



JOHANNES GUTENBERG
UNIVERSITÄT MAINZ

Commissioning and First Results of the SUPAX Experiment and Prospects for Detecting High-Frequency Gravitational Waves with GravNet

Dissertation

zur Erlangung des Grades “Doktor der Naturwissenschaften” am Fachbereich
Physik, Mathematik und Informatik der Johannes Gutenberg-Universität in Mainz

Tim Schneemann

geb. in Gotha
Mainz, den 07.07.2026

1. Gutachter: Univ.-Prof. Dr. Matthias Schott
2. Gutachter: Univ.-Prof. Dr. Michael Wurm
Datum der mündlichen Prüfung: 31.08.2026

Dieses Werk ist lizenziert unter einer Creative Commons Namensnennung 4.0 International Lizenz (CC BY 4.0).

Abstract

The axion is a hypothetical particle originally proposed to solve the strong CP problem. Incidentally, its properties also make it a very promising cold dark matter (CDM) candidate. In haloscope experiments, CDM axions are searched for via their conversion into photons in the presence of a strong magnetic field, where the resulting signal is enhanced by a resonant cavity in a cryogenic environment and detected using a low-noise readout chain.

This thesis reports on the first operation and results of the SUPAX experiment, a superconducting axion search experiment, and the proposed GRAVNET detector network, which intends to utilize existing axion haloscope infrastructure to search for High-Frequency Gravitational Waves (HFGW) originating from another dark matter candidate: Primordial Black Holes (PBH).

During several experimental runs, SUPAX has studied the feasibility of Niobium Nitride (NbN) as a superconducting coating with and without the application of an external magnetic field, and provided an exclusion limit for an axion mass range of ± 0.72 MHz around 8.26548 GHz ($m_a \approx 34.18 \mu\text{eV}$) down to $g_{a\gamma\gamma} > 1.67 \cdot 10^{-13} \text{ GeV}^{-1}$. This limit is recast into a limit on dark photons (DP), another CDM candidate, yielding an exclusion down to $\chi > 1.97 \cdot 10^{-14}$. Another experimental run without an external magnetic field resulted in a limit on the DP coupling down to $\chi > 8.78 \cdot 10^{-15}$ around 8.4712 GHz ($m_{\text{DP}} \approx 35 \mu\text{eV}$).

Finally, a software package has been developed to simulate the output of the current data acquisition chain of GRAVNET for HFGW signals originating from PBH mergers. This way, analysis techniques can be tested and developed before the experimental setup is completed.

Table of contents

1. Introduction	1
2. The Case for Dark Matter	4
2.1. Dynamical Evidence for Dark Matter	4
2.2. Gravitational Lensing Observations	6
2.3. Cosmic Microwave Background and the Λ CDM Model	7
2.4. Emerging Limitations of Λ CDM and Alternative Ideas	9
2.5. Candidates for Dark Matter	10
2.6. Types of Dark Matter Detection Experiments	14
3. Theory	18
3.1. Dark Matter Halo	18
3.2. Axion Physics	19
3.3. Dark Photon Physics	31
3.4. RF-Cavity Based Haloscopes	36
3.5. Study of Superconducting RF-Cavities	44
4. The SUPAX Experiment	47
4.1. Experimental Setup	47
4.2. Data Acquisition	49
4.3. The Cavity	49
4.4. Tracking of Environmental Conditions during Data Taking	50
4.5. Temperature and Cooling Management	51
5. Cavity Characterisation	55
5.1. Simulations in CST Studio	55
5.2. Methodology of Characterisation	56
5.3. Performance of a NbN Coated Cavity in a Strong Magnetic Field	59
6. Axion and Dark Photon Analysis	63
6.1. System Noise Temperature	63
6.2. Signal attenuation	65
6.3. Analysis Overview	68
6.4. Constructing the Grand Unified Spectrum	70
6.5. Savitzky-Golay Filter Optimisation	72
6.6. Limit Setting	76
6.7. Results	84
7. High-Frequency Gravitational Waves	97
7.1. Theoretical Background	97

7.2. Gravnet	99
7.3. Detecting HFGW in RF-Cavities	101
8. Future Prospects	114
8.1. Further SRF Cavity Development	114
8.2. That Single-Bin Power Excess in the Dedicated Dark Photon Measurement .	114
8.3. Achievable Axion Limits with One Year of Data Acquisition	115
9. Conclusion	119
Bibliography	121
Appendix	132
A. Normalisation of Continuous and Discretised Lineshape	132
B. Computationally Efficient Limit Setting	134
C. Characterisation Data for the NbN Superconducting S-Curve	137
D. Cryostat Operation	138

1. Introduction

The field of particle physics is one of the, if not the, most significant and researched fields of physics of the 21st century. Over the past decades, the **Standard Model** of particle physics (**SM**), developed in the 1960s and 70s, has changed very little and been repeatedly shown to be a precise description of the majority of fundamental physics. It explains three of the four fundamental forces of physics: the electromagnetic (EM), the weak and the strong force. These forces are described in quantum electrodynamics (QED), which merges with the weak force into one single electroweak interaction above a unification energy, and quantum chromodynamics (QCD). The fourth fundamental force, gravity, remains unexplained by the Standard Model. The observation of the Higgs boson in 2012 at the Large Hadron Collider (LHC) constitutes the experimental confirmation of the existence of all particles that make up the Standard Model. This also qualified the LHC to be the main contributor to the field of particle physics in modern times. The constituents of the SM can be shown in a single image (see Fig. 1.1).




Figure 1.1.: The Standard Model of particle physics.¹

With the content of the presented elementary particles all ordinary matter can be described. However, cosmology has presented a problem which the SM cannot solve. For nearly a century, evidence has been mounting (starting with Fritz Zwicky's paper in 1933²) that 'ordinary matter', everything explained in the SM, most likely only accounts for a

1. Introduction

fraction of everything that is in our universe. The dominant theory combining (nearly) all observations is called **Lambda Cold Dark Matter** (Λ CDM).




Figure 1.2.: Constituents of the universe according to Planck results on Λ CDM.³

While the Λ , representing dark energy and said to be responsible for the expansion of the universe in this theory, is the subject of a whole field of research, we will concentrate on the CDM part - cold dark matter. Cold in this context means that the particles move at non-relativistic speeds while other dark matter theories also allow for warm or hot dark matter, so particles moving at moderate or highly relativistic velocities. As can be seen in Fig. 1.2, dark matter accounts for roughly 27% of the universe, more than five times the amount of what is described by the Standard Model. The problem of Λ CDM is that it reveals the lack of understanding we have of what $\approx 95\%$ of the universe is actually made of. While the composition seen in Fig. 1.2 is relatively clear the question of how we can describe the constituents of the two sectors outside of ordinary matter is still unanswered, hence they are named ‘dark’. Dark in this context means that neither dark matter nor dark energy emit, absorb or reflect any light in the way ordinary matter does. This makes experimental detection of such particles very challenging. Nevertheless, theorists have come up with several possibilities of how to describe dark matter, many including new particles that are not part of the Standard Model. Theories including such hypothesised particles are hence known as BSM - *Beyond Standard Model* - theories. Given the constant inflow of new ideas and theories from theoretical physics, it is the task of experimental physics to test if these theories can be proven right or wrong by experiment. One such experiment is SUPAX, a superconducting axion search experiment and the central subject of this thesis.

The work presented here investigates the potential of a new superconducting microwave-cavity haloscope to search for dark matter particles and related new-physics signals in the GHz regime. Its central goal is to establish the SUPAX experiment as a contributor to coupling limits set for axion and dark-photon candidates with rest mass in the μeV regime. The relevant cavity and cryogenic measurement systems are characterised, and data from several experimental runs are analysed to derive exclusion limits on axion-photon couplings and dark-photon kinetic mixing, and the performance of NbN is studied as a possible superconducting coating in this context. In addition, the thesis explores the broader use of haloscope infrastructure for the detection of high-frequency gravitational waves. To this end,

it develops a simulation framework for gravitational-wave signals and enables first steps toward signal analysis for the proposed GRAVNET network. The following chapters introduce the theoretical motivation for these searches, describe the experimental setup and its characterisation, present the resulting dark matter limits, and discuss the prospects for future haloscope-based searches.

2. The Case for Dark Matter

This chapter discusses the most compelling evidence for the existence of dark matter. It begins with the discovery of the need of dark matter due to observations of the velocity of rotating galaxies and additional hints toward dark matter such as gravitational lensing. It then turns to Λ CDM, the standard model of cosmology, the observations on which it is based, as well as its limitations and alternatives. The most relevant dark matter candidates are introduced, and finally the different types of detection experiments are discussed.

2.1. Dynamical Evidence for Dark Matter

The kickoff into the search for dark matter as we know it today was Fritz Zwicky's paper 'Die Rotverschiebung von extragalaktischen Nebeln'² ('The redshift of extragalactic nebulae'). In his paper, Zwicky studied the Coma cluster, a cluster of more than 1,000 galaxies in approximately 100 Mpc distance from Earth. He observed that the luminous (one might call it visible) mass was by far insufficient to cause the gravitational effects observable in the cluster. Specifically, the galaxies within the cluster were moving much faster than could be explained by the visible mass alone. He concluded that there must be a significant amount of unseen - *dark* - matter contributing to the total mass of the cluster. He reached this conclusion by applying the Virial theorem to the cluster. For this, he assumed that the cluster reached a mechanically stationary state, so the Virial theorem states

$$\bar{\epsilon}_k = -\frac{1}{2}\bar{\epsilon}_p \quad (2.1)$$

where $\bar{\epsilon}_k$ and $\bar{\epsilon}_p$ describe the average kinetic and potential energy of the system. Assuming that all mass of the cluster, consisting of 800 individual nebulae of mass $\approx 10^9 M_\odot$ ($M_\odot$ = solar mass) and a radius of the entire cluster of $\approx 10^6$ ly (ly = light years), is distributed uniformly, this theorem results in

$$\sqrt{v^2} = 80 \frac{\text{km}}{\text{s}} \quad (2.2)$$

while he observed Doppler speeds of $\approx 1000 \frac{\text{km}}{\text{s}}$. This leads to the conclusion that the mass density must be roughly 400 times larger than observed. Even when dropping the assumption of a stationary equilibrium and therefore assuming that all the potential energy appears as kinetic this only accounts for a factor of 1/2 in equation (2.1) "[...] and the need for an enormous density of dark matter remains."²

While there had been talk about dark matter before this work, the earliest explicit mention of "dark matter" has been made by Poincaré in 1906,⁴ most astronomers of the time believed dark matter was likely to be made from faint stars, i.e. ordinary matter, rather than a completely new type of matter. Also, Zwicky was the first to claim that there is significantly

2.1. Dynamical Evidence for Dark Matter

more dark than luminous matter as at the time there was a consensus that there “is probably less [dark matter] than the total mass of visible stars, possibly much less.”⁵

In 1937, Zwicky published another paper in which he built on his previous findings.⁶ Besides continued calculation using the Virial theorem, confirming his prior work, he extended his analysis of the mass-to-light ratio, comparing the total mass of galaxies and clusters to their luminosity. He found that for clusters this ratio was much higher than for individual galaxies, implying a significant amount of dark matter. Following his findings, Zwicky introduced the concept of using gravitational lensing as a method to detect dark matter (see sec. 2.2). He proposed that massive objects like galaxy clusters could bend light from background objects, a phenomenon predicted by general relativity and thereby reveal the presence of unseen mass. Though not observed at the time, this idea was pioneering and has since become a crucial tool in modern astrophysics, as we shall see in the next section.

However, before we get to more recent evidence of dark matter - the observation of gravitational lensing - we shall see another example of galactic dynamics suggesting the presence of dark matter: the rotation speeds of the Andromeda nebula.




Figure 2.1.: Rubin and Ford’s findings on the flat rotation curve of the Andromeda nebula updated with recent data.^{7,8} To reach the observed rotational velocities at large radii the need for a dark matter halo is modeled. Figure taken from [9].

In 1970 Vera Rubin and Kent Ford found that the rotation curve of the Andromeda Galaxy remained flat, i.e. constant, at large radii, see Figure 2.1.¹⁰ Instead of decreasing with distance from the centre, as expected from Keplerian motion (where the velocity should drop off inversely with the square root of the distance, $v_{\text{rot}} \propto 1/\sqrt{R}$), the rotational velocity stayed approximately constant beyond the visible edge of the galaxy. The flat rotation curve suggested that the mass of the Andromeda Galaxy continued to increase well beyond the extent of the visible matter (mainly stars and gas). This implied that a significant amount of mass was present in the outer regions of the galaxy that did not emit light, contributing

2. The Case for Dark Matter

to the observed rotational velocities. This result provided strong observational evidence supporting Fritz Zwicky's earlier hypothesis about dark matter. While Zwicky had inferred the presence of dark matter in galaxy clusters, Rubin and Ford's observations confirmed that individual galaxies also contained significant amounts of dark matter.

2.2. Gravitational Lensing Observations

Let us now come to the observations of what Zwicky proposed as a direct consequence of general relativity (GR) and dark matter: gravitational lensing. In GR, light does not travel in a straight line in the presence of mass. Instead, mass curves spacetime, and light follows these curved geodesics. Any mass concentration between an observer and a background galaxy therefore acts like a lens, slightly deflecting the light. The amount and pattern of distortion depend only on the total gravitational potential along the line of sight, not on how much of that mass happens to be luminous or *dark*. By comparing the observed shapes, positions and brightness of background galaxies with what they would look like without the lens, one can reconstruct the projected mass distribution of the lens and directly test whether luminous baryons are sufficient to explain the observed gravitational effect.

