototype setup we used a square coil with a side length of one meter, which is connected to a waveform generator to provide the external sinusoidal field. A sketch of the positions of the excitation coil relative to the AMS for the different field directions, as well as picture of the setup in x-direction is shown in Figure 4.6. The magnetic sensor to measure the shielding is placed in the middle of the prototype at the centre of the profile structure.

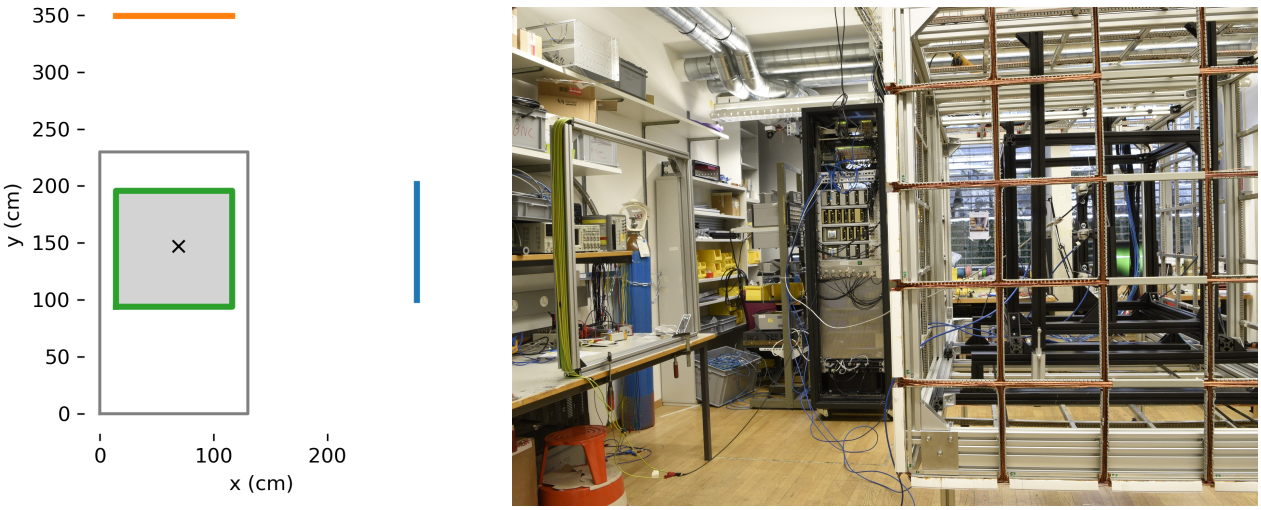

(a) (b)

Figure 4.6: Left side: Schematic depiction of the shielding measurement, where the position of the excitation coil is coloured in blue, orange and green for the x-, y- and z-direction measurement. The mu-metal cube is shaded in light grey, and the outline of the prototype is black. The cross in the middle of the setup corresponds to the central sensor used in the evaluation of the measurement. Right side: Picture of the shielding measurement setup. The excitation coil is positioned to produce a field in x-direction. The profile structure inside the AMS holds the central magnetic sensor.

The coil was placed in different positions depicted in Figure 4.6a in order to create fields in a dominant direction, either x, y, and z or a combination of the latter. The positioning of the coil was restricted by the geometry of the lab. For each created field direction, the corresponding B field component readout at the magnetic sensor was investigated. This procedure makes the measurement comparable. The x-direction was tested on all configurations, whereas the other setups were tested only on configuration 3 which has proven to give the best results.

The position of the coil and the current in each setup was tuned such that the uncompensated amplitude of the oscillation was approximately 8 µT. This tuning of the amplitude is needed as the shielding factor depends on the amplitude and is bigger for higher amplitudes.

The shielding, calculated by the amplitude of the uncompensated divided by the compensated field, was measured for different frequencies over several orders of magnitude. The comparison of the measurement with outside fields in x-, y- and z-direction in configuration 3 is shown in Figure 4.7. It can be seen, that the curves for the shielding factor in all three directions follow the same form. For

low frequencies there is a plateau, where the shielding factor is stable, up to around 100 mHz and then starts decreasing for higher frequencies. This feature can be explained by the inductance of the coil system, which causes that the currents can not be changed fast enough for high frequencies. Around the frequency of interest, 5.6 mHz, the shielding factor is stable. In the following, the shielding factor at these frequencies will be referred to as the shielding factor of the system in a certain setup.

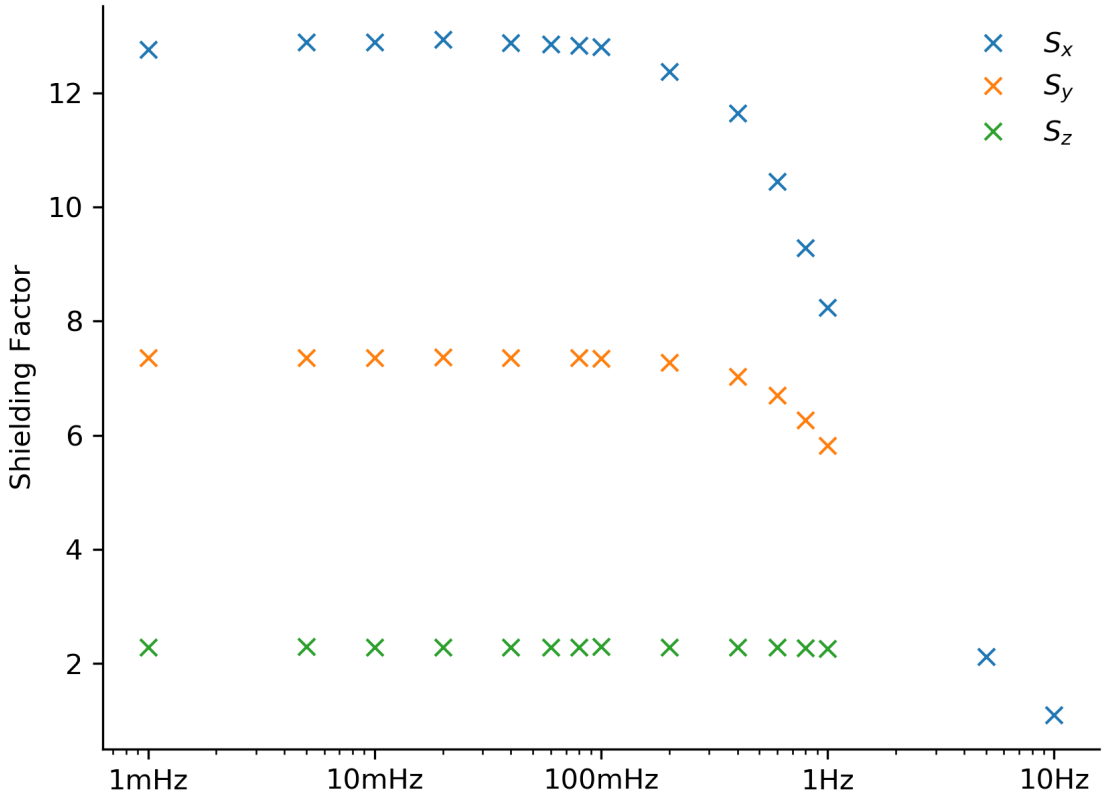

Figure 4.7: The shielding factor in x-, y- and z-direction plotted at different frequencies.

Most prominently, it can be seen, that the quasi static shielding factors for the different directions differ up to a factor of 6 between the z- and x-direction. This behaviour can be explained by the position of the excitation coil in the different setups. Depending on the alignment and the distance of the coil to the centre magnetic sensor, the fraction of higher order fields, which can not be compensated by the system, changes. To support this hypothesis, a simulation analysing the static shielding factor for these three setups was conducted. In the simulation the field was implemented using the Biot-Savart law and the field at the feedback sensor was analysed for residual higher order terms. It was found that the static shielding factor in x-direction is 2.1 times bigger than in y-direction and 7.0 times bigger than in z-direction. Taking into account that for the simulations a perfect prototype system was implemented, the simulation results are in agreement with the measured values.

The shielding in x-direction was measured for all three configurations of sensor positions. The summary is given Table 4.4. From this table it can be seen that the shielding factor is an interplay between a good condition number on one hand while on the other hand ensuring that sensors are placed only in regions where the compensation coil field is homogeneous enough. Especially when compensating the magnetic field to zero, the offset of the attenuated oscillation from 0 µT is relevant. This offset called $B_{\mathrm{res},x}$ is minimised when the magnetic sensors are placed in a more homogeneous

region. It has to be noted that the specified sensor accuracy is $\pm 0.5\,\mu T$.

	config. 1	config. 2	config. 3
S_x	8.45	6.69	12.86
$B_{\mathrm{res},x}$	$0.51\,\mu T$	$-0.18\,\mu T$	$-0.03\,\mu T$

Table 4.4: Comparison of the shielding factors in x-direction S_x and the residual offset in x-direction $B_{\mathrm{res},x}$ for the three configurations of sensor positions.

For the prototype setup it was not possible to place the magnetic sensors in a region with higher homogeneity, as the volume was already very constrained. The positions of configuration 3 showed the best performance and were therefore used in further measurements.

4.5 Stability Measurements

The second performance criterion, apart from the shielding, is the stability of the system. The n2EDM experiment will run continuously over several years. Therefore, especially the long-term stability of the AMS must be assessed. In the following, the tools and results of such measurements are discussed.

To quantify the stability of a system over several frequencies the standard tool is to use the Allan-Standard-Deviation [60]. For a time series the standard deviation can be estimated for different time periods. The Allan-Standard-Deviation is given as

$$\sigma_{\mathrm{Allan}}(\tau) = \sqrt{\frac{1}{2(N-1)} \sum_{i=1}^{N-1} \left(B_{i+1}^{\tau} - B_i^{\tau}\right)^2}, \tag{4.5}$$

where τ is the integration time, $N = T/\tau$ the number of subdata samples of length τ and i the number of the subdata sample. B_{i+1}^{τ} is then the integration of the magnetic field over τ for the i-th subdata sample. The Allan deviation was calculated for a measurement taken over night with configuration 3 which has shown the best performance. In Figure 4.8, it is plotted for all nine magnetic sensors, eight for the feedback algorithm and one to monitor the field in the centre of the prototype. Even in a quiet environment, such as during the night at ETH, the prototype is able to improve the field stability over the range of 10 s and 1000 s.

4.6 Interplay with the Mu-metal Cube

In the surrounding field compensation system (SFC) of the old nEDM experiment, it was shown that the placement of the feedback sensors is of great importance, especially if the coil system itself is not orthogonal [53]. Even though the AMS for n2EDM is well defined, the impact of the fluxgate positioning on the stability of the system can be investigated. The main difference to the case described in the previous sections, is the fact that a mu-metal cube inside the volume of the AMS distorts the field produced by the AMS. Therefore any simulation which investigates the feedback sensor positioning needs an underlying simulation taking into account the effect of the mu-metal.

4.6.1 COMSOL simulations

The tool which was used to perform the simulations is COMSOL Multiphysics. The simulation made use of the "Magnetic Fields, no Currents" interface, in which the background fields of all the eight coils was read into. The mu-metal was simulated by a 5 cm thick cube in order to reduce meshing problems. The relative magnetic permeability of the cube is set to 2'000, but it was found that the shape of the magnetic field outside the cube did not depend on the relative magnetic permeability, if they are high enough.

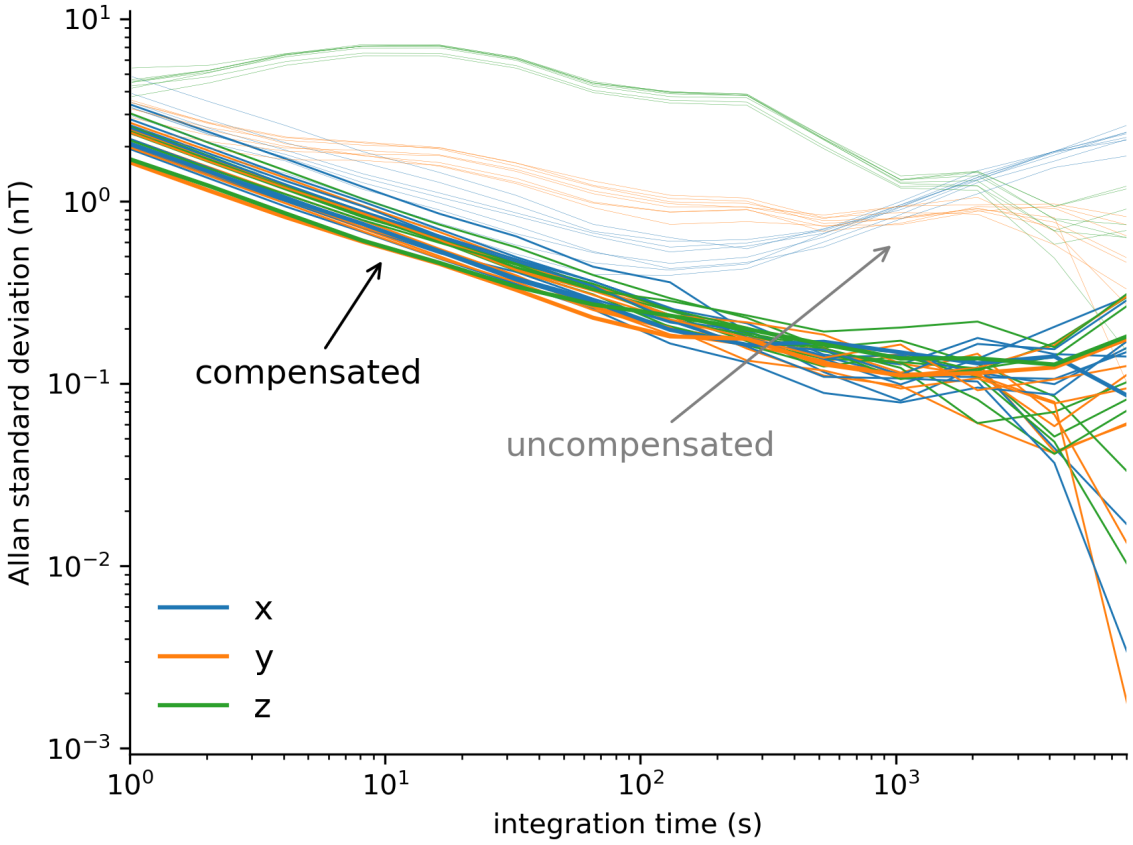

Figure 4.8: Allan deviation for a long time measurement taken during night using configuration 3. It can be seen that, even during a quiet night, the prototype AMS can improve the stability for integration times between 10 s and 1000 s.

In a first step, the simulation had to be validated with magnetic field measurements. For this reason a profile was inserted in the prototype AMS volume along the x-axis and the magnetic field was measured at several positions along the profile. Additionally a measurement of the background field was taken at every position, such that the background field could be subtracted. For every measurement point and every coil the x-, y- and z-component of magnetic field could be compared to the value of the COMSOL simulation. Unfortunately, the mapper could not be used, as the support structure of the mu-metal cube is not compatible with the mapper rail. In Figure 4.9 this comparison is shown for the y-coil. It can be seen that simulation and measurements match over the whole range of the AMS width in x-direction.

4.6.2 Shielding

With the previously validated COMSOL simulation, the magnetic field values in the volume taking into account the distortion by the mu-metal cube were calculated for all eight coils. These simulations served as input to the feedback sensor position optimisation. As in the case without mu-metal the algorithm was trying to minimise the condition number of the system.

The algorithm was used to calculate various sets of feedback sensor positions on different inputs which were compared to each other. The first figure of merit is how much the measured condition number diverges from the calculated one. Similarly to the sensor positioning study without mu-metal

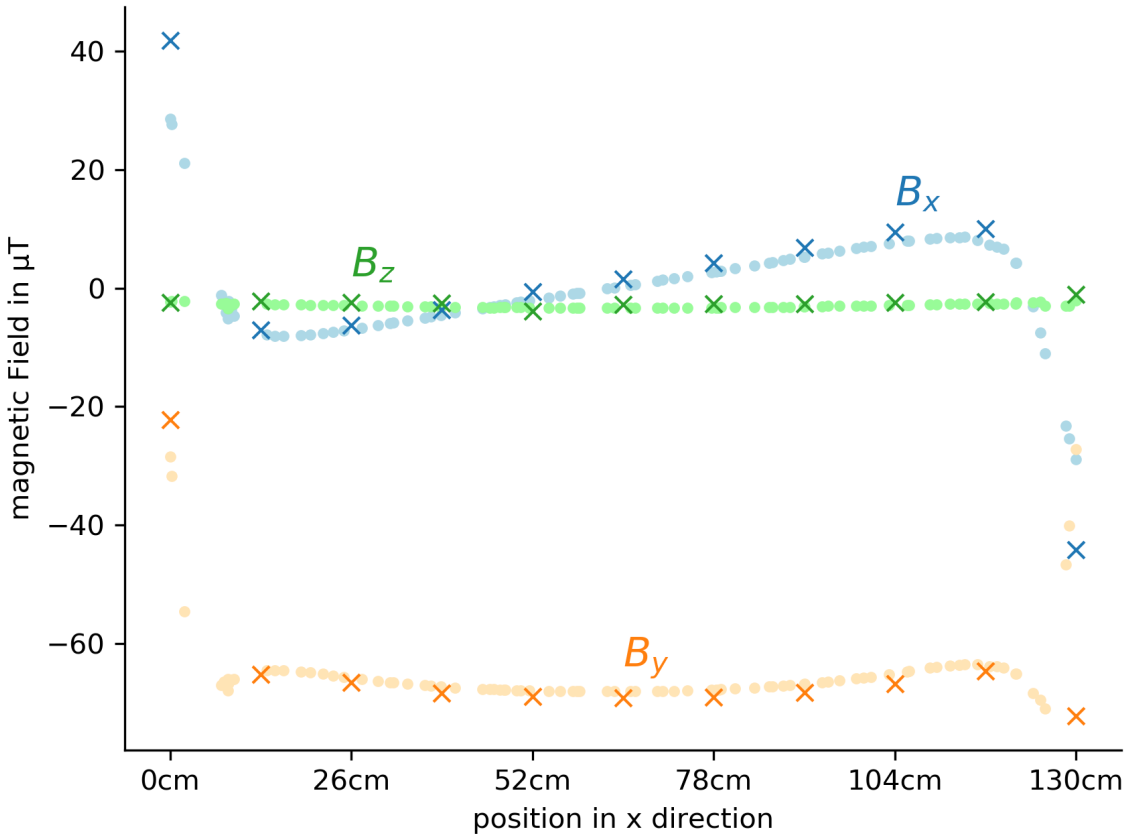

Figure 4.9: Comparison between the simulated magnetic field (dotted line) and the measured one (indicated with crosses) for B_x, B_y and B_z produced by the y-coil. The measurement was taken along a line in x-direction in front of the mu-metal.

cube, configurations 5 and 6 were calculated with the optimisation algorithm such that they lie in the region which is initially homogeneous to 0.5 µT. In the second case a method called the scaling of the gradient was used to look for the minimal condition number. The scaling of the gradients accounts for the fact that the goal fields for the homogeneous and gradient fields do not have the same strength (50 µT vs. 20 µT/m over the characteristic length of 1 m). Therefore, in the matrix calculated during the optimisation algorithm the matrix is scaled by the goal field strength. This should make the coupling to the homogeneous and gradient coils equally important. For configuration 7 the optimisation algorithm was run, without constraints, on the homogenous region, again employing the scaling of the gradients.

In Table 4.5 the condition numbers are listed for all the three examined configurations. The calculated condition number is the value directly coming from the optimisation algorithm. For the consecutive measurement the magnetic sensor could not always be perfectly placed on the suggested positions due to geometrical reasons. Therefore, the simulated condition number, re-evaluates the value with the positions used in the corresponding shielding measurement. The deviations between the suggested and implemented positions are normally less than 1 cm. The last row contains the measured condition number as it can be extracted from the measured matrix.

Clearly the unconstrained configuration 7 yields the highest condition number throughout all categories. This is mainly due to the fact, that those positions are closest to the coils and therefore the deviation from the simulated field is the highest. For the other two configurations the condition

	config. 5	config. 6	config. 7
calculated	9	/	/
simulated	9.2	8.42	14.39
measured	9.6	10.1	10.7

Table 4.5: Comparison of the theoretical, simulated and measured condition numbers for the three configurations of sensor positions.

numbers yield similar values, nevertheless the scaled approach produces a smaller simulated condition number.

For all the three configurations the shielding in x-direction was compared. The measurement was done in the same way as described in the previous section about the shielding without mu-metal cube.

Again the amplitudes of the compensated field and the offset of the residual field from zero was looked at. The comparison of those values for shielding in x-direction can be found in Table 4.6. The values of the compensated magnetic fields correspond to an additional shielding factor of approximately 2. Compared to the previously measured shielding factors those values are rather low.

	config. 4	config. 5	config. 6	config. 7
B_x[nT]	57	53	55	42
$B_{\mathrm{res},x}$[nT]	-108	-200	-167	-210

Table 4.6: Comparison of the compensated magnetic field B_x and the residual field $B_{\mathrm{res},x}$ for configurations 4-7.

Nevertheless, several facts need to be considered. First, when looking at the uncompensated measurement only the mu-metal will provide shielding. It is known that the shielding factor of the mu-metal is higher, when the outside disturbances are bigger. Thus in the case with the AMS, a combination of a lower mu-metal shielding factor and the additional shielding by the AMS is measured. Therefore the AMS shielding factor is underestimated. It was measured that the shielding factor of the mu-metal drops around 40% for a decrease of the outside amplitude of a factor of 10.

Second, as discussed previously the fields produced by the excitation coils contain a certain amount of higher order fields, this is especially true for the fields in x- and z-direction. In the case with the mu-metal the fraction of higher order fields gets larger due to the bending of the field lines and makes it harder to be compensated by the AMS. This can be verified by looking at the measurements in y-direction conducted with configuration 6, where the disturbance could be attenuated to 6 nT which corresponds to a shielding factor of 79.

In Table 4.6 it can also be observed that the residual fields $B_{\mathrm{res},x}$, are an order of magnitude higher than in the measurements without the mu-metal. The reason for this is again the distorted field lines, which show a different field at the magnetic sensors, than the field which causes this distorted field lines and thus should be compensated.

Because it is assumed that the effect of the mu-metal vanishes if the surrounding field is compensated close enough to 0, a shielding measurement was attempted with the best configuration coming from the optimisation without mu-metal. Surprisingly configuration 3 performed well at attenuating the amplitude of the disturbance, but still revealed an offset in the residual field of a few hundred nT.

The details and values of all measurement conducted can be found in [61].

4.7 Conclusion

In this chapter it was demonstrated, that coils constructed with the "method of simple coil design" are able to achieve the required 1 µT homogeneity in the region of interest. The positions of the feedback

sensors were optimised for with an algorithm making use of the gradient descent method. It was shown that the thereby produced condition numbers can also be reproduced by a measurement with the prototype system. Several configurations of feedback sensor positions were tested for their shielding factor in the case without mu-metal and a preferred configuration could be chosen. Additionally, it was shown that the AMS performs well in regards of stability in the range of 10 s to 1000 s. The COMSOL simulation including a mu-metal cube in the AMS could be validated by measurements with the prototype system. The output of the simulation was used to optimise the condition number of the feedback sensor positions. It was shown that if the outside field is consisting of homogeneous and first order gradient fields, the amplitude of a disturbance can be attenuated. Nevertheless an offset in the residual magnetic field of several hundreds of nT has to be expected. The presented approach can be used in n2EDM to produce an optimised set of sensor positions. Those positions can then serve as a starting point in the optimisation of the AMS performance.

Chapter 5

Mapping

In Section 3.3, it was shown that one of the input parameters to the algorithm of simple coil design, is the target fields one wants to create. Therefore, the external fields expected on the experimental site have to be known. After the predecessor experiment nEDM had been dismantled in October 2017, the area was empty for some months until the first construction works for n2EDM started. During this period of time, two mapping campaigns were conducted, where the external magnetic field environment was investigated more closely.

In this chapter we are going to describe the tools with which the magnetic fields were measured and what external fields are expected for n2EDM data taking. The goal was to quantify these magnetic fields to a precision of a few µT, to give a limit what the magnetic field range of the next generation AMS must at least be.

5.1 The Mapping Tower

The outline of the empty experimental area where n2EDM is situated was approximately 10 m by 10 m. The MSR itself, having a height of 5 m, will be raised 2 m above the ground. The general idea of the mapping is to gain magnetic field information in as many points as possible in the described volume. In order to reach the 8 m in height, the mapper was designed to be a tower-like structure, called the mapping tower. It has a modular design, consisting of 5 identical modules being stacked on top of each other. The triangular holding structure, which forms the base of the mapping tower, is normally used for building portable event stages. Each module is 2 m tall and has three 3-axis fluxgate magnetic sensors attached to it.

It was important in the design phase as well, that the footprint of the mapper is as small as possible, such that narrow spots can be mapped. For this reason the tower was mounted on a heavy aluminium ground plate, which prevented the tower from tilting over. To allow the construction to move across the experimental area three wheels were attached.

Additionally the mapping tower is made of non magnetic materials, such that the measurements were not biased. A picture of the mapping tower is shown in Figure 5.1.

5.1.1 Positioning System

Once the magnetic field information is collected, it needs to be connected to a position in space for later analysis. Hence an accurate positioning system is needed. Two strong aluminium profiles of approximately 4 m length each, screwed together in a right angle served as local coordinate system. Two string potentiometers were installed at the end of the coordinate system (point A and C), and one at the point where the aluminium profiles are connected (point B).

The two string potentiometers A and B were attached to a screw in the middle of the mapper. The length of the strings can be calculated from the measured output voltages. By intersecting two circles, centred at the location of the string potentiometer with radii equating the string length, the

Figure 5.1: The mapping tower in action during the January 2018 mapping campaign.

respective local positions of the mapping tower can be inferred. The additional string potentiometer C has been used to calculate the direction in which the axis of the mapping tower is pointing. A visualisation of the mapper geometry is shown in Figure 5.2.

Together with the absolute height of the fluxgate above ground, all necessary spatial information is available in order to put the measured magnetic fields in context. During the first mapping campaign, it was noticed that certain points which are in line with one of the coordinate's system arms had very poor precision with regards to position. Furthermore, the connection between the arms was not absolutely rigid and could thus introduce an additional imprecision to the measurement. Consequently, a slightly different configuration was used in the Summer 2018 mapping campaign: Instead of an L-shape piece just one aluminium profile was used as coordinate system. All three string potentiometers were attached in line on this bar. For the second campaign the 5 m by 5 m platform, which served as support for the n2EDM apparatus, was already installed. Thus, the coordinate system was rigidly attached to this structure and would not be moved any more.

5.1.2 Calibration

Following the procedure described above, the positioning information can only be measured with respect to the local coordinate system e.g. the angle made of aluminium profiles. However all the relevant information, such as the position of the MSR, is given with respect to a global PSI system. Due to the lack of accessible benchmarks, it is not possible to measure the location of our local coordinate system within the global framework with sufficient accuracy. For the n2EDM experiment, the axes of the experimental coordinate system will be parallel to the global PSI ones, with the centre of the precession chamber serving as origin. In order to calibrate our local systems, benchmarks points were used, which were drawn on the floor. Subsequently, these points were precisely measured by the

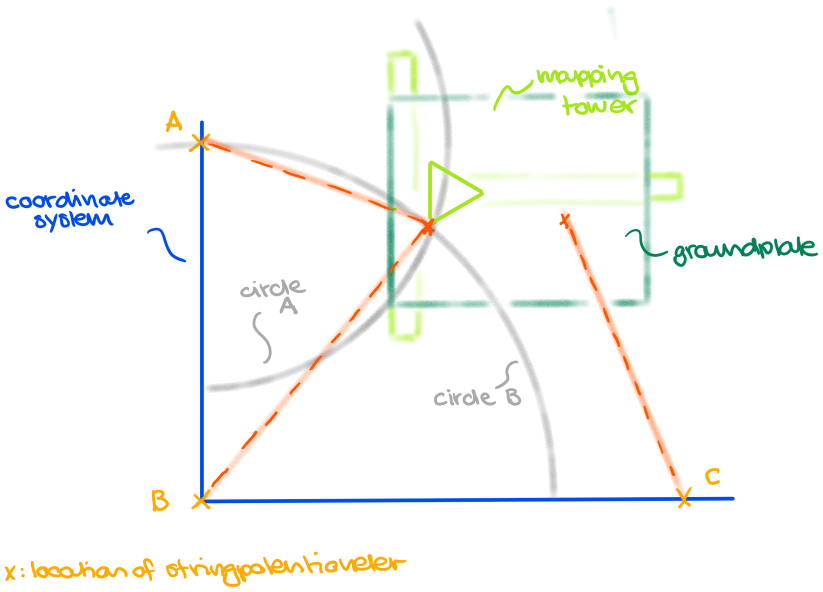

Figure 5.2: Top view of the mapper positioning system: The position of the mapper is inferred by intersecting the two circles centred at the string potentiometer positions (A,B). The string potentiometer at position C is used to calculate the rotation of the mapper (green).

point	Δx[cm]	Δy[cm]
δ	2.50	-3.45
η	0.87	-0.17
ζ	-0.38	2.40
γ	-2.52	-1.48
β	-0.48	2.69

Table 5.1: Deviation in x- and y-direction for every calibration point of the fit. Positions are illustrated in Figure 5.3

surveying group at PSI.

In order to establish calibration maps, the mapping tower was accurately centred onto these benchmark points and the voltage was then read out. The three unknowns (x-position, y-position and angle of the local system with respect to the global coordinates) were then fitted to the data. An example of such a calibration map, which was taken during the January 2018 mapping campaign can be seen in Figure 5.3.

In this figure it can also be seen that one point in the centre of the area and one in the lower left corner have been rejected as calibration points. The point in the centre, has been drawn on movable wooden plates, the fact that those plates can be easily shifted disqualified it from the use in the calibration. For the point in the lower left corner, it can be seen, that it is almost in line with the aluminium profile, meaning that the system is very insensitive to its position. In the January mapping campaign, several such maps were taken. Hereby also the relative position of the local coordinates system within the global coordinate system has been significantly altered.

The only means to quantify the quality of a calibration, is to compare the global coordinates of the calibration points to the mapped calibration points. In Table 5.1 the spatial differences is shown for one pair of calibration maps. It can be seen that the differences between the global and mapped coordinates are less than a few cm.

point	Δx[cm]	Δy[cm]
AA	1.52	-1.47
BB	-0.51	-1.65
CC	-0.65	1.09
DD	-1.14	0.20
EE	0.78	1.83

Table 5.2: Deviation in x- and y-direction for every calibration point to the fit, in the June 2018 mapping campaign.