One typically distinguishes between weak and strong lensing. In the weak lensing regime, typical for large-scale structures like galaxy clusters, the distortions of individual background galaxies are tiny, but statistically coherent over large areas of the sky. By averaging over many galaxies one can obtain a map of projected mass. This method has been applied to merging clusters such as 1E0657-558 (the Bullet Cluster) in a paper by Douglas Clowe et al.¹¹ There, optical observations (see Fig. 2.2) show two concentrations of collisionless galaxies, while X-Ray data reveal that the hot intracluster plasma, the dominant baryonic component, has been slowed down and left behind during the collision. The weak lensing reconstruction, however, shows that the main mass peaks are centred on the galaxies and significantly offset from the X-Ray gas.

In a purely baryonic universe the gravitational potential would be expected to follow the dominant visible mass (the plasma), which it clearly does not. This spatial separation between the concentrations of baryons and gravitational potential is therefore difficult to reconcile with modifications to gravity* alone and strongly supports the presence of an additional collisionless dark matter component.

In contrast to these weak lensing effects, strong gravitational lensing occurs in regions where the projected mass density is high enough that light rays are bent by large angles, producing multiple images, giant arcs or even nearly complete Einstein rings of background galaxies. In such systems, the observed image geometry tightly constrains the enclosed mass on scales of a few kiloparsecs. Both individual galaxies and cluster cores that act as strong lenses systematically require far more mass than is present in stars and hot gas alone, and this excess mass must be extended and roughly collisionless. Together with weak lensing in merging clusters, as discussed above, the universe's gravitational potentials are dominated by non-luminous matter, consistent with the proposition of dark matter.

*such as Modified Newtonian Dynamics (MOND) for example




Figure 2.2.: Shown is the merging Bullet Cluster, Figure taken from [11]. The green contours can be interpreted as the mass distribution, obtained by gravitational lensing analysis. The right image shows that the mass distribution does not coincide with the centre and the distribution of the plasma, which constitutes the majority of the observable matter. In the image on the left one can see how the centre of the mass, indicated by white contours corresponding to 1, 2 and 3 σ confidence and the location of the plasma centre, indicated by turquoise + symbols, are clearly not aligned.

2.3. Cosmic Microwave Background and the Λ CDM Model

In 1964 two relevant parts of the same story played out no more than 60 km apart from each other. Arno Penzias and Robert Wilson were working with a very sensitive horn antenna at a Bell Labs facility in Holmdel, New Jersey. In order to get a proper measurement they needed to get rid of any noise in their signal, including cooling the radio receiver itself to 4 K to reduce thermal noise. Once everything was set up they detected a noise signal 100 times larger than expected. It was detected at 4080 MHz and “isotropic, unpolarised, and free from seasonal variations”.¹² Any efforts to get rid of this background signal, including the relocation of some pigeons that had nested in the antenna, were unsuccessful, leaving the two physicists quite puzzled. Unbeknownst to them, a group of physicists at the nearby Princeton University reasoned that the Big Bang must have released tremendous amounts of radiation that should now be red-shifted into the microwave regime. Once Penzias heard that someone had independently found a theory that predicted exactly the spurious signal he had observed, he contacted the researchers at Princeton. In 1965 they published their results and were rewarded the Nobel Prize in 1978 for the discovery of what is known today as the Cosmic Microwave Background (CMB).

Since its discovery, the CMB has been studied extensively with the most recent results came from the Planck satellite in 2018. As previously mentioned, around the turn of the millennium the Λ CDM model became the Standard Cosmological Model and the CMB (both its existence and what it reveals) has been by far the major contributor of experimental evidence to the model. The model interprets the CMB as relic radiation from around 380,000 years after the Big Bang. Prior to this point, baryons and photons formed a dense and hot plasma, a particle soup in which photons were continuously scattered by free electrons. As the universe expanded and cooled to ≈ 3000 K, protons and electrons combined to form

2. The Case for Dark Matter

hydrogen, a process known today as *recombination*. The sudden drop in free electrons meant that free photons were no longer scattered, effectively decoupling radiation from matter. The universe had become transparent and these photons have since propagated nearly unimpeded, stretching their wavelengths from the infrared to the microwave regime due to cosmic expansion, also called redshift. Today, its spectrum is an almost perfect blackbody at 2.7 K, with anisotropies at the 10^{-5} level, shown in Fig. 2.3. These anisotropies




Figure 2.3.: Angular power spectrum showing the anisotropies of the CMB with data taken by Planck¹³ in red and the best fit to the theoretical spectrum predicted by Λ CDM in blue (top). Below, the residuals of data and best fit are shown. Note the scale change at $\ell = 30$. Plot taken from [13].

encode the density and velocity fluctuations of the primordial plasma, giving us the earliest image of the universe at an age of approximately 380,000 years. Studying the anisotropies reveals information about the densities of baryonic matter, dark matter and dark energy. According to Λ CDM, before recombination the photons and baryons behaved as a single, tightly coupled fluid (the aforementioned particle soup) executing acoustic oscillations inside gravitational potentials seeded mostly by dark matter. The pattern of these oscillations is preserved in the CMB as a sequence of acoustic peaks. Baryons add inertia to the plasma which enhances the amplitude of compressional phases of the oscillations relative to the rarefaction phases. This odd-even peak asymmetry effectively measures the baryon loading in the fluid. By fitting the asymmetry with the Λ CDM model predictions, one determines the baryon density today to be $\Omega_b \approx 0.05$, meaning that only about five percent of the universe is made of ordinary matter. Dark matter, on the other hand, shapes the gravitational potentials in which the plasma oscillates. Increasing the dark matter content deepens these potentials, alters the epoch when matter first dominated over radiation, and shifts the relative heights and locations of the acoustic peaks in a calculable way. Matching these features to observation shows that the dark matter density must be about $\Omega_c \approx 0.25$, or roughly a quarter of the universe. In addition to determining the matter content, the CMB also provides strong evidence for dark energy through its constraints on the geometry and total energy density of the universe. The angular positions of the acoustic peaks indicate that

the universe is extremely close to spatially flat. In a flat universe the total mass-energy density must equal the critical density, yet the combined contribution of baryonic and dark matter amounts to only about thirty percent. This discrepancy points to the presence of an additional, smoothly distributed component that does not cluster like matter. Introducing such a component - dark energy - raises the total density to the critical value implied by the observed geometry. This yields a dark energy density of $\Omega_\Lambda \approx 0.70$ as per the results of the Planck satellite.¹³

2.4. Emerging Limitations of Λ CDM and Alternative Ideas

Recent years have uncovered increasing inconsistencies in the theory of Λ CDM in its current form. Here we will discuss some of the most prominent ones, possible solutions, and alternative ideas that do not involve dark matter such as Modified Newtonian Dynamics (MOND).

Although Λ CDM continues to provide an excellent description of many key cosmological observations, including the CMB and the distribution of large-scale structures, it is becoming clear that it may not be the final word on the dynamics of the cosmos. Although not unusual in physics, one of the main issues is that neither dark matter nor dark energy has been identified in the laboratory. Various dark matter searches have pushed to extremely small interaction strengths without finding a conclusive signal and the dark energy observed by Planck in the CMB is many orders of magnitude smaller than predicted by QFT. This discrepancy is commonly known as Vacuum Catastrophe.

Alongside these conceptual concerns, experimental observations have emerged over the past few years that put Λ CDM under pressure. The most widely discussed is the Hubble tension. It describes the discrepancy between the value obtained for the Hubble constant H_0 by analysing the early universe, primarily through the CMB, and the value obtained by studying late-universe phenomena such as supernovae and Cepheids. With the constant development in sensitivity there are now two camps, which obtain different values with a greater than 4σ significance and non-overlapping error margins. The most prominent representative of the early-universe, Λ CDM supporting camp is Planck¹³ which obtains $H_0 = (67.4 \pm 0.5) \text{ km s}^{-1} \text{ Mpc}^{-1}$ while the opposing camp is spearheaded by analyses of the Hubble Space Telescope¹⁴ claiming $H_0 = (73.2 \pm 1.3) \text{ km s}^{-1} \text{ Mpc}^{-1}$. Whether this reflects unrecognised systematics or genuinely new physics is still an open question, but the persistence of the discrepancy across different observational strategies has made it difficult to dismiss.

These issues, among others, have stimulated a broad landscape of theoretical proposals. One of the most ambitious is the previously mentioned Modified Newtonian Dynamics, altering the gravitational theory itself. MOND, along with a multitude¹⁵ of modified theories of gravity (MOG), seeks to explain galactic dynamics without the need for dark matter, while other MOG theories attempt to account for cosmic acceleration without invoking a cosmological constant, such as the Λ in Λ CDM. While these ideas can address certain specific discrepancies, none has yet matched the comprehensive success of Λ CDM across all available data, particularly the CMB and gravitational lensing. Importantly, even if Λ CDM were eventually replaced, the need for some form of non-luminous matter would almost certainly remain. The previously mentioned phenomena such as galaxy rotation curves and lensing,

2. The Case for Dark Matter

but also other effects not discussed in detail here such as the formation and dynamics of galaxies^{16,17} require mass beyond what is visible, regardless of the fate of the Λ term.

2.5. Candidates for Dark Matter

Having established the overwhelming astrophysical and cosmological evidence for the existence of dark matter, the natural next question is what it is made of. While the observations discussed in the previous section point toward a non-baryonic component that interacts predominantly through gravity, they place relatively weak constraints on what dark matter is actually made of. Aside from requiring non-zero mass, stability on cosmological timescales, and the absence of strong interactions with the SM, a wide range of possibilities remains open. As a result, dark matter does not correspond to a single well-defined theoretical prediction but rather to a large landscape of candidate particles and compact objects arising from different *new physics* theories.




Figure 2.4.: Schematic overview of various dark matter models. Figure from T. Tait¹⁸

This theoretical diversity is illustrated in Fig. 2.4, which highlights the sheer multitude of ideas that have been proposed over the past decades. Dark matter candidates span many orders of magnitude in mass, interaction strength, and production mechanism, ranging from astrophysical objects such as massive compact halo objects or primordial black holes to elementary particles motivated by extensions to the Standard Model. In the following, a selection of the most prominent dark matter candidates is reviewed. The discussion begins with neutrinos as representative of hot (or warm) DM which helps clarify the physical requirements for a successful DM candidate. Subsequently, other classes of candidates are introduced including MACHOs, WIMPs, primordial black holes, dark photons and finally axions and axion-like particles, the latter three being studied in detail in this thesis.

2.5.1. Neutrinos

Neutrinos were (and in some sense still are) an attractive dark matter candidate for a very simple reason: they already have many of the right qualitative features, without the need to introduce new particles or extend the Standard Model. They are electrically neutral, extremely long-lived on cosmological timescales, and interact only weakly with ordinary matter. Once neutrino oscillations established that neutrinos are nonzero in mass, it also became clear that they must contribute some non-baryonic mass density to the universe.

However, the same facts that make Standard Model (also called *active*) neutrinos compelling also make them unable to account for the dominant dark matter component in Λ CDM. The key issue is not whether neutrinos gravitate, they do, but when they slow down enough to behave like pressureless matter during the era when cosmic structure grows. A useful dividing line is whether a DM candidate is still relativistic at the time when perturbations on galactic scales first enter the horizon. Active neutrinos decouple very early (around the MeV epoch, or roughly 1 s after the Big Bang) and then essentially free-stream with a thermal momentum distribution whose temperature redshifts, just like the CMB discussed in the previous section. For realistic neutrino masses, their momenta remain far above their rest mass at early times relevant for small-scale structure, so they behave as hot dark matter: high-momentum, relativistic particles.

This becomes a concrete problem once one tries to connect it to the Λ CDM picture of structure formation. As discussed in sec. 2.3, cold dark matter perturbations provide gravitational potential wells into which baryons can later fall once photon pressure no longer prevents collapse. If instead the dominant dark matter were hot, these observed small-scale fluctuations would be erased before they ever had a chance to grow, suppressing early galaxy formation and shifting structure formation toward a “top-down” sequence (large objects collapse first, then fragment). Numerical studies and observed small-scale structure favor the opposite, “bottom-up” hierarchy that comes naturally from cold dark matter.

Even leaving early universe dynamics aside, there is an additional, more local argument against active neutrinos making up galactic halos. Because neutrinos are fermions, packing enough of them into the Milky Way halo (more on the DM Halo model in sec. 3.1) to reproduce the measured local dark matter density implies a phase-space density that forces the typical neutrino velocities to be very large given their tiny masses. Put differently: if you try to explain the halo with $\sim$ eV-mass fermions, the required velocity dispersion comes out comparable to (and exceeding) galactic escape speeds, so the halo would not remain gravitationally bound in the way we observe.

Finally, cosmology constrains the total active neutrino mass density rather directly, because the cosmic relic neutrino number density is tightly fixed by thermal history while neutrinos are relativistic. Planck’s CMB analysis limits the sum of neutrino masses to well below an eV ([13] specifies $\sum m_\nu < 0.12$ eV), implying that active neutrinos are contributing $\lesssim 1\%$ of the total matter density.

As we have emphasised throughout, the above arguments apply specifically to *active* neutrinos it is only logical that the failure of these active neutrinos as a main dark matter candidate motivates a new neutrino-like state, called *sterile* neutrinos. A sterile neutrino is a neutral fermion, which does not participate in the weak interaction the way active neutrinos do. Its only generic connection to the SM can be through mixing with active neutrinos.¹⁹

From a theoretical perspective, such states arise naturally once one extends the SM to

2. The Case for Dark Matter

account for the observed neutrino masses. In many frameworks, sterile (often right-handed) neutrinos play a central role in generating the small masses of active neutrinos, making their existence well-motivated even independently of DM considerations. What distinguishes sterile neutrinos from their active counterparts in a cosmological context is that their interaction history and abundance are not fixed by thermal equilibrium in the early universe. As a result, both their mass scale and their phase-space distribution can differ significantly from those of active neutrinos.[20, 21]

This opens up the possibility that sterile neutrinos could behave very differently with respect to structure formation. In particular, sterile neutrinos with masses in the keV range can become non-relativistic, i.e. cold, much earlier than active neutrinos. Depending on how they are produced in the early universe, their momentum distribution can be sufficiently cold that they cluster on galactic and sub-galactic scales.¹⁹ Such scenarios demonstrate that neutrino-like particles are not categorically excluded as dark matter, provided they lie beyond the Standard Model, though they are subject to strong cosmological and astrophysical constraints.

2.5.2. MACHOs

Massive compact halo objects (MACHOs) have been proposed to make up DM in the 1990s. The name MACHO includes several types of objects such as brown, white, and red dwarfs, neutron stars and even black holes in some contexts. MACHO searches mainly use gravitational microlensing to detect potential candidates. The idea is that of a MACHO passing by a star in such a way that the gravity around the MACHO distorts the light of the star, causing the star to appear brighter - *lensed* - on Earth. Since this is a phenomenon purely based on gravity, it does not need the MACHO to emit or reflect any light at all which makes it perfectly suited for the search for DM. Around the turn of the millennium, the two main groups searching for MACHOs via gravitational lensing nearly completely ruled out MACHOs as a possible candidate to explain DM. Although the MACHO collaboration claimed to have found a signal which may account for 20% of DM in our galaxy,²² this claim was not confirmed by the EROS-2 collaboration²³ which did not find such a signal, even with a higher sensitivity than the aforementioned. Furthermore, later Hubble Space Telescope searches for such a signal have turned out to be fruitless for the claim made by the MACHO collaboration.

2.5.3. WIMPs

With MACHOs more or less ruled out* as a prime candidate, the scientific community shifted its attention towards other candidates, one of the main ones being Weakly Interacting Massive Particles (WIMPs). These particles appear in many BSM theories such as SUSY (supersymmetric) theories or in theories with extra dimensions. The mass range of WIMPs is generally assumed to be on the GeV - TeV scale. While they have been produced in the early universe they are now non-relativistic which qualifies them as CDM. Although

*Naturally there still is a minority standing against the trend and continuing their search for MACHOs. However, the overwhelming majority of DM-searching physicists have moved on to different candidates. This trend can also be observed in the shift of funding since the 90s with the bulk going to large WIMP experiments (partly because of their size, see XENONnT or LUX-ZEPLIN) and a rapidly growing share attributed to axion searches such as ADMX.

the so-called ‘WIMP miracle’ correctly predicts the dark matter abundance of $\approx 25\%$ by its natural freeze-out mechanism it has recently lost its appeal by the lack of evidence for WIMPs by large experiments such as XENON,²⁴ to name just one of many. While they have found a rare, never before observed, isotope decay - that of ^{124}Xe with an estimated half life of 1.8×10^{22} years²⁵ - the WIMP, in the words of XENON Editorial Board member Marc Schumann, “despite of many years of effort, was not yet discovered.”²⁶

2.5.4. Primordial Black Holes and Similar Sources of High-Frequency Gravitational Waves

Mergers of Primordial Black Holes (PBH*), as well as other exotic compact objects such as boson and fermion stars,^{27–30} gravitino stars³¹ or dark matter blobs,³² can result in gravitational waves (GW), just like larger objects, such as ordinary black holes, can. Gravitational waves from large objects have been observed in large detectors such as LIGO with frequencies in the range of Hz – kHz. However, much lighter objects, such as those listed above, are expected to emit gravitational waves in the GHz regime. It has been shown³³ that, in the presence of a strong magnetic field, these High Frequency Gravitational Waves (HFGW) behave like an effective current in resonant cavities, allowing for the search for them with existing, dark matter searching, RF-cavity setups. Besides PBH mergers there might be more sources of GWs in the GHz regime, such as axion super-radiance and even a stochastic background, comparable to the CMB. This will be discussed in more detail in section 7.1.

2.5.5. Dark photons

The dark photon (DP) is a vector particle and considered the boson of an additional U(1) symmetry to the Standard Model. It is hard to pinpoint the first work or person who came up with the idea since many papers refer to roughly the same idea under several different names such as the paraphoton, hidden-sector photon, secluded photon, U-boson or, as it will be referred to in this thesis, dark photon. The name originates, as can be expected, from the fact that, being considered dark matter, it is *dark*, i.e. not charged under SM gauge groups. The *photon* part comes from the fact that the only expected interaction is via kinetic mixing with the SM photon, since it has the same quantum numbers. While there is extensive theory behind massless and massive dark photons we will only consider massive dark photons here. This is due to the fact that the massless case does not couple directly to any of the SM currents through operators of dimension four or lower and in the massive case the mass of the DP is a very good estimate of the energy (or frequency) of the resulting SM photon, potentially detected via a cavity haloscope. The theory of dark photons will be discussed in further detail in section 3.3.

2.5.6. Axions and ALPs

Having discussed some other candidates, let us get to the most promising candidate for DM at the moment - the axion. Taking a step back from the actual hands-on physics, the main difference between the axion and the rest of the candidates is that the theory behind the ‘original’ axion - the so-called QCD axion - was not crafted in an attempt to create a

*More detailed introduction of PBHs in chapter 7

2. The Case for Dark Matter

DM candidate. Quite to the contrary the axion fell out of a possible solution of the strong CP-problem (more on that in section 3.2.1) as a byproduct and it did not take long for physicists to realise that this hypothesised particle constitutes a marvellous solution to not just one but two of the biggest riddles left in particle physics simultaneously. Shortly after that many models were published involving the axion or very similar particles, so-called ALPs (axion-like particles).

They showed production and detection mechanisms for axions and ALPs with the rest mass of the new particles ranging over several orders of magnitude across theories. This also allowed collider physics to get involved in the search for a particle which at first seemed to be only possible with sub-eV mass. Naturally, no PhD thesis involving axions would be complete without mentioning the rather unusual origin of the name. In 1978 Frank Wilczek published his famous paper *Problem of Strong P and T Invariance in the Presence of Instantons*³⁴ in which he coined the name ‘axion’ for the first time. In his 2004 Nobel lecture he explained the origin of the name: “I named them axions after a laundry detergent, because they clean up a problem with an axial current.”³⁵

2.6. Types of Dark Matter Detection Experiments

There are four main types of searches for DM: astrophysical observation, direct and indirect detection as well as producing DM in colliders. Astrophysical observations of the space around the Earth sparked the idea of DM as previously explained. However, it fails to give an explicit description of how DM is structured or what it is made of exactly. Cosmology and astrophysics have delivered a multitude of reasons why DM should exist. The task for theoretical and experimental particle physicists now is to figure out in which ways one might be able to detect and thereby explain what DM actually is.

2.6.1. Collider Experiments

One way is to look for DM in events recorded in colliders. The idea is that by colliding high energy SM particles one can produce dark matter particles, which either promptly decay into detectable SM particles, or escape the detector resulting in a missing energy signature. Now one can search the collision results for patterns consistent with those predicted by the theory of DM decays. The great advantage this method has over most others is that it is independent of the presence of naturally occurring DM since the particles are produced in the lab. Regarding axions, colliders are not searching for QCD axions*, but for axion-like particles (ALPs). This is because realistic QCD axions are expected to be extremely light, typically with sub-eV mass, far below the energy reach of colliders, and their couplings are fixed by the axion-decay constant, making them inaccessible to direct production at collider energies. For ALPs this relation does not necessarily hold, so these could have high masses, but still small couplings. In collider experiments, ALPs are most commonly searched for via their coupling to photons, gluons, and electroweak bosons. Another important production mechanism involves coupling to the Higgs boson. If such couplings exist, exotic Higgs decays like $h \rightarrow aa$, where a denotes an ALP, become possible. These decays can result in various final states, such as four photons or combinations of photons and leptons. In addition to

*those which would solve the strong CP problem, see section 3.2.1

2.6. Types of Dark Matter Detection Experiments

prompt decays, ALPs may also be long-lived, leading to displaced vertices or missing energy signatures if they escape the detector before decaying. Both hadron and lepton colliders contribute to ALP searches. Hadron colliders such as the LHC provide access to very high energies and large production rates, making them sensitive to heavy ALPs and rare processes. General-purpose experiments like ATLAS and CMS perform a broad range of searches that can be reinterpreted in terms of ALP models. Lepton colliders, on the other hand, offer a cleaner environment as they typically have much lower background in comparison to high-energy hadron colliders. Experiments such as Belle II are well suited for probing low-mass ALPs with very weak couplings.

2.6.2. Indirect Searches

Indirect searches aim at detecting SM particles produced by dark matter annihilation, decay, or self-interactions in astrophysical environments. In contrast to collider and direct searches, these experiments do not attempt to observe the interaction of DM with a detector directly, but rather look for excesses or characteristic spectral features in cosmic radiation that can be attributed to DM processes. Typical targets include gamma rays, charged cosmic rays, and neutrinos originating from regions with high expected DM density. A prominent example is the observation of high-energy gamma rays. Annihilation or decay of DM particles can produce gamma rays either directly or through intermediate SM states that subsequently radiate photons. Space-based instruments such as the Fermi Gamma-ray Space Telescope survey the sky with high precision and search for deviations from known astrophysical backgrounds. Another important class of indirect probes involves neutrinos. If DM accumulates in massive bodies like the Sun through scattering with SM particles, subsequent annihilation may produce high-energy neutrinos that can escape and be detected on Earth. Large-volume neutrino observatories such as the IceCube Neutrino Observatory search for an excess of neutrinos from the direction of the Sun or the galactic centre. While neutrino signals are challenging to detect due to small interaction cross sections, they offer a comparatively background-free detection channel. Overall, indirect searches provide a complementary approach to collider and direct searches. Their main limitation lies in the dependence on astrophysical modelling and background estimation, which introduces significant uncertainties, a typical problem of experimental astrophysics.

2.6.3. Direct Searches

The third method is the ‘direct search’, for which the SUPAX experiment is one example. This thesis is therefore going to put the emphasis on such searches. The ‘direct search’ strategy focuses on the detection of processes in which DM particles interact directly with SM particles. One such example is the transfer of momentum and energy in scattering processes. Experiments like the aforementioned XENON or CRESST have been researching such processes for many years and continue to do so.

While the previously mentioned experiments try to exploit the nuclear recoil of DM-SM interactions, the following concepts are based on the idea of a DM particle converting into SM particles, mainly axions converting into photons by the inverse Primakoff effect.³⁶

One such idea is helioscope experiments. These experiments rely on the idea of axions being produced in great quantities inside our Sun, hence the name. One main downside

2. The Case for Dark Matter




Figure 2.5.: The current limits on axion-photon coupling $g_{a\gamma\gamma}$ plotted against the axion mass m_a . There is a visible difference between the wide ranging broad-band helioscopes with lower absolute sensitivity in green and the high-sensitivity, but narrow-band haloscopes in blue. Figure made with [37].

of helioscopes is the amount of assumptions necessary to justify the approach. Helioscopes depend on an existing flux of axions from the sun which is assumed to be large enough to be detected with reasonable sensitivity on Earth. This solar axion flux has a very broad energy distribution, allowing for scanning sweeps over large areas of axion masses. Helioscope experiments like (Baby)IAXO use high magnetic fields to convert solar axions into detectable X-rays, utilizing the inverse Primakoff effect. In Figure 2.5 one can clearly see that although helioscopes (green) span over a much wider range of masses than the soon to be discussed haloscopes (blue), their sensitivity rarely reaches near the yellow band representing the QCD axions. Alternative experiments requiring fewer necessary assumptions include light-shining-through-wall (LSW) experiments. As with colliders, these experiments do not depend on naturally occurring DM. LSW experiments use two areas of strong magnetic field, separated by an optical barrier. The idea is that one or more photons of a laser will convert into axions via the Primakoff effect on one side of the optical barrier, travel through it, and turn into a detectable photon on the other side via the inverse of the same process. This would create the impression the photon had passed the optical barrier or to put it more simply: light shining through a wall. One such example is the ALPS II³⁸ experiment. While the independence of parameter assumptions and models may be attractive these experiments have not reached the sensitivity of modern halo- or helioscope experiments.

2.6. *Types of Dark Matter Detection Experiments*

The last and main type of direct search however is yet to come. Haloscopes, such as SUPAX, are based on the idea of a DM halo around our galaxy, which can be modelled as an isotropic, isothermal sphere. Note that helioscopes, colliders and LSW experiments do not rely on the existence of a galactic dark matter halo and hence only probe for the principal existence of axions, but not their natural abundance as dark matter. A detection in one such experiments does not automatically constitute evidence for dark matter. Haloscopes on the other hand explicitly probe for dark matter axions. Section 3.1 will explain the galactic halo in greater detail.

3. Theory

This chapter will cover the theoretical grounding of the SUPAX experiment. First the concept of the DM halo is discussed, followed by the hypothesised interaction of axions and dark photons with the Standard Model and how one can utilise this to detect them with cavities. Thereafter, the characterisation of such a haloscope cavity and the basics of superconductors in this context are discussed.