Because several calibration runs with the coordinate system in the same place are available in the January 2018 campaign, the reproducibility of the mapping was investigated as well. This is done by transforming a grid of point in the local coordinate system with two calibrations, and to look at the distance between the global mapped points and the measured points. The fitted parameters of these two calibrations are given below:

- Calibration 1: Δx = 1216.5 cm, Δy = 966.7 cm, $\phi = 107.9°$
- Calibration 2: Δx = 1216.7 cm, Δy = 966.8 cm, $\phi = 108.0°$

Here Δx, and Δy denote the offset in x- and y-direction respectively, where ϕ is the angle between the coordinate systems. The above described analysis reveals that the two calibrations taken, within 10 minutes, on average only produce a 0.12 cm difference in the global coordinate system. This very good compatibility is not surprising, looking at the high level of agreement which is already given in the calibration parameters.

To explore the influence of an error in the calibration parameters on the globally mapped points, the analysis was repeated with a larger disagreement between the two calibrations. Values of Δx and Δy of maximally 10 cm were considered. It was found, that this shift linearly propagates to a global misplacement of the same magnitude. This means that a 10 cm difference in Δx directly corresponds to a on average 10 cm misplacement of the mapped points. Looking at differences in angles, the effect is even stronger, as a small angle mismatch of $\pm$ 1° results in a discrepancy of the mapped points of approximately 12 cm.

As mentioned previously, the geometry in the June 2018 mapping campaign was slightly different. This did not change anything in the manner the calibration has been conducted. As the platform itself has been very well positioned, one could directly relate the coordinates of the calibration points to some benchmark points on the platform. The five points used for calibration are labelled AA to EE. In Table 5.2 the difference between the global measured and the global mapped calibration points is shown. Compared to the January 2018 calibrations, the discrepancy was even smaller.

The visualisation of the only June 2018 calibration map can be seen in Figure 5.4. The calibration parameters used for all the June 2018 maps are the following:

- offset x-direction: 533.42 cm
- offset y-direction: 1009.38 cm
- angle: 89.18°

Undoubtedly, it is surprising that the rotation angle of the aluminium profile is not exactly 90° even though it is fixed to the platform. Probably this additional rotation comes from the fact, that the string potentiometer have different offsets which have not been fitted perfectly in the string potentiometer calibration. Generally it is assumed, that the June 2018 calibration is even more reproducible than the January 2018 calibration as additional uncertainties, such as the angle between the profiles were eliminated. Furthermore, the distances on the platform were much smaller than in the mapping on the ground, which implies a smaller error on the readout voltages.

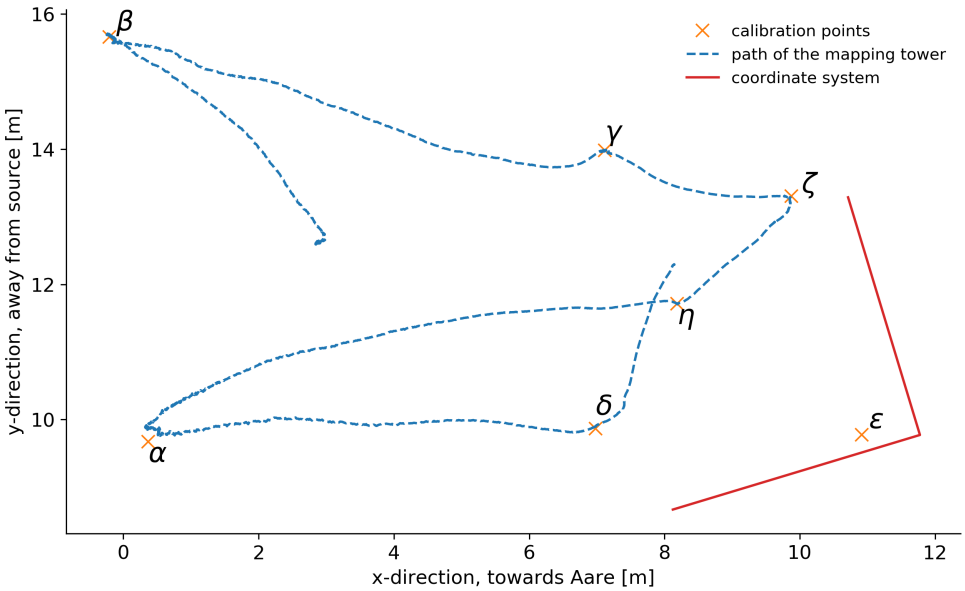

Figure 5.3: Plotted in orange is the position of the benchmark points used for calibration in area south. The blue line depicts the path of the mapping tower during the calibration map. The red L-shape is the local coordinate system, where the string potentiometers were attached.

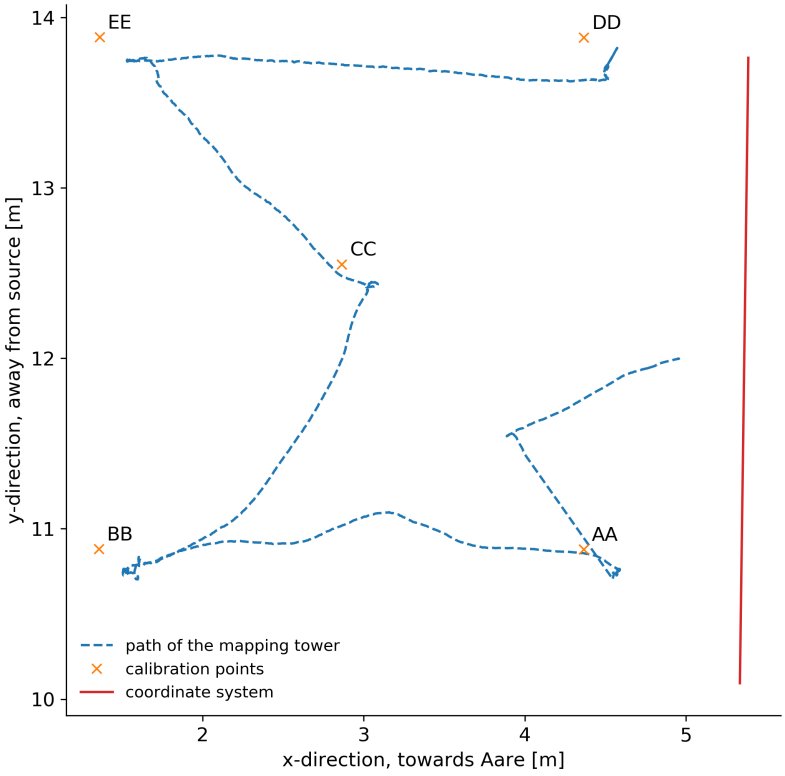

Figure 5.4: Depiction of the June 2018 calibration map. The red line on the right side represents the aluminium profile which served as coordinate system. The mismatch of the path of the tower mapper and the calibration points, is due to the fact that the relative position of the screw to the positioning device has not been corrected for.

5.2 Mapping Campaigns

To record a magnetic field map, the mapping tower was moved over the whole floor of the area to be mapped. For a good quality of the maps it is important to map quickly on one hand, but on the other hand to also cover the area as densely as possible. For the construction of the active magnetic shielding we especially rely on a good knowledge of the time dependent part of the field. In our case these are the SULTAN and COMET magnets which create most of the field changes on the experimental site. In order to retrieve the contribution of these magnets, the approach was to map the area once with the magnet on and another time with the magnet off. Ideally these two maps should be taken within a short period of time. As these maps had then to be subtracted from each other and we always took different paths in the different maps, the data had to be binned. Naturally the binning destroys some of the positioning information, which in this case is not a problem as the positioning accuracy of the mapping tower is estimated to be in the cm range. The main differences in the setup of the two mapping campaigns were already highlighted in the preceding part. Apart from those, the tower mapper has been used for the first time during the January 2018 mapping campaign. Naturally this led to a numerous amount of hardware and software start up problems. In consequence, it was not always possible to conduct the mappings in the most preferable way, meaning that background maps could not always be taken directly after magnet maps. Additionally it was encountered that the operation of the surrounding main magnets did not always adhere to the announced schedule, thus it occurred that maps were taken during ramps or magnet maps were taken after the magnet was already turned off. For the subsequent analysis, it was often not possible to pair maps which were taken one after another within a short timescale.

In Table 5.3, a list of the most important maps during the January 2018 campaign can be seen.

As the time between the magnet map and the background map can be seen as a quality criterium, the January 2018 campaign is considered to be less reliable, and has therefore more the character of a test measurement. Nevertheless, in some cases where the quality of the January 2018 maps is reliable it is interesting to compare the results of the two campaigns.

Additionally, the intrinsic performance of the two mapping campaigns can be investigated, by looking at the Cartesian harmonic decomposition of two maps with the same magnet configuration which were taken shortly one after another. For the January 2018 mapping campaign we took the two COMET maps taken at 17:01 and 17:17 respectively on 04/01/2018. When comparing the harmonics decomposition of the two maps they do not differ more than 1.2 µT in any of the components. Performing the same test on the two COMET background maps from the June 2018 mapping campaign taken at 00:34 and 00:54 on 16/06/2018 the maximum difference in any component of the Cartesian harmonics decomposition yields 0.13 µT. Even though there is marginally less time between the January 2018 maps, the stability of the June 2018 maps is approximately one order of magnitude better.

In the June 2018 mapping campaign we relied less on the announced operation schedule of the

Maps January 2018 Mapping Campaign

date	time	COMET	SULTAN	COBRA
2018-01-04	15.41	on	off	off
2018-01-04	17.01	on	off	off
2018-01-04	17.17	on	off	off
2018-01-05	16.12	off	off	off
2017-12-21	10.04	on	9.5 T	on
2017-12-19	18.42	on	off	on
2017-12-21	14.23	on	2T	on
2017-12-21	17.35	on	off	on

Table 5.3: List of the maps taken during the January 2018 mapping campaign.

surrounding facilities, but rather consulted the system of remote magnetometers. These are around a dozen magnetic sensors which were installed in December 2017 not only in the UCN area South, but in the whole experimental hall. The surrounding magnetic field can be monitored and large field changes can be spotted. As the remote magnetometers are distributed in a large area, small field changes, for example from the COBRA magnet, can not be seen, therefore it is not ideal for the extraction of gradients due to its large distance to the magnetic sources. It is also thanks to the remote magnetometers that in hindsight some of the January 2018 maps could be identified as not relevant. Ramps especially of SULTAN and COMET can be very well seen in those data, an example is given in Figure 5.5.

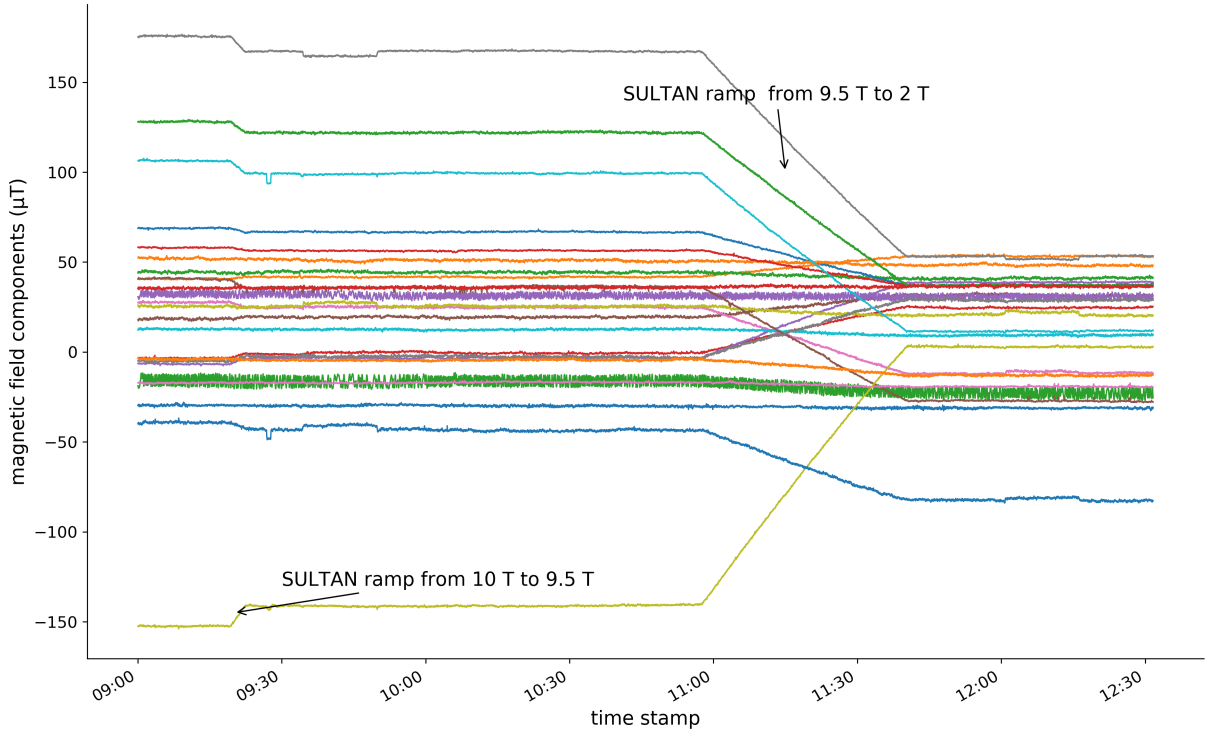

Figure 5.5: In the data of the remote magnetometers, large magnetic field changes can be identified.

In Table 5.4 the list of the maps from the June 2018 mapping campaign which were used in the analysis is given.

The general approach was to fit Cartesian Harmonic polynomials up to 3rd order to the field. In the next section we only show the decompositions up to 1st order gradients. First, by comparing the 0th to 3rd order decomposition, it could be seen that the standard deviation of the residuals drops several µT to the region of around 1 µT between only using homogenous field and using up to 1st order polynomials. The difference between 1st and 2nd order is often less than 0.1 µT and between 2nd and 3rd order much less than this. Weighing these numbers up to the fact that being able to create 2nd order fields doubles the construction effort from eight to fifteen coils we can conclude eight coils have a good cost performance ratio. Furthermore, if really needed, additional coils can always be added on top of the existing ones.

In the following we are going to investigate the varying fields of the surrounding magnet facilities

Maps June 2018 Mapping Campaign

date	time	COMET	SULTAN	COBRA
2018-06-15	17.45	on	off	off
2018-06-15	23.51	on	off	off
2018-06-16	00.34	off	off	off
2018-06-16	00.55	off	off	off
2018-06-13	15.33	on	9.5 T	off
2018-06-13	17.12	on	off	off
2018-06-18	10.12	on	1 T	off
2018-06-18	10.59	on	1 T	off
2018-06-18	11.19	on	2.5 T	off
2018-06-18	18.09	on	off	off
2018-06-15	14.36	on	off	on
2018-06-15	14.42	on	off	off
2018-06-15	15.10	on	off	off

Table 5.4: List of the relevant maps taken during the June 2018 mapping campaign.

independently.

5.2.1 COMET

The COMET magnet is normally running during most of the data taking and will only be occasionally ramped down for maintenance. Therefore we are not only interested in the pure COMET field but also in the total magnetic field in our area. For the COMET magnet we are in the comfortable situation that during the June 2018 mapping campaign we have two maps were COMET is on and two maps with COMET off, where three of them have been taken within approximately one hour, meaning directly after the ramps down (see maps 1-4 in Table 5.4).

As can be seen from Table 5.3 a total of three maps was collected within one and a half hours, while COMET was running. However, the corresponding background map could only be taken the day after, because of operation schedule mismatches. Also the same background map had to be subtracted for all three COMET on maps. The three resulting pure COMET decompositions, are compatible in each component within 2 µT, where the gradients were scaled with the characteristic length of 5 m.

Comparing now the January 2018 to the June 2018 decompositions, the homogenous contributions still agree within 2 µT, whereas the linear gradient contributions agree within 4 µT on the characteristic length of 5 m. This can be explained by the fact that the linear gradients describe the field more locally, as gradients tend to fall off faster than the homogeneous parts of the field. As the background map in January 2018 has been taken almost a day apart, this finding expresses the amount of variation which needs to be expected in the gradients with identical magnet configuration. Additionally it was helpful that the COMET January 2018 maps were taken during the extended Christmas holiday, where only a small number of magnet activities had to be expected.

5.2.2 SULTAN

SULTAN belongs to the EDIPO superconductor test facilities [62] and can be operated up to 11 T and in different modes. It is typically ramped up every workday in the morning and ramped down every evening. Throughout the day it is often operated at different field strengths and therefore will cause most of the dynamic fields during n2EDM data taking. The ramp-up in the morning up to 10 T takes approximately an hour and is therefore certainly in the frequency range which needs to be compensated by the AMS. It is advantageous for us that, due to the strong fields, the ramps can be very well seen in the remote magnetometer data. During the June 2018 mapping campaign we were

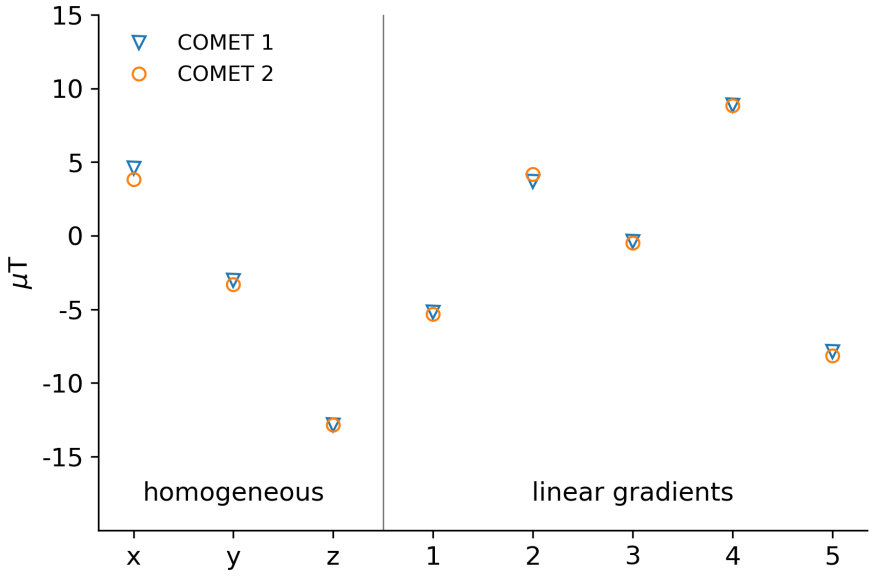

Figure 5.6: Decompositions of the two COMET pure maps into Cartesian harmonics. The gradients are intrinsically given in µT/m but were scaled by 5 m which is the side length of the MSR. The linear gradients 1-5 correspond to the Cartesian harmonic polynomials 4-8 in Table 3.1.

able to map three different field strengths of SULTAN, one at almost full field and two at lower fields. Due to its layout, SULTAN is expected to produce mainly fields in x- and y-direction, which has been confirmed as shown in 5.7.

When scaling up the homogeneous components to full field strength at 11 T, one ends up with $-34\,\mu T$ and $41\,\mu T$ in the x- and y-component respectively.

During the January 2018 mapping campaign almost no useful data could be taken, as in one case SULTAN was already in ramp down during the mapping and in the other case the background map has been taken four days apart from the SULTAN on map. When trying to decompose this map, around $20\,\mu T$ are missing in the y-component. It has to be assumed that the background was substantially differing in these two maps.

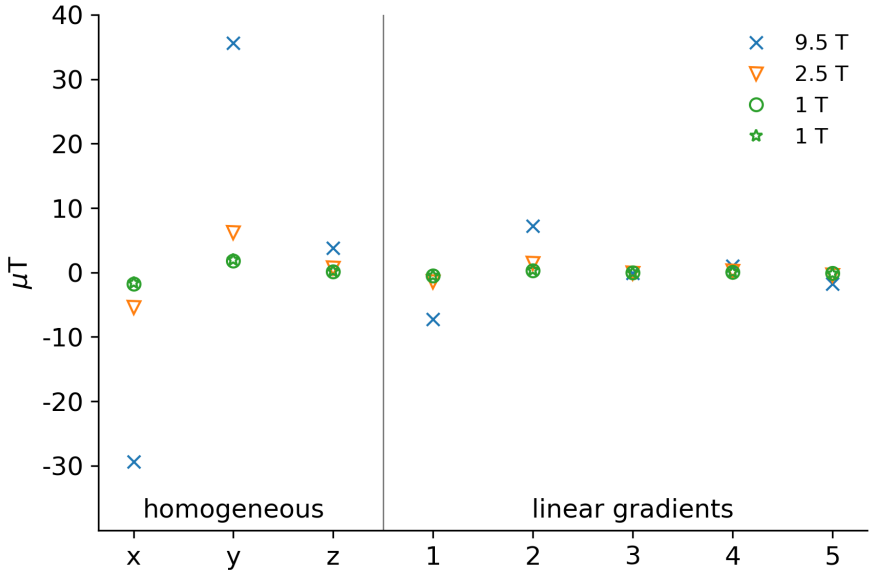

Figure 5.7: Decompositions of the two SULTAN pure maps into Cartesian harmonics. The gradients are intrinsically given in µT/m but were scaled by 5 m which is the side length of the MSR. The linear gradients 1-5 correspond to the Cartesian harmonic polynomials 4-8 in Table 3.1.

5.2.3 COBRA

COBRA is a superconducting solenoid magnet and part of the MEG experiment. It produces a central field of 1.23 T [63] and is operated irregularly. Among the discussed sources of dynamic magnetic fields, COBRA is the weakest. Nevertheless it produces a few µT mainly in x- and y-direction on the n2EDM experimental site. The COBRA contribution could only be extracted in the June 2018 mapping campaign. In Figure 5.8 two successful decompositions of the pure COBRA field into Cartesian harmonics polynomials are shown.

Even though the homogeneous contributions are rather small, the linear gradient fields are large in comparison. This is because COBRA was designed to produce gradient fields, such that emitted positrons at any emission angle can be guided efficiently.

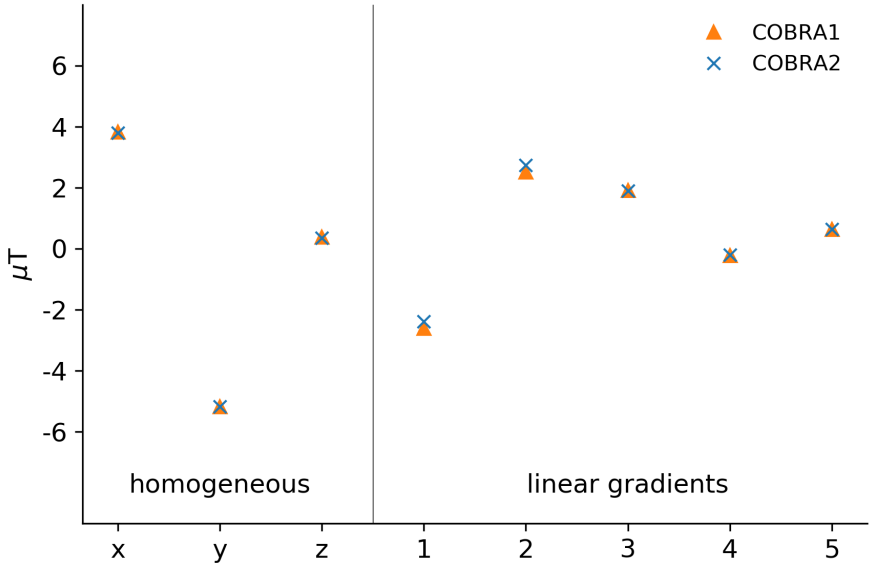

Figure 5.8: Decompositions of the two COBRA pure maps into Cartesian Harmonics. The gradients are intrinsically given in µT/m but were scaled by 5 m which is the side length of the MSR. The linear gradients 1-5 correspond to the Cartesian harmonic polynomials 4-8 in Table 3.1.

5.3 Mapping Summary

Utilising the maps from the two mapping campaigns, it was possible to decompose the main sources of dynamic fields into their Cartesian harmonic components. The thereby following requirements for the AMS performance are important inputs for the design algorithm. The most relevant fields created on the n2EDM experimental site are summarised in Table 5.5.

Additionally it has been shown that the results are reproducible among different maps and even mapping campaigns. However, this procedure is not a precision measurement and therefore adequate margins will be included in the AMS design. The analysis suggests that target fields of the AMS of 50 µT homogeneous fields and 5 µT/m 1st order gradient fields will be able to compensate any occurring magnetic field configuration.

C	S	c	x	y	z	1	2	3	4	5
x			4.2	-3.2	-12.5	-5.3	-3.9	-0.4	8.9	-8.0
	x		-34.0	41.2	4.4	-8.4	8.4	-0.2	1.2	-2.1
		x	3.8	-5.2	0.4	-2.5	2.6	1.9	-0.2	0.6

Table 5.5: Decomposition of relevant background magnetic fields into Cartesian harmonics. In cases were several maps could be decomposed, the average is cited. The values always refer to the full field of the corresponding magnet and are given in µT. The linear gradient values were scaled with a characteristic length of 5 m C: COMET S: SULTAN c: COBRA

Chapter 6

Design

In this chapter we are going to discuss the underlying ideas of the design for the n2EDM active magnetic shielding. The AMS was developed utilizing the method of simple coil design as discussed in Chapter 3, but there are deviations and additional considerations which will be highlighted in this chapter.

Most importantly we will discuss each step of the design procedure in context of the later construction of the system. Many deviations from an ideal system had to be accepted in order to ensure feasibility and practicability. One main issue which appears is the amount of total current, which had to be kept small, as it directly translates into weight and thus installation effort.

From the mapping described in Chapter 5 it can be seen that the fields were only reproducible down to a few µT. Therefore the system was specified to provide a field homogeneity of $\pm 1\,\mu T$ around the target fields. This value would then directly translate into a static shielding factor of 30 - 100 which enables data taking also during SULTAN ramps.

In Section 3.3 it was pointed out that the algorithm of simple coil design requires three main ingredients which influence the result: the target magnetic field(s), the grid and the set of points on which the optimisation will be conducted. In the following we will call the latter points of interest (POI's).

6.1 Target Fields

The choice of target fields for the AMS corresponds equally to a choice of coils and shapes of magnetic fields that can be produced and thus compensated. For the design of the prototype the Cartesian harmonics served as a basis for the coil system. Even though we can very well model the field of e.g. a SULTAN or COMET, this generic basis of coils has many advantages. For instance it provides a large flexibility in adapting to changing magnetic field environments, which for example would happen if a new facility were to be built nearby. In an diverse scientific institution such as PSI this can never be excluded. Additionally, with the mapping it was shown that with a decomposition of the various configurations the background magnetic field can be modelled very well. For these above mentioned reasons it has been decided to pursue the same approach as with the prototype and design coils which each correspond to a Cartesian harmonic polynomial as listed in Table 3.1.

The two questions which consequently needed to be answered, are: Up to which order of Cartesian harmonics should the system be built, and what should be the values of the fields? From the mapping it could be seen, that after a first order subtraction of the compensated fields, the residuals in the sub-µT range did not decrease any more. Consequently, on the µT resolution level which the mapping provides, there are no significant contributions of higher than 1st order gradients. This is probably due to the fact, that the experimental area is sufficiently far away from strong magnetic sources, such that higher order field components are already drastically decreased.

Looking at Table 3.1, it can be seen that in order to compensate up to first order gradients eight coils are required; one for each spatial direction plus five for the first order gradients. Considering the

subsequent construction and cost of eventual power supplies, eight is a practical number. In order to expand the system such that also 2nd order gradients could be compensated another eight coils would be needed. Contrasting this information to the fact that the contribution of 2nd order gradients are in the sub-µT range, led to the decision that only eight coils will be implemented to compensate up to 1st order gradients.

After analysing the comparison of decomposition for the various background fields and possible magnet contributions in Table 5.5 and adding a generous 20 % margin led to the following values of magnetic fields that need to be produced (and thus compensated) by the system:

- 50 µT for the three zeroth order homogeneous coils
- 5 µT/m for the five first order linear gradient coils

These two values plus the Cartesian harmonic polynomials up to first order provide the first input required for the algorithm of simple coil design.

6.2 Grid Design

As one of the goals of the AMS is to provide the most homogeneous field on the MSR surface, it is advantageous to place the grid structure as far away as possible. However, the space around the MSR is very limited and the experimental area is surrounded by massive biological shielding. Around the whole experiment a wooden shell is built to ensure a thermal stability of 1° C over the whole volume. This wooden structure will in the following be referred to as the thermo-house. As the coils of the AMS need later on to be supported, it has been decided, that the thermo-house should serve as holding structure. First of all, it makes an additional holding structure redundant and additionally the wood facilitates the later installation work. Equally important, the thermo-house fills the available space around the MSR efficiently and allows the AMS to be installed at a reasonable distance.

It is clear that the number of tiles in the grid has a big influence on the homogeneity of the magnetic field. With a finer meshing, the currents can be placed more accurately where they are needed. However also more holding structure and wires need to be installed resulting in a increased installation time. To estimate the approximate impact of the number of tiles, a regular grid structure with different meshing has been simulated.

The number of tiles needed to achieve a certain homogeneity can only be used as a guideline for the grid design, as it is not possible to implement a regular grid. Firstly, there are supports for the platform and the MSR which intercept a regular grid, but additionally there are big openings in the thermo-house itself. The main openings which break the regularity of the grid are shown in Figure 6.2. Most of these openings are used to access the thermo-house or to transport equipment in and out, such as the big front opening, the front door and the craning access. The magnet opening on the ceiling has a special importance, as it is used to crane out the super-conducting magnet in case of malfunction.

Obviously, the depicted openings are a threat to field homogeneity, as there are areas on the thermo-house where no current is flowing. The method to counter this current-free space are the so called connected doors. A connected door is a grid structure placed on the removable part of the thermo-house, which is topologically separated from the rest of the grid. When implementing this correctly into the algorithm of simple coil design the calculated simple loops do not connect the grid covered opening and the main grid. A visualisation of this procedure is shown in Figure 6.3.

To connect the simple loops on the connected door with the other simple loops on the main grid, a twisted pair cable going to the separated grid and back will be installed. Consequently, the currents will cancel and no additional magnetic field will be created, thus the cable can be as long as needed to remove the thermo-house cap. This makes it possible to completely detach the connected door from the thermo-house and therefore making the opening accessible. The main disadvantage of this scheme is the fact that the currents surrounding the opening and the ones on the connected door can not be cancelled, thus resulting in higher currents in the grid and thicker cable bundles around the openings.