3.1. Dark Matter Halo

Observations of galactic rotation curves (recall sec. 2.1) suggest that dark matter in galaxies is distributed in extended, approximately spherical halos that dominate the gravitational potential. To directly detect dark matter particles in this halo, we must make at least two assumptions about the general structure of the halo. The first is regarding the mass distribution and the other is the velocity distribution of the particles. Since the 1930s, the main assumption was, and continues to be, that the galactic halo follows the Virial theorem, i.e. it is virialised. It is common for haloscope experiments to assume a fully virialised halo and this thesis is no exception. The flat rotation curves imply that the total enclosed mass scales linearly with the distance, the radius as we assume a spherical halo, from the galactic centre $M(r) \propto r$ over a wide range of distances. From this follows a density profile of the form

$$\rho(r) \propto \frac{1}{r^2} \quad (3.1)$$

commonly referred to as an isothermal sphere. This introduces some flaws: the density approaches infinity as the radius approaches the galactic centre and for $r \rightarrow \infty$ the total mass of the halo also approaches infinity. Characterizing the dark matter halo in a way that eliminates or at least softens these two problems has led to a multitude of models, many of which can be expressed in a double power law of form

$$\rho(r) = \rho_0 \left(\frac{r}{r_0} \right)^{-\gamma} \left[1 + \left(\frac{r}{r_0} \right)^\alpha \right]^{(\gamma-\beta)/\alpha} \quad (3.2)$$

where ρ_0 is a characteristic density, r_0 a scale radius, and the parameters α, β, γ depend on the model. This family includes both cored ($\gamma = 0$) profiles, often referred to as pseudo-isothermal spheres, and cusped ($\gamma > 0$) profiles such as the Navarro-Frenk-White³⁹ profile with parameters $(\alpha, \beta, \gamma) = (1, 3, 1)$, one of the most commonly used models. Finding the correct model is still the subject of a very open debate, however, since lab-bound dark matter experiments are sensitive only to local properties of the halo, the behaviour at very large and very small radii has little practical impact. At the position of the Earth, approximately

8 kpc from the galactic centre, all commonly used halo models predict similar densities. For axion haloscope searches it is standard to assume a local dark matter density of

$$\rho_{\text{DM}} \simeq 0.45 \text{ GeV cm}^{-3} \quad (3.3)$$

which lies within the range inferred from both local dynamical measurements and global fits to the Milky Way rotation curve.⁴⁰ In a virialised halo, the dark matter velocity distribution in the Galactic rest frame is expected to be approximately Maxwellian. The characteristic velocity scale is set by the local circular velocity $v_c \simeq 220 \text{ km/s}$, which reflects the depth of the galactic gravitational potential. The corresponding root-mean-square velocity

$$\langle v^2 \rangle^{1/2} = \sqrt{3/2 \cdot v_c^2} \simeq 270 \text{ km/s} \quad (3.4)$$

is often referred to as the virial velocity.

For axion haloscope experiments, this velocity distribution plays a central role because it determines the axions' velocity dispersion. Since we assume cold, i.e. non-relativistic, dark matter axions, we can write their total energy as $E \simeq m_a c^2 + \frac{1}{2} m_a v^2$, so that the halo's velocity dispersion produces a narrow spectral line with a fractional width of order $\Delta E/E \sim \langle v^2 \rangle / c^2 \approx 10^{-6}$. Since energy and frequency are related by $E = h\nu$, the same relation applies to the frequency spectrum of the axion-induced signal

$$\frac{\Delta\nu}{\nu} \approx 10^{-6} \quad (3.5)$$

where $\nu_a = m_a c^2 / h$ is the axion Compton frequency. It is convenient to characterise this intrinsic linewidth by defining an effective axion quality factor

$$Q_a \equiv \frac{\nu_a}{\Delta\nu} \approx 10^6. \quad (3.6)$$

This quantity is not a property of the detector but rather a consequence of the velocity dispersion of axions in the galactic halo. As for any quality factor, Q_a is a measure of how peaked the axion signal is.

In haloscope experiments it is usual to assume a homogeneously distributed galactic halo that follows some form of eq. (3.1) or (3.2), allowing for an estimated local dark matter density as in eq. (3.3). The same assumption is made in this entire thesis. However, it should be noted that this is not undisputed. There is ample room for discussion that a DM halo may not be homogeneously distributed,⁴¹ in which case any axion-photon* coupling limit set with haloscopes is to be treated with caution.

3.2. Axion Physics

Although the previous section has used the terms dark matter and axion synonymously, all relations also hold if you assume other light dark matter candidates, such as massive dark photons, instead of axions. The following section will discuss the strong CP problem and how the absence of an electric dipole moment (EDM) of the neutron hints towards the existence of axions specifically. Furthermore, the interaction of the axion field and the electromagnetic

*or dark photon-photon

3. Theory

field as an expansion of the Standard Model Lagrangian will be used to derive the Maxwell equations including the axion field. Ultimately the boundary conditions of a cavity will be applied to these Maxwell equations to derive the signal power one can expect from an axion conversion in a resonant cavity.

3.2.1. The Strong CP Problem

Modern quantum field theory (QFT) is built around the concept of symmetry. A symmetry is defined as a transformation acting on a physical system that leaves the action, and therefore the equations of motion, invariant. Symmetries play a central role in the SM, determining both the allowed interactions and the conservation laws obeyed by physical processes.

Noether’s theorem⁴² states that for every continuous symmetry of the action there exists a conserved current, and hence a conserved quantity, and vice versa. Well-known examples include translational invariance resulting in conservation of momentum, or rotational invariance resulting in conservation of angular momentum.

In addition to continuous symmetries QFTs may also respect discrete symmetries. Of particular importance are charge conjugation (**C**), which exchanges particles with their antiparticles, parity (**P**), which corresponds to spatial inversion $\vec{x} \rightarrow -\vec{x}$, and time reversal (**T**), which reverses the direction of time.

The first hint of one of these symmetries being violated was discovered in 1956 when Lee and Yang⁴³ proposed that parity might not be conserved in weak interactions, a hypothesis that was soon confirmed experimentally by Wu⁴⁴ in her study of β decay in cobalt nuclei. The observed angular distribution of emitted electrons demonstrated maximal parity violation in the weak force, earning Lee and Yang, but not Wu, the Nobel Prize in 1957. Further insight came from the Goldhaber experiment,⁴⁵ which showed that neutrinos are exclusively left-handed, while anti-neutrinos are exclusively right-handed. Here, “left-handed” refers to the projection of the particle’s spin along its direction of motion: for a left-handed particle the spin points opposite to the momentum, while for a right-handed particle it points along the momentum. If **C** symmetry were conserved, left-handed anti-neutrinos would also exist, but they do not. Charge conjugation is therefore also violated in the weak interaction. Nevertheless, the combined **CP** symmetry could still have remained intact.

This possibility was ruled out in 1964 by the observation of **CP** violation in neutral kaon decays.⁴⁶ The weak interaction thus violates **CP** symmetry, a feature that has been incorporated into the SM through the complex phase of the Cabibbo-Kobayashi-Maskawa (CKM) matrix, which describes the probability of quarks changing flavour via weak interactions. While this 3×3 matrix consists of nine real parameters, they are not all independent from each other such that it can also be written with 3 angles and the just mentioned **CP**-violating phase, conventionally denoted by δ_{CKM} .

In classical theories, the conserved current J^μ can be calculated via Noether’s theorem $\partial_\mu J^\mu = 0$, but at the quantum level symmetries that are exact classically may be broken. In the SM, the divergence of a classically conserved axial current J^μ takes the form

$$\partial_\mu J^\mu = \frac{g^2}{32\pi^2} F_{\mu\nu} \tilde{F}^{\mu\nu} \quad (3.7)$$

where g is the relevant gauge coupling, $F_{\mu\nu}$ the (generic) gauge field strength tensor, and

$\tilde{F}^{\mu\nu} = \frac{1}{2}\epsilon^{\mu\nu\alpha\beta}F_{\alpha\beta}$ its dual, connected via the Levi-Civita tensor.

$$\epsilon_{\mu\nu\alpha\beta} = \begin{cases} +1 & \text{if } (\mu, \nu, \alpha, \beta) \text{ is an even permutation of } (0, 1, 2, 3) \\ -1 & \text{if } (\mu, \nu, \alpha, \beta) \text{ is an odd permutation of } (0, 1, 2, 3) \\ 0 & \text{otherwise} \end{cases} \quad (3.8)$$

A term in the form of the r.h.s. of eq. (3.7) appears for every coupling between fermions and gauge symmetries. As a consequence, the Lagrangian gains terms of the form

$$\mathcal{L}_\theta = \frac{g^2}{32\pi^2}\theta F_{\mu\nu}\tilde{F}^{\mu\nu} \quad (3.9)$$

where θ_{EW} and θ_{QCD} are the vacuum angles of the electroweak and strong force, with $\theta \in [0, 2\pi)$. The presence of the antisymmetric $\epsilon_{\mu\nu\alpha\beta}$ implies that these terms generally violate P and T symmetries. As the combined symmetry CPT is generally conserved, the T violation implies that these terms hence violate CP as well. The θ angles therefore provide a measure of the intrinsic CP violation in the corresponding gauge groups.

Experimentally we have seen that CP violation in the electroweak sector is sizeable and there is no a priori theoretical reason for θ_{QCD} to take any particular value. One way to probe CP violation in the strong force is the electric dipole moment (EDM) of the neutron d_n . The neutron has a non-zero magnetic dipole moment μ_n and applying the Wigner-Eckart theorem, we know that both d_n and μ_n are along the spin axis, i.e. they are either parallel or anti-parallel. Consider a neutron for which d_n and μ_n are parallel, so pointing in the same direction. Under parity reversal d_n changes sign while μ_n does not and under time reversal μ_n changes sign while d_n does not. In other words, either a P or a T inversion creates a particle that is observably different from before. A non-zero d_n therefore violates P and T symmetry, where a T symmetry violation again implies a CP violation, as before.




Figure 3.1.: The nEDM generating mechanism. The black vertices are CP-conserving while the gray vertex is CP-violating. Figure taken from [47].

3. Theory

Figure 3.1 shows the mechanism that generates the EDM of the neutron. The calculation of this interaction has been presented in [48, 49], the relevant result for our case is

$$d_n \sim \theta_{\text{QCD}} \frac{e}{m_N^2} \frac{m_u m_d}{m_u + m_d} \sim \theta_{\text{QCD}} \times 10^{-16} e \text{ cm} \quad (3.10)$$

where m_N is the nucleon mass and m_u, m_d are the up- and down-quark masses. The current experimental upper limit $d_n \simeq 10^{-26} e \text{ cm}$ implies the remarkably strong constraint

$$\theta_{\text{QCD}} < 10^{-10}. \quad (3.11)$$

The extreme smallness of this bound, compared to the natural expectation $\theta_{\text{QCD}} \sim \mathcal{O}(1)$, constitutes the strong CP problem.

From a pragmatic point of view, a value $\theta_{\text{QCD}} \approx 0$ is not logically inconsistent, just as $\theta_{\text{EW}} \approx 30^\circ$ is perfectly acceptable. However, the history of physics (and mathematics) has shown repeatedly that when a dimensionless parameter is extraordinarily close to a special value, this often signals the presence of an underlying mechanism or symmetry. A number found to be close to 3.1415 naturally invites a geometric explanation enforcing π . Likewise, a number as close to zero as θ_{QCD} motivates the search for a dynamical or symmetry-based suppression mechanism.

The situation becomes even more dramatic once one takes into account that quarks not only couple to the strong interaction but also to the electroweak sector via Yukawa interactions with the Higgs field. In this case, the relevant CP-violating parameter is not θ_{QCD} alone, but

$$\bar{\theta} = \theta_{\text{QCD}} + \theta_Y \quad \text{with} \quad \theta_Y = \arg(\det[Y_d Y_u]) \quad (3.12)$$

where Y_u and Y_d are the Yukawa coupling matrices responsible for the up- and down-type quark masses.⁵⁰ Even if a justification for a vanishing θ_{QCD} were found, one would still expect a non-zero contribution from θ_Y .

There is no reason for θ_Y to be small. On the contrary, since the CP violation phase of the CKM matrix mentioned earlier has been determined⁵¹ as $\delta_{\text{CKM}} = 1.147 \pm 0.026$, far from zero, it is natural to expect θ_Y to be of comparable magnitude. The experimental constraint of $\bar{\theta} < 10^{-10}$ would then require θ_Y and θ_{QCD} to be individually of order 1 with opposite sign, but equal in magnitude to a precision of at least one in 10^{10} . Because the strong and electroweak interactions are independent sectors of the SM, there is no fundamental relation enforcing such a cancellation. Achieving it would therefore require an extreme fine-tuning of two a priori unrelated parameters.

Fine-tuning of this severity is generally frowned upon in a fundamental theory as the odds of these two independent parameters being nearly exactly the same are astronomically small.* This observation strongly motivates the search for a dynamical mechanism that naturally drives $\bar{\theta}$ to zero, rather than relying on an implausible numerical coincidence. As will be discussed in the following section, the Peccei-Quinn mechanism provides precisely such a solution.

*Picking two random numbers between 0 and 2π and having them disagree no more than 10^{-10} is roughly as likely as picking the ace of spades from a shuffled deck of cards six times in a row.

3.2.2. Solution of the CP Problem: The Peccei-Quinn Mechanism and the Axion

The most promising solution to the CP problem has been proposed by Peccei and Quinn nearly half a century ago in 1977.^{52,53} The Peccei-Quinn (PQ) mechanism offers a dynamical explanation for the smallness of $\bar{\theta}$ by promoting it from a fixed parameter to a field whose vacuum expectation value is determined by the dynamics of the theory. The key ingredients are the introduction of a new global chiral symmetry and its spontaneous breaking, which in turn gives rise to a new pseudoscalar degree of freedom: the axion.

The basic mechanism underlying the PQ solution is spontaneous symmetry breaking (SSB). Consider a complex scalar field $\phi(x)$ with a potential of the form

$$V(|\phi|) = \frac{1}{2}\mu^2|\phi|^2 + \frac{1}{4}\lambda|\phi|^4 \quad (3.13)$$

and its Lagrangian

$$\mathcal{L} = \frac{1}{2}(\partial_\nu \phi^*)(\partial^\nu \phi) - V(|\phi|) \quad (3.14)$$

which is invariant under phase rotation $\phi \rightarrow e^{i\alpha}\phi$, i.e. a U(1) symmetry. For positive μ^2 , the scalar field has mass μ . Consequently, in case $\mu^2 = 0$, the field is massless. However, for $\mu^2 < 0$, the minimum of the potential is not at $|\phi| = 0$ but on a circle with $|\phi| = \sqrt{-\mu^2/\lambda} =: v$. The SSB can be realised through a phase transition from the $\mu^2 > 0$ case to the $\mu^2 < 0$ case as shown in Fig. 3.2.




Figure 3.2.: At high energies, i.e. high temperatures T , (left) the potential has its minimum at the origin, $|\phi| = 0$. As the system cools down, the symmetry spontaneously breaks (middle) and the minimum moves to a non-zero field value $|\phi| = v > 0$.

The course of such a phase transition is commonly assumed to be a temperature dependence of μ^2 such that $\mu^2(T = 0) < 0$ and $\mu^2(T) > 0$ for sufficiently high temperature with some critical temperature $T_c > 0$ at which $\mu^2(T_c) = 0$. Below this critical temperature, the point $|\phi| = 0$ becomes a local maximum and fluctuations will drive the field away from the origin toward one of the degenerate minima with $|\phi| = v$, giving rise to the name of spontaneous symmetry breaking. Assuming that the universe has cooled down far enough that we are in the interesting $\mu^2 < 0$ regime, we can expand the field around the minimum, writing

$$\phi(x) = [v + \rho(x)] e^{ia(x)/f_a} \quad (3.15)$$

3. Theory

where $\rho(x)$ and $a(x)$ are real scalar fields parametrising the radial and angular excitations of the complex field. Here f_a is a constant with mass dimension one that sets the normalisation of the angular field. Its physical interpretation will become clear in the context of the PQ mechanism below. Using eq. (3.15) we can write out the Lagrangian of (3.14) as

$$\mathcal{L} = \frac{1}{2}(\partial_\nu \rho)^2 + \frac{v^2}{2f_a^2}(\partial_\nu a)^2 + \left(\frac{\rho^2}{2} + v\rho\right) \frac{1}{f_a^2}(\partial_\nu a)^2 - V(\rho(x) + v) \quad (3.16)$$

where the first and second terms are the kinetic energy of the two fields respectively and the third describes their interaction. The field $a(x)$ is only present in the Lagrangian through derivatives, notably lacking a mass term. Each such massless mode corresponds to a Goldstone boson.




Figure 3.3.: “Sombrero” potential illustrating the SSB of a global $U(1)$ symmetry. Figure taken from [54].

If the symmetry were exact, the Goldstone boson would remain massless. However, if the symmetry is only approximate, it becomes a pseudo-Goldstone boson and acquires a small mass. This is precisely the situation in the PQ mechanism. It introduces a new global chiral symmetry $U(1)_{\text{PQ}}$ with a corresponding complex scalar field ϕ_{PQ} . The field develops a vacuum expectation value f_a , spontaneously breaking $U(1)_{\text{PQ}}$. At the classical level, this symmetry can be arranged such that the CP-violating parameter $\bar{\theta}$ is unphysical. Crucially, $U(1)_{\text{PQ}}$ is not an exact symmetry at the quantum level, and this chiral anomaly generates an interaction term

$$\mathcal{L}_{agg} = \frac{g^2}{32\pi^2 f_a} a G_{\mu\nu}^a \tilde{G}^{a\mu\nu} \quad (3.17)$$

where a is the angular mode of the PQ field, identified as the axion by Weinberg and Wilczek. Comparing this to the θ term in QCD (see eq. (3.9), where we now write the gluon field strength explicitly as $G_{\mu\nu}^a$ instead of the generic notation $F_{\mu\nu}$) shows that the physically relevant combination is $\bar{\theta} + a(x)/f_a$. In other words, the PQ mechanism promotes $\bar{\theta}$ to a dynamical field variable. While the PQ symmetry is broken at the scale f_a , the anomaly induces a small explicit breaking that lifts the degeneracy of the vacuum along the angular

direction. In the intuitive picture of the “sombbrero” potential of Fig. 3.3, the originally flat trough is slightly tilted by QCD, selecting a unique minimum. The axion field relaxes into this minimum, located at

$$\bar{\theta} + \frac{a(x)}{f_a} = 0 \quad (3.18)$$

and dynamically cancels out the CP-violating parameter in QCD, irrespective of the initial value of $\bar{\theta}$. The strong CP problem is therefore solved dynamically rather than by fine-tuning: CP violation in the strong interaction is driven to zero by the dynamics of the axion field.

At low energies, the axion-gluon interaction of eq. (3.17) is dominated by mixing between axions and neutral pions, giving a relation for the axion mass.

$$m_a^2 = \frac{m_\pi^2 f_\pi^2}{f_a^2} \frac{m_u m_d}{(m_u + m_d)^2} \quad (3.19)$$

The axion inherits the pion’s coupling to several particles, especially its two-photon decay, shown in Fig. 3.4, which is the key interaction exploited in haloscope-type detectors.




Figure 3.4.: Axion decay to two photons (left) and stimulated axion to photon conversion by interaction with an external magnetic field $\mathbf{B}_0$ (right).

The axion-photon interaction is given by

$$\mathcal{L}_{a\gamma\gamma} = g_\gamma \left(\frac{e}{4\pi} \right)^2 \frac{a}{f_a} F_{\mu\nu} \tilde{F}^{\mu\nu} \quad (3.20)$$

or the more commonly used form

$$\mathcal{L}_{a\gamma\gamma} = -\frac{1}{4} g_{a\gamma\gamma} a F_{\mu\nu} \tilde{F}^{\mu\nu} \quad (3.21)$$

where the model-dependent dimensionless coefficient is connected to the interaction strength $g_{a\gamma\gamma}$, usually simply called coupling, of the axion with two photons via

$$g_{a\gamma\gamma} = \frac{\alpha g_\gamma}{\pi f_a} = \frac{\alpha g_\gamma}{\pi \sqrt{\chi}} m_a \quad (3.22)$$

where α is the fine-structure constant. At zero temperature, lattice QCD calculations give a topological susceptibility of order $\chi \simeq [77.6 \text{ MeV}]^4$.

3. Theory

3.2.3. Interaction with the Electromagnetic Field

In this section we will see how the axion interacts with the electromagnetic (EM) field. Of special interest with regard to the SUPAX experiment (or cavity haloscopes in general) is the resulting electric field $\mathbf{E}_a$, mostly expressed in terms of RF power P , from the interaction of the axion field with an external magnetic field $\mathbf{B}_e$. Since the axions considered in these types of experiments are generally assumed to be cold dark matter and therefore have very low velocity, the calculations will not consider any relativistic effects. The derivation of the Maxwell equations follows the *Theoretical Foundations* of Millar et al.,⁵⁵ with additional steps clarified in the adaptation presented in [56].

As shown in the previous section, the Lagrangian describing EM currents and photons can be extended by the axion field a and its interaction with the EM field (see eq. (3.21)) yielding the EM-Lagrangian

$$\mathcal{L} = \frac{1}{2}(\partial_\mu a)^2 - \frac{1}{2}m_a^2 a^2 - \frac{1}{4}F_{\mu\nu}F^{\mu\nu} - J^\mu A_\mu - \underbrace{\frac{1}{4}g_{a\gamma\gamma}aF_{\mu\nu}\tilde{F}^{\mu\nu}}_{\text{interaction term}} \quad (3.23)$$

where

$$F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu \quad (3.24)$$

$$\tilde{F}^{\mu\nu} = \frac{1}{2}\epsilon^{\mu\nu\alpha\beta}F_{\alpha\beta} \quad (3.25)$$

are the EM field tensor and its dual, as before. As usual, the EM field-strength tensors are derived from the vector potential A^μ and the electric 4-current J^μ .

$$A^\mu = (A_0, \mathbf{A}) \quad (3.26)$$

$$J^\mu = (\rho, \mathbf{J}) \quad (3.27)$$

For the axion extension, a is the scalar axion field, m_a is the axion mass and $g_{a\gamma\gamma}$ is the coupling constant which determines the strength of the interaction between the axion and the EM fields. For convenience the following calculations will be done in natural units, so $\hbar = c = 1$. Starting from the Lagrangian in (3.23) the next step is to solve the Euler-Lagrange equations

$$\frac{\partial \mathcal{L}}{\partial a} - \partial_\lambda \left[\frac{\partial \mathcal{L}}{\partial (\partial_\lambda a)} \right] = 0 \quad (3.28)$$

$$\frac{\partial \mathcal{L}}{\partial A_\kappa} - \partial_\lambda \left[\frac{\partial \mathcal{L}}{\partial (\partial_\lambda A_\kappa)} \right] = 0 \quad (3.29)$$

in order to derive the Maxwell equations including the scalar axion field. While (3.28) is used to derive the equations of motion (EOM) of the axion field it is eq. (3.29) which will be used to derive Maxwell's equations with the axion field as additional source. The first term does not contribute, the second term is

$$\frac{\partial \mathcal{L}}{\partial (\partial_\lambda A_\kappa)} = - \left[\partial^\lambda A^\kappa - \partial^\kappa A^\lambda \right] + \frac{1}{2}g_{a\gamma\gamma}\epsilon^{\mu\nu\lambda\kappa} [\partial_\mu A_\nu - \partial_\nu A_\mu] a \quad (3.30)$$

$$= -F^{\lambda\kappa} + \frac{1}{2}g_{a\gamma\gamma}\epsilon^{\lambda\kappa\mu\nu}F_{\mu\nu}a \quad (3.31)$$

and inserting this into (3.29) one gets the EOM

$$\partial_\alpha F^{\alpha\beta} = \frac{1}{2} g_{a\gamma\gamma} \epsilon^{\alpha\beta\mu\nu} \partial_\alpha (F_{\mu\nu} a). \quad (3.32)$$

These equations can be simplified by applying Maxwell's constraint equations

$$\partial_\alpha F_{\mu\nu} + \partial_\mu F_{\nu\alpha} + \partial_\nu F_{\alpha\mu} = 0. \quad (3.33)$$

By extending (3.32)

$$\begin{aligned} \partial_\alpha F^{\alpha\beta} &= \frac{1}{2} g_{a\gamma\gamma} \epsilon^{\alpha\beta\mu\nu} \partial_\alpha (F_{\mu\nu} a) \\ &= \frac{1}{2} g_{a\gamma\gamma} \left[\epsilon^{\alpha\beta\mu\nu} F_{\mu\nu} (\partial_\alpha a) + \epsilon^{\alpha\beta\mu\nu} a (\partial_\alpha F_{\mu\nu}) \right] \end{aligned} \quad (3.34)$$

one can eliminate the last term by first permuting the indices and then applying (3.33).

$$\begin{aligned} \epsilon^{\alpha\beta\mu\nu} \partial_\alpha F_{\mu\nu} &= \frac{1}{3} \left[\epsilon^{\alpha\beta\mu\nu} \partial_\alpha F_{\mu\nu} + \epsilon^{\mu\beta\nu\alpha} \partial_\mu F_{\nu\alpha} + \epsilon^{\nu\beta\alpha\mu} \partial_\nu F_{\alpha\mu} \right] \\ &= \frac{1}{3} \epsilon^{\alpha\beta\mu\nu} [\partial_\alpha F_{\mu\nu} + \partial_\mu F_{\nu\alpha} + \partial_\nu F_{\alpha\mu}] = 0 \end{aligned} \quad (3.35)$$

Therefore (3.32) simplifies to

$$\partial_\alpha F^{\alpha\beta} = \frac{1}{2} g_{a\gamma\gamma} \epsilon^{\alpha\beta\mu\nu} F_{\mu\nu} \partial_\alpha a. \quad (3.36)$$

This can now be written in terms of $\mathbf{E}$ and $\mathbf{B}$ using

$$F^{0i} = -E_i \Rightarrow F_{0i} = E_i, \quad (3.37)$$

$$F^{ij} = -\epsilon_{ijk} B_k \Rightarrow F_{ij} = -\epsilon_{ijk} B_k \quad (3.38)$$

which reveals that the interaction term takes the form

$$\frac{1}{4} g_{a\gamma\gamma} a F_{\mu\nu} \tilde{F}^{\mu\nu} = -g_{a\gamma\gamma} a \mathbf{E} \cdot \mathbf{B}. \quad (3.39)$$

Let us now consider the EOM by splitting them into the two cases $\beta = 0$ and $\beta = 1, 2, 3$. $\beta = 0$ yields

$$\begin{aligned} \partial_i F^{i0} &= \frac{1}{2} g_{a\gamma\gamma} \epsilon^{i0\mu\nu} F_{\mu\nu} \partial_i a \\ &= -g_{a\gamma\gamma} [F_{23} \partial_1 a + F_{31} \partial_2 a + F_{12} \partial_3 a] \\ &= g_{a\gamma\gamma} [\epsilon_{231} B_1 \partial_1 a + \epsilon_{312} B_2 \partial_2 a + \epsilon_{123} B_3 \partial_3 a] \\ &\Rightarrow \partial_i E_i = g_{a\gamma\gamma} B_i \partial_i a \end{aligned} \quad (3.40)$$

while $\beta = 1, 2, 3$ will only be written out for $\beta = 1$ since the remaining two cases are just permutations of $\beta = 1$.