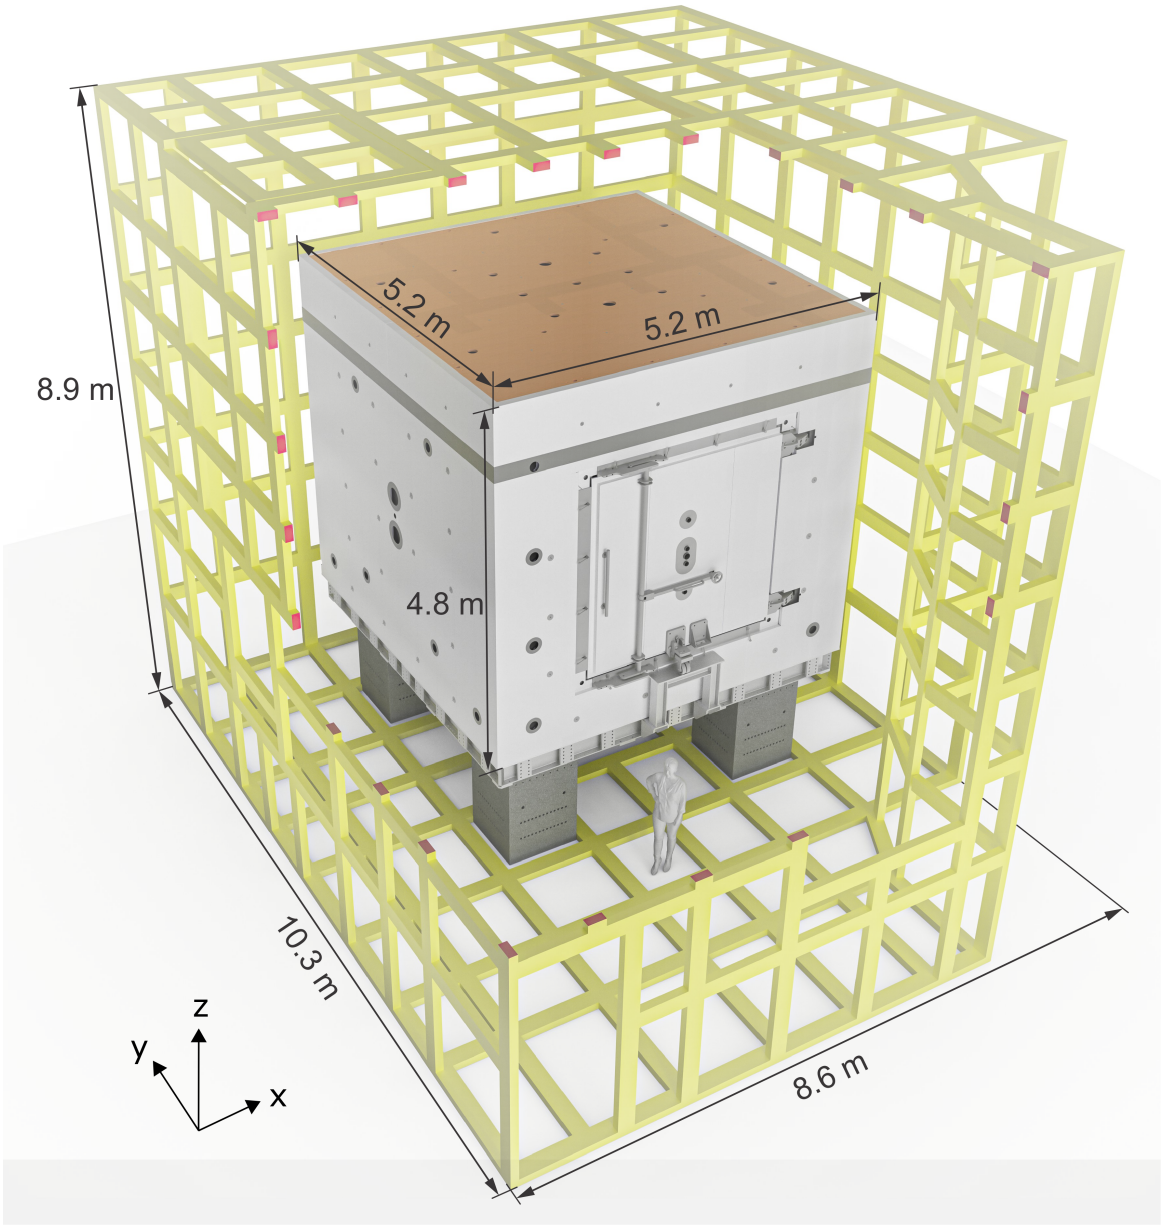

Figure 6.1: CAD model of the AMS grid structure of n2EDM including the MSR and a human for reference. The origin of the coordinate system is in the very centre of the MSR. Courtesy of Georg Bison.

Therefore, the impact of the openings with or without connected door on the homogeneity of the magnetic field was investigated thoroughly. Wherever possible, connected doors were omitted. It was shown that most of the openings can be left open as they are either small enough or far enough away from the MSR. The only opening which has to be covered is the top magnet opening. The connected door structure can be seen in Figure 6.2.

The process of grid design was mainly heuristic, starting off from a first grid version provided by the technician taking into account all interfaces to other parts of the apparatus. In several iterations of field homogeneity calculations, edges were added to the grid. Because a too coarse meshing of the grid led to local magnetic field residuals on the MSR, it was possible to evaluate efficiently where additional connections had to be included. Another important maxim was to use as many through-going connections around the whole thermo-house as possible as this later facilitates the construction.

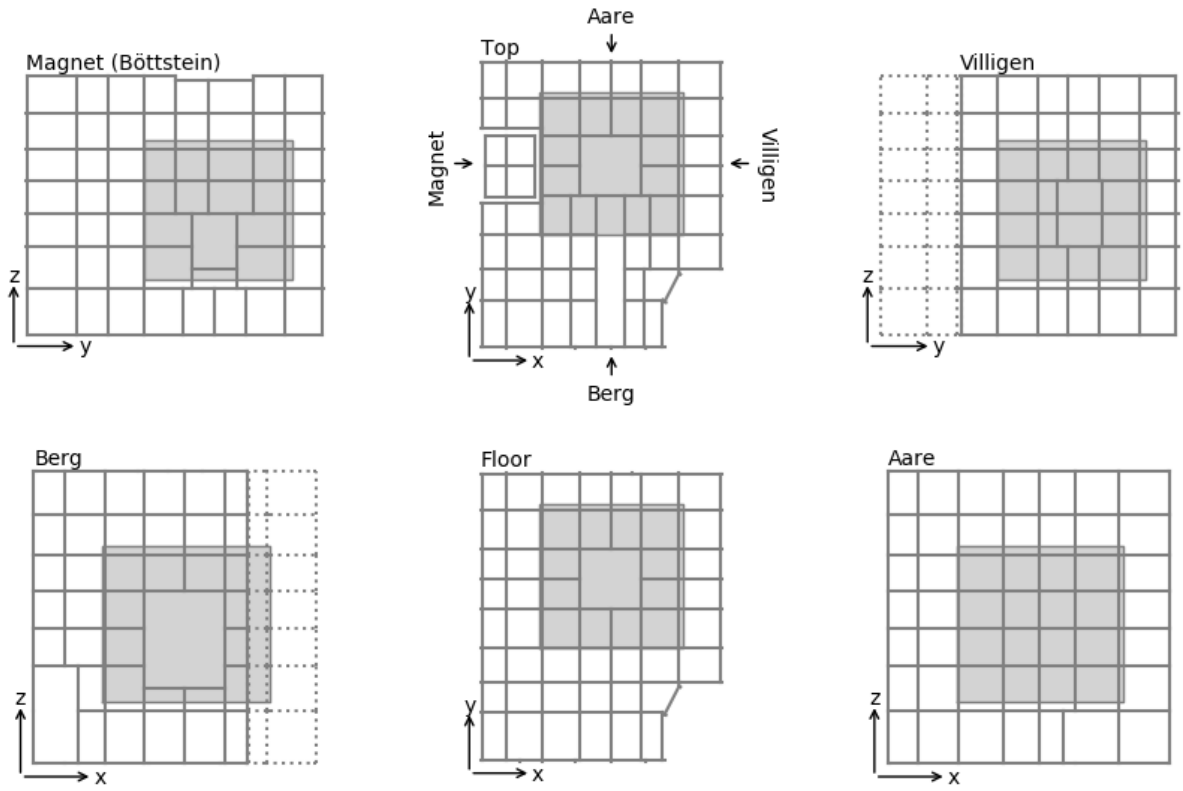

Figure 6.2: This is a dice view plot of the final AMS structures, the main openings and irregularities can be well seen.

For the same reason, only right angles were allowed at the vertices even though any shape would be possible from an algorithmic point of view.

The final grid structure for the n2EDM AMS features the impressive numbers of 778 edges, 473 vertices and 308 tiles.

6.3 Choice of Optimisation Volume

The final input required by the algorithm of simple coil design is the volume on which the magnetic field will be optimised. As the calculation is done numerically, a uniformly distributed throughout the volume chosen set of 1600 points, the points of interest (POI), which sample the volume, will be provided. Intuitively, it is clear that the larger the given volume, the worse the overall performance is. The performance is evaluated by calculating the produced field in validation points on the surface of the MSR, different from the optimisation points and then subtracting the target field in each component separately. It was decided that the minimum and maximum value of these residuals is a good measure for the homogeneity on the shield. For the performance evaluation a set of 9600 randomly placed points is used. It was checked that these points sample the MSR surface densely enough.

Because later on a volume with a good homogeneity will be needed to position the feedback sensors, it was decided that the POI will lie in a shell around the MSR, which is parametrised by an inner and an outer padding thickness of the shell, which corresponds to the inner and outer distance to the MSR. Due to the nature of the Maxwell equations, it is expected that the inner padding thickness has no significant effect on the field quality, as there are no sources enclosed. Therefore, it has been set to 10 cm. However, the larger the outer diameter is, the closer the optimisation region will be to the

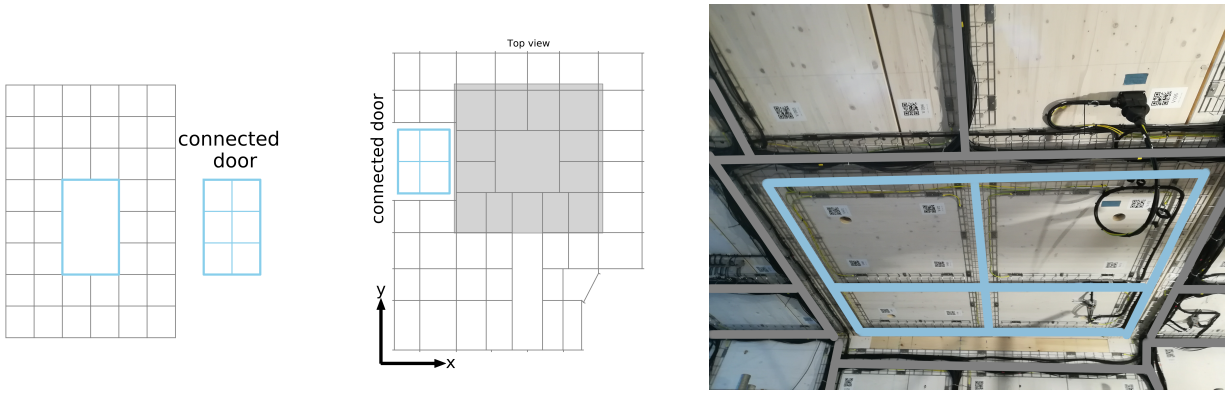

(a) Schematic view of the principle of a connected door.

(b) Location of the connected door in the n2EDM AMS.

(c) Picture of the connected door during the construction process.

Figure 6.3: Visualisation of the working principle of a connected door.

grid and thus more current is needed to maintain the field. This effect can very well be seen in Figure 6.4. Therefore small currents in the grid have to be balanced against homogenous volume where the magnetic sensors can be placed. The optimal outer thickness was taken to be 20 cm.

Another main problem in the optimisation procedure, is the irregular outline of the thermo-house walls. In Figure 6.5 it can be seen that there is a kink in the grid and some of the tiles are far away from the POI's depicted in blue. In consequence, some of those tiles carried very high currents while only improving the field quality on the MSR insignificantly. As mentioned high currents in the grid are a threat to the constructibility of the system. The solution to the problem is the introduction of regularisation to the optimisation problem. Expressed mathematically:

$$\begin{pmatrix} \vec{\mathbf{B}} \\ \vec{\mathbf{0}} \end{pmatrix} = \begin{pmatrix} \mathbf{M} \\ \lambda \cdot \mathbb{1} \end{pmatrix} \cdot \vec{\mathbf{I}}, \tag{6.1}$$

where M is the correlation matrix, $\vec{\mathbf{B}}$ the target magnetic field in POI's, $\vec{\mathbf{I}}$ the currents in the tiles and λ regularisation parameter. It follows from this equation, that additionally to the optimisation of the target magnetic fields there is also a certain penalty attributed to large currents. Consequently, current is traded against field quality. The higher the λ parameter, the more important it is to have small currents and vice versa. A vanishing λ parameter would correspond to the solution without regularisation.

The value of λ is not a priori known and has to be set for each coil and system individually in a process we call λ-optimisation. During this process the optimisation will be conducted for a set of different λ values. The total amount of current in the system and the quality of the magnetic field produced, expressed as minimum and maximum deviation to the target field, will be plotted against the λ values. An example of such a plot is shown in Figure 6.6 for the x-coil.

It can be seen, that minimum/maximum of the deviations start to stagnate below a value of 10^{-8} T/A for λ T/A. At the same time, the mean current in the edges, which directly corresponds to the total current in the system has a minimum at around $\lambda = 5 \cdot 10^{-9}$ T/A. For the x-coil $\lambda = 6 \cdot 10^{-9}$ T/A has been chosen. Also for all the other coils the appropriate value for λ has been chosen manually. In Table 6.1 the chosen λ values for all the coils are given.

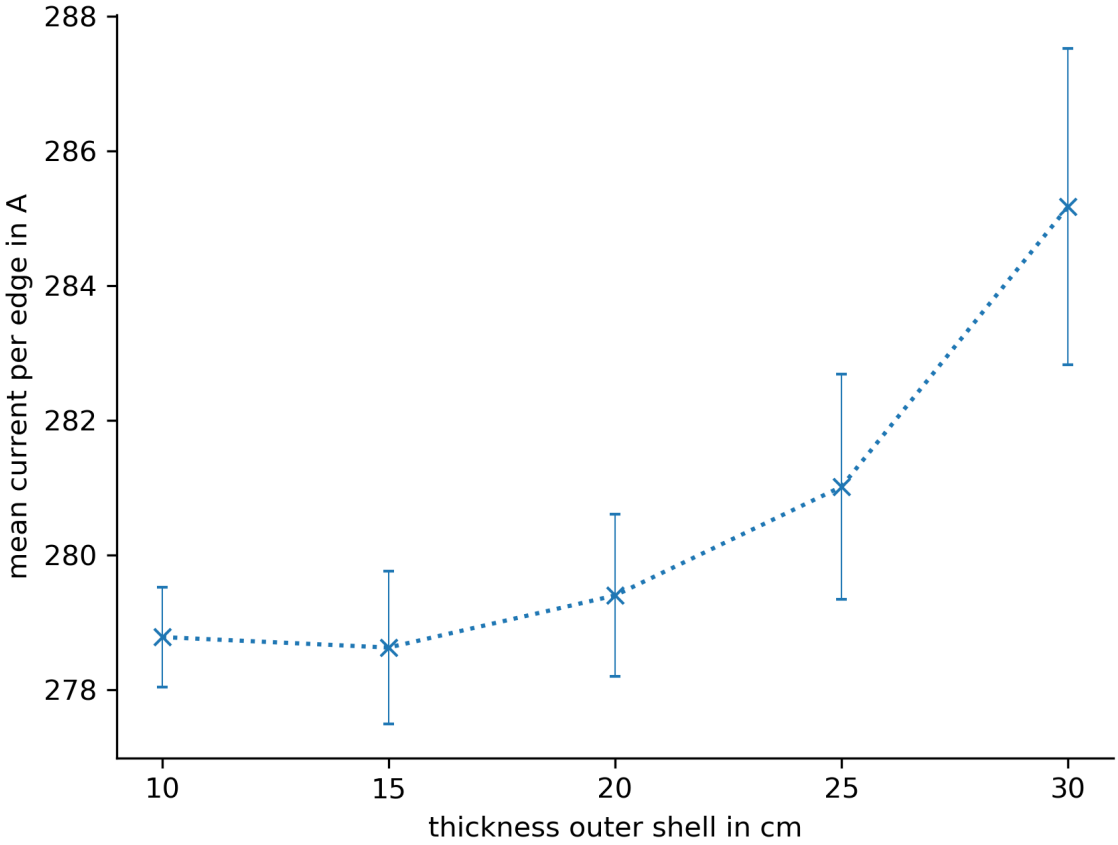

Figure 6.4: The size of the outer thickness of the optimisation shell on the x-axis against the total amount of current in the grid on the y-axis. The error bars correspond to the standard deviation after conducting the calculation ten times with random sets of points for a given value of the outer thickness. The increasing standard deviations are coming from the increasing volume, where the points of interest can be placed.

coil	λ [T/A]
x-coil	$6 \cdot 10^{-9}$
y-coil	$4 \cdot 10^{-9}$
z-coil	$4 \cdot 10^{-9}$
1st gradient	$1 \cdot 10^{-8}$
2nd gradient	$1 \cdot 10^{-8}$
3rd gradient	$8 \cdot 10^{-9}$
4th gradient	$8 \cdot 10^{-9}$
5th gradient	$1 \cdot 10^{-8}$

Table 6.1: List of the λ value for all coils constituting the n2EDM AMS. Where the linear gradients 1-5 correspond to the Cartesian harmonica polynomials 4-8 in Table 3.1.

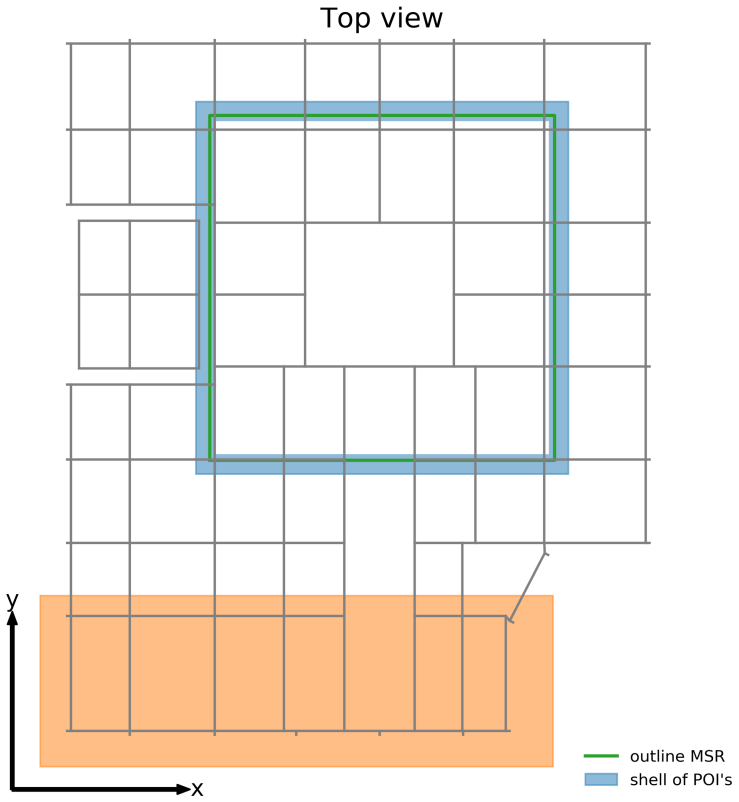

Figure 6.5: Schematic view of the outline of the AMS including the MSR in green. Drawn in light blue is a possible set of POI's. Tiles which do not have significant impact on the field on the MSR are shaded in orange.

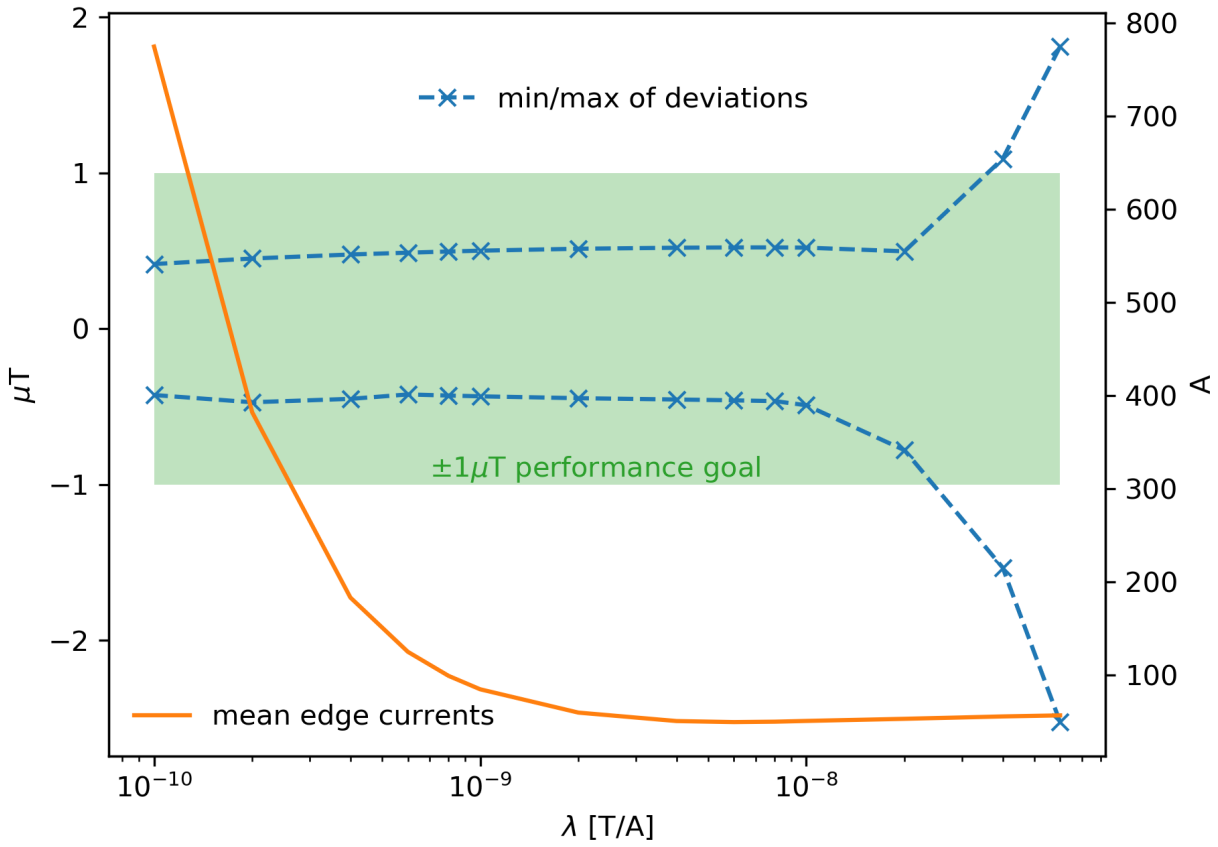

Figure 6.6: The minimal and maximal deviation of the target field on the MSR is plotted against the lambda parameter for the optimisation procedure of the x-coil. The $\pm 1\,\mu$T performance goal is indicated by the green box. On the second axis the mean currents in the edges of the grid are plotted for the set of lambda parameters.

6.4 Additional Adjustments to the Algorithm

One of the heuristic parts of the algorithm used is the search for simple loops in the grid of currents. To facilitate this procedure, a precision parameter has been introduced. This parameter specifies how much the currents in two edges need to coincide in order to be able to connect them with a simple loop. This significantly speeds up the extraction of the simple loops, but also causes currents that will be left in the grid after all simple loops are found. Additionally, the algorithm stops with looking for simple loops if the currents in all edges are less than half the minimal current. This is for the simple reason that this is the smallest current which will be rounded up to the minimal current. Indeed does the minimal current have an influence on the achieved homogeneity. As it also increases the winding work, it is subject to an optimisation which is presented in Chapter 7. For the above mentioned argumentation the field produced after simplification is of lower quality than when taking the prior exact values. Additionally, all of the eight coils consist of more than a hundred simple loops, where it is not possible to understand how much each simple loop improves the quality. Therefore a procedure to re-optimise the currents and reduce the number of simple loops was introduced for the sake of simplifying the construction.

All the simple loops were ordered by importance, which is defined as the amount of total field they produce on the MSR. Afterwards, for different numbers of most important simple loops, the current was re-fitted to achieve best performance with the reduced number of simple loops. These results were then compared to each other. In Figure 6.7 such a comparison is shown for the 4th gradient coil, which originally consisted of 123 simple loops.

There is no strict cut-off for the number of simple loops which must be chosen, and the optimisation for each coil is done manually. By looking at the plot for the 4th gradient coil, it can be seen that the maximum field deviation on the MSR significantly differs between using the first 60 or 70 simple loops, but almost no difference in field quality can be seen between 70 and 80 simple loops. Therefore, it has been decided to wind only the most important 70 simple loops for the 4th gradient coil, this kind of optimisation reduced the number of simple loops by more than 40%.

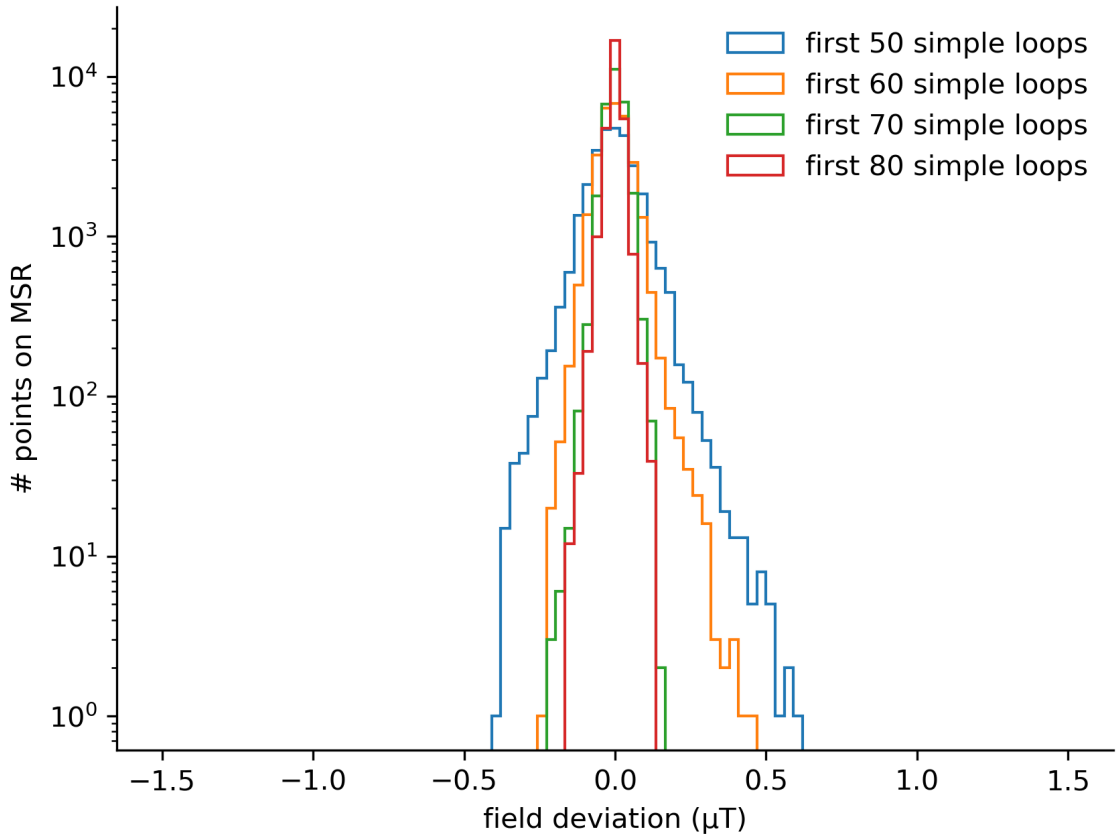

Figure 6.7: Histogram of the residual fields on the MSR for different numbers of most important simple loops used in the 4th gradient coil of the system.

6.5 Static Performance

Up to now, the performance of the coils was only treated separately. However, in order to check that the specifications are met, the performance of the system as a whole needs to be examined. This means that the performance of all the coils needs to be evaluated at the same time. As the paths of the simple loops are bound to a grid with finite tile size, it is evident that the goal field can only be approximated and never be perfectly achieved. Imagining a system of perfect coils creating homogeneous fields in x-,y- and z-direction, the correlation matrix would look the following:

$$M = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \tag{6.2}$$

For the above mentioned reasons, this ideal correlation matrix, where turning on all the currents to their maximum current value produces perfectly homogeneous fields in all three spatial directions, is not achievable. The correlation matrix of a real world system would rather look the following:

$$M = \begin{pmatrix} 0.9 & 0.02 & 0.05 \\ 0.01 & 1.1 & 0.02 \\ 0.1 & 0.05 & 0.9 \end{pmatrix} \tag{6.3}$$

Where the on-diagonal contributions are close to one but not exactly, and the off-diagonal contributions are close to zero. This means for example considering the x-coil, that turning on the full current would

produce only 90% of the target field in x-direction, but also produces small fields in y- and z-direction. This non-uniformity is a part of what was classified earlier on as the deviations of the field from the target field, and might be seen as a problem at first glance. However when the system will act as a whole in feedback mode, then the y- and z-contributions of the x-coil will be compensated by the other two coils. This mechanism of the other coils cancelling the residuals of one coil will be referred to as cross compensation.

The failure of the x-coil to reach 100% of the target field in full operation only implies that the system needs to be able to drive more current in order to achieve the full target fields. Thus, it can be concluded that non-perfection of the correlation matrix is not a problem insofar as it can be diagonalised, and will be automatically taken care of by the system in operation. Therefore, when presenting the static performance of the system, cross compensation effects are included.

Moreover, it is clear that the performance of the coils is evaluated on the exact same points on the MSR. Therefore the residuals of all the eight coils in one point will be added. This is what naturally happens when operating the system and in most cases this will reduce the residuals on the MSR as the coils do not have coinciding spots of maximum deviation. Rather will the residuals compensate each other over the area of the MSR.

Last, to estimate the static performance of the system it is important to know in which configuration the system is running. The surrounding fields that need to be compensated have to be known. Up to now, the residuals were considered at full current, and as the system's residuals scale linearly with the amount of current this is the worst case assumption. In the following the residuals will be given in COMET, SULTAN and full current mode, where COMET and SULTAN represent the realistic state in which the AMS will be most of the time. The residuals in these two modes are typically smaller, as the current in the coils will not be ramped to their maximum.