$$\begin{aligned} \partial_0 F^{01} + \partial_2 F^{21} + \partial_3 F^{31} &= \frac{1}{2} g_{a\gamma\gamma} [\epsilon^{01\mu\nu} F_{\mu\nu} \partial_0 a + \epsilon^{21\mu\nu} F_{\mu\nu} \partial_2 a + \epsilon^{31\mu\nu} F_{\mu\nu} \partial_3 a] \\ &= g_{a\gamma\gamma} [F_{23} \partial_0 a + F_{30} \partial_2 a + F_{02} \partial_3 a] \\ &\Rightarrow -\partial_0 E_1 - \epsilon_{213} \partial_2 B_3 - \epsilon_{312} \partial_3 B_2 = g_{a\gamma\gamma} [-\epsilon_{231} B_1 \partial_0 a - E_3 \partial_2 a + E_2 \partial_3 a] \\ &\Rightarrow -\partial_0 E_1 + (\partial_2 B_3 - \partial_3 B_2) = g_{a\gamma\gamma} [-B_1 \partial_0 a + (E_2 \partial_3 a - E_3 \partial_2 a)] \end{aligned} \quad (3.41)$$

3. Theory

Writing this in terms of $\mathbf{E}$ and $\mathbf{B}$, one finally arrives at the Maxwell equations including the axion field.

$$\nabla \cdot \mathbf{E} = \rho - g_{a\gamma\gamma} \mathbf{B} \cdot \nabla a \quad (3.42a)$$

$$\nabla \times \mathbf{B} - \partial_t \mathbf{E} = \mathbf{J} + g_{a\gamma\gamma} (\mathbf{B} \partial_t a - \mathbf{E} \times \nabla a) \quad (3.42b)$$

$$\nabla \cdot \mathbf{B} = 0 \quad (3.42c)$$

$$\nabla \times \mathbf{E} + \partial_t \mathbf{B} = 0 \quad (3.42d)$$

3.2.4. Detection in Cavities

Having derived the Maxwell equations, including the axion field, we will now impose the boundary conditions inherent to the experiment. Since the axions of interest here are considered cold, i.e. non-relativistic, dark matter, their de Broglie wavelength

$$\lambda_{\text{dB}} = \frac{h}{\gamma m v} \quad (3.43)$$

is of the order of 50 m assuming the dark matter velocity of $230 \frac{\text{km}}{\text{s}}$ mentioned in sec. 3.1 and an axion rest mass of $33 \mu\text{eV} \approx 8 \text{ GHz}$. Hence, it is safe to assume spatial homogeneity on our detector scale. The axion field

$$a(t) = a_0 e^{-i\omega t} \quad (3.44)$$

can therefore be written purely as a function of time. Its gradient $\nabla a(t) = 0$ in equations (3.42a) and (3.42b) vanishes. In the context of a cavity haloscope, there are no electromagnetic charges or currents, so eq. (3.27) simply reads $J^\mu = (0, 0, 0, 0)$, simplifying (3.42b) further. The Maxwell equations in the context of cavity haloscope experiments therefore read

$$\nabla \cdot \mathbf{E} = 0 \quad (3.45a)$$

$$\nabla \times \mathbf{B} - \dot{\mathbf{E}} = g_{a\gamma\gamma} \mathbf{B}_e \dot{a} \quad (3.45b)$$

$$\nabla \cdot \mathbf{B} = 0 \quad (3.45c)$$

$$\nabla \times \mathbf{E} + \dot{\mathbf{B}} = 0 \quad (3.45d)$$

where $\mathbf{B}_e$ denotes the externally applied magnetic field along the z -axis

$$\mathbf{B}_e = B_0 \hat{\mathbf{z}}. \quad (3.46)$$

The derivation of the signal power one can expect to detect in an RF-cavity in this context will now follow [57] and [58], which present the original work of Sikivie⁵⁹ in a more comprehensible way. To calculate the signal power that an axion can deposit in a cavity as shown in Fig. 3.5, we start by taking the time derivative of eq. (3.45b)

$$\nabla \times \dot{\mathbf{B}} - \ddot{\mathbf{E}} = g_{a\gamma\gamma} \ddot{a} \mathbf{B}_e \quad (3.47)$$




Figure 3.5.: Concept of a cavity-based haloscope. A (possibly superconducting) cavity is fitted with a pickup antenna and is placed in an external magnetic field in a cryogenic environment to limit thermal noise. The antenna is connected to a spectrum analyser with an intermediate cryogenic amplifier.

and substituting $\mathbf{B}_e$ and $\dot{\mathbf{B}}$ with (3.46) and (3.45d) respectively.

$$-\nabla \times (\nabla \times \mathbf{E}) - \ddot{\mathbf{E}} = g_{a\gamma\gamma} \ddot{a} B_0 \hat{\mathbf{z}} \quad (3.48)$$

We can use Lagrange's formula for the vector triple product

$$\nabla \times (\nabla \times \mathbf{E}) = \nabla (\underbrace{\nabla \cdot \mathbf{E}}_{=0}) - \Delta \mathbf{E} = -\Delta \mathbf{E} \quad (3.49)$$

with the Laplace operator

$$\Delta \mathbf{x} = \nabla^2 \mathbf{x} = (\nabla \cdot \nabla) \mathbf{x} = \left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + \frac{\partial^2}{\partial z^2} \right) \mathbf{x} \quad (3.50)$$

and eq. (3.45a) to write

$$\Delta \mathbf{E} - \ddot{\mathbf{E}} = g_{a\gamma\gamma} \ddot{a} B_0 \hat{\mathbf{z}}. \quad (3.51)$$

The EM fields are split into a sum of orthonormal eigenmodes $\mathcal{E}_m(\mathbf{x})$ which solve the Poisson equation

$$\nabla^2 \mathcal{E}_m(\mathbf{x}) = -\omega_m^2 \mathcal{E}_m(\mathbf{x}) \quad (3.52)$$

with their respective eigenfrequencies ω_m . Stating that the modes are orthonormal, the normality is ensured by

$$\int_V d^3\mathbf{x} \mathcal{E}_m \cdot \mathcal{E}_n = V \delta_{mn} \quad (3.53)$$

3. Theory

where V is the total volume of the cavity and δ_{mn} is the Kronecker-Delta. Splitting the electric field $\mathbf{E}$ into time- and space-dependent components

$$\mathbf{E} = \sum_m E_m(t) \boldsymbol{\mathcal{E}}_m(\mathbf{x}) \quad (3.54)$$

and inserting this in (3.51) yields

$$\ddot{E}_m + \omega_m^2 E_m + \Gamma_m \dot{E}_m = -g_{a\gamma\gamma} B_0 \ddot{a} C_m. \quad (3.55)$$

This describes the time evolution of the amplitude in the m th mode where energy losses have been parametrised by a decay rate Γ_m and a mode-dependent geometric form factor

$$C_m = \frac{1}{B_0 V} \int_V d^3\mathbf{x} B_0 \hat{\mathbf{z}} \cdot \boldsymbol{\mathcal{E}}_m \quad (3.56)$$

has been defined. Writing the form factor in this way shows that only the overlap of each mode's E -field in the direction of the external magnetic field can be excited by the axion. Fourier transforming eq. (3.55), we can write it as

$$(\omega_m^2 - \omega^2 + i\omega\Gamma_m)E_m(\omega) = -g_{a\gamma\gamma} B_0 C_m \omega^2 a(\omega) \quad (3.57)$$

$$\Leftrightarrow E_m(\omega) = -g_{a\gamma\gamma} B_0 C_m \frac{\omega^2 a(\omega)}{\omega_m^2 - \omega^2 + i\omega\Gamma_m}. \quad (3.58)$$

The total energy stored in the cavity, U , can be expressed as the sum over the energy of all modes U_m .

$$U = \int d^3\mathbf{x} \frac{1}{2} (|\mathbf{E}|^2 + |\mathbf{B}|^2) = \sum_m \frac{1}{2} |E_m|^2 \left(\frac{\omega^2 + \omega_m^2}{2\omega^2} \right) = \sum_m U_m \quad (3.59)$$

The energy in a mode can be measured with a small antenna, but the power extracted contributes to the losses

$$\Gamma_m = \Gamma_m^c + \Gamma_m^s \quad (3.60)$$

where Γ_m^c denotes any losses connected to the cavity itself, such as currents in the walls - a significant loss source for superconducting cavities (see sec. 3.5.2). Γ_m^s denotes the ‘losses’ which are recorded as output signal. The power of that signal therefore is

$$P = \Gamma_m^s U_m = \Gamma_m^s \frac{|g_{a\gamma\gamma} B_0 C_m a_0|^2 V}{2} \frac{\omega^4}{(\omega_m^2 - \omega^2)^2 + (\omega\Gamma_m)^2}. \quad (3.61)$$

Here we introduce the quality factor Q , a quantity describing the energy loss per oscillation of the mode m as

$$Q_m = \frac{\omega_m}{\Gamma_m} \quad (3.62)$$

and assume that we are only interested in frequencies corresponding to axion masses near the resonance frequency of the cavity ω_c , hence $\omega \approx \omega_c \approx m_a$ in natural units. Since we are only considering this one mode we can relabel Q_m here to Q_L , the *loaded* quality factor.

$$\frac{1}{Q_L} = \frac{1}{Q_0} + \frac{1}{Q_s} \quad (3.63)$$

3.3. Dark Photon Physics

It can be expressed as the reciprocal sum of two components. The *unloaded* quality Q_0 measures the energy lost per cycle in a cavity without antennas, i.e. surface resistance in the walls, and Q_s which parametrises the signal ‘lost’ through the antenna. From here we can define the coupling of the antenna

$$\beta = \frac{Q_0}{Q_s} \quad (3.64)$$

and rewrite eq. (3.63) in a rather common form for RF systems.

$$Q_L = \frac{Q_0}{1 + \beta} \quad (3.65)$$

Lastly, we define the local dark matter density $\rho_{\text{DM}} = m_a^2 a_0^2/2$ and assume that axions constitute all of it. Now we can write the axion-induced power

$$P_{\text{sig}} = g_{a\gamma\gamma}^2 V_{\text{eff}} \rho_{\text{DM}} B_0^2 Q_L \frac{1}{m_a} \frac{\beta}{\beta + 1} \quad (3.66)$$

of the signal received by an antenna in an RF-cavity with an effective volume $V_{\text{eff}} = VC^2$.

3.3. Dark Photon Physics

The dark photon and its characteristics have been mentioned in the introduction. Now we shall investigate the signal power of the dark photon and how it is similar to the axion case, with the one big difference stemming from the fact that the DP is a vector boson rather than a scalar.

3.3.1. Dark Photon Electrodynamics

The derivation of the the signal power for a DP-induced signal follows nearly the same procedure as for an axion-induced signal. For the full derivation, see [60]. Instead of introducing a scalar field, one introduces the vector field (and with it a vector gauge boson X) which kinetically mixes with the visible photon. The resulting Lagrangian can be written as

$$\mathcal{L} = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} - \frac{1}{4}X_{\mu\nu}X^{\mu\nu} + J^\mu A_\mu + \frac{m_X^2}{2}(X^\mu X_\mu + 2\chi X_\mu A^\mu) \quad (3.67)$$

where χ is the kinetic mixing, i.e. the interaction strength of the dark and SM photon. A is the photon produced in EM interactions, and X the DP sterile state. From there, one can derive the DP-induced electric field to be

$$|\mathbf{E}_0| = \left| \frac{\chi m_X}{\epsilon} \mathbf{X}_0 \right| \quad (3.68)$$

where ϵ is the dielectric constant of the medium ($\epsilon = 1$ in vacuum). But as one can only access one mode of the cavity at a time, usually defined as the z -axis (more on this in sec. 3.4.2), the resulting accessible field is

$$|\mathbf{E}_0| = \left| \frac{\chi m_X}{\epsilon} \mathbf{X}_0 \cos \theta \right| \quad (3.69)$$

3. Theory

where $\cos \theta = \hat{\mathbf{z}} \cdot \mathbf{X}_0/|\mathbf{X}_0|$. Thus, the DP signal depends not only on the mode overlap as before, but also on the orientation of the DP field relative to the cavity's orientation. From here one can derive the signal power analogously to before.

$$P_{\text{sig}} = \chi^2 V_{\text{eff}} \rho_{\text{DM}} Q_L m_{\text{DP}} \frac{\beta}{\beta + 1} |\cos \theta|^2 \quad (3.70)$$

3.3.2. The Vector Field Nature of the Dark Photon

Since the dark photon is a vector boson, as opposed to the scalar axion, we must consider the overlap of the dark photon field and the axis of the mode the experiment is aligned with. In SUPAX the cavity is zenith-pointing, meaning its axis (and thus the resonant mode) is aligned with the local vertical: the line from the Earth's centre through the cavity to the zenith. Due to Earth's rotation this axis is time-dependent and can be written as the unit vector

$$\hat{\mathbf{Z}}(t) = \begin{pmatrix} \cos \lambda \cos(\omega t) \\ \cos \lambda \sin(\omega t) \\ \sin \lambda \end{pmatrix} \quad (3.71)$$

in the equatorial coordinate system $(\mathbf{x}_e, \mathbf{y}_e, \mathbf{z}_e)$ where λ denotes the latitude of the experiment, for the setup used in Mainz $\lambda = 50.0^\circ\text{N}$, and $\omega = 2\pi/(1 \text{ day})^*$. For the dark photon field, its polarisation $\hat{\mathbf{X}}$ is either randomly distributed or fixed along an axis. As shown in (3.69), the electric field induced by kinetic mixing of the dark photon field scales with $\mathbf{E} \propto \cos \theta$ where θ is the angle between the cavity's z -axis and the field. For the signal power in the cavity we know that it scales with

$$P_{\text{sig}} \propto \mathbf{E}^2 \propto \cos^2(\theta) = (\hat{\mathbf{Z}} \cdot \hat{\mathbf{X}})^2. \quad (3.72)$$

We can therefore derive that only the fraction

$$\cos^2(\theta) = \frac{P_{\text{sig}}}{P_X} \quad (3.73)$$

of the produced power by dark photon conversion via kinetic mixing, P_X , can be recorded as a signal using the resonant mode along the cavity's axis.