In Figure 6.8 the theoretical residuals expected on the surface of the MSR are shown for the full field configuration. It can be clearly seen that the surfaces which are closest to the AMS, the Aare and the Villigen side, encounter the highest residuals. In the case of the Aare side even the imprint of the grid structure on the residuals can be seen. This suggests, that the system exploits its potential and could only be improved by a finer meshing or a larger distance to the MSR. The summary of the expected theoretical static performance of the AMS system is given in Table 6.2 also including the most relevant COMET and SULTAN configuration.

mode	minimum [μT]	maximum [μT]
full current	-1.5	1.1
COMET	-0.6	0.9
SULTAN	-0.4	0.7

Table 6.2: Static performance of the AMS system in full current, COMET and SULTAN mode as per design, specified as minimum and maximum of the residuals on the MSR surface.

After the above explanation the consequent question to ask is: Why are there still residuals on the MSR left, if using cross compensation? The answer is fairly simple: The problem is that any sort of imperfection does not only produce homogeneous and 1st order gradient fields, but also higher order gradients. These are the fields which can not be compensated with the coils of the system. In the next section we are going to investigate the possible reasons for those higher order imperfections when considering the real construction of the system.

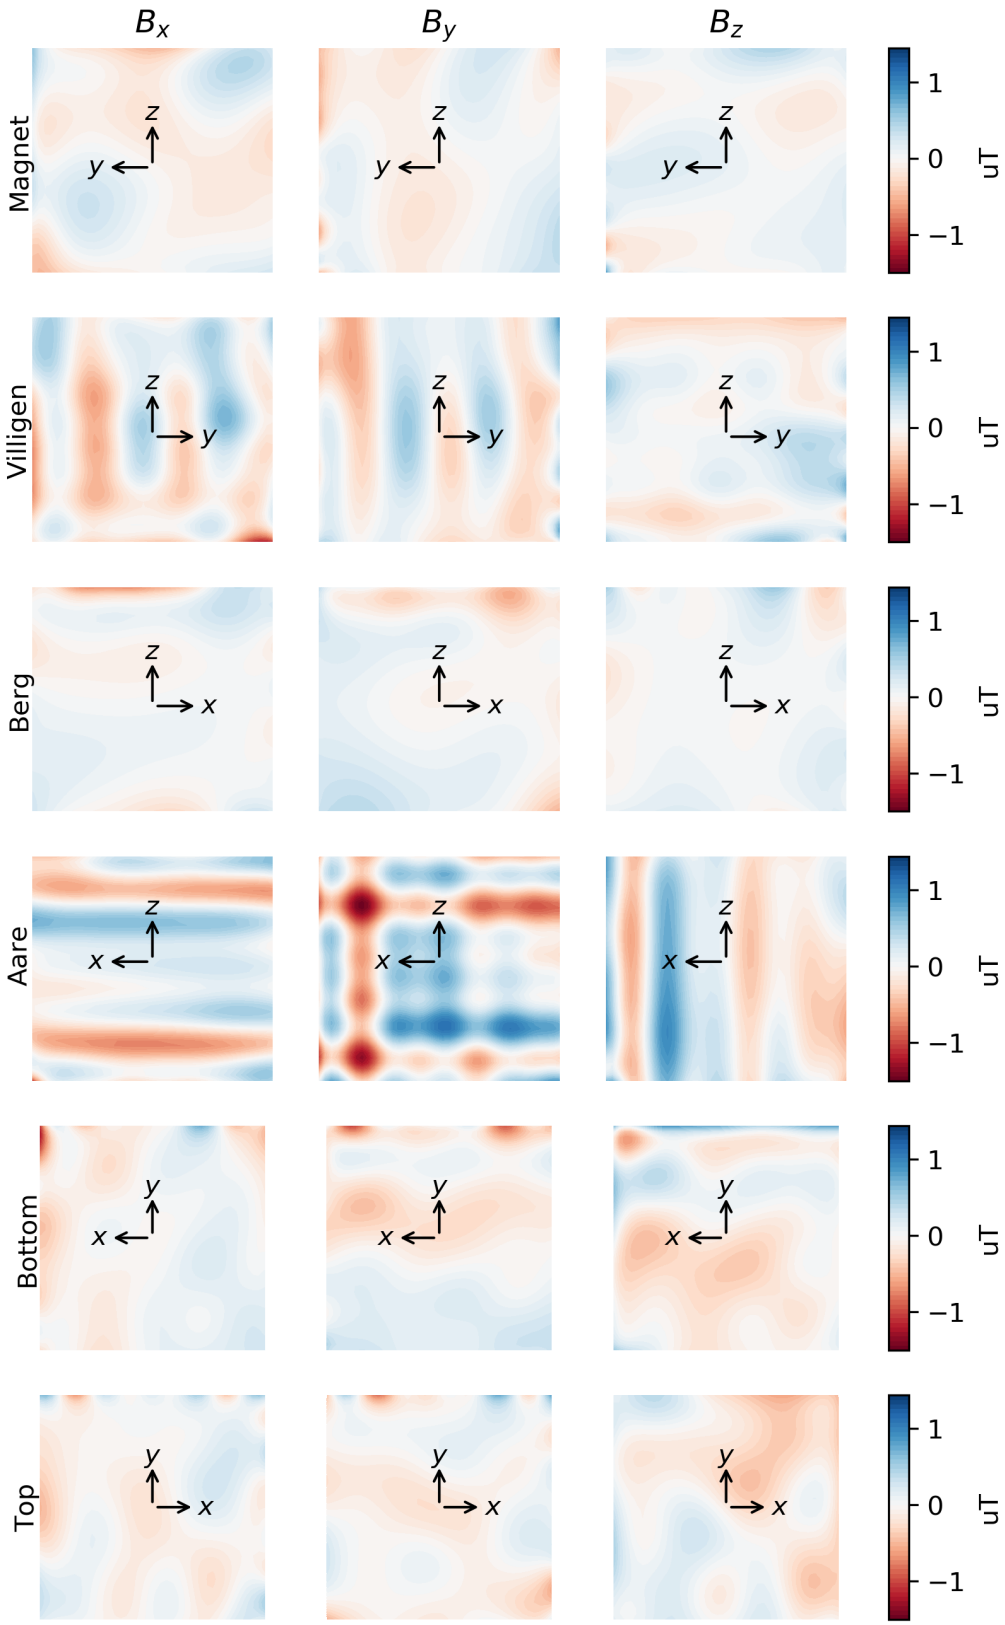

Figure 6.8: Plot showing the residuals on the MSR for full field configuration. The columns corresponds to the x-, y- and z-component of the magnetic field respectively. The MSR surface is looked at from the outside, the coordinate system is indicating the orientation.

6.6 Deviations in the Real World Implementation

For the installation many wires taking the same path e.g. simple loop will be merged to a bundle. When doing an estimate of the bundle cross section after the simplification procedure, it is clear that the bundles will become fairly thick. This has effects on the accuracy of the winding process, and directly implies that not all of the wires will be following the theoretically perfect path used in the calculations. Hence, this might produce deviations from the calculated ideal fields. In the next two paragraphs we are going to quantify the effects which those unavoidable imperfections cause.

6.6.1 Bending

In Figure 6.9 a thick wire bundle running around a corner is shown. The main issue can be spotted on the inner side of the bundle, where the wires bend much earlier around the corner. Ideally, the bundles take a path in a quarter circle around the corner, where its radius is a measure for the deviation from the sharp edge.

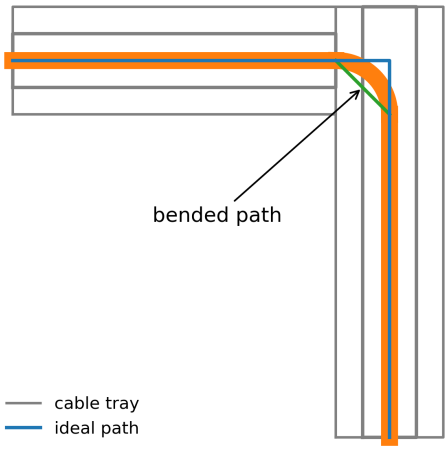

(a) Sketch of the situation when a thick wire bundle runs around a corner.

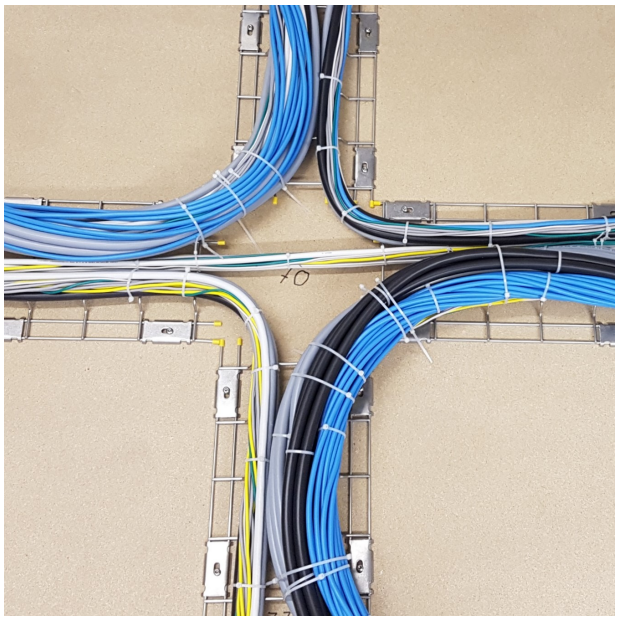

(b) Prototype construction of implementing the corner with the highest number of wires in the full system running through.

Figure 6.9: Illustration and prototyping of the bending problem.

This bending radius can not be calculated theoretically, as it depends a lot on the respective mounting quality of the bundles. Nevertheless, the junction containing the most amount of cables was constructed beforehand in the lab. This resulted in an estimate of the worst bending radii being between 10 cm and 15 cm. These numbers give a first idea of realistic inputs for the calculations.

Instead of a quarter circle a straight line has been used in the calculation procedure, as it was shown that the two calculations result in a field difference of less than 0.01 µT. There are two versions of the algorithm, one where the bending radius was fixed to a certain value in every corner and a second one where the bending radius depended on the amount of current in the bundle. Clearly the first approach is a worst case scenario. Therefore the results of the former calculation are presented here. Fixed bending radii between 4 cm and 15 cm have been evaluated.

For the evaluation, cross compensation and scaling to the magnet configurations was used as explained in the preceding section. In Figure 6.10 it can be seen that the residuals on the MSR do not increase noticeably with larger bending radii. It was verified that only for bending radii of several dozens of cm an effect can be seen. The reason for this behaviour is that changing the bending radius

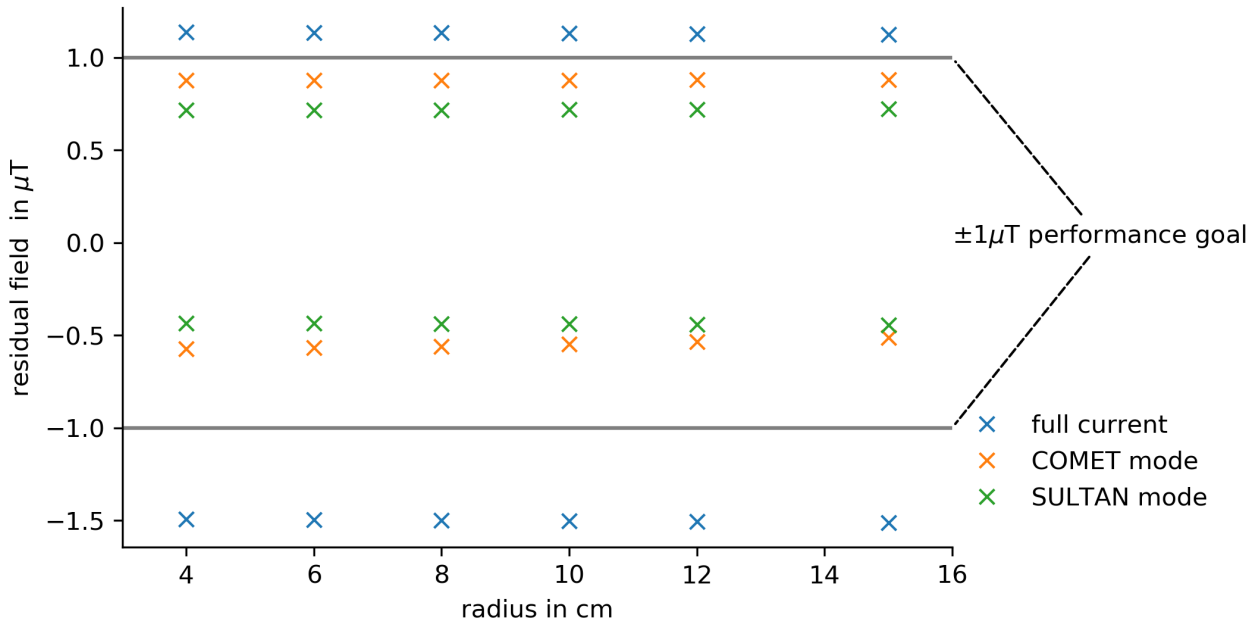

Figure 6.10: Results of the bending calculations using the same fixed bending radius in all corners.

is a smooth transformation, which does not influence the ratios between the magnetic fields produced by each coil. This implies that more current is needed to be fed into the system in order to achieve full field, but the distortion is not significant.

In conclusion, it can be said that in our case the bending radius caused by the thickness of the wire bundle does not present a threat to the performance.

6.6.2 Displacement

The ideal path used in the optimisation procedure is running exactly in the middle of the cable trays. When having several hundreds of cables running on the same cable tray, it is evident that most of them will be displaced in the plane normal to the cable tray. This situation is depicted in Figure 6.11.

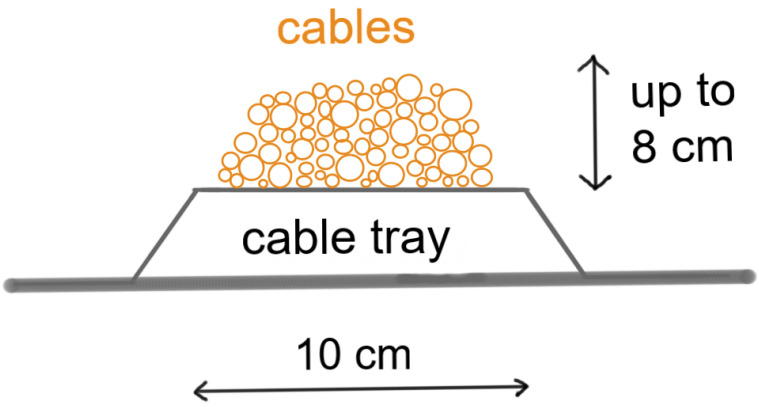

Figure 6.11: Sketch of the situation with a lot of cables running on the same tray, making it impossible for all the wires to take the ideal path.

The largest calculated cross section on one cable tray was estimated to be approximately $70\,\mathrm{cm}^2$. Therefore we conducted calculations, where the simple loops were randomly shifted. For the horizontal displacement on the cable tray a random number from a flat random distribution from $-5\,\mathrm{cm}$ to $5\,\mathrm{cm}$ was drawn, and from $-4\,\mathrm{cm}$ to $4\,\mathrm{cm}$ for the vertical displacement. This made an area of $80\,\mathrm{cm}^2$ where the simple loops could be placed.

Because the shifts were drawn at random, several runs had to be performed and the mean was calculated. The results are presented in Table 6.3 below. Again cross compensation was taken into account.

mode	minimum [µT]	maximum [µT]	σ min/max[µT]
full current	-1.7	1.5	0.25/0.17
COMET	-0.8	1.0	0.11/0.17
SULTAN	-0.6	0.9	0.10/0.14

Table 6.3: Results of the displacement calculations for full current, COMET and SULTAN mode, after averaging 10 runs.

It can be seen that the residuals increase clearly compared to the theoretical static performance. However, they are still within the performance goal. Contrarily to the bending, the displacement does change the ratios between the magnetic fields produced by the simple loops and therefore also induces 2nd order residuals.

Again, this calculation is clearly a worst case, as the $80\,\mathrm{cm}^2$ area is only occupied in the case of a large amount of cables running on the same cable tray. If less cables run on the same cable tray they will follow the ideal path more closely.

Nevertheless, this is an important result for the construction of the system, as it implies that it does not really matter where on the cable trays the bundles need to be mounted. This gives a certain freedom to the people installing the system and significantly facilitates their work.

6.7 Conclusion

Using the method of simple coil design an active magnetic shielding for n2EDM was designed considering spatial constraints as well as general feasibility. It was shown that the residuals on the MSR in the relevant SULTAN and COMET magnetic mode do not exceed $\pm 1\,\mu\mathrm{T}$. As a summary the final theoretical residuals for the full current magnetic field is shown on a cut through the $z = 0$ plane of the experimental area. Additionally the goal performance is met as well when construction imperfections are considered. In the next chapter the construction process is discussed in details.

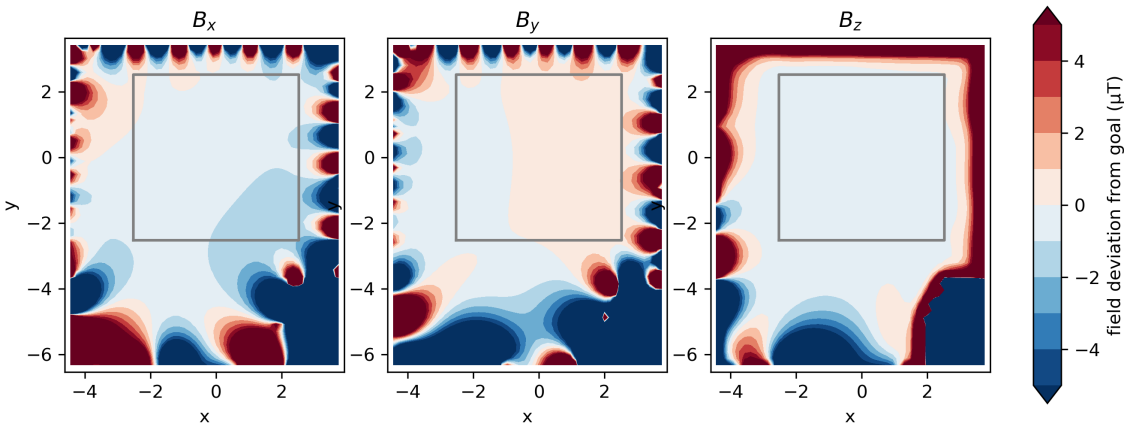

Figure 6.12: Cut through the experimental area at $z = 0$, where the final theoretical residuals are plotted for the full field in all coils including cross compensation. Due to the design process, the residuals are minimised around the MSR which is drawn in grey.

Chapter 7

Construction

In the next section we are going to move on from theoretical questions to questions about how to bring such a system to life. We are translating the simple loops into a system that is eventually going to compensate magnetic fields. The first question that needs to be answered is the choice of currents.

7.1 Choice of Currents

After the theoretical design section the active magnetic shielding consists of a grid and simple loops with corresponding currents. First it is investigated what needs to be considered in order to power the system with current.

When looking at a specific coil we are given between 50 - 80 simple loops which need to be operated in series in order to achieve the target fields. Each of these simple loops have a different nominal current and needless to say it is not feasible to employ one power supply for each simple loop. It is also not practical to use a minimal current and wind each simple loop up to 100 times.

The solution lies somewhere in between, by decomposing the total current in each simple loop into a set of elementary currents. This method was already employed in the construction of the prototype at ETH. In the case of the n2EDM AMS it was decided to use three elementary currents by minimising both number of power supplies and winding effort. Those elementary currents can be different for each simple loop. In total 24 (8 coils x 3 elementary currents) power supplies are needed to power the AMS.

As an example we look at a certain simple loop, which needs a current of 93 A and the set of elementary currents is [20 A, 5 A, 1 A]. Then, the 93 A are decomposed in the following way:

$$93\,\mathrm{A} = 4 \times 20\,\mathrm{A} + 2 \times 5\,\mathrm{A} + 3 \times 1\,\mathrm{A} \tag{7.1}$$

This means that the simple loop will need to be wound four times with the 20 A-wire, twice with the 5 A-wire and three times with the 1 A-wire.

The set of elementary currents was subject to an optimisation due to two main reasons. First, the choice of the minimal elementary current influences the homogeneity at the maximal field. The smaller the minimal elementary current, the better the quality of the achieved field. Second, if the minimal elementary current is small, there is a tendency for more windings and thus less efficiency in the upcoming winding process. It was chosen that 1 A will be used as minimal elementary current, as the field distortions are still insignificant with this value.

Second, different ratios between the three elementary currents lead to different number of windings and thus weight. The optimisation of the elementary currents uses the weight of the system as a figure of merit as it also includes the number of windings.

The above mentioned optimisations were conducted separately for every coil. Apart from the minimal elementary current, the set of elementary currents varies from coil to coil. In the summary Table 7.1 the elementary currents and resulting coil weights are listed.

coil	E1	E2	E3	weight
x	15 A	5 A	1 A	248 kg
y	15 A	5 A	1 A	327 kg
z	15 A	5 A	1 A	306 kg
1G	10 A	3 A	1 A	65 kg
2G	12 A	5 A	1 A	80 kg
3G	15 A	4 A	1 A	109 kg
4G	8 A	3 A	1 A	70 kg
5G	12 A	5 A	1 A	112 kg
				1.3 t

Table 7.1: Table summarising the optimal set of elementary currents (E1-E3) for each coil. The weight is only considering pure copper, not taking into account rounding to commercially available sizes and insulation. 1G-5G are the working names of the gradient coils during construction and correspond to what was previously called 1st-5th linear gradient.

The calculated weight is a minimal estimate, as only the copper weight for the minimal wire diameter was taken into account in order to fulfil Ohm's law while staying below 50 V. The voltage of 50 V was chosen due to the availability of power supplies, the total power dissipation and most importantly due to safety regulations at PSI. The detailed lists of expected resistances and wire diameters for each coil can be found in Appendix B. These numbers significantly influence the final weight. The total amount of cable needed sums up to at least 51 km. The cable tray holding the highest number of cables will support 227 single wires.

Another important aspect which had to be investigated prior to the installation was the expected heat load. Summing up the thermal power of all the eight coils yields 7.25 kW, which is half of the power of the air conditioning. Normally the coils will run in either SULTAN and COMET mode, which reduces the thermal power output to 2.9 kW and 3.2 kW respectively.

7.2 Installation Procedure

In the first step, the cable holding structure was mounted on the inside of the thermo-house. The positions were measured with a precision of 1 cm with the help of a laser meter. The cable trays were tested for their magnetic properties beforehand and did not show any significant magnetisation.

The cable holding structure then had to be grounded for electrical safety. The grounding was done in a way, that there were no closed loops, such that eddy currents could flow in the trays.

Four professional electricians were working on the installation of the AMS for several months. For all of the approximately 500 simple loops that had to be installed, the procedure always looked the same and is described in the following with the help of the photographs in Figure 7.1.

The installation process starts with the bundle preparation. As the current in all simple loops is decomposed into the three elementary currents as discussed previously, typically a bundle consists of several windings of different circuits. The length of the simple loops was calculated beforehand and the right amount of cable was cut as seen in Figure 7.1a. To prevent wasting too much cable a cable meter as depicted in Figure 7.1b was used.

In Figure 7.1c it can be seen that each individual cable was labelled with a barcode including a human readable identifier of the cable. The barcodes are used for the mobile winding application as described in Appendix C to make sure that the correct wire/bundle is used in the wiring process. Afterwards the bundles were stored tidily until later use as seen in Figure 7.1d.

The next step is then to install the cable along the path of the simple loop. For this purpose every vertex was labelled with a number and a QR code. The electricians were provided with a wiring plan, where the path which the bundle had to take was indicated together with the vertex numbers the

bundle will pass. This part is depicted in Figure 7.1e. The bundles were attached to the cable trays with cable ties.

This part of the installation was taking most of the time, as the simple loops can include up to several dozens of vertices and the bundles can be very heavy. Even the use of cable ties was optimised in order to speed up the installation. For the z-coil which is mainly running in the xy-planes wooden hooks were installed. This increased the installation speed as in the horizontal sections the bundles did not have to be fastened with a cable tie but were hold together by the hook.

It is easily imaginable that wrong paths could be taken. To fix these mistakes after the installation would be almost impossible, as there will be several layers of bundles on top. Therefore a smartphone application was programmed to prohibit such mistakes. As depicted in Figure 7.1f, with the smartphone application, one can choose which coil needs to be checked by scanning the barcodes on the bundles. The person checking the path follows the path of the bundle and scans the QR codes on every vertex the bundle passes, as seen in Figure 7.1g. If there is a deviation from the foreseen path, the person is informed immediately and can correct the mistake.

After a simple loop is installed correctly, the windings belonging to the same elementary current need to be crimped together, as in Figure 7.1h. In Figure 7.1i the start and end of the circuit for a simple loop are connected to terminals on DIN rails.

As a simple loop is a closed path, the start and end position can be chosen freely. Because in the end the simple loops have to be connected by so called interconnection wires, the position of the DIN rails was optimised such that the total path connecting all the simple loops of a coil was minimised. This is important in order to minimise the resistance of the coil, which directly relates to the wire thickness and thus weight for a fixed voltage.

In a last step, shown in Figure 7.1j, the resistance for each simple loop and circuit is measured and compared to the calculations. During the crimping it can very easily happen, that instead of connecting the windings, one of them is crimped to itself. With a resistance test such mistakes can be identified.

Because the coils are designed to produce fields in all three spatial directions, it is clear that there will be many crossings between the simple loops. If too many crossings occur by turn at the same place, then the wires will stack up which can induce errors in the magnetic fields due to the position deviations. Thus the order in which the simple loops within a coil are installed was optimised, such that crossings are minimised.

The only check which makes to almost 100% sure that no other mistakes were made, is to power the coils and measure the current in every edge with an Ampère-meter in a selection of 10-20 edges of the simple loop. The measured currents can then be compared to the model. This test was conducted at the end of the construction and mainly discovered simple loops running in the wrong direction.

During the writing of this thesis the AMS was still under construction. A picture of it during installation can be seen in Figure 7.2c.

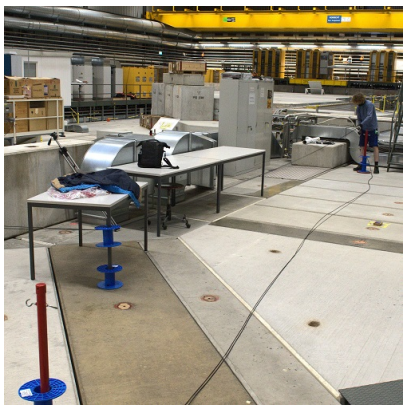

(a)

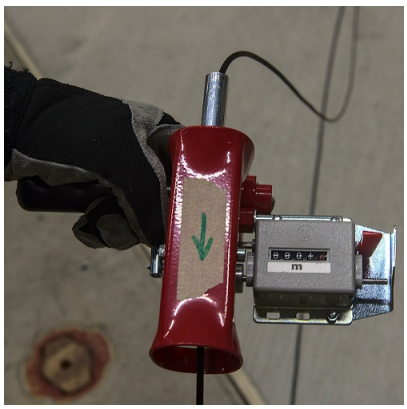

(b)

(c)

(d)

(e)

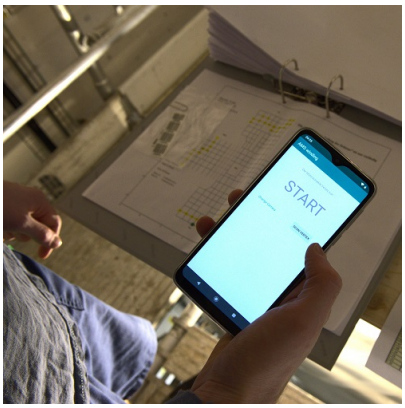

(f)

(g)

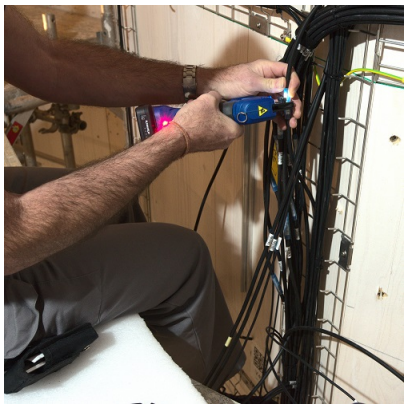

(h)

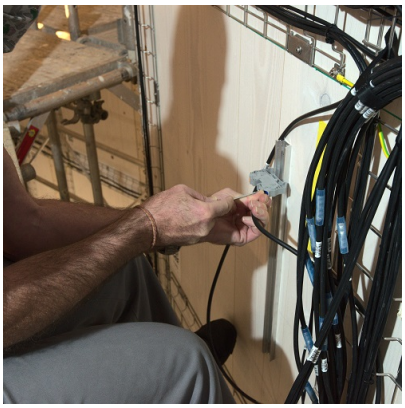

(i)

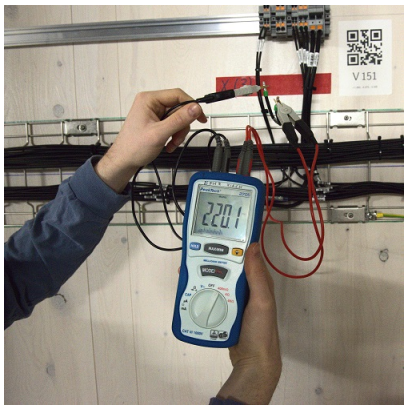

(j)

Figure 7.1: Photographic representation of the bundle installation procedure. Follow description in the text.

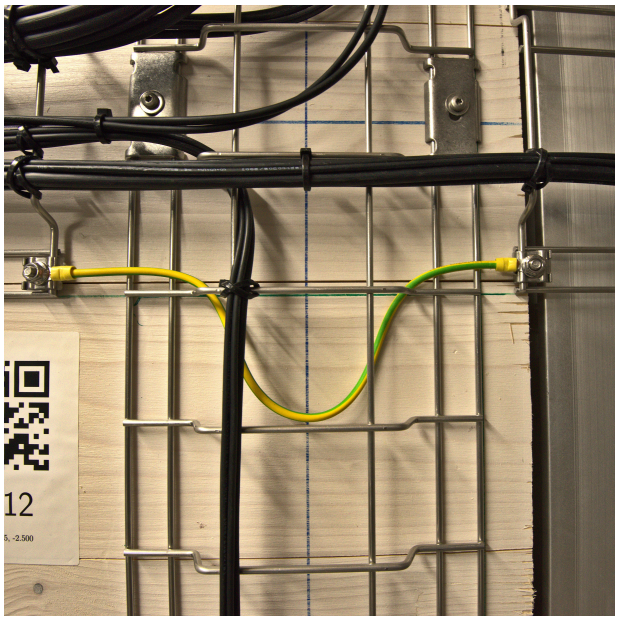

(a) Close-up of the grounding connecting two horizontal cable trays.

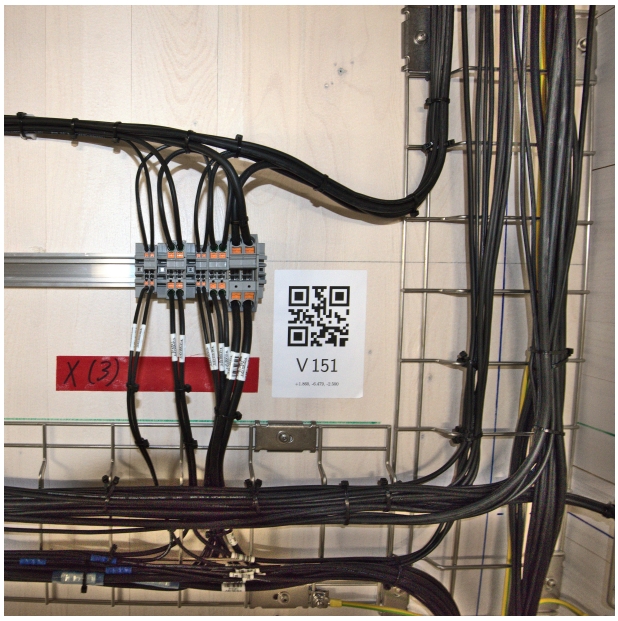

(b) Close-up of a DIN rail, where 5 circuits from two different simple loops are terminated.

(c) Photo of the AMS while still under construction, taken from top of the MSR. At this point in time less than half of the coils were installed.

Figure 7.2: Impressions of the AMS under construction.

Chapter 8

Performance

During the last year of the Doctorate underlying this PhD thesis the coil system was under construction. Even though solid procedures were put in place to ensure the correctness of the system, the final validation must be the measurement of the field produced around the MSR. In the following section the static performance measurements are presented.

Despite the fact that the feedback system has not yet been implemented, tentative positions for the positioning of the feedback sensor are provided to ensure best possible performance. Those positions can serve as a starting point in the upcoming optimisation process once the system is running in feedback mode.

8.1 Validation of the Fields

The main problem for the validation of the fields produced by the AMS coils is the fact that the AMS was installed after the magnetically shielded room which can not be removed. Therefore the fields produced by the AMS can not be measured separately and the effect of the MSR must be taken into account. To calculate how the magnetic field lines are bent in presence of the MSR a COMSOL simulation similar to the one for the prototype (described in Section 4.6.1) was set up. This way the measured field can be validated against the simulated field.

Because the AMS is a very complicated system with around 500 different wire paths, it was not feasible to implement the complex coil structure in a timely manner. Therefore it was decided to use background fields and investigate the changes the MSR forces on the background fields. For this reason the "Magnetic Fields, no Currents" interface of COMSOL was used. The MSR model featured in the simulation is 10 cm thick and has a relative magnetic permeability of 10'000. This simplification had to be employed, as the real MSR is a complex structure, consisting of several layers of mu-metal with different relative magnetic permeabilities. It has been shown that for magnetic permeabilities above 2'000 the field outside the MSR did not significantly change any more.

Apart from validating the measured fields, the simulation aims at predicting the best position of the magnetic feedback sensors. Therefore, it is important that the fields close to the MSR are well described, as this is the most homogeneous region. Here, the sensors are planned to be installed. This allows for a simplistic design of the simulation, where the calculated imperfect fields produced by the MSR can be replaced by the continuous perfect fields directly representing the Cartesian harmonic functions. This procedure has the advantage, that every point of magnetic field information needed in the simulation can be calculated analytically.

The downside of the above mentioned simplification is the fact that the field had to be validated in positions were the AMS is expected to produce fields not deviating from the target fields. These positions can be mainly found close to the MSR on the Villigen and Berg side of the thermo-house. These spots are also farthest away from irregularities of the grid, such as big openings and the kink.

In Figure 8.1 the magnetic field components produced by the y-coil in an xy-plane at $z = 0$, which correspond to the centre of the MSR, are shown. The z-component of the magnetic field remains at

zero, whereas the B_x and B_y components are bent. Because the field produced by the y-coil is mainly pointing in the y-direction in front of the MSR, the MSR will bend those lines such that they are guided through the mu-metal. Therefore the B_y components are reduced along the sides of the MSR which are parallel to the y-direction. However, because of the bending of the field lines, magnetic field components in B_x are acquired.

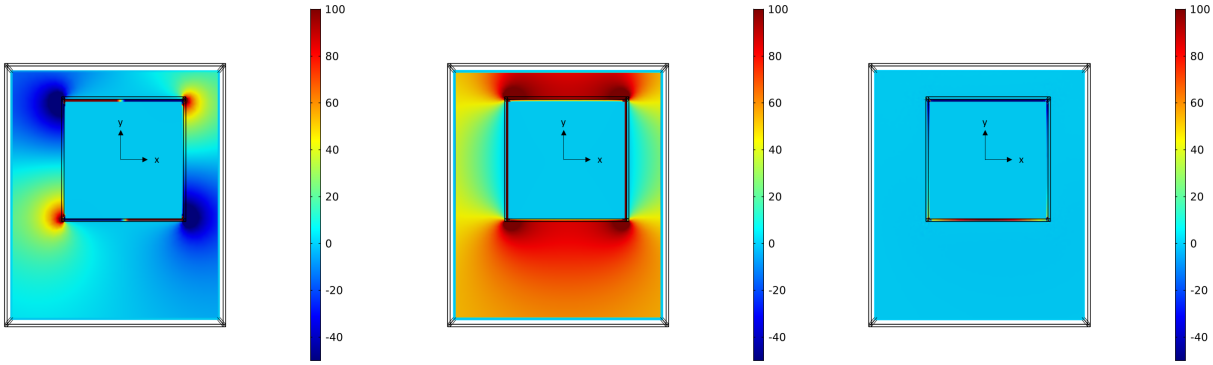

(a) B_x (b) B_y (c) B_z

Figure 8.1: COMSOL simulation of the field produced by the y-coil at $z = 0$.

Conducting this performance measurement was mainly a question of the available space. Contrary to the mapping campaign prior to the start of the n2EDM construction, the area was at this point already to crowded to perform a mapping. Thus, the important positions could not be reached with a movable mapping structure. Therefore the measurements were taken with Stefan Meyer flux gates attached to a 3.6 m long profile. Ten positions were marked along the profile, with a distance of 40 cm in between. The profile was installed at two positions to map the x- and y-direction respectively, the positions are sketched in Figure 8.2. The exact global positions were determined relative to the thermo-house walls with the help of a laser distance meter. A spirit level was used to align the profiles parallel to the ground.

Each of the eight coils was measured separately in order to minimise the impact of background fields on the data, the coils were operated with a square pulse of a few seconds length such that the amplitude could be extracted. The measurement for every position on the item profile took approximately 8 s. Because of the range of the used magnetic sensors, the coils were only operated at half of the maximum current.

In Figure 8.3 such a measurement is plotted on top of the simulation. This example features the y-coil of which the COMSOL fields are shown in Figure 8.1. The measurement of all three components was taken parallel to the y-axis between the MSR and the Villigen wall. It can clearly be seen that the simulation describes the measurements very well for all three spatial directions. This was validated for all measurement directions and coils of the AMS.

In order to quantitatively compare the measurements to the simulation the deviation to the simulation was calculated, for every measurement point, resulting in an RMS estimate for all coils. These numbers are summarised in Table 8.1. Even though the deviations are higher for the homogeneous x-, y- and z-coils, it can be stated that the discrepancies are mainly caused by shifts between the measurements and the simulations, due to imprecise position measurements. Because all the measurements agree very well with the simulation regarding the shape, the simulation could be validated.

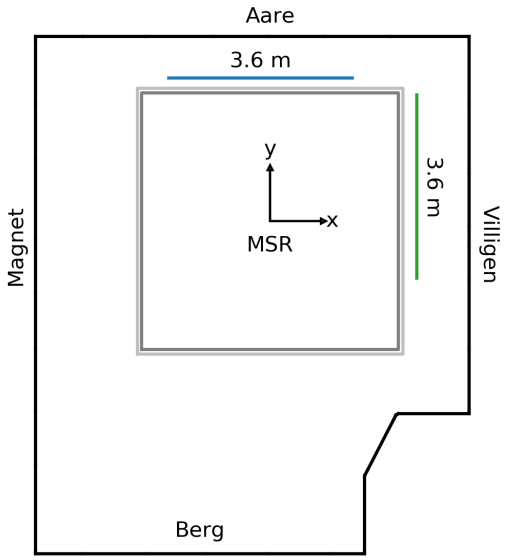

Figure 8.2: Position of the item profile for the validation measurement in x-direction (blue) and y-direction (green), relative to the MSR and AMS outline. The height above ground was approximately 3.2 m, corresponding to approximately $z = -1$ m.

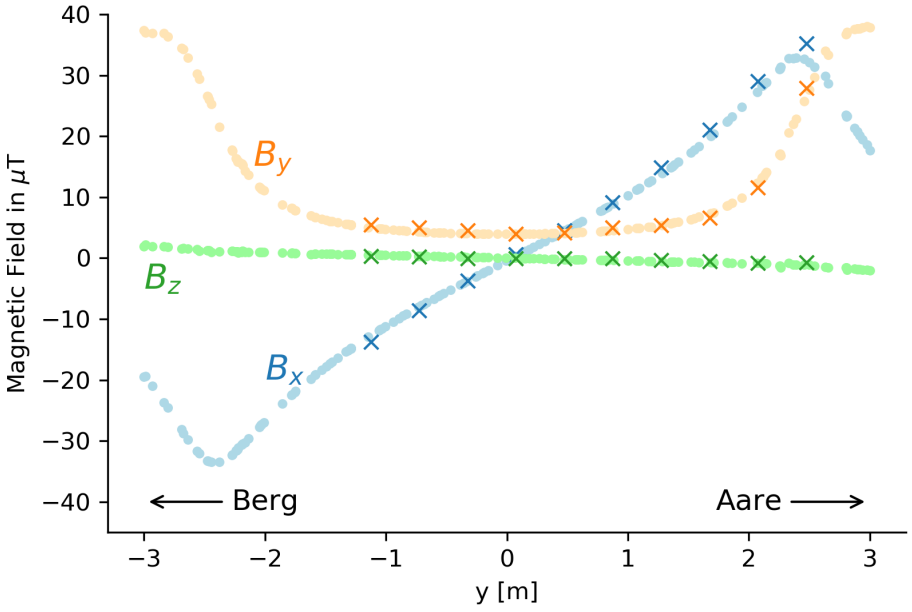

Figure 8.3: This plot features the magnetic field values measured on the line in y-direction with the y-coil at 50% power. Crosses represent the measurement points, whereas the pale dotted line represents the simulation values.

	x-direction			y-direction		
	σ_{B_x}[µT]	σ_{B_y}[µT]	σ_{B_z}[µT]	σ_{B_x}[µT]	σ_{B_y}[µT]	σ_{B_z}[µT]
x-coil	0.5	2.6	0.7	2.2	1.6	0.7
y-coil	0.8	0.9	0.8	1.1	1.1	0.2
z-coil	0.2	0.7	0.6	0.9	0.4	0.6
1G-coil	0.1	0.6	0.1	0.1	0.2	0.6
2G-coil	0.1	0.6	0.2	0.3	0.4	0.5
3G-coil	0.2	0.5	0.8	0.4	0.2	0.3
4G-coil	0.1	0.1	0.1	0.4	0.2	0.3
5G-coil	0.1	0.3	0.2	0.1	0.2	0.1

Table 8.1: RMS error between the simulation and measurement for each coil, magnetic field direction and mapping axis.

8.2 Feedback Sensor Positioning

The validated COMSOL simulation was then used to calculate the magnetic field for each coil at a set of approximately 50'000 points in a volume around the MSR. These magnetic field values served as input to the magnetic sensor position optimisation algorithm as already used in 4.3. The possible magnetic field positions were constrained to a shell which is exceeding the MSR by 20 cm on all sides.

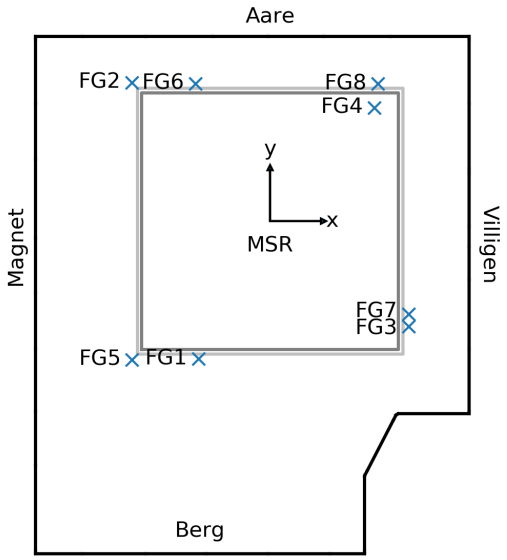

Figure 8.4: Position of the magnetic sensors as given in Table 8.2, relative to the MSR and the thermo-house walls. The z-coordinates of the sensor positions can be found in Table 8.2. Note that FG4 is not penetrating the MSR, it is positioned below it.

The positions of the eight magnetic sensors illustrated in Figure 8.4 yield a scaled condition number (see Section 4.6.2) of 12. Even though this scaled condition number will not be achieved when measuring with the running system, it was shown in the prototype that those positions achieve better shielding performance. As the plot in Figure 8.4 only shows a projection onto the xy-plane, the full set of coordinates is given in Table 8.2.

	x[m]	y[m]	z[m]
FG1	-1.4	-2.7	-1.1
FG2	-2.7	2.7	-2.1
FG3	2.7	-2.1	-1.9
FG4	2.0	2.2	-2.4
FG5	-2.7	-2.7	2.7
FG6	-1.5	2.7	2.7
FG7	2.7	-1.8	2.7
FG8	2.1	2.7	2.7

Table 8.2: Coordinates of the optimal positions for eight magnetic feedback sensors. The MSR has a height of 4.8 m but the centre is offset by 0.14 m, therefore all positions lie outside the cube.

Because the AMS of n2EDM was not yet operational when these positions were calculated, it could not be verified if the simulated condition number will be achieved. Even though a study on the optimal number of feedback sensors in terms of the condition number was performed, the final decision on the number of magnetic sensors was not yet taken. The above presented positions must be considered as starting point of the performance optimisation once the AMS is fully operational. Considering the respective tests in the prototype the calculated positions are promising.

Part III

Axion-Like Particles

Chapter 9

Introduction

The axion was originally postulated by Peccei and Quinn [64] to provide a solution to the so called strong CP problem, which can be explained in the following way. In the Standard Model (SM) of particle physics any term that does not violate the symmetry rules can be written down. This is also true for the following term in Quantum Chromodynamics (QCD) [65]:

$$\mathcal{L}_{QCD} = \frac{g_s^2}{32\pi^2}\theta G^a_{\mu\nu}\widetilde{G}^{a\mu\nu}, \tag{9.1}$$

where g_s is the strong coupling constant, G and $\widetilde{G}$ the gluon field and its dual respectively and θ a CP-violating phase. With the current best upper limit on the neutron electric dipole moment [16] it can be found that $\theta < 10^{-10}$ [66]. It is not clear why this parameter is not of order one and therefore constitutes the strong CP-problem, a fine-tuning problem.

Peccei's and Quinn's solution of the problem is the introduction of an additional U(1) symmetry which is spontaneously broken at a higher scale. The corresponding Nambu-Goldstone boson is called the axion [67, 68] and replaces θ by the dynamic axion field $a(x)$ [69]. The θ angle is then given by: $\theta = a(t)/f_a$. The axion field is oscillating around the minimum of its potential at $\theta = 0$ and naturally explains the smallness of θ.

The scale f_a at which this newly introduced symmetry is broken was initially set to the electroweak scale, however this was soon after ruled out experimentally. Nevertheless, the Peccei-Quinn mechanism is underlying other axion-like particle models [70–75], where the symmetry is broken at higher scales. Even though some of those particles can not solve the strong CP problem, they are good cold dark matter (CDM) candidates [76, 77].

The phenomenology of the axion, mainly depends on the value of the symmetry breaking scale f_a and is expected to lie close to one of the fundamental scales of particles physics M_{GUT} and M_{Planck}.

Low-mass (sub-eV) axion like particles can be created via non thermal mechanisms at the beginning of our universe through vacuum misalignment [78–80]. They subsequently form a coherently oscillating classical field

$$a(t) = a_0 \cos(m_a c^2 t/\hbar), \tag{9.2}$$

where the frequency depends only on the axion mass m_a. Because these particles get virialised by gravitational interactions during structure formation they have a finite coherence time [81]

$$\tau_{\mathrm{coh}} \sim \frac{2\pi}{m_\phi v_{\mathrm{vir}}} \sim 10^3 \left(\frac{2\pi}{m_\phi}\right), \tag{9.3}$$

with $v_{\mathrm{vir}}^{\mathrm{local}} \sim 300\,\mathrm{km/s}$ and corresponds to $\frac{\Delta\omega}{\omega} \sim 10^{-6}$.

Under the assumption that the axion field makes up for the whole dark matter (DM) content of the universe [82], which is

$$\rho_{CDM}^{local} \approx 0.4\,\frac{\mathrm{GeV}}{\mathrm{cm}^3}, \tag{9.4}$$

then an upper limit on the axion mass can be set to be $m_a < 0.1\text{eV}$. From the argument that the de Broglie wavelength of the axion must not be larger than the smallest dwarf galaxies in order to take part in structure formation, a lower limit on the axion mass can be derived, which yields: $m_a > 10^{-22}\text{eV}$ [75, 83–85].

The axion can interact with the standard model in the following way:

$$\mathcal{L}_{int} = \frac{C_g}{f_a}\frac{g_s^2}{32\pi^2} a G^a_{\mu\nu}\widetilde{G}^{a\mu\nu} + \frac{C_\gamma}{f_a}\frac{e^2}{32\pi^2} a F_{\mu\nu}\widetilde{F}^{\mu\nu} - \sum_{f=n,p,e}\frac{C_f}{2f_a}\partial_\mu a f\gamma^\mu\gamma^5 f \tag{9.5}$$

The first term describes the interaction with the gluon fields, whereas the second term gives the interaction with the electromagnetic field. In the last term the coupling of the fermions in the Standard Model is represented. The parameters C_g, C_γ and C_f are dimensionless model dependant parameters.

Axions can be sought for by using any of the above mentioned interactions. Previously, experiments looking for axions focused on the interaction with the electromagnetic field. These experiments can be grouped into three different categories, in parenthesis are given example experiments: helioscopes (CAST), microwave cavities (ADMX) and shining light through wall experiments (ALPS).

In ADMX the inverse Primakoff effect [86], where axions are converted to photons in the presence of a strong magnetic field, is used [87]. The power is enhanced in a microwave cavity which can be tuned to the axion energy. Similarly, in the CAST experiment [88] the same mechanism is used to look for axions produced by the sun. As the energies of those axions is different they will be converted to X-rays instead of microwave photons. In both experiments the drawback is that a power, which is proportional to the square of the coupling constants, is measured. This makes it difficult to cover the low axion-photon-coupling parameter space.

In the ALPS experiment, a powerful coherent laser is shone at a wall. By collecting the photons on the other side of the wall it can be measured how many photons converted to axions, passed the wall and converted back to photons on the other side [89]. Unfortunately, this scheme scales quarticly in the interaction constant and is even less sensitive to small values of the coupling.

A few years ago it was suggested to use the axion-gluon coupling to search for axions [69]. This coupling generates an oscillating electric dipole moment in any nucleon, as for example the neutron:

$$d_n(t) \approx +2.4\times 10^{-16}\frac{C_g a_0}{f_a}\cos(m_a t)e\text{cm}, \tag{9.6}$$

where $\frac{a}{f_a}$ is the previously discussed θ angle. The axion gluon field also induces an oscillating electric dipole moment in atoms[81]. Here we give the example of the ^{199}Hg atom.

$$d_{Hg}(t) \approx +1.3\times 10^{-19}\frac{C_g a_0}{f_a}\cos(m_a t)e\text{cm} \tag{9.7}$$

Because of the Schiff screening theorem [23], the induced oscillating electric dipole moment is intrinsically three orders of magnitude smaller than the one of the neutron.

Additionally to the lower frequency bound due to galaxy formation, further cosmological constraints can be set on the axion-gluon coupling. First, if axions exist they would also be produced in the core of stars. As they only weakly interact with standard matter, they would contribute a lot to the energy loss of stars. This would alter the known evolutions of stars, therefore, with the supernova SN1987A a limit for the coupling strength can be set [90–92]. Second, the axion interaction would also have an influence on the primordial ^{4}He abundance [93], which can be compared to measurement and theoretical calculation. Thus, a limit on the interaction strength can be set. In Figure 9.1 the above mentioned astrophysical and cosmological constraints are plotted.

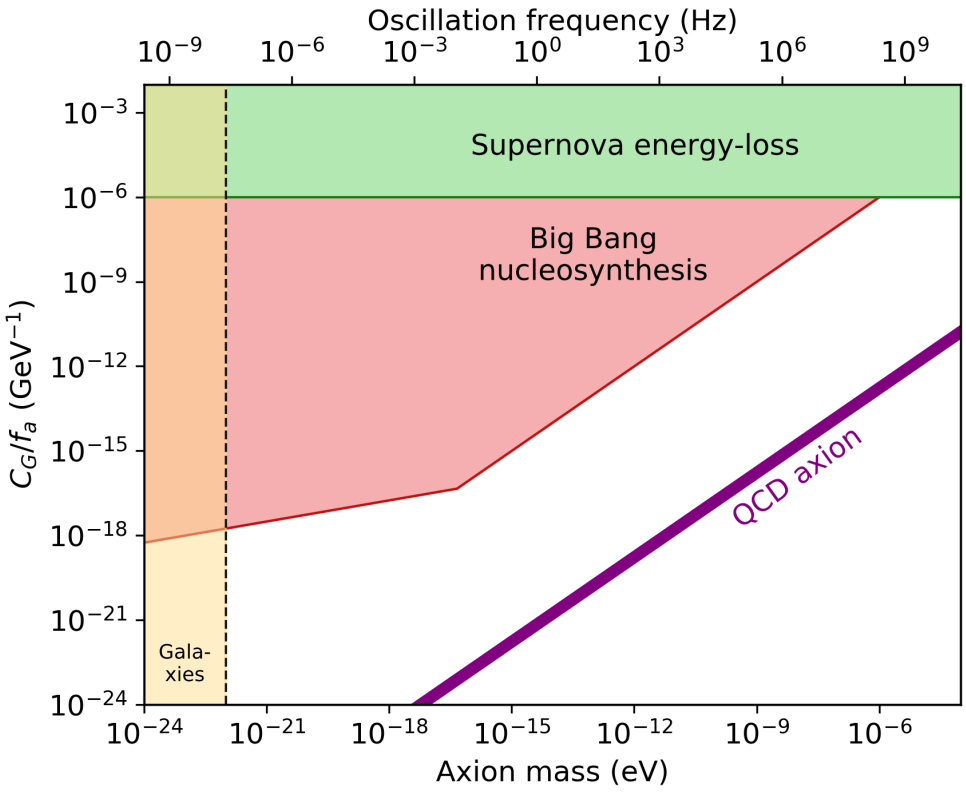

Figure 9.1: Visualisation of the constraints of the gluon-axion coupling due to cosmology. Limits coming from Supernova energy-loss, Big Bang nucleosynthesis and galaxy formation as explained in the text. Courtesy of Michał Rawlik.

Chapter 10

Reanalysis

In 2017, our collaboration published the first exclusions on the axion-gluon coupling coming from a laboratory experiment, using data from nEDM [50]. At this time the nEDM analysis had not yet been finalised, which had an impact on the published exclusions. In this chapter we are going to outline the reanalysis which led to improvement of approximately a factor two at frequencies from $2 \cdot 10^{-8}$ Hz to $5 \cdot 10^{-5}$ Hz compared to the previously published result. For this reason we will in the first section outline the general analysis procedure, followed by a discussion of the 2017 results and its limitations. Finally the improved result of the reanalysis will be presented.

10.1 Analysis Procedure

The analysis methods used for the reanalysis of the axion search are similar to the ones of the initial search. In the following the analysis steps are explained.

10.1.1 Periodograms

From the Equation 9.6 in Chapter 9 it can be seen that the aim of the axion analysis is to look for oscillations of the neutron electric dipole moment across a wide range of frequencies. Such a spectral analysis is normally performed immediately by using a Fourier transform. The result of applying a Fourier transform on a continuous function is the power spectrum. However, in most cases only a limited amount of data points are registered due to sampling. Therefore, the Fourier transform can not be applied and only an estimator of the power spectrum, called the periodogram, can be calculated. The classical periodogram is called the Schuster periodogram [94]. It is calculated in the following way:

$$P_S(f) = \frac{1}{N} \left| \sum_{n=1}^{N} g_n e^{-2\pi i f t_n} \right|^2 , \tag{10.1}$$

where g_n is a sequence of N measured values at equidistantly sampled time points t_n. Clearly, even sampling does not apply to the nEDM dataset, where normally data points were taken every five minutes, sometimes interrupted by longer breaks due to maintenance work on the apparatus. Generally the classical periodogram can be extended to the non-uniform sampling case by allowing the use of window functions which reflect the structure of the underlying time series. However, this extension distorts some of the important statistical properties of the classical periodogram, for example the result that equidistantly sampled Gaussian white noise is chi-squared distributed in the periodogram.

In 1982 Scargle addressed this issue with the introduction of the Lomb-Scargle periodogram [95]. It can be written in the following way:

$$P_{LS}(f) = \frac{A^2}{2}(\sum_n g_n \cos(2\pi f(t_n - \tau_n)))^2 + \frac{B^2}{2}(\sum_n g_n \sin(2\pi f(t_n - \tau_n)))^2, \tag{10.2}$$

where τ_n an arbitrary function. Scargle showed that it fulfils the following properties.

i In case of even spacing it reduces to the classical periodogram

ii It is equivalent to least-square fitting of sines and cosines to the data

iii It is invariant to time shifts, i.e. transforms $t \to t + \Delta t$

Especially the second point is of interest for this analysis, it reduces the procedure to least-squares fitting at the frequencies of interest. Normally, the time series to be analysed have a zero baseline. In the case of the nEDM data, this can not be assumed for the whole data set, therefore the Ansatz can be generalised to be including an offset into the fit function.

$$A\cos(\omega t) + B\sin(\omega t) + C, \tag{10.3}$$

where A, B and C are the fit parameters at each frequency. The estimated power can then be written as

$$P(\omega) = \frac{N}{4}\left(A^2 + B^2\right), \tag{10.4}$$

which reflects the normalisation used by Scargle [95], which retains the χ^2-distribution at each frequency in the case of white noise.

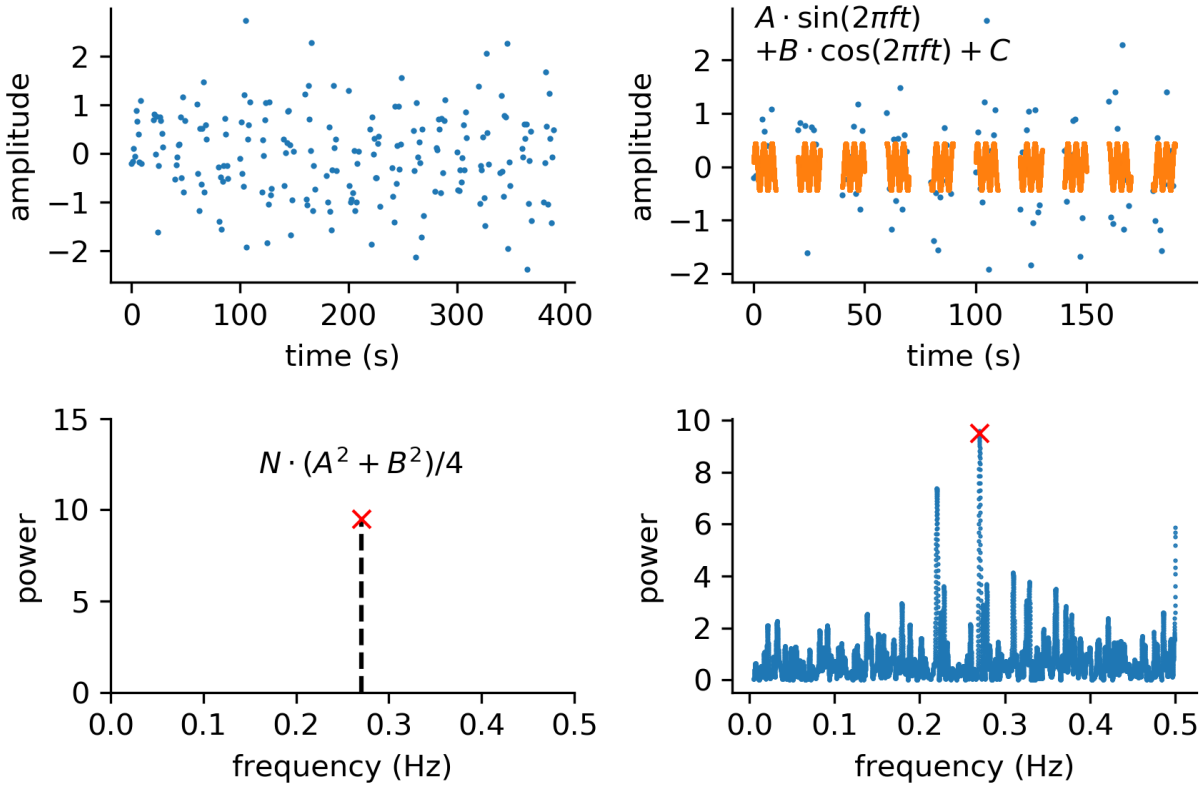

Figure 10.1: Top left: The time series constructed as described in the text. Amplitude in arbitrary units. Top right: Results of the fit plotted on top of the time series. Bottom left: Power calculated from the fit plotted against the investigated frequency. Bottom right: Full periodogram, where the power at the frequency investigated in the first three panels in highlighted.