Random polarisation case

If we assume $\hat{\mathbf{X}}$ to be randomly polarised, θ is not time-dependent since it is randomly polarised for *any* time t . Therefore, we can find the expected overlap $\langle \cos^2(\theta) \rangle$ by simply integrating over the solid angle $d\Omega = \sin \theta d\theta d\phi$ of a unit sphere with surface area $\int d\Omega = 4\pi$.

$$\langle \cos^2(\theta) \rangle = \frac{1}{4\pi} \int d\Omega \cos^2(\theta) \quad (3.74)$$

$$= \frac{1}{4\pi} \underbrace{\int_0^{2\pi} d\phi}_{=2\pi} \int_0^\pi d\theta \cos^2(\theta) \sin(\theta) \quad (3.75)$$

*In this context 1 day refers to a sidereal day, which is 23:56:04 h as opposed to the 24 h solar day. Due to the Earth circling the Sun it rotates $\approx 361^\circ$ in a solar day with respect to the galactic rest frame, not 360° .

For the θ integration we can substitute

$$\mu = -\cos \theta \quad (3.76)$$

$$\frac{d\mu}{d\theta} = \sin(\theta) \Leftrightarrow d\theta = \frac{d\mu}{\sin(\theta)} \quad (3.77)$$

and can now calculate the overlap.

$$\langle \cos^2(\theta) \rangle = \frac{2\pi}{4\pi} \int_{\mu(0)=1}^{\mu(\pi)=-1} d\mu \mu^2 \quad (3.78)$$

$$= \frac{1}{2} \left. \frac{\mu^3}{3} \right|_{-1}^1 = \frac{1}{2} \left(\frac{1}{3} - \left(-\frac{1}{3}\right) \right) = \frac{1}{2} \cdot \frac{2}{3} = \frac{1}{3} \quad (3.79)$$

Thus, the signal power is to be reduced by a factor of $1/3$ in the randomised polarisation scenario, independently of when or how long data was taken.

Fixed polarisation case

Conversely, if we assume $\hat{\mathbf{X}}$ to be polarised along a fixed axis

$$\hat{\mathbf{X}} = \begin{pmatrix} \sin \theta_X \cos \phi_X \\ \sin \theta_X \sin \phi_X \\ \cos \theta_X \end{pmatrix} \quad (3.80)$$

the angle $\theta(t)$ is indeed time-dependent, since the overlap of

$$\cos^2 \theta(t) = (\hat{\mathbf{X}} \cdot \hat{\mathbf{Z}}(t))^2 \quad (3.81)$$

is not constant for any time t . We define the time-independent $A = \sin \theta_X \cos \lambda$ and $B = \cos \theta_X \sin \lambda$ to simplify the scalar product.

$$\cos \theta(t) = \hat{\mathbf{X}} \cdot \hat{\mathbf{Z}}(t) = A \underbrace{(\cos(\omega t) \cos \phi_X + \sin(\omega t))}_{=\cos(\omega t - \phi_X)} \sin \phi_X + B \quad (3.82)$$

$$= A \cos(\omega t - \phi_X) + B \quad (3.83)$$

We find the expected overlap for a finite data acquisition time T as the time average over this time interval T .

$$\langle \cos^2 \theta \rangle_T = \frac{1}{T} \int_0^T \cos^2 \theta(t) dt \quad (3.84)$$

$$= \frac{1}{T} \int_0^T [A \cos(\omega t - \phi_X) + B]^2 dt \quad (3.85)$$

$$= \frac{A^2}{2} + B^2 + \frac{1}{\omega T} \left[\frac{A^2}{4} (\sin(2(\omega T - \phi_X)) + \sin(2\phi_X)) + 2AB (\sin(\omega T - \phi_X) + \sin(\phi_X)) \right] \quad (3.86)$$

Since θ_X and ϕ_X are unknown, a first approach in approximating $\langle \cos^2 \theta \rangle_T$ is to sample over θ_X and ϕ_X and take the 5th percentile as a cutoff value for setting a limit with a

3. Theory

95% confidence level. This, however, introduces a statistical impurity into the limit setting, resulting in artificially weakening the true 95% CL limit. In order to construct a statistically robust exclusion limit we first consider the measured power P as a sum of a fully measurable dark matter signal $P_{\text{sig}} = P_X$ and noise P_N .

$$P = P_{\text{sig}} + P_N \quad (3.87)$$

The noise is assumed to be Gaussian distributed around 0 with standard deviation σ_N , which is a fair assumption as will be shown later in the analysis section. An exclusion limit at 95% CL is then obtained by demanding

$$\frac{P_S}{P_N} > \Phi^{-1}(0.95) \simeq 1.64 \quad (3.88)$$

where Φ^{-1} denotes the inverse cumulative distribution function of the normal distribution. When including the vector nature of the dark photon, the situation changes and, as shown in eq. (3.73), the signal power is reduced by the factor $\langle \cos^2 \theta \rangle_T$. Therefore, the total readout power is

$$P = P_{\text{sig}} + P_N = P_X \langle \cos^2 \theta \rangle_T + P_N \quad (3.89)$$

where $\langle \cos^2 \theta \rangle_T$ is the time-averaged overlap defined in eq. (3.84). Consequently, the signal power P_{sig} is not a fixed value but rather a random variable drawn from the distribution $f(\langle \cos^2 \theta \rangle_T)$ and therefore the total power P is drawn from the joint distribution of P_X and the Gaussian P_N . The relevant CL condition follows from the probability that the measured power falls below the expectation from noise alone. Since $\langle \cos^2 \theta \rangle_T$ is distributed according to $f(\langle \cos^2 \theta \rangle_T)$ on the interval $[0, 1]$ (see eq. (3.84), the integral is normalised) and P_N follows a Gaussian distribution $g(P_N)$ the joint probability distribution is simply the product

$$f_{\text{joint}}(x, N) = f(x) g(P_N) \quad (3.90)$$

since the overlap and the noise are independent, where, from here on, we define $x \equiv \langle \cos^2 \theta \rangle_T$ for ease of reading. The probability that the observed total power P is smaller than zero is therefore

$$\mathbb{P}(P < 0) = \int_0^1 dx \int_{-\infty}^{-P_X x} dP_N f(x) g(P_N). \quad (3.91)$$

This expression makes explicit that for each possible overlap x , we must integrate the noise up to the value $P_N = -P_X x$, i.e. the boundary where the total power vanishes as per eq. (3.89). The inner integral is the cumulative distribution Φ of the Gaussian noise evaluated at $-P_X x$. For this we generally know

$$\Phi\left(\frac{x - \mu}{\sigma}\right) = \frac{1}{2} \left[1 + \text{erf}\left(\frac{x - \mu}{\sigma\sqrt{2}}\right) \right] \quad (3.92)$$

and can therefore write

$$\int_0^1 dx f(x) \frac{1}{2} \left[1 + \text{erf}\left(\frac{P_X}{\sigma_N\sqrt{2}} x\right) \right] = \mathbb{P}(\text{miss}) \quad (3.93)$$

where $\mathbb{P}(\text{miss})$ denotes the probability of missing a true signal, for example 5% for 95% CL.

Solving eq. (3.93), given a free choice of CL, for P_X/σ_N gives the critical signal-to-noise ratio (SNR) for the vector case. To relate this to the case ignoring the vector nature of the DP, i.e. $P_{\text{sig}} = P_X$ and therefore eq. (3.88) holds, we define the exclusion factor

$$\langle \cos^2 \theta \rangle_T^{\text{excl.}} = \frac{\Phi^{-1}(\text{CL}) \sigma_N}{P_X} \quad (3.94)$$

which, for the common choice of CL = 95%, is

$$\langle \cos^2 \theta \rangle_T^{\text{excl.}} (\text{CL} = 95\%) \simeq \frac{1.64 \sigma_N}{P_X}. \quad (3.95)$$

This factor quantifies how much the required power for a 95% exclusion is degraded once the fixed polarisation of the DP is included and is thereby equivalent to the 1/3 obtained in eq. (3.79) for the random polarisation case.

In practice the distribution of $f(x)$ is not known analytically* and must be obtained numerically by Monte Carlo simulation, shown in Fig. 3.6.




Figure 3.6.: Distribution $f(\langle \cos^2 \theta \rangle_T)$ for latitude $\lambda = 50.0^\circ\text{N}$ (Mainz) and total data acquisition times of $T = 8438\text{ s}$ (dedicated dark photon run without magnet) and $T = 10318\text{ s}$ (axion run with magnet). The resulting exclusion factors $\langle \cos^2 \theta \rangle_T^{\text{excl.}}$ are significantly higher than the trivial 5th percentile cutoff. They fall into the 19th and 20th percentiles, well within the typical 15-30th percentile range given in [60].

We draw n random orientations (θ_X, ϕ_X) , compute the corresponding $\langle \cos^2 \theta \rangle_T$ and collect these values as a sample from $f(x)$. With this sample, eq. (3.93) can be approximated as

$$\mathbb{P}(\text{miss}) = \int_0^1 dx f(x) \Phi\left(\frac{P_X}{\sigma_N} x\right) \approx \frac{1}{n} \sum_{i=1}^n \Phi\left(\frac{P_X}{\sigma_N} x_i\right) \quad (3.96)$$

for sufficiently large n , where x_i are the simulated values of $\langle \cos^2 \theta \rangle_T$. We then solve for the value P_X/σ_N that yields $\mathbb{P}(\text{miss}) = 1 - \text{CL} = 0.05$ and thereby $\langle \cos^2 \theta \rangle_T^{\text{excl.}}$ via eq. (3.95). The authors of [60] note that the value for $\langle \cos^2 \theta \rangle_T^{\text{excl.}}$ obtained in this way is typically in

*since one cannot know the parameters (θ_X, ϕ_X) , characterizing the polarisation axis of a field of which one does not know if it even exists

3. Theory

the 15-30th percentile of $f(\langle \cos^2 \theta \rangle_T)$ - far above the 5th percentile one might erroneously use, giving a significant enhancement for the limit. This estimation can be run for several integration times, and as one can see in Fig. 3.7 the exclusion factor approaches a limit for increasing integration time, depending on the geographical latitude of the experimental setup. For $\lambda \approx \pm 35.26^\circ$ this approaches the factor 1/3 of the random polarisation case, as shown in [60].




Figure 3.7.: Distribution width of $f(\langle \cos^2 \theta \rangle_T)$, resulting $\langle \cos^2 \theta \rangle_T^{\text{excl.}}$ and comparison to $\langle \cos^2 \theta \rangle_T = 1/3$ for random polarisation at latitude $\lambda = 50.0^\circ\text{N}$ (Mainz) and various total data acquisition times between 0 and 48 h. The exclusion factor $\langle \cos^2 \theta \rangle_T^{\text{excl.}}$ for the dedicated DP run and the axion searching run are marked in green and orange respectively.

3.4. RF-Cavity Based Haloscopes

Haloscopes, especially those in form of RF-cavities (schematically shown in Fig. 3.5), have great potential to reach the QCD-axion-photon coupling for masses below the meV scale. Just like helioscopes, they rely on the assumption of a naturally occurring axion flux in contrast to the LSW approach of stimulating axion production in the lab. While helioscopes assume the Sun as the axion source, haloscopes assume a galactic halo of dark matter (see sec. 3.1). The idea in these experiments is that in a resonant microwave cavity, with a strong external magnetic field inside, an axion converts into a photon via the inverse Primakoff effect. The sensitivity is enhanced in the region near the resonance frequency of the cavity, as derived in section 3.2.4.

Fig. 2.5 already showed the trade-off between bandwidth and sensitivity. By focusing on an extremely narrow frequency range, haloscopes achieve much higher sensitivity than

broadband approaches such as helioscopes or light-shining-through-wall experiments. The drawback is that a conventional metal cavity can only probe one frequency at a time, though this limitation can be mitigated by tuning the cavity. Tuning in this context refers to the possibility of shifting the resonant frequency of a cavity. The most common way of doing this is by inserting a tuning rod. In this setup, a metal rod is inserted in the cavity and by changing the position of the tuning rod, e.g. by rotating it, the resonance frequency changes. This allows for a sweep of a broader frequency range rather than probing just one fixed frequency. This, however, constantly changes the effective volume of the cavity as well as the EM field distribution inside the cavity. Another approach is to shift the actual size of the cavity. For this, a cavity made from two halves is used and these are slowly separated from each other, changing the geometry of the cavity and therefore also the resonant frequency. This comes with issues of its own such as by opening the cavity it is not trivial to make sure that the resonant nature of the cavity still holds true and the EM field is not leaking out.⁶¹

In SUPAX, tuning is conducted via pressure. Although in theory we assume that the experiment is conducted in vacuum we will see in sec. 4.1 that the cooling of the cavity is not provided by a coldhead or dilution fridge in vacuum but a liquid helium cryostat and therefore low pressure (50 – 100 mbar) helium is present in and around the cavity. This allows for tuning by changing the pressure since the resonance frequency depends on the dielectric inside the cavity. The pressure can only be varied in a certain range for cooling reasons, but it allows very granular control over the resonance frequency over a span of close to 1 MHz.

Some leading experiments in this area of research are ADMX,^{62,63} HAYSTAC,⁶⁴ RADES⁶⁵ and CAPP.⁶⁶ CAPP and RADES are of special interest regarding our experiment since they have shown that superconducting cavities are very much a viable option to improve the sensitivity of the measurement, with CAPP even maintaining superconductivity in high magnetic fields.⁶⁷

Haloscope experiments mainly concentrate on the mass regions of $0.4 - 40 \mu\text{eV}$ ($\approx 0.1 - 10 \text{ GHz}$). This is no coincidence but deliberately chosen since many theories⁶⁸⁻⁷¹ predict the axion mass in exactly this range.