In Figure 10.1 the creation of a Lomb-Scargle periodogram for a toy example is described visually. The structure of the time series is such, that data is taken in 20 bunches of 10 s with 10 s pauses in-between. The bunch structure is a feature, which will also be encountered similarly in the nEDM

experimental data. During the bunch a data point is taken every second with a 0.2 s jitter, drawn from a flat random distribution. An oscillating signal of amplitude 0.2 in an arbitrary unit and a frequency of 0.27 Hz was folded into Gaussian noise with amplitude 0.8 in the same arbitrary unit. The data time series is plotted in the top left panel. The next step is to fit a signal according to Eq. 10.3 to the data, in the top right panel the fit at the signal frequency 0.27 Hz is plotted on top of the data for the first 200 s. From the amplitude of the sine and cosine contribution of this fit, the power (Eq. 10.4) can be calculated and plotted against the examined frequency: with the evaluation at 0.27 Hz shown in the bottom left panel. The last two steps can now be repeated to cover the whole frequency range to build up the periodogram. It can be seen in the bottom right panel that even if a signal is injected, it does not mean that it constitutes the absolutely highest peak. These spurious peaks come from both the timing structure, as well as just random fluctuations. How to separate real signals from spurious peaks will be examined in the next section.

Before, we mentioned that in the last step the powers need to be calculated for the whole range of frequencies, which is not a clear statement at all. The choice of the evaluation frequencies in a Lomb-Scargle periodogram is of high importance. The lower frequency end is fairly easy to define as it is given by the inverse of the time spanned by the whole data set and can be complemented with the evaluation at $f = 0$, which corresponds to a constant offset. For the upper end of the frequency range there are many different approaches such as using the mean of the sampling interval or the minimal sample spacing [96] or basing the limit on the expectation of signals [96]. Most importantly, the spacing of the frequencies needs to be decided. This is fairly important, because by using a too widely spread frequency grid it can be that the height of certain peaks is underestimated, or even that whole peaks are missed. In the literature a frequency grid of

$$\Delta f = 1/mT \tag{10.5}$$

where T is the total observation time and m a value between 5 and 10, is considered a good choice [97, 98].

10.1.2 Null Hypothesis Testing

In the above described way, the periodogram for the data time series can be calculated. As noted earlier, the periodogram will contain many peaks, some being false peaks induced by statistical noise, some coming from aliasing and others being real peaks of a spectral signal. The question is now how to quantify the relevant peaks in the spectrum as opposed to spectral features only coming from the time structure.

One way to treat this problem is the use of Monte Carlo simulations. The power distribution at each frequency can be estimated by simulating the non-signal case, the so called null hypothesis, many times. The known analytic form of the signal-less case as calculated by Scargle [95] is exponentially distributed in power. If the power distribution at a certain frequency under the null hypothesis is given by

$$Z = P_{H_0}(\omega), \tag{10.6}$$

then the cumulative distribution function (CDF) of this power distribution can be written as:

$$F_Z(z) = 1 - e^{-z}, \tag{10.7}$$

which describes the probability of measuring a higher power than z under the null hypothesis. The CDF instead of the probability density function (PDF) is used in this case because binning, which always brings a certain information loss, can be omitted. For a certain power the p-value, which describes the probability to achieve this or a higher power, is then $1 - F_Z(z)$.

If we assume now a case where a high power is measured at a certain frequency, many Monte Carlo simulations would be needed to estimate the probability to encounter such a high power at this frequency, because the smallest probability that can be probed is given by the inverse number

of Monte Carlo simulations. Therefore the approach is to simulate a reasonable amount of data and based on 10.7 using

$$F_\omega(z) = 1 - A_\omega e^{-B_\omega z}, \tag{10.8}$$

as a fit function at each frequency, where A_ω and B_ω are the fit parameters and do not relate to the parameters used in Eq. 10.4. Those parameters can then be used to calculate the probability that the peaks are arising from random fluctuations. A reasonable number of simulations then means, to throw as many Monte Carlo simulations as needed to perform the fit within a bearable computation time.

Nevertheless, a detection can not immediately be reported if the probability of a peak exceeds a certain predefined p-value threshold, two things need to be considered additionally. First, even though a peak is detected at a certain frequency it has to be considered, that when looking at many frequencies, as in a periodogram, it is very likely that a rare event is detected. Or expressed differently, the power of the most significant peak also follows a random distribution. This so called"false alarm probability" is given in [96] as

$$p_{\text{global}} = 1 - F_Z(z)^N, \tag{10.9}$$

where N is the number of investigated events that is data points. This directly leads to the second consideration, that the above equation only holds true in the case of independent events. Due to the uneven sampling of the time series underlying the Lomb-Scargle periodogram and the arbitrary choice of frequencies, this does not hold true. Therefore N needs to be replaced with an effective number N_{eff} of independent frequencies. This number can be inferred, by fitting the right hand side of Equation 10.9 to the distribution of the most significant peaks in each periodogram.

10.1.3 Signal Hypothesis

If no detection can be claimed, the next step is to look what signals can be excluded. The question to ask is then which axion signals would be seen in the data and because they are not detected can thus be excluded? For the purpose of finding an answer to this question many axion signals with different parameters need to be tested and simulated. As an axion signal is characterised by two parameters, the amplitude A and the frequency ω, the exclusion parameter space is 2D.

What needs to be evaluated is the probability that an axion signal with certain parameters produces less power at the axion frequency than what is observed in the data. To be consistent with prior work [54], this can be written as:

$$\Pr(P^{H(\omega,A)}(\omega) < P^D(\omega)), \tag{10.10}$$

where $H(\omega, A)$ is the axion hypothesis under investigation characterised by ω and A.

The procedure to calculate this probability and the exclusion limit is the following:

- Generate 1000 time series with an axion signal with a pair of parameters (A, ω).
- Calculate the periodogram for each of the time series
- From this data the CDF can be calculated, which is then used to infer the p-value at $P^D(\omega)$ by interpolation.
- For every frequency there are many amplitude/p-value pairs. Again, by interpolation, the amplitude corresponding to a p-value of 0.5 can be found.

This set of 5% amplitude and frequency then forms the exclusion limit. Nevertheless, with this procedure one would also reject parts of the parameter space where the experiment is not sensitive to. A possible solution is the use of the CLs method as described in [82]. The idea is to scale the exclusion p-value by the probability to also reject the null hypothesis. It is defined as

$$\text{CLs} = \frac{\Pr(P^{H(\omega,A)}(\omega) < P^D(\omega))}{\Pr(P^{H_0}(\omega) < P^D(\omega))}. \tag{10.11}$$

Now if the H_0 as well as the axion hypothesis $H(\omega, A)$ are very low, the CLs will still be of order one and thus not be excluded. Eventually Equation 10.11 will be used to claim an exclusion.

10.2 Discussion of the 2017 Result

To elaborate what improvements can be made to the 2017 analysis, it is crucial to understand the data taking procedure for nEDM and the relevant parts of the subsequent analysis.

10.2.1 Data Taking

In the following we are only going to highlight the procedures relevant to the axion analysis. A more detailed description can be found in [99, 100].

In Section 2.1, it was discussed that in order to retrieve a neutron electric dipole moment, the Larmor precession frequency with an electric and magnetic field in a parallel and antiparallel state needs to be measured:

$$d_n = h\frac{f_{\mathrm{n},\uparrow\downarrow} - f_{\mathrm{n},\uparrow\downarrow}}{4E}. \tag{10.12}$$

Because the two measurements in the parallel and antiparallel state are taken at different times, magnetic field fluctuations can be mistaken as a signal. Therefore, the neutron Larmor precession frequency is corrected for with a Hg-comagnetometer.

The Hg-comagnetometer works such that the precession chamber is additionally filled with polarised Hg-atoms, hence the name comagnetometer. By reading out their Larmor precession frequency with a discharge lamp the magnetic field in the precession chamber can be measured.

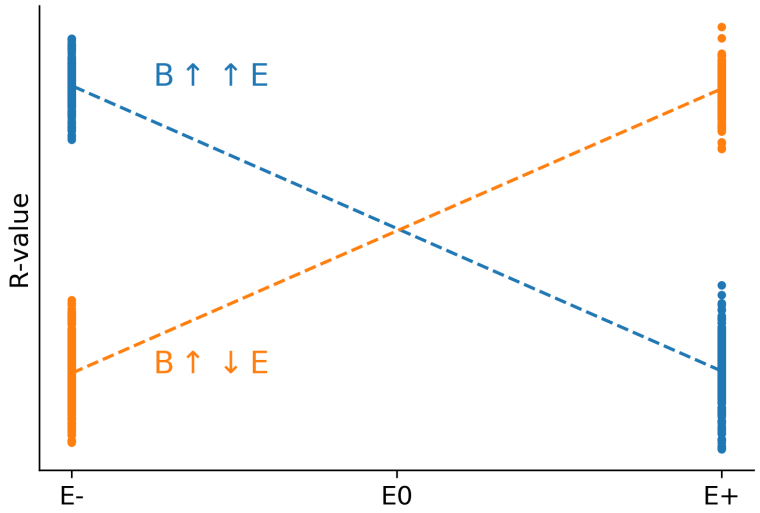

Figure 10.2: Depiction of the E vs $\mathcal{R}$ fit, which is performed for every of the 112 cycle long subdataset separately.

The ration between the neutron and mercury Larmor precession frequency, which is to first order independent of the magnetic field, is called $\mathcal{R}$. If $\mathcal{R}$ is measured in a perfect field, then $\mathcal{R} = \frac{\gamma_\mathrm{n}}{\gamma_\mathrm{Hg}}$ holds true, where γ_n and γ_Hg are the gyromagnetic ration of the neutron and the mercury respectively. But in our experimental environment some deviations need to be considered and therefore $\mathcal{R}$ must be written as

$$\mathcal{R} = \frac{f_n}{f_{Hg}} = \frac{\gamma_n}{\gamma_{Hg}}\left(1 + \delta_\mathrm{grav} + \delta_T + \delta_\mathrm{other}\right). \tag{10.13}$$

δ_T considers deviations coming from transverse magnetic fields, whereas subdominant effects such as the Ramsey-Bloch-Siegert shift and the earth rotation correction are considered in δ_{other}[100].

The most interesting term contributing to the understanding of the data taking procedure is the gravitational correction δ_{grav}. This term is coming from the fact that the neutrons and the mercury atoms do not sample the magnetic field in the exact same region due to their different heights of the centre of mass. δ_{grav} is found to be

$$\delta_{\text{grav}} = \frac{G_{\text{grav}} z}{B_0}, \tag{10.14}$$

where z is the centre of mass offset, B_0 the strength of the holding field and G_{grav} a combination of the relevant higher order modes of the magnetic field [100].

The neutron electric dipole moment is now measured by using the so called E vs. $\mathcal{R}$ fit. The $\mathcal{R}$ values for the $E\uparrow$ and $E\downarrow$ are fitted for the two different electric field directions, where B is parallel to E or antiparallel to E. The $\mathcal{R}$ value of the crossing points of the two fits corresponds to the $\mathcal{R}$ value in this run and the slope is proportional to d_{n}. With this procedure which is depicted in Figure 10.2 a pair ($\mathcal{R}$, d_{n}) is collected for each run.

Unfortunately, the measured dipole moment $d_{\text{n}}^{\text{meas}}$ is not equal to the true dipole moment $d_{\text{n}}^{\text{true}}$ due to systematic effects contributing to a false nEDM. This correction is mainly coming from the geometric phase effect[101]. As this effect is proportional to the electric field direction and strength it directly mimics the signature of an electric dipole moment.

The geometric phase effect systematic applies to both species, but is larger for the Hg-atoms. As the frequencies of both species are needed to calculate $d_{\text{n}}^{\text{meas}}$ the false electric dipole moment of the Hg-atoms contributes the largest fraction to the false EDM correction. The measured electric dipole moment must then be written as

$$d_{\text{n}}^{\text{meas}} = d_{\text{n}}^{\text{true}} + \frac{\hbar\gamma_{\text{n}}\gamma_{\text{Hg}}}{8c^2} R^2 \left(G_{\text{grav}} + G_{3,0} \left(\frac{\text{R}^2}{4} + \frac{H^2}{10} \right) \right) \tag{10.15}$$

in a cylindrical precession chamber, where R is its radius and H it is height. $G_{3,0}$ is a higher order contribution to the magnetic field, similar to G_{grav} used in the gravitational correction of $\mathcal{R}$ [100].

The higher mode $G_{3,0}$ as well as the transversal magnetic field B_T can be very well extracted from the offline magnetic mapping of the measurement volume [102]. Therefore the $\mathcal{R}$-value and the measured d_{n} can be corrected for using all terms including values which can be measured. This is not possible for G_{grav} as a unachievable accuracy of $< 1\,\text{pT/cm}$ would be needed. The corrected electric dipole moment and the corrected $\mathcal{R}$-value can then be connected through [100]

$$d_{\text{n}}^{\text{corr}} = d_{\text{n}}^{\text{true}} + B \frac{\hbar\gamma_{\text{Hg}}^2}{8c^2 z} R^2 \left(\mathcal{R}^{\text{corr}} - \frac{\gamma_{\text{n}}}{\gamma_{\text{Hg}}} \right). \tag{10.16}$$

To get a handle on the G_{grav} systematic effect the so-called crossing-point analysis was established. As G_{grav} mainly depends on the gradient of the magnetic field in the z-direction $\frac{dB_z}{dz}$ it can be measured by applying different gradients throughout the data taking process. A line can then be fitted through the ($d_{\text{n}}^{\text{corr}}$, $\mathcal{R}^{\text{corr}}$) data points for the different gradients. If the sequence of gradients is applied to both magnetic field directions, then the true value of d_{n} is found at the crossing of the two fits for opposite magnetic field directions. A visualisation of the procedure can be found in Figure 10.3.

To follow the described analysis procedure data needs to be taken with different B and E field directions and additionally, for each combination different vertical magnetic field gradients.

The magnetic field was reversed every four weeks during the 253-day long data taking campaign. During the periods with same magnetic field direction typically five different gradients were applied and changed on a daily basis. Finally, the electric field was changed in a pattern following [8,48,8,48] cycles corresponding to electric field states [zero, up, zero down]. The whole sequence took 112 cycles, which corresponds to a period of around nine hours.

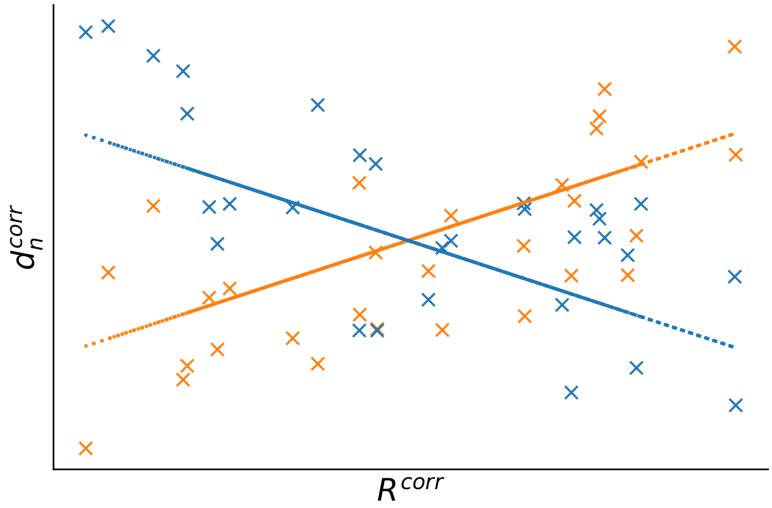

Figure 10.3: Depiction of the crossing point analysis, where the pairs ($d_{\mathrm{n}}^{\mathrm{corr}}$, $\mathcal{R}^{\mathrm{corr}}$) are plotted. For each relative magnetic and electric field direction a linear fit is performed. At the intersection of the two fit $d_{\mathrm{n}}^{\mathrm{true}}$ can then be extracted and via Eq.10.15 G_{grav} can be calculated.

10.2.2 Result and Limitations

In 2017 when the first search for axion-like particles was published the nEDM analysis was not yet finalised. Therefore unknown offsets, which are different for each part of the data with the same magnetic field direction and vertical gradient, were present in the data. This led to the decision that an offset for each combination of magnetic field direction and vertical gradient needed to be introduced in the fit as described in Section 10.1. Introducing these offsets and the decision to fit the two electric field directions separately caused that only data of approximately four hours length could be fitted continuously.

In Figure 10.4, it can clearly be seen that the sensitivity of the data starts dropping for frequencies lower than the corresponding two to three day interval, in this so called short time-base. The long time-base analysis is using data from the ILL-Sussex-RAL experiment, fitting time periods longer than a year.

Even though the long time-base analysis is compensating for the loss at lower frequencies, it is evident that the exclusion of the long time-base is intrinsically smaller. Therefore an overall improvement of the exclusion can be achieved by expanding the short base analysis to lower frequencies.

In the next section we are going to explore how the 2017 result can be improved and present the newly achieved result.

10.3 Reanalysis

By the time the reanalysis started, the nEDM analysis was finalised, unblinded and the result was published. Therefore the correction terms needed could be extracted. For completeness the versions of the analysis code which were used to extract the time series are listed below:

- Run level: 2019-12-05_09deb6f08806c9ab5cb997a8a885975d3cab9e94
- Cycle level: 2019-11-24_38ad7b7007ef93dc11ea6f229ca293d54b4c02b7
- Import: 2019-11-21_038f1a92f1d0a1fce6aca62426ffa3d324b8c64b

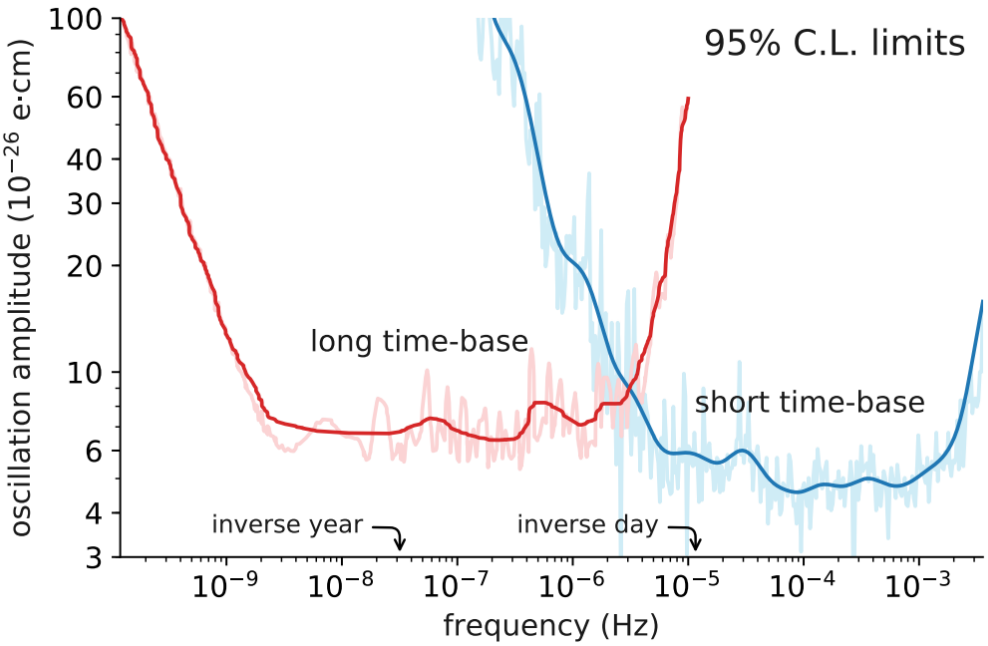

Figure 10.4: Reproduction of the result plot of the 2017 analysis from [50].

- Maps: 2019-08-14_6df996f6e296c2fd17578dfcbb2502315e88cdd9

In the following it is explained in detail what corrections are applied to the time series and it is illustrated with a toy example how the analysis strategy influences the exclusion limits. Finally, the improved limit on the axion-gluon coupling is presented.

10.3.1 Improvements

In the final step of the nEDM analysis, the crossing point fit is performed. Because from the analysis $d_{\mathrm{n}}^{\mathrm{true}}$ and $d_{\mathrm{n}}^{\mathrm{measured}}$ are known, with Equation 10.15, G_{grav} can be extracted and used in the time series to correct $d_{\mathrm{n}}^{\mathrm{false}}$. The false EDM correction is the only correction which depends on the relative electric and magnetic field direction. Because the corrections which depend on the magnetic field only, would be needed to remove the offset completely, are not known well enough, it was found that a better sensitivity can be achieved by keeping the offset. Additionally due to the offsets all magnetic field related systematics not causing frequency shifts correlated with the E-field are absorbed. With the applied false EDM correction the related systematic error is accurate to within the determination of the crossing point.

The main improvement which allows for better exclusion at low frequencies is the choice to fit the two relative electric and magnetic field directions at the same time. This allows the offsets to be determined without losing sensitivity. We will illustrate the impact of a combined relative direction fit with the help of the following example. The time structure of the example has similarities with the real time series. We imagine having a time series of 16 s which is equidistantly sampled every 0.1 s with a jitter of 0.03 s. Every 4 s the magnetic field gradients are changed which results in a random offset taken from a normal distribution with width 4 in arbitrary units for each block of data. Within this segment of same magnetic field gradient, an electric field direction change is performed every second.

In a first step, it is important to understand how an axion signal would look under the above described circumstances. In Figure 10.5 the underlying time structure is visualised in all four panels. The blue and orange areas on the timeline represent the parallel and antiparallel relative electric and magnetic field direction respectively. As there is no change of sign of the magnetic field in our imaginary time series, the parallel and antiparallel relative direction can be used interchangeably with

electric field up and down. The change in the magnetic field gradient, and thus the indication of a new offset is visually given by the red line.

In the first panel of Figure 10.5, an axion signal with frequency 25 mHz and an amplitude of one in arbitrary units is plotted for the first 16 s. In the second and third panel only electric field up and down respectively are plotted, but additionally considering the offsets for same magnetic field gradients. It can also be seen that for the electric field up direction the axion signal is inverted, as it represents a modulation of the $\mathcal{R}$-value which depends on the relative electric and magnetic field direction. When fitting the offsets in the two separate cases, the fitted offsets can almost fully absorb the axion signal.

The situation when both electric field directions are taken into consideration at the same time, is plotted in the fourth panel. The above mentioned inversion of the axion signal with the electric field direction can now be seen more clearly. It is important to notice that in case of an underlying axion signal a separation between the electric field up and down direction data can be spotted.

In Figure 10.6 the case of white noise, hence the H0 hypothesis is examined. Again in all five panels the underlying time structure is given by the blue and orange shades as well as the red lines separating different magnetic field gradients. In the first subplot white noise data, with a width of 0.8 in an arbitrary unit and offsets with a distribution width of 4, are plotted. The offset between the magnetic gradient changes can be spotted.

Because we want to see how a simultaneous fit of both electric field directions can improve the exclusion at frequencies slower than the period of same offset data (4 s), we try to fit an axion signal with $f_a = 25\,\text{mHz}$ to the white noise. This signal is depicted in the second panel.

We then use a least squares optimisation to fit this sinusoidal axion signal to the electric field up (third panel) part and electric field down (fourth panel) part of the noise data. It is evident that the fit yields some amplitude and thus power for the fits for separate electric field directions. This non-zero power is due to the offsets, which can be adjusted to describe the axion signal for oscillations significantly slower than the magnetic field gradient change period. In the real analysis this will be the case for all frequencies slower than approximately the inverse of a day and therefore results in a rise of power in the H0 periodogram for those frequencies.

Contrastingly, when looking at the combined fit in the fifth panel of Figure 10.6, the fitted amplitude is close to zero. The underlying reason is the observation made from Figure 10.5, that if there is an axion signal there is also a separation in the electric field up and electric field down data, which is clearly not the case for the white noise data as plotted in the first panel. Thus, the H0 periodogram produces low power estimates also for the low frequencies.

Clearly the above example depicts the situation for a single simulation of H0 data. Therefore the analysis on the toy data was repeated 1000 times to calculate the mean power at the frequency of 25 mHz for both the noise and the signal case. The outcome is presented in Table 10.1. It is noticeable that in the H0 case the mean power is on average 10 times lower and can therefore directly translate to a factor ten improvement of the exclusion.

	calculated power		
hypothesis	up	down	combined
H0	6.40	6.15	0.66
signal	16.5	15.7	20.7

Table 10.1: Summary of the toy example. The power produced by the time series when using either the whole time series or only the up/down relative direction part under the H0 and signal hypothesis. In the H0 case when only using up/down data, ten times more power is produced for low frequencies than for the combined dataset. For the signal case, the axion amplitude can be reconstructed equally well in both situations.

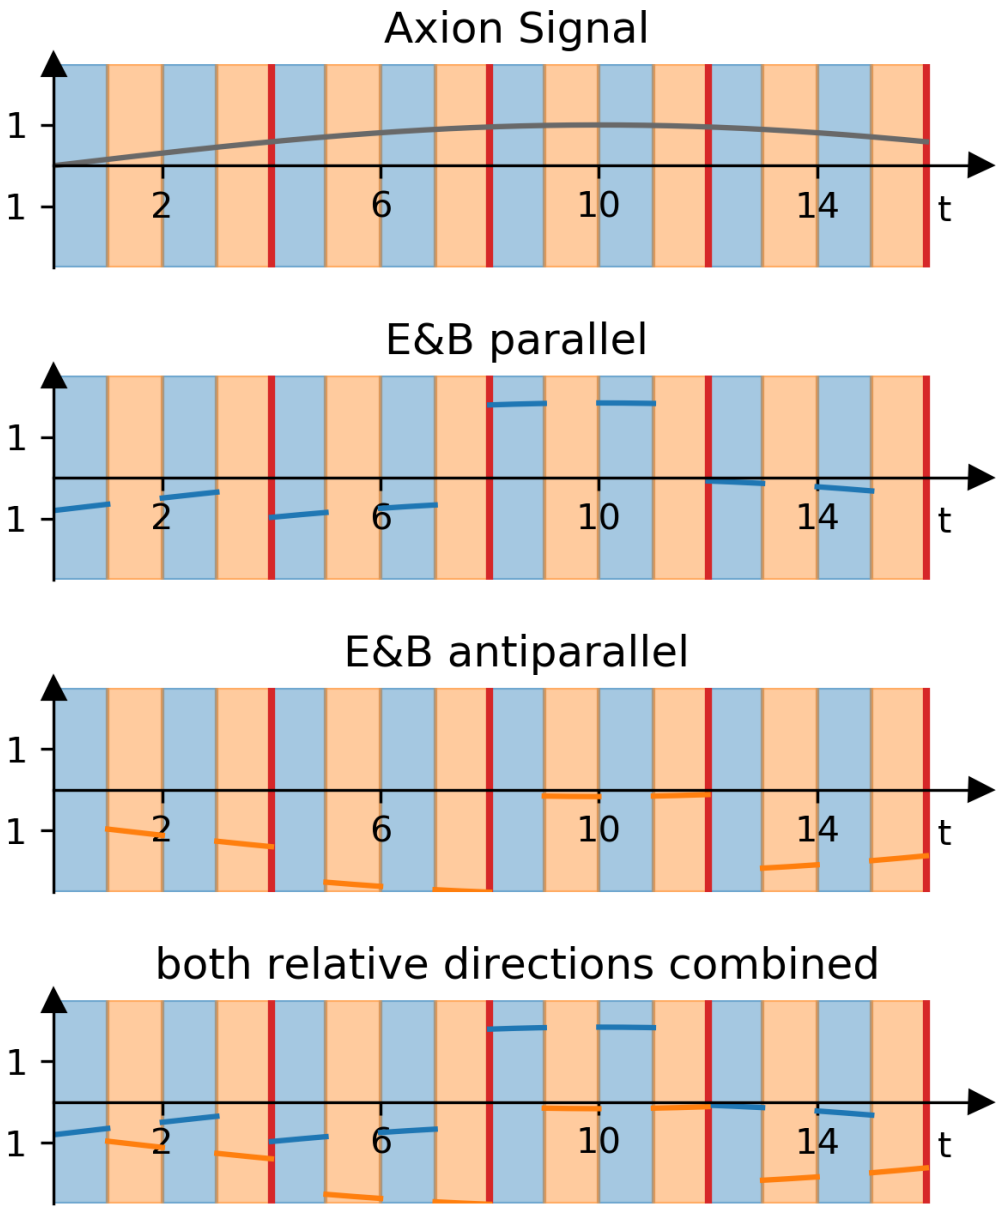

Figure 10.5: Depiction of a low frequency axion signal. For the EB parallel and EB antiparallel dataset only the part of the axion signal in the corresponding interval is considered. The axion signal is mirrored at the x-axis comparing EB-parallel to EB-antiparallel. When both relative directions are combined a separation between the EB-parallel and EB-antiparallel part is seen, which unveils an axion signal.

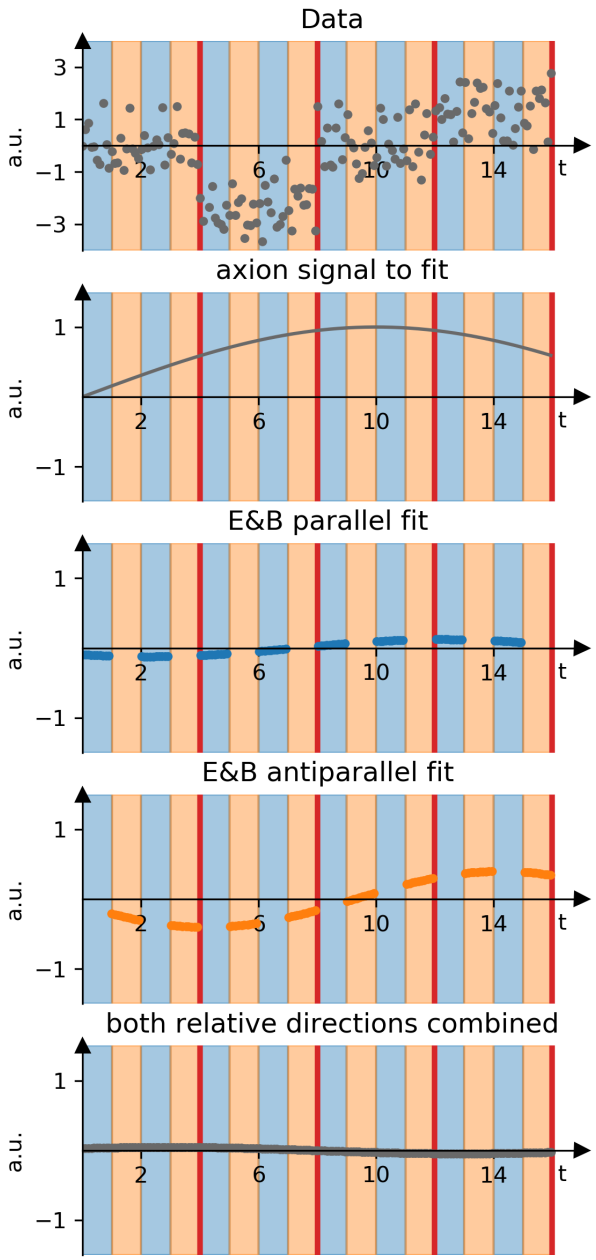

Figure 10.6: In the first panel, toy data consisting only of Gaussian noise and the offset for the gradient changes is plotted. The axion model which should be included is plotted in the second panel. In the third and fourth panel the fit of the axion model to the noise data is plotted for EB parallel and EB antiparallel respectively. It can be seen that some large signal amplitude is reconstructed, even if the data only includes white noise. In the fifth panel both relative directions are included for the fit. It can be seen that the fit yields only an invisibly small amplitude which correctly reproduces the white noise case.

10.3.2 Results

For the reanalysis of the 2017 result, the whole dataset consisting of 52'427 cycles taken over the course of approximately 500 days was used. For the H0 analysis a total of 43'737 linearly spaced frequencies, spanning from 10^{-9} Hz to 10^{-4} Hz were tested. Thus expanding the previous short time-base analysis by at least two orders of magnitudes. The frequency spacing corresponded to

$$\Delta f = \frac{1}{10 \cdot T} = 2.3\,\text{nHz}, \tag{10.17}$$

where T is the total timespan of the dataset.

Two points can be noted when looking at the periodogram of the data treated identically as in the nEDM analysis in Figure 10.7. First, towards higher frequencies, the periodogram seems to become more noisy. In reality this is an artefact from the linear frequency sampling. Because the x-axis is in logarithmic scale, many more individual frequencies are plotted per unit length of the plot. Therefore, it is also more probable to encounter over- or under-fluctuating amplitudes for the higher frequencies.

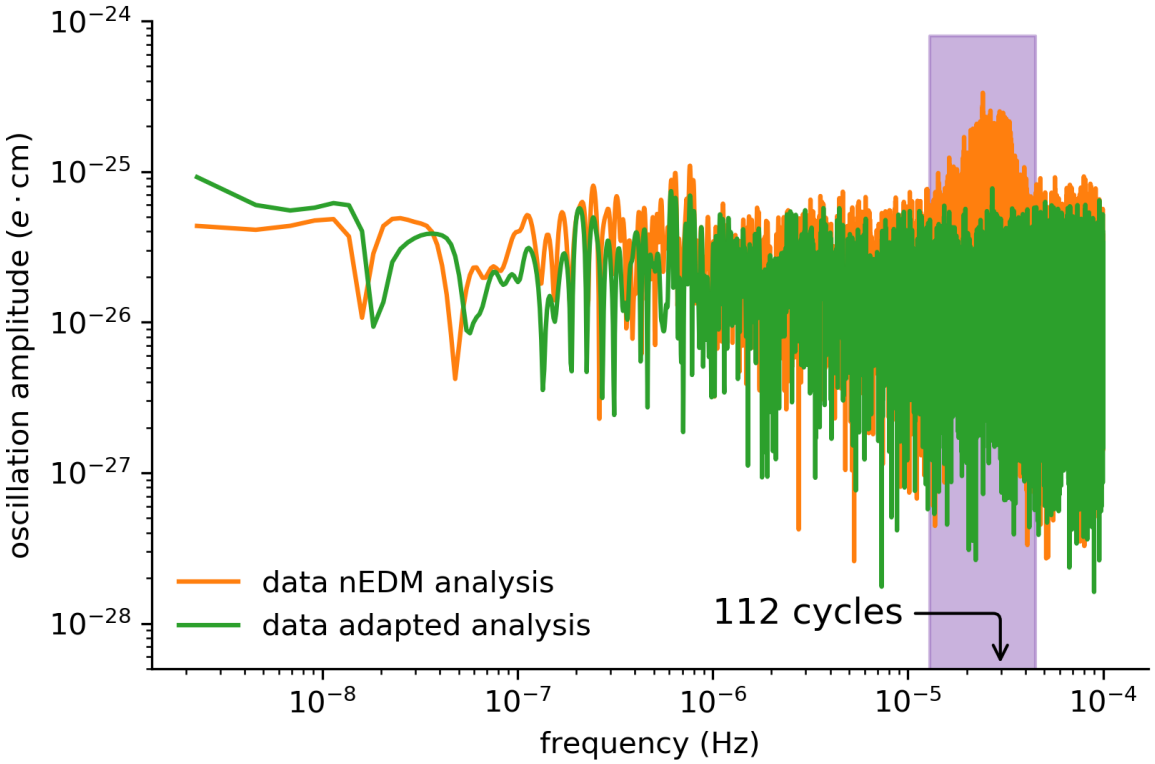

Figure 10.7: Periodogram comparison of the data sets once using an analysis identical to the one for nEDM and once using an adapted analysis where the G_{grav} correction was applied to each subdataset of 112 cycles. Due to the alternative analysis the excess in amplitude around the 112 cycles frequency can be reduced.

Second, there is a rise above the general noise level around frequencies between $1.3{\cdot}10^{-5}$ Hz and $4.5{\cdot}10^{-5}$ Hz (shaded in light purple). Also the frequency corresponding to the timespan of 112 cycles falls into this region. 112 is a special number in terms of cycles, because it is the length of a full pattern of electric field inversion which was applied during data taking. In Figure 10.8 a periodogram of the electric field direction was calculated. It is visible that the same structure as for the original periodogram of Figure 10.7 appears. Because the duration of the 112 cycles pattern was not always

the same, due to breaks inbetween the electric field changes or after such a sequence, the feature in the spectra is a broad bump instead of a sharp peak.

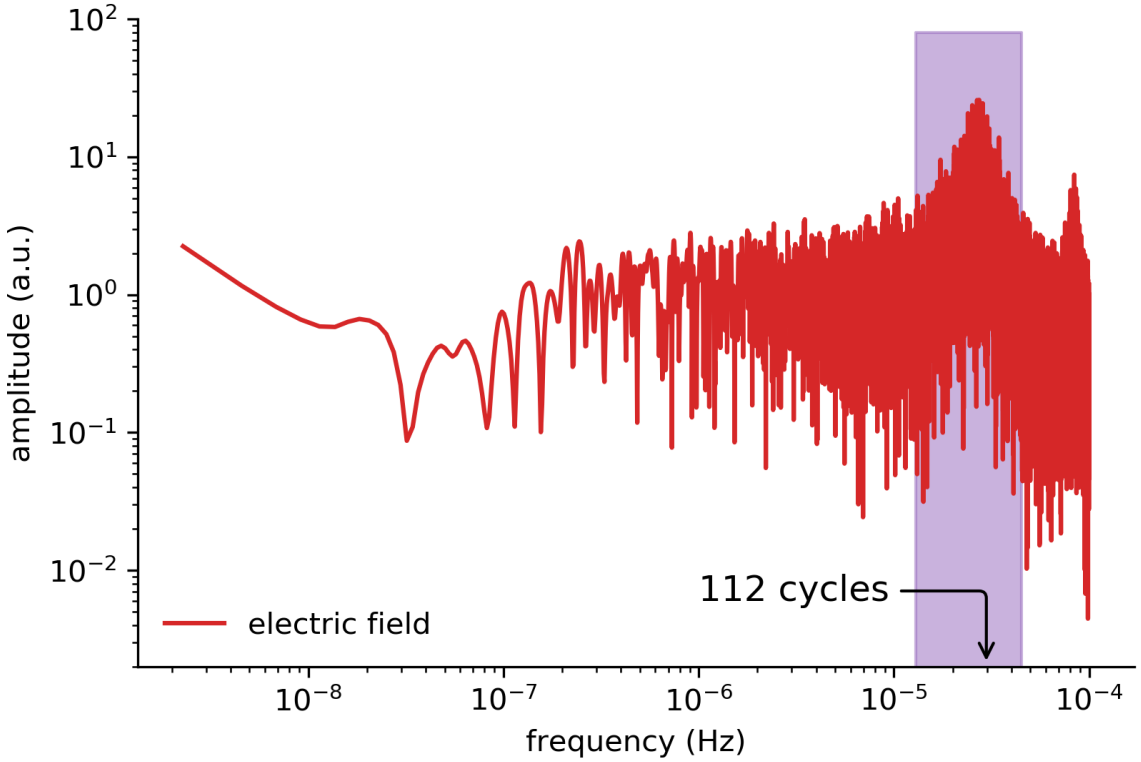

Figure 10.8: Periodogram of the electric field amplitude revealing the 112 cycle structure of the electric field reversal. The second peak at around $8{\cdot}10^{-5}\,\mathrm{Hz}$ corresponds to 48 cycles which is the duration of one electric field state.

The reason for the 112 cycles pattern in the dataset using the original nEDM analysis is the correction of G_{grav}. This is only corrected for with one mean value over the period of a whole data set which lasted approximately 2-3 days. Because it is very common that small but clearly visible gradient drifts occur over the course of a week the common correction induced an offset for each of the 112 cycles patterns. For the second data periodogram in Figure 10.7 the G_{grav} correction was therefore applied for every subdataset which led to the disappearance of the bump. The analysis searching for a static nEDM is not affected as d_{n} is calculated for blocks of 112 cycles and then combined.

For the detection analysis the data periodogram with the adapted analysis was compared to the null hypothesis. A thousand simulations only containing Gaussian noise were conducted and the global 1σ to 5σ thresholds were calculated. This comparison is plotted in Figure 10.9. Clearly, there is still a bump around the 112 cycles frequency visible. At this point this is only due to the timing structure of the data points, which clearly still reveal a 112 cycle periodicity.

Additionally, no peak exceeds the 5σ threshold which is required to claim a detection. Therefore it can be proceeded to the exclusion of axion signals in the whole frequency region.

For the exclusion analysis also one thousand simulations were performed for each pixel of the amplitude-frequency parameter space. As there are no resonant structures expected in the frequency domain, 6000 logarithmically spaced frequencies from $1{\cdot}10^{-9}\,\mathrm{Hz}$ to $1{\cdot}10^{-3}\,\mathrm{Hz}$ were tested. The examined amplitudes spanned 30 values from $1.5{\cdot}10^{-27}\,e\,\mathrm{cm}$ to $1.2{\cdot}10^{-24}\,e\,\mathrm{cm}$ which were also logarithmically spaced.

In Figure 10.10 the newly achieved axion exclusion limit is plotted against the limits achieved in

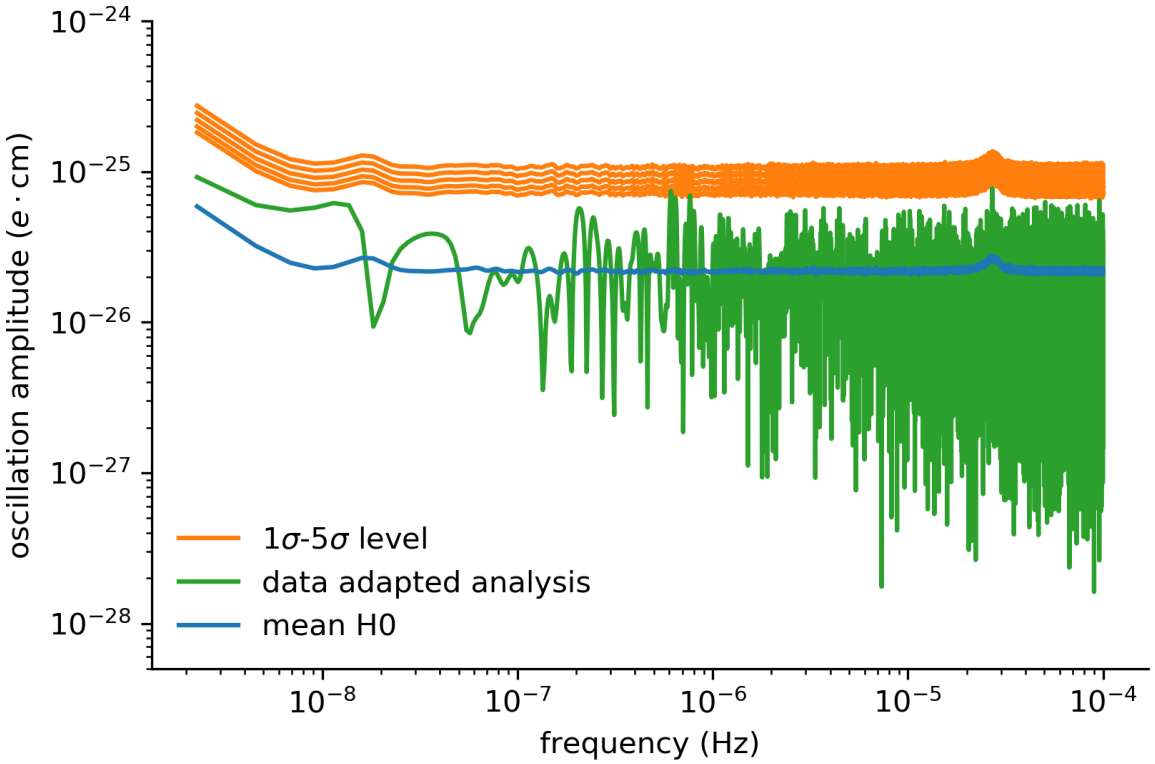

Figure 10.9: For the detection analysis the data periodogram is plotted against the 1σ to 5σ thresholds. To claim a detection the 5σ limit needs to be exceeded. The imprint of the electric field reversal pattern of 112 cycles can still be seen in the averaged H0 hypothesis plot.

the 2017 analysis with the short time-base and the long time-base analysis. It can be seen that the new result extends the limit of the short time-base analysis further down to low frequencies. From frequencies below $1 \cdot 10^{-8}$ Hz the sensitivity is decreasing as the total temporal length of the dataset is exceeded.

Additionally, in Figure 10.11 the exclusion is plotted in terms of the axion-gluon coupling C_G. The analyses using the nEDM dataset can thus expand the limits down to a region of $m = 1 \cdot 10^{-22}$ eV excluded by galaxy formation. The old limits are thus improved with the new analysis.

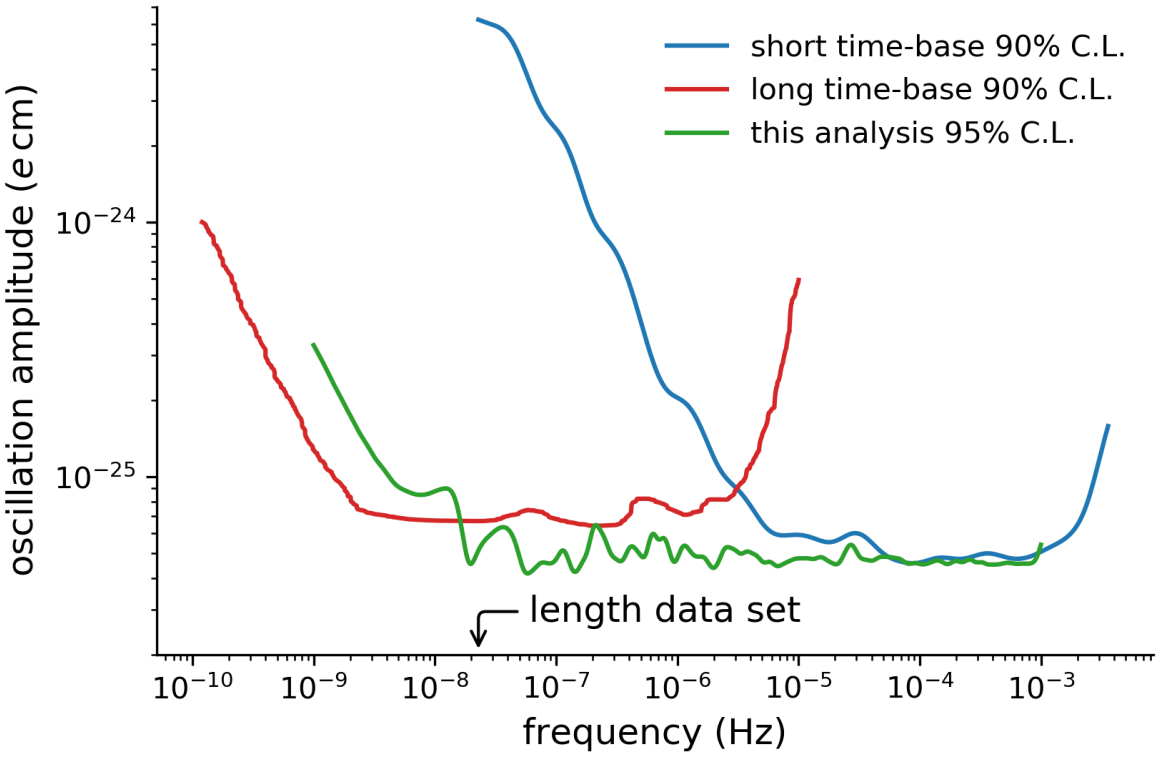

Figure 10.10: The exclusion limits achieved with this analysis compared to the short time-base and long time-base analyses of 2017. The short time-base analysis was also performed on the nEDM dataset. It can be seen that the new analysis can expand the short time-base analysis to lower frequencies, only dropping off at approximately $2 \cdot 10^{-8}$ Hz which corresponds to the inverse length of the dataset.

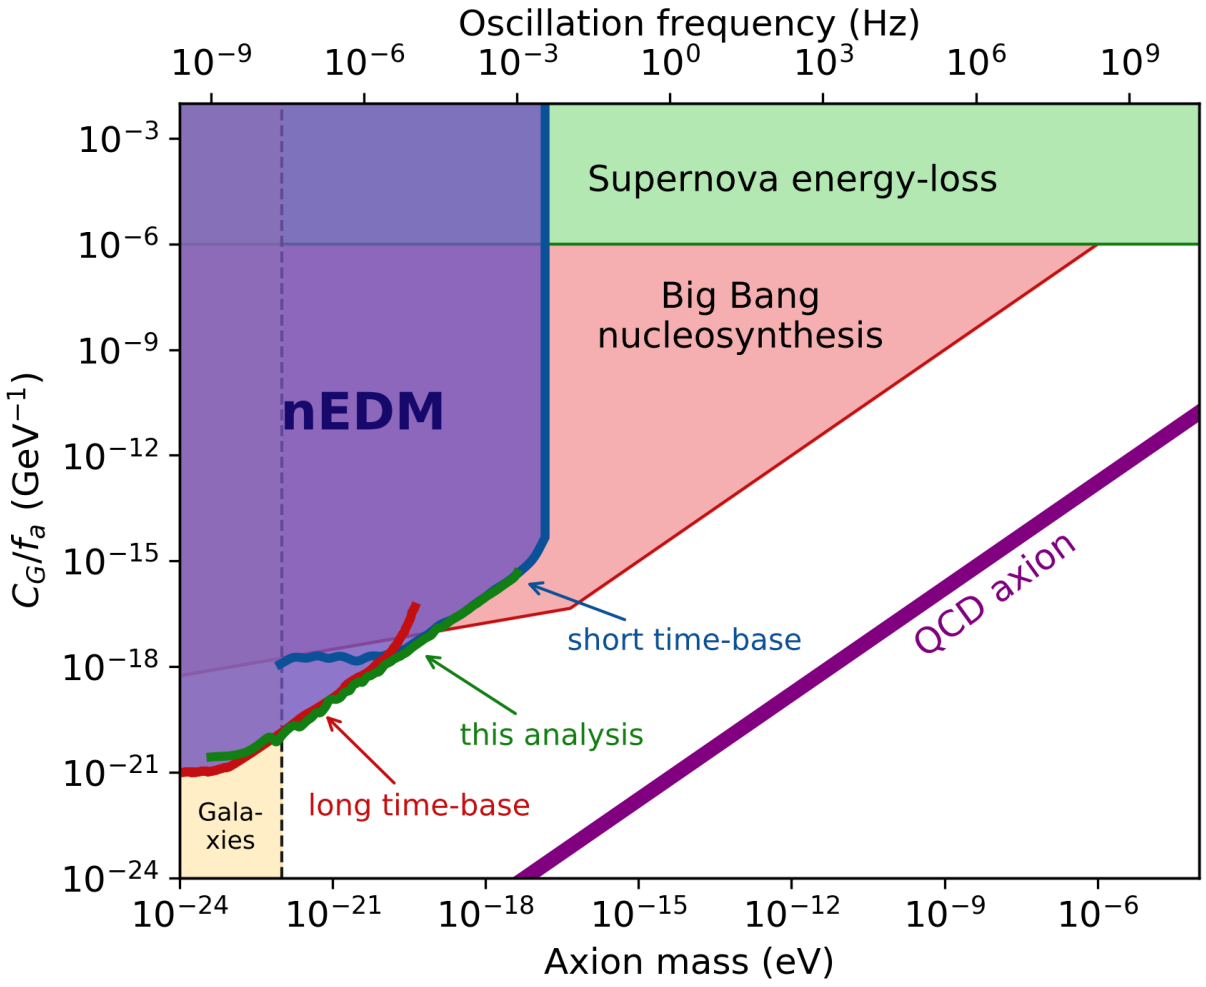

Figure 10.11: This reanalyis and the 2017 analysis are plotted as exclusions of the axion-gluon coupling C_G. The explanations for the cosmological limits are given in Chapter 10.

Acknowledgements

First I want to thank Prof. Klaus Kirch for the opportunity to pursue a PhD in his group. I appreciated very much the responsibility and freedom I was entrusted with. He always managed to look at the problems I encountered from a different point of view, which either made them disappear or at least made them less scary. Thank you for being a role model and caring about the human factor. I could definitely grow during my time here.

I would like to acknowledge the effort by Prof. Stéphanie Roccia and Prof. Hans-Arno Synal to serve as my co-examiners and also the time they invested in reading my thesis.

Next I would like to express my gratitude to Nick Ayres, especially during the last part of my thesis you were always there when I needed your advice or opinion. I enjoyed very much our collaboration and your experience on the axion analysis, without you everything would have been much harder. I also want to thank you for the enormous investment of proof-reading my thesis.

I also collaborated closely with the other members of nEDM at ETH group without whom the impressive AMS could never have been built. Thank you Jochen for always having an open ear for problems of any sort and making sure I was caring enough about the details. Thank you Vira for making the women power in our group complete and always being positive. Merci Luke Noorda for making sure the AMS wires are taking the right paths.

PSI not only hosts world-leading facilities and our n2EDM experiment, but also world class scientists. I am very grateful that I could profit from your huge knowledge and experience. I owe many thanks to Bernhard Lauss, Philipp Schmidt-Wellenburg, Georg Bison and Geza Zsigmond for sharing their insights with me. Special thanks goes to the best mechanic, Michi Meier, when conducting my small experiments at PSI I could profit his expertise and abilities.

The most memorable events during my PhD are certainly the many real-life collaboration meetings I could join. Not only did it train my ability in giving talks, but there was always a friendly atmosphere of curiosity where one could learn something new. Many thanks to the whole nEDM collaboration.

Even though not all of the people in the ETH group were working on the coolest of all experiments (or in particle physics at all), I always enjoyed the company of Aldo Antognini, Karsten Schuhmann, Manuel Zeyen, Ryoto Iwai, Ivana Belosevic and Michał Rawlik. The diversity of different topics in the group definitely enlarged my horizon as did the discussions over lunch with the above mentioned people.

Special thanks goes to Michał for entrusting me his "babySFC" as it used to be called. Even though I was sometimes being annoyed about your baby I could always count on your support and knowledge even years later. Apart from this, thanks for being a true friend.

I am convinced that science needs more female role models in visible positions. You, Anna Soter are one of those for me. Thank you for giving me the trust that I can make it, your support and precious advice.

In German you would say: "Geteiltes Leid ist halbes Leid", this is certainly also true for the faith of PhD students. It is always special to share experiences with fellow PhD students. I'm happy I could share with Pin-Jung Chiu, Kim Ross and Duarte Pais. The nEDM summer school in 2019 is still a lively memory of mine.

During my PhD I had the opportunity to supervise many semester theses. I appreciated the dedication of the students I was taking care of: Tim Röthlisberger, Wilke Grosche, Florian Meier and Rebecca Hampp. Additionally, I was lucky enough to be allowed to work with several apprentices.

Thank you Cornel Andreoli for entrusting these young people to me. I enjoyed observing their evolution and could always count on the work of Cyril Strässle, Jonas Oertli, Michel Philippin and Simon Hauri. Big thanks goes to Simon for taking professional pictures of the prototype and the AMS at PSI:

Even though physicist are doing important science someone needs to take care of the administrative work. At ETH these are Caroline Keufer-Platz, Bettina Lareida, Gabriela Amstutz and Jennifer Zollinger and at PSI Anita Govaerts Van Loon. I appreciated your patience and organisational skills.

Last but not least I owe many thanks to my family:

Ich möchte mich bei meinen Eltern bedanken für die Unterstützung und die Liebe, die ich auf meinem Weg bis hierhin erfahren durfte. Danke Papi, dass du meine Interesse an der Physik geweckt und geförderst hast. Merci Mami, dass du immer an mich glaubst und mir hilfst die wichtigen Dinge im Leben zu sehen. Ich bin auch sehr dankbar dafür, dass mein Mémé und Pépe immer hinter mir stehen und sich dafür interessieren was ich tue.

Lieber Valentin, ohne dich hätte ich wohl den letzten und härtesten Teil meiner Doktorarbeit nicht so schadlos überstanden. Merci für deine Zuversicht, Ruhe und bedingungslose Liebe. Ich bin unendlich dankbar dich in meinem Leben zu haben.

Appendix A

Construction Prototype

As elaborated in Chapter 4, the prototype consists of eight coils, out of which three produce homogeneous fields in each spatial direction and five are first order gradient coils. Each of the coils consists of a number of simple loops, characterised by the corresponding total current. Similarly to the construction of the n2EDM AMS, the currents were decomposed into three elementary currents to decrease the number of windings. As the power supplies were already available, the decomposition into elementary current followed the current output of the power supply. In the case of the prototype the elementary currents were: 5 A, 1 A and 0.2 A. Consequently, there was a maximum of four windings which needed to be effectuated with the 1 A- and 0.2 A-wire.

The decomposition, and thus number of windings with each wire in order to achieve total current in the simple loop was given at the top of each wiring plan as seen in Figure A.2. Additionally, the path the wire should take for the respective simple loop was plotted in a 3D plot of the prototype. Every vertex in the grid has a number, and the simple loop was also given as a sequence of vertex numbers. An example of such an installation plan is given in Figure A.2.

In Table A.1 the number of total simple loops for each coil is given. Not all of the simple loops had to be wound with all elementary currents. As the simple loops are ordered by total current, the bigger the simple loop number, the more probable it is that it does not need to be wound with the 5 A- and 1 A-wire. All the wiring plans for the eight coils can be found on the group drive: Z:\projects\nEDM\SFC\prototype.

For each coil and elementary currents the simple loops that need to be wound need to be selected. In the prototype those simple loops belonging together will be wound one after another with the same continuous wire. Because of this, the order and the starting points of the simple loops had to be chosen, such that the winding process could directly be continued. If there were simple loops which did not have a common vertex, they had to be connected through a short path. In the end this short path had to also be taken on the way back, such that the currents flowing in counteracting directions could be cancelled.

	sl
x-coil	50
y-coil	19
z-coil	50
1grad	40
2grad	61
3grad	37
4grad	32
5grad	59

Table A.1: Total number of simple loops for each coil of the prototype.

The wire used was an enamelled copper wire, with diameters between 0.8 mm and 1.2 mm. As the coating of the wires consisted only of a thin layer of varnish, it has to paid attention to not scratch the surface which would have resulted in a short between the coils.

Even though, the volume taken up by the wires was estimated prior to the installation, towards the end of the construction the cable channels became very crowded. At certain vertices they had to be reinforced with some mechanical structure as can be seen in Figure A.1.

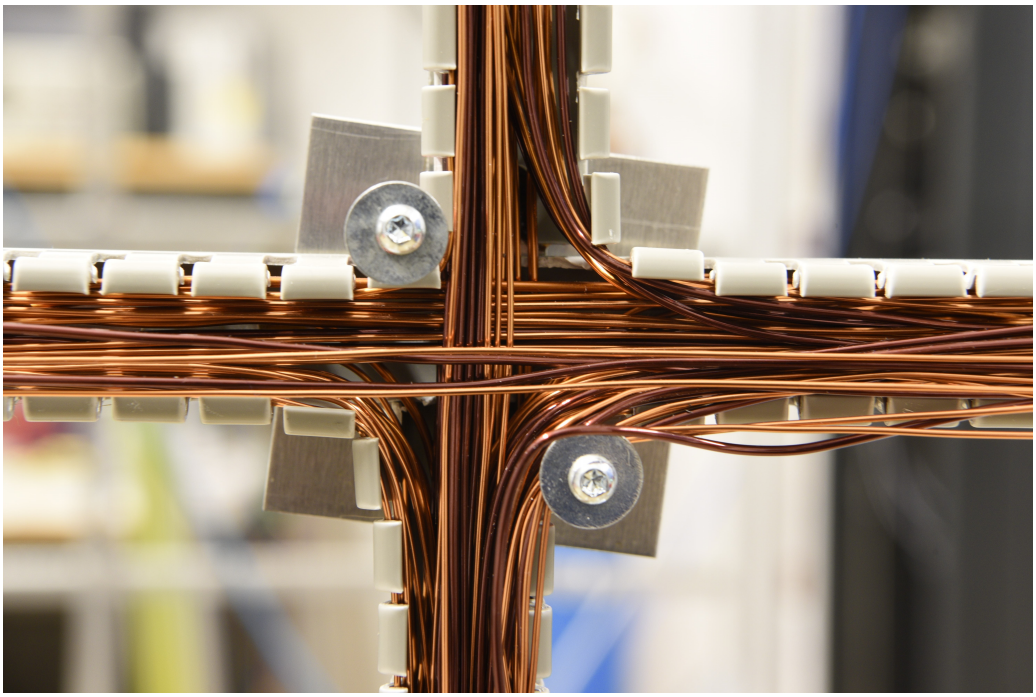

Figure A.1: Close-up of a cable channel towards the end of the construction, where certain vertices had to be reinforced with some mechanical structure.

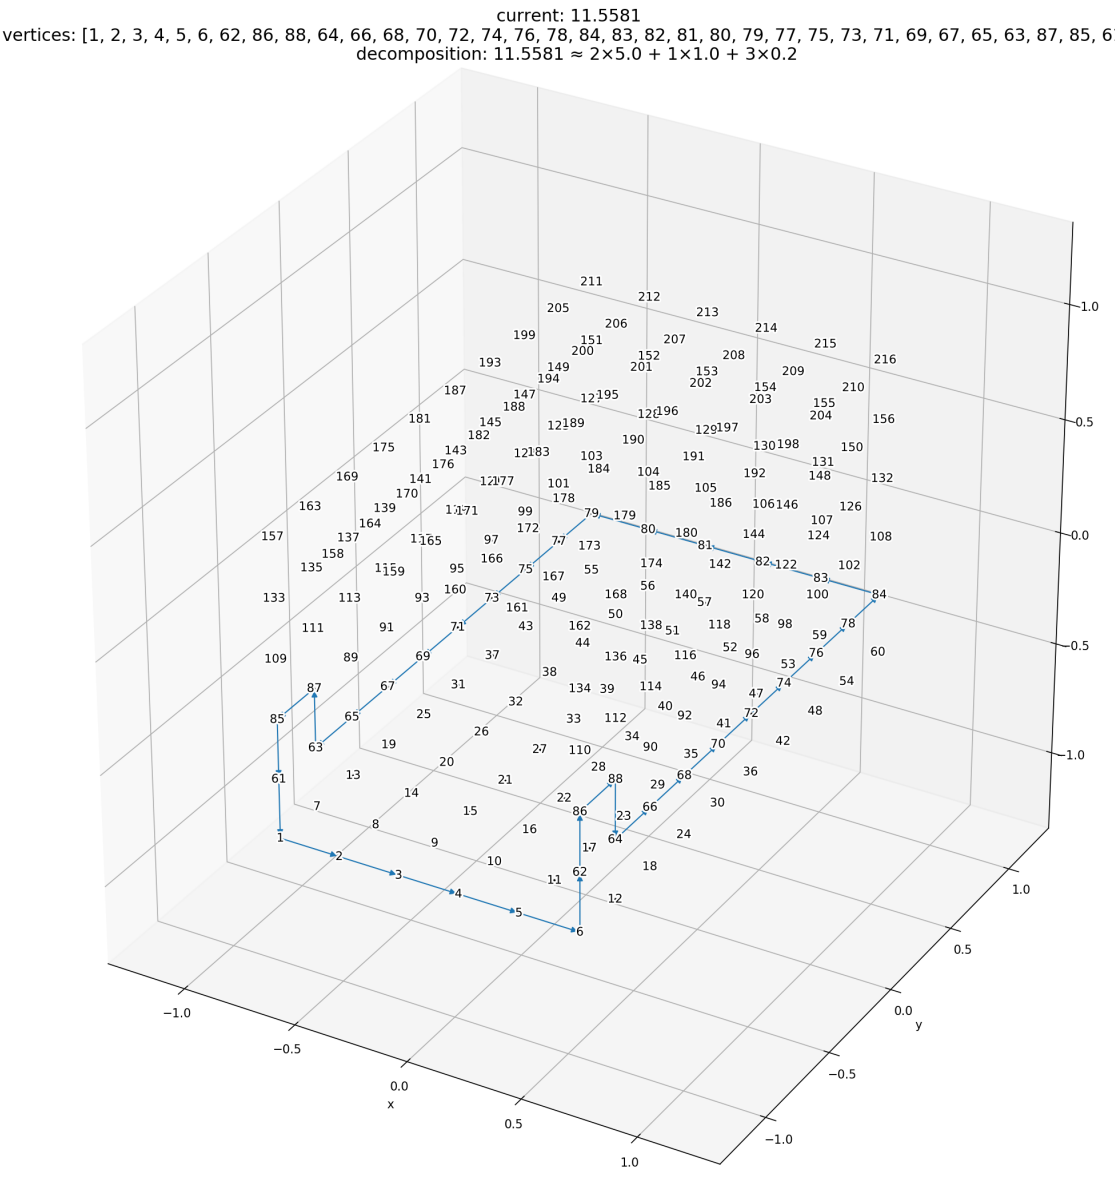

Figure A.2: Installation instruction for the simple loop No. 5 of the z-coil. The outline of the prototype is given in terms of the vertex numbers. On top the wire path is given as sequence of vertex number. Additionally, the decomposition in elementary currents provides the information on how many windings with each wire are needed. The direction of the current flow is indicated by small arrows.

Appendix B

Details n2EDM AMS

For reasons of electrical safety the power supplies driving the coils will be maximally operating at 50 V. Therefore, in the choice of the cables for each coil and circuit it has to be made sure, that the resistance is not exceeding $R_{\max} = 50\,\mathrm{V}/I$.

For the cross section of the cables only commercially available solutions were considered. Whenever the resistance was too close to maximal allowed one, the next bigger cross section was used. All choices of wires still provide some margin, such that even larger currents could be driven.

coil	$I_{\max}$[A]	A [mm^2]	L [km]	R[Ω]	$R_{\max}$[Ω]
x	15	16	1.7	2.1	3.3
	5	6	2.2	7.2	10
	1	2.5	4.2	20.9	50
y	15	16	1.9	2.3	3.3
	5	6	2.3	7.7	10
	1	4	5.1	25.3	50
z	15	16	1.9	2.3	3.3
	5	6	2.2	7.2	10
	1	4	5.1	25.4	50
G1	10	6	1.0	3.5	5
	3	4	1.8	8.7	16.7
	1	1.5	2.1	28.8	50
G2	12	6	0.9	2.9	4.2
	5	4	1.3	6.3	10
	1	2.5	3.6	28.4	50
G3	15	10	0.9	1.8	3.3
	4	4	2.0	9.8	12.5
	1	2.5	3.4	26.8	50
G4	8	6	1.1	3.7	6.25
	3	4	2.0	9.6	16.7
	1	1.5	2.2	29.0	50
G5	12	6	0.9	3.0	4.2
	5	6	1.7	5.6	10
	1	2.5	4.0	32.2	50

Table B.1: Maximal current $I_{\max}$, wire length L, wire cross section A, total resistance R and maximal resistance $R_{\max}$ for every circuit of each coil.

Appendix C

AMS Winding Application

Already during the winding process of the prototype, which has been accompanied by visualisations of the simple loops paths, we noticed that mistakes can happen very quickly even when working cautiously. Even though the winding process of the n2EDM differs from the one of the prototype, it was clear that due to our high performance goal on the field homogeneity such mistakes need to be prevented, ideally during the winding process or immediately after.

For his reason an Android application was developed, which helps preventing and finding such winding mistakes. Additionally it provides modes to check the current and composition of each bundle.

Every vertex of the grid was labelled with a QR code (a 2D barcode) sticker also indicating the coordinates and the name of the vertex. Additionally, every cable and cable bundle also has a barcode attached to it, containing the coil and simple loop information. An example of both a vertex QR code and a cable label can be seen in Fig. C.1.

When opening the application, a screen appears (Fig. C.2a) on which one can choose the task which needs to be completed. Regardless of which mode is used, the user is asked to provide the coil name and the simple loop number s/he is working with. The choice of the coil name can be seen in Fig. C.2b. An identical scrolling list is provided to afterwards choose the simple loop number.

To make sure that the input provided by the user matches the bundle s/he is working with, the user is asked to scan the barcode on the bundle. The generic scanner view can be seen in Fig. C.2c. If this information matches the user is forwarded to the main window, which is different for each mode.

In the "Winding" mode the user is installing a bundle. In the main window C.2d the next vertex through which the bundle needs to pass is displayed in bright pink. By clicking on the "Scan Vertex" button the generic scanner view is opened. Every QR code of the vertices which are passed during installation need to be scanned. If the QR code matches the given vertex, the name of the next vertex is displayed. When all vertices are scanned and the bundle is installed correctly, the user is asked to provide her/his credentials.

When using the "Check Path" mode, the task of the user is to follow a previously installed wire and scan the QR codes of the vertices along the way, without knowing which is the correct path. The instead of displaying the next vertex, the application states which was the last vertex which has been scanned correctly, thus providing orientation after eventual breaks or problems. Both the winding and the check path mode have a "Change Camera" button, which switches between front and rear camera, which is convenient when working with vertices attached to the ceiling.

The third mode is the "Check Bundle" mode. It is used to check that the correct wires are in the bundle. After scanning the barcode on the whole bundle, one can proceed to scan all the individual wires. In the main window (C.2e), the application displays the number of the wires which have not been scanned yet. When all wires are successfully scanned the user can continue with installing the wire.

In a final step when all simple loops of a coil are interconnected, the "Check Current" mode is used to test this connection. For each simple loop a current measurement can be submitted, if it matches

Figure C.1: QR codes used to identify the vertices.

the deposited value the measurement is stored. Else, the user is informed that the measurement diverges from the correct values and some verification of the connections has to be made.

The application works such that all the information regarding bundle paths, decomposition and expected currents are stored externally on an Amazon Web Service (AWS) bucket. This has the advantage that the application does not need to be reinstalled when changes to the files are made. Also, all the installation information, which bundle was installed when by who is uploaded. The AWS bucket is kindly provided by Prime Office AG.

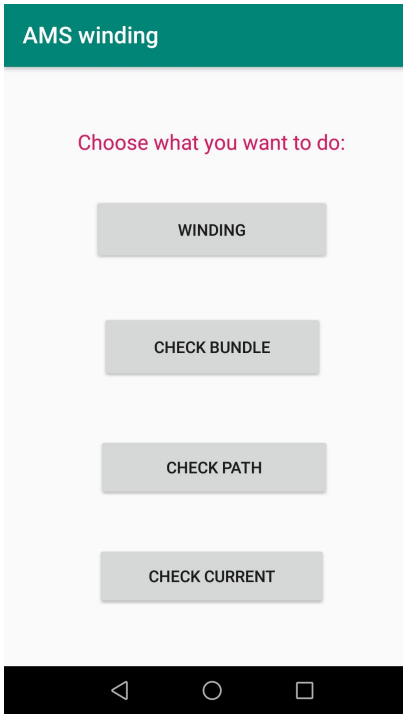

(a) Start window.

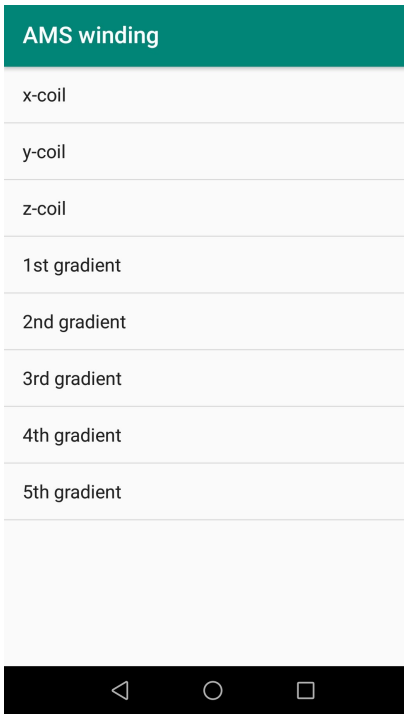

(b) Choice of the coil name.

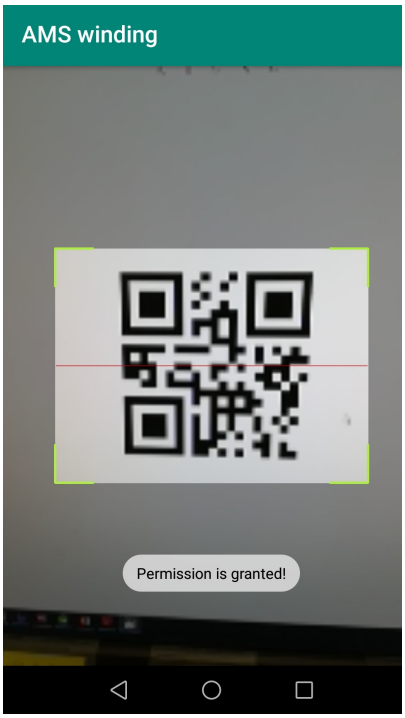

(c) Generic scanner view.

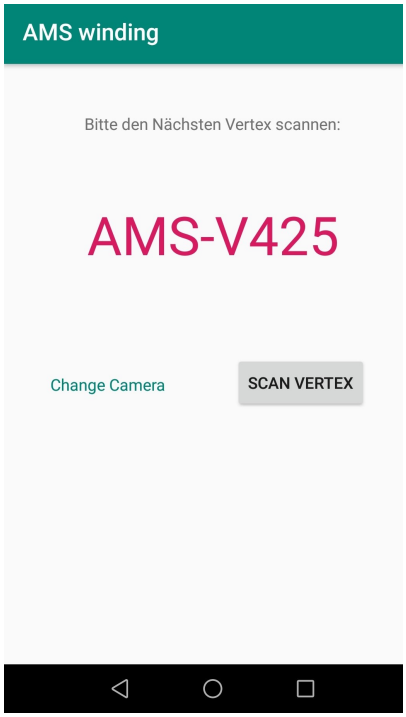

(d) Winding mode.

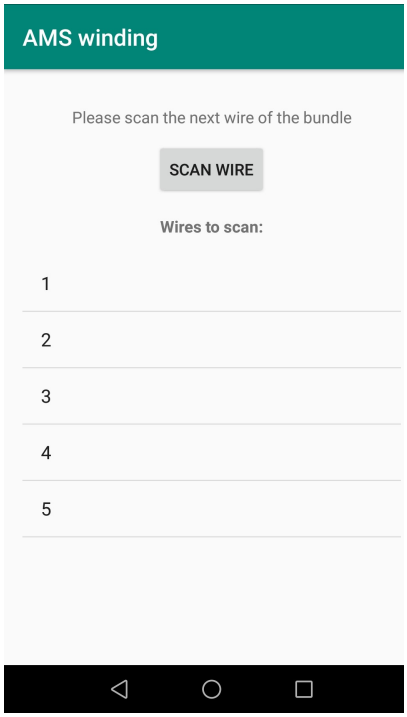

(e) Current checking mode.

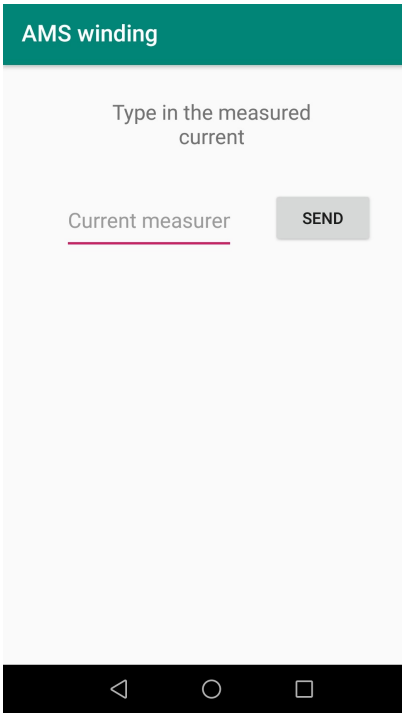

(f) Bundle composition mode.

Figure C.2: Screenshots of the main windows as seen in the AMS installation application.

Bibliography

1. N. Aghanim *et al.* Planck 2018 results - VI. Cosmological parameters. *Astronomy & Astrophysics* **641,** A6 (2020).
2. A. D. Sakharov. Violation of CP invariance, C asymmetry, and baryon asymmetry of the universe. *Soviet Physics Uspekhi* **34,** 392 (1991).
3. C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes & R. P. Hudson. Experimental Test of Parity Conservation in Beta Decay. *Physical Review* **105,** 1413 (1957).
4. B. Povh, K. Rith, C. Scholz, F. Zetsche & W. Rodejohann. Teilchen und Kerne. *Springer-Lehrbuch* (2014).
5. J. H. Christenson, J. W. Cronin, V. L. Fitch & R. Turlay. Evidence for the 2π decay of the K_2^0 meson. *Physical Review Letters* **13,** 138–140 (1964).
6. H. Burkhardt *et al.* First evidence for direct CP violation. *Physics Letters B* **206,** 169–176 (1988).
7. J. R. Batley *et al.* A precision measurement of direct CP violation in the decay of neutral kaons into two pions. *Physics Letters B* **544,** 97–112 (2002).
8. B. Aubert *et al.* Observation of the decay $B^+ \to K^+K^-\pi^+$. *Physical Review Letters* **99** (2007).
9. J. P. Lees *et al.* Study of CP asymmetry in $B^0\bar{B}^0$ mixing with inclusive dilepton events. *Physical Review Letters* **114** (2015).
10. S. W. Lin *et al.* Difference in direct charge-parity violation between charged and neutral B meson decays. *Nature* **452,** 332–335 (2008).
11. R. Aaij *et al.* Observation of CP Violation in Charm Decays. *Physical Review Letters* **122** (2019).
12. R. Aaij *et al.* Updated measurement of decay-time-dependent CP asymmetries in $D^0 \to K^+K^-$ and $D^0 \to \pi^+\pi^-$ decays. *Physical Review D* **101** (2020).
13. M. Starič *et al.* Measurement of D^0–$\bar{D}^0$ mixing and search for CP violation in $D^0 \to K^+K^{\pi^+\pi}$ decays with the full Belle data set. *Physics Letters B* **753,** 412–418 (2016).
14. M. B. Gavela, P. Hernandez, J. Orloff & O. Pène. Standard Model CP-Violation and Baryon Asymmetry. *Modern Physics Letters A* **09,** 795–809 (1994).
15. P. Huet & E. Sather. Electroweak baryogenesis and standard model CP violation. *Phys. Rev. D* **51,** 15 (1995).
16. C. Abel *et al.* Measurement of the Permanent Electric Dipole Moment of the Neutron. *Physical Review Letters* **124,** 81803 (2020).
17. B. Graner, Y. Chen, E. G. Lindahl & B. R. Heckel. Reduced Limit on the Permanent Electric Dipole Moment of Hg 199. *Physical Review Letters* **116** (2016).
18. K. Blum, C. Delaunay, M. Losada, Y. Nir & S. Tulin. CP violation beyond the MSSM: baryogenesis and electric dipole moments. *Journal of High Energy Physics* **2010,** 1–29 (2010).

19. S. J. Huber, T. Konstandin, T. Prokopec & M. G. Schmidt. Electroweak phase transition and baryogenesis in the nMSSM. *Nuclear Physics B* **757,** 172–196 (2006).

20. S. J. Huber, T. Konstandin, T. Prokopec & M. G. Schmidt. Baryogenesis in the MSSM, nMSSM and NMSSM. *Nuclear Physics A* **785,** 206–209 (2007).

21. J. Kozaczuk, S. Profumo, M. J. Ramsey-Musolf & C. L. Wainwright. Supersymmetric electroweak baryogenesis via resonant sfermion sources. *Physical Review D* **86,** 096001 (2012).

22. M. Thomson. *Modern Particle Physics* (Cambridge University Press, 2013).

23. L. I. Schiff. Measurability of nuclear electric dipole moments. *Physical Review* **132,** 2194–2200 (1963).

24. C.-P. Liu & J. Engel. Schiff screening of relativistic nucleon electric-dipole moments by electrons. *Physical Review C* **76,** 028501 (2007).

25. V. V. Flambaum & A. Kozlov. Screening and finite-size corrections to the octupole and Schiff moments. *Physical Review C* **85,** 068502 (2012).

26. W. B. Cairncross *et al.* Precision Measurement of the Electron's Electric Dipole Moment Using Trapped Molecular Ions. *Physical Review Letters* **119,** 153001 (2017).

27. V. Andreev *et al.* Improved limit on the electric dipole moment of the electron. *Nature* **562,** 355–360 (2018).

28. J. J. Hudson *et al.* Improved measurement of the shape of the electron. *Nature* **473,** 493–496 (2011).

29. G. W. Bennett *et al.* Improved limit on the muon electric dipole moment. *Physical Review D* **80,** 052008 (2009).

30. F. Rathmann & N. N. Nikolaev. Electric dipole moment searches using storage rings. arXiv: `1904.13166` (2019).

31. I. B. Khriplovich & S. K. Lamoreaux. *CP Violation without Strangeness* (Springer-Verlag, 1997).

32. R. Golub & S. K. Lamoreaux. Neutron electric-dipole moment, ultracold neutrons and polarized 3He. *Physics Reports* **237,** 1–62 (1994).

33. N. F. Ramsey. A molecular beam resonance method with separated oscillating fields. *Physical Review* **78,** 695–699 (1950).

34. J. H. Smith *et al.* Experimental Limit to the Electric Dipole Moment of the Neutron. *Physical Review* **108,** 120 (1957).

35. J. C. Peng. Neutron electric dipole moment experiments. *Modern Physics Letters A* **23,** 1397–1408 (2008).

36. C. A. Baker *et al.* Improved Experimental Limit on the Electric Dipole Moment of the Neutron. *Physical Review D - Particles, Fields, Gravitation and Cosmology* **98,** 149102 (2006).

37. F. Shapiro. Electric dipole moments of elementary particles. *Uspekhi Fizicheskih Nauk* **95,** 145–158 (1968).

38. N. F. Ramsey. Electric-dipole moments of elementary particles. *Reports on Progress in Physics* **45,** 95–113 (1982).

39. J. M. Pendlebury *et al.* Revised experimental upper limit on the electric dipole moment of the neutron. *Physical Review D - Particles, Fields, Gravitation and Cosmology* **92,** 1–22 (2015).

40. G. Bison *et al.* Comparison of ultracold neutron sources for fundamental physics measurements. *Phys. Rev. C* **95,** 045503 (4 2017).

41. B. Lauss. Ultracold Neutron Production at the Second Spallation Target of the Paul Scherrer Institute. *Physics Procedia* **51.** ESS Science Symposium on Neutron Particle Physics at Long Pulse Spallation Sources, NPPatLPS 2013, 98–101 (2014).

42. C. A. Baker *et al.* Apparatus for measurement of the electric dipole moment of the neutron using a cohabiting atomic-mercury magnetometer. *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **736,** 184–203 (2014).

43. C. Abel *et al.* Optically pumped Cs magnetometers enabling a high-sensitivity search for the neutron electric dipole moment. *Phys. Rev. A* **101,** 53419 (2020).

44. S. Afach *et al.* Dynamic stabilization of the magnetic field surrounding the neutron electric dipole moment spectrometer at the Paul Scherrer Institute. *Journal of Applied Physics* **116,** 84510 (2014).

45. I. Altarev *et al.* Test of Lorentz Invariance with Spin Precession of Ultracold Neutrons. *Physical Review Letters* **103,** 81602 (2009).

46. I. Altarev *et al.* New constraints on Lorentz invariance violation from the neutron electric dipole moment. *EPL* **92,** 51001 (2010).

47. S. Afach *et al.* Constraining interactions mediated by axion-like particles with ultracold neutrons. *Physics Letters, Section B: Nuclear, Elementary Particle and High-Energy Physics* **745,** 58–63 (2015).

48. P.-J. Chiu. *Doctoral Thesis (in preparation)* PhD thesis (2021).

49. C. Abel *et al.* Statistical sensitivity of the nEDM apparatus at PSI to $n - n'$ oscillations. *EPJ Web of Conferences* **219,** 07001 (2019).

50. C. Abel *et al.* Search for axion-like dark matter through nuclear spin precession in electric and magnetic fields. *Phys. Rev. X* **7,** 1–8 (2017).

51. N. J. Ayres *et al.* The design of the n2EDM experiment. *The European Physics Journal C* **81** (2021).

52. M. Rawlik *et al.* A simple method of coil design. *American Journal of Physics* **86,** 1–7 (2018).

53. B. Franke. *Investigations of the Internal and External Magnetic Fields of the Neutron Electric Dipole Moment Experiment at the Paul Scherrer Institute* PhD thesis (2013).

54. M. Rawlik. *Active magnetic shielding and Axion Dark Matter search* PhD thesis (2018).

55. G. Golub & W. Kahan. Calculating the Singular Values and Pseudo-Inverse of a Matrix. *Journal of the Society for Industrial and Applied Mathematics Series B Numerical Analysis* **2,** 205–224 (1965).

56. W. Grosche. *Optimisation of Fluxgate Positions for the n2EDM Experiment at PSI* 2020.

57. R. Hampp. *Investigation on the Active Magnetic Shielding Prototype for the n2EDM at PSI* 2021.

58. J. Bork, H. Hahlbohm, R. Klein & A. Schnabel. The 8-layered magnetically shielded room of the PTB: Design and construction (2001).

59. K. I. Harakawa, G. Kajiwara, K. Kazami, H. Ogata & H. Kado. Evaluation of a high-performance magnetically shielded room for biomagnetic measurement. *IEEE Transactions on Magnetics* **32,** 5256–5260 (1996).

60. D. W. Allan. Statistics of Atomic Frequency Standards. *Proceedings of the IEEE* **54** (1966).

61. R. Hampp. *Investigation on the Active Magnetic Shielding System for the n2EDM Experiment at PSI - Part 2* 2021.

62. EPFL. *SULTAN facility* `https://www.epfl.ch/research/domains/swiss-plasma-center/research/superconductivity/page-97675-en-html/` (2021).

63. M. collaboration. *Development of the COBRA Magnet for the MEG experiment*

64. R. D. Peccei & H. R. Quinn. CP conservation in the presence of pseudoparticles. *Physical Review Letters* **38,** 1440–1443 (1977).

65. R. D. Peccei. The Strong CP Problem and Axions. *Lecture Notes in Physics* **741,** 3–17 (2006).

66. M. Pospelov & A. Ritz. Electric dipole moments as probes of new physics. *Annals of Physics* **318,** 119–169 (2005).

67. S. Weinberg. A New Light Boson? *Phys. Rev. Lett.* **40,** 223–226 (4 1978).

68. F. Wilczek. Problem of Strong P and T Invariance in the Presence of Instantons. *Phys. Rev. Lett.* **40,** 279–282 (5 1978).

69. P. W. Graham & S. Rajendran. Axion dark matter detection with cold molecules. *Physical Review D - Particles, Fields, Gravitation and Cosmology* **84,** 1–13 (2011).

70. E. Witten. Some properties of O(32) superstrings. *Physics Letters B* **149,** 351–356 (1984).

71. J. P. Conlon. The QCD axion and moduli stabilisation. *Journal of High Energy Physics* **2006,** 078 (2006).

72. P. Svrcek & E. Witten. Axions in string theory. *Journal of High Energy Physics* **2006,** 051 (2006).

73. A. Arvanitaki, S. Dimopoulos, S. Dubovsky, N. Kaloper & J. March-Russell. String axiverse. *Physical Review D* **81,** 123530 (2010).

74. P. Arias *et al.* WISPy cold dark matter. *Journal of Cosmology and Astroparticle Physics* **2012,** 013 (2012).

75. D. J. E. Marsh. Axion Cosmology. *Physics Reports* **643,** 1 (2016).

76. W. Hu, R. Barkana & A. Gruzinov. Fuzzy Cold Dark Matter: The Wave Properties of Ultralight Particles. *Physical Review Letters* **85,** 1158 (2000).

77. D. J. E. Marsh & J. Silk. A model for halo formation with axion mixed dark matter. *Monthly Notices of the Royal Astronomical Society* **437,** 2652–2663 (2014).

78. L. F. Abbott & P. Sikivie. A cosmological bound on the invisible axion. *Physics Letters B* **120,** 133–136 (1983).

79. M. Dine & W. Fischler. The not-so-harmless axion. *Physics Letters B* **120,** 137–141 (1983).

80. J. Preskill, M. B. Wise & F. Wilczek. Cosmology of the invisible axion. *Physics Letters B* **120,** 127–132 (1983).

81. Y. V. Stadnik. *Manifestations of Dark Matter and Variations of the Fundamental Constants in Atoms and Astrophysical Phenomena* (Springer International Publishing, 2017).

82. P. A. Zyla *et al.* Review of particle physics. *Progress of Theoretical and Experimental Physics* **2020,** 1–2093 (2020).

83. B. Bozek, D. J. E. Marsh, J. Silk & R. F. G. Wyse. Galaxy UV-luminosity function and reionization constraints on axion dark matter. *Monthly Notices of the Royal Astronomical Society* **450,** 209–222 (2015).

84. H.-Y. Schive, T. Chiueh, T. Broadhurst & K.-W. Huang. Contrasting Galaxy Formation from Quantum Wave Dark Matter, ψDM, with ΛCDM, using Planck and Hubble Data. *The Astrophysical Journal* **818,** 89 (2016).

85. P. S. Corasaniti, S. Agarwal, D. J. E. Marsh & S. Das. Constraints on dark matter scenarios from measurements of the galaxy luminosity function at high redshifts. *Physical Review D* **95,** 083512 (2017).

86. P. Sikivie. Experimental Tests of the "Invisible" Axion. *Phys. Rev. Lett.* **51** (1983).

87. S. J. Asztalos *et al.* SQUID-Based Microwave Cavity Search for Dark-Matter Axions. *Phys. Rev. Lett.* **104** (2010).

88. V. Anastassopoulos *et al.* New CAST limit on the axion-photon interaction. *Nature Physics* **13,** 584–590 (2017).

89. K. Ehret *et al.* Resonant laser power build-up in ALPS-A "light shining through a wall" experiment. *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **612,** 83–96 (2009).

90. P. W. Graham & S. Rajendran. New observables for direct detection of axion dark matter. *Physical Review D - Particles, Fields, Gravitation and Cosmology* **88,** 1–13 (2013).

91. G. G. Raffelt. Astrophysical methods to constrain axions and other novel particle phenomena. *Physics Reports* **198,** 1–113 (1990).

92. G. G. Raffelt. Astrophysical Axion Bounds. *Lecture Notes in Physics* **741,** 51–71 (2008).

93. K. Blum, R. T. D'Agnolo, M. Lisanti & B. R. Safdi. Constraining axion dark matter with Big Bang Nucleosynthesis. *Physics Letters, Section B: Nuclear, Elementary Particle and High-Energy Physics* **737,** 30–33. eprint: 1401.6460 (2014).

94. A. Schuster. On the investigation of hidden periodicities with application to a supposed 26 day period of meteorological phenomena. *Journal of Geophysical Research* **3,** 13 (1898).

95. J. D. Scargle. Studies in astronomical time series analysis. II - Statistical aspects of spectral analysis of unevenly spaced data. *The Astrophysical Journal* **263,** 835 (1982).

96. J. T. VanderPlas. Understanding the Lomb-Scargle Periodogram. *The Astrophysical Journal* **236,** 16 (2017).

97. J. T. VanderPlas & Ž. Ivezić. Periodograms for Multiband Astronomical Time Series. *The Astrophysical Journal* **812,** 18 (2015).

98. J. W. Richards *et al.* Construction of a Calibrated Probabilistic Classification Catalog: Application to 50k Variable Sources in the All-Sky Automated Survey. *The Astrophysical Journal* **203,** 32 (2012).

99. C. Abel *et al.* nEDM experiment at PSI: data-taking strategy and sensitivity of the dataset. *EPJ Web Conf.* **219,** 02001 (2018).

100. C. Abel *et al.* Magnetic-field uniformity in neutron electric-dipole-moment experiments. *Physical Review A* **99,** 042112 (2019).

101. J. M. Pendlebury *et al.* Geometric-phase-induced false electric dipole moment signals for particles in traps. *Physical Review A - Atomic, Molecular, and Optical Physics* **70,** 1–20 (2004).

102. C. Abel *et al.* Mapping of the magnetic field to correct systematic effects in a neutron electric dipole moment experiment. arXiv: 2103.09039 (2021).