3.4.1. Characterisation of an RF System

A haloscope cavity such as in the SUPAX experiment can be characterised as an RF system with two ports. The characterisation is usually performed using a Vector Network analyser (VNA) and yields three fundamental figures: The resonance frequency f_0 , the loaded quality factor Q_L and the coupling coefficient β . With these and eq. (3.65) we can calculate the unloaded quality factor $Q_0 = Q_L(1 + \beta)$, which is inherent to the resonating system, in this case the cavity. Physically it is the ratio of the energy stored in the system to the energy lost per cycle

$$Q_0 = 2\pi \frac{E_{\text{stored}}}{E_{\text{loss per cycle}}} \quad (3.97)$$

where the factor 2π originates from the per-cycle scaling. It depends on the material the cavity is made of, the geometry of the resonator and the surface roughness inside the cavity. For example, a simple electropolishing treatment of the inner walls can improve a cavity's

3. Theory

Q_0 by several percent. The dominant loss mechanism in an RF cavity is ohmic dissipation in the cavity walls. RF currents flow only within the skin depth $\delta = \sqrt{2/(\mu_0\omega\sigma)}$, so the effective surface resistance is

$$R_s = \sqrt{\frac{\omega\mu_0}{2\sigma}} \quad (3.98)$$

where σ is the electrical conductivity of the wall material, e.g. copper. The unloaded quality factor scales inversely with the surface resistance.

$$Q_0 \propto \frac{1}{R_s} \propto \sqrt{\sigma} \quad (3.99)$$

Surface roughness increases the effective surface area and forces the RF surface currents to follow longer paths. This increases the effective surface resistance $R_s^{\text{eff}} > R_s$, thereby increasing dissipation losses and reducing Q_0 . Polishing reduces the microscopic roughness on length scales comparable to the skin depth, lowering R_s^{eff} , and thus increasing Q_0 .

The determination of β is done by analysing the scattering parameters matrix

$$S = \begin{pmatrix} S_{11} & S_{12} \\ S_{21} & S_{22} \end{pmatrix} \quad (3.100)$$

obtained via a Vector Network analyser (VNA). A VNA is both a coherent RF source and a receiver. It generates a continuous-wave RF signal with known amplitude and phase at one port and measures the complex amplitude and phase of the waves reflected from and transmitted through the cavity at all ports. Directional couplers inside the VNA separate the incident (forward) and reflected waves at each port. By sweeping over the resonance frequency, the complex scattering parameters S_{ij} of eq. (3.100) are obtained. In this context “radiated” power refers to power launched by the VNA into a given port of the cavity, “reflected” power is the fraction of that incident wave returning to the same port, and “transmitted” power is the fraction emerging from the opposite port.

In the two-port case of the cavity, the scattering parameters relate the complex amplitudes of the outgoing waves b_i to the incident waves a_j .

$$b_i = \sum_{j=1}^2 S_{ij} a_j \quad (3.101)$$

The reflection coefficients at the ports are

$$S_{11} = \left. \frac{b_1}{a_1} \right|_{a_2=0} \quad S_{22} = \left. \frac{b_2}{a_2} \right|_{a_1=0}, \quad (3.102)$$

i.e. the ratio of the reflected wave to the incident wave at the same port when the other port is matched. The transmission coefficients are

$$S_{21} = \left. \frac{b_2}{a_1} \right|_{a_2=0} \quad S_{12} = \left. \frac{b_1}{a_2} \right|_{a_1=0}, \quad (3.103)$$

describing the fraction of power transmitted from port 1 to port 2 and vice versa. The reflected and transmitted power fractions are given by $|S_{ij}|^2$. At resonance, $|S_{11}|$ and $|S_{22}|$ are directly related to the coupling of the antennas to the cavity mode.

With these parameters, which are measured via the VNA, one can determine the couplings β_1 and β_2 of the respective ports P1 and P2 via

$$\beta_1 = \frac{1 - |S_{11}|}{|S_{11}| + |S_{22}|} \quad (3.104)$$

$$\beta_2 = \frac{1 - |S_{22}|}{|S_{11}| + |S_{22}|}. \quad (3.105)$$

In practice, the coupling factor of the readout antenna is determined from the minimum of the corresponding reflection coefficient (S_{11} for β_1) at resonance, as described in sec. 5.2. The two port expressions of equations (3.104) and (3.105) are given only for completeness. They are not used for the experimental analysis as one port is deliberately severely undercoupled and terminated by design.

The S-parameters are adjusted by lengths L_1 and L_2 of the readout antennas. Port 2 will, in our case, be weakly coupled and used to insert a signal during the characterisation process, as shown in Figure 3.8. In Figure 3.9 one can see a simple simulation of four different antenna




Figure 3.8.: The Vector Network analyser (blue box) generates a sweep over several frequencies. The signal is transmitted through a directional coupler into the cavity. The reflected signal at antenna #1 is redirected by the directional coupler into the readout, measuring S_{11} . The transmitted signal through the cavity is picked up by antenna #2, measuring S_{21} . The same procedure is done in reverse to determine S_{22} and S_{12} . For a given cavity (fixed Q_0) all S-parameters are determined by the length of the antennas L_1 and L_2 .

lengths L_1 while L_2 is held constant. This effect gets more pronounced, the higher the quality factor of the cavity Q_0 is. This means that when improving the experimental setup by, e.g., switching from a simple copper cavity to a superconducting cavity the sensitivity to the

3. Theory

antenna length increases dramatically to the point where 0.1 mm can make a difference in coupling of the order of $\approx 30\%$. This of course makes the process of optimising the cavity's coupling increasingly difficult with increasing Q_0 . An additional factor is that especially for superconductors the surface resistance in the superconducting regime is *very* sensitive to the temperature. An ideally coupled antenna installed in a superconducting cavity at 3.8 K could already be slightly overcoupled at 3.7 K and will certainly be massively overcoupled at 3.5 K. This then creates a trade-off problem between the coupling - affecting the signal strength of a possible signal - and ambient temperature - affecting the system noise - which will play a role in the signal-to-noise ratio.




Figure 3.9.: S_{11} for different antenna lengths L_1 where $l_1 < l_2 < l_3 < l_4$. For $L_1 = l_1$ the antenna is strongly undercoupled, meaning there is significant signal loss in the readout at that antenna. For $L_1 = l_2$ the antenna is less undercoupled, visible by the sharper and deeper peak. For $L_1 = l_3$ the antenna is critically coupled. This is the desirable length for the antenna to minimise signal loss near the resonance frequency. If the antenna is too long ($L_1 = l_4$), i.e. overcoupled, the peak widens as well.

3.4.2. Modes in RF Cavities

As discussed in section 3.2.4, in a cavity there are infinitely many ways in which an EM field can oscillate, which are called modes. Like most axion searching cavity haloscopes, SUPAX is using the so-called TM_{110} mode. For simplicity we assume a perfectly conducting rectangular cavity occupying the region $0 < x < a$, $0 < y < b$, $0 < z < d$ filled with a homogeneous, source-free medium of permittivity ε and permeability μ . All fields are assumed to vary harmonically in time as $e^{j\omega t}$. With this assumption, in the absence of

sources, Maxwell's equations are

$$\nabla \times \mathbf{E} = -j\omega\mu \mathbf{H} \quad (3.106)$$

$$\nabla \times \mathbf{H} = j\omega\varepsilon \mathbf{E} \quad (3.107)$$

$$\nabla \cdot \mathbf{E} = 0 \quad (3.108)$$

$$\nabla \cdot \mathbf{H} = 0. \quad (3.109)$$

We curl (3.106) and insert (3.107). Once more using (3.49), this yields the vector Helmholtz equation

$$(\nabla^2 + k^2) \mathbf{E} = 0 \quad (3.110)$$

$$k = \omega\sqrt{\mu\varepsilon}. \quad (3.111)$$

Each component of $\mathbf{E}$ therefore satisfies the scalar Helmholtz equation. For the transverse-magnetic (TM) modes with respect to the z -direction, the longitudinal magnetic field vanishes.

$$H_z = 0, \quad E_z \neq 0 \quad (3.112)$$

The longitudinal electric field satisfies

$$(\nabla^2 + k^2) E_z = 0 \quad (3.113)$$

and once E_z is known, the remaining field components follow from Maxwell's equation. For a perfect electric conductor, the tangential electric field must vanish at all cavity walls. At the side walls ($x = 0, a$ and $y = 0, b$), E_z is tangential and must therefore satisfy

$$E_z(0, y, z) = E_z(a, y, z) = 0 \quad (3.114)$$

$$E_z(x, 0, z) = E_z(x, b, z) = 0 \quad (3.115)$$

and at the end walls ($z = 0, d$) the tangential field components are E_x and E_y , which must also vanish.

$$E_x(x, y, 0) = E_y(x, y, 0) = 0 \quad (3.116)$$

$$E_x(x, y, d) = E_y(x, y, d) = 0 \quad (3.117)$$

We seek a separable solution of the form

$$E_z(x, y, z) = X(x)Y(y)Z(z). \quad (3.118)$$

Substituting this into (3.113) and dividing by XYZ yields

$$\frac{1}{X} \frac{d^2 X}{dx^2} + \frac{1}{Y} \frac{d^2 Y}{dy^2} + \frac{1}{Z} \frac{d^2 Z}{dz^2} + k^2 = 0. \quad (3.119)$$

We can solve this by separating the variables

$$\frac{d^2 X}{dx^2} + k_x^2 X = 0 \quad (3.120)$$

3. Theory

$$\frac{d^2 Y}{dy^2} + k_y^2 Y = 0 \quad (3.121)$$

$$\frac{d^2 Z}{dz^2} + k_z^2 Z = 0 \quad (3.122)$$

with

$$k^2 = k_x^2 + k_y^2 + k_z^2. \quad (3.123)$$

Applying the boundary conditions (3.114) and (3.115) gives

$$X(x) = \sin\left(\frac{m\pi x}{a}\right) \quad m = 1, 2, 3, \dots \quad (3.124)$$

$$Y(y) = \sin\left(\frac{n\pi y}{b}\right) \quad n = 1, 2, 3, \dots \quad (3.125)$$

with

$$k_x = \frac{m\pi}{a}, \quad k_y = \frac{n\pi}{b} \quad (3.126)$$

where we only consider non-trivial ($E_z \neq 0$) solutions as modes, so $m, n \geq 1$. For TM modes, the transverse electric field components are related to E_z by

$$E_x = -\frac{1}{k_t^2} \frac{\partial^2 E_z}{\partial x \partial z} \quad (3.127)$$

$$E_y = -\frac{1}{k_t^2} \frac{\partial^2 E_z}{\partial y \partial z} \quad (3.128)$$

where

$$k_t^2 = k_x^2 + k_y^2. \quad (3.129)$$

Now using the boundary conditions (3.116) and (3.117) we get

$$\left. \frac{dZ}{dz} \right|_{z=0} = \left. \frac{dZ}{dz} \right|_{z=d} = 0 \quad (3.130)$$

and a solution of (3.122) satisfying this is

$$Z(z) = \cos\left(\frac{p\pi z}{d}\right) \quad p = 0, 1, 2, \dots \quad (3.131)$$

with

$$k_z = \frac{p\pi}{d}. \quad (3.132)$$

The longitudinal electric field is therefore

$$E_z = E_0 \sin\left(\frac{m\pi x}{a}\right) \sin\left(\frac{n\pi y}{b}\right) \sin\left(\frac{p\pi z}{d}\right) \quad (3.133)$$

and the transverse electric field components follow as

$$E_x = -E_0 \frac{k_x k_z}{k_t^2} \cos\left(\frac{m\pi x}{a}\right) \sin\left(\frac{n\pi y}{b}\right) \sin\left(\frac{p\pi z}{d}\right) \quad (3.134)$$

$$E_y = -E_0 \frac{k_y k_z}{k_t^2} \sin\left(\frac{m\pi x}{a}\right) \cos\left(\frac{n\pi y}{b}\right) \sin\left(\frac{p\pi z}{d}\right) \quad (3.135)$$

From the resonance condition (eq. (3.111)) we can also calculate the resonance frequency $f_{mnp} = \omega_{mnp}/2\pi$ just from the side lengths a, b, d as

$$f_{mnp} = \frac{1}{2\sqrt{\mu\epsilon}} \sqrt{\left(\frac{m}{a}\right)^2 + \left(\frac{n}{b}\right)^2 + \left(\frac{p}{d}\right)^2}. \quad (3.136)$$

For the TM_{110} mode used in the SUPAX cavity this yields a resonance frequency of

$$f_{110} = 8.26 \text{ GHz} \quad (3.137)$$

in vacuum, which is 60 MHz away from the next higher mode f_{111} . We can also estimate the overlap of the mode with the magnetic field for the SUPAX cavity by applying

$$C_{mn}^2 = \frac{64}{\pi^4} \frac{1}{m^2 n^2} \quad (3.138)$$

$$C_{11}^2 = \frac{64}{\pi^4} \approx 0.657 \quad (3.139)$$

known from [59] for a cylindrical cavity with a rectangular cross section. Due to the rounded edges in the cavity, the form factor acquired via simulation in CST Studio is slightly lower at $C^2 = 0.629$. A graphical simulation of the mode used can be seen in Fig. 3.10. This mode was chosen since it is the base mode and therefore has the highest quality factor and the setup of the experiment demands a mode aligned with the axis of the external magnetic field, which is along the z -axis of the cavity.




Figure 3.10.: The TM_{110} mode in the SUPAX cavity, visualised in the CST Studio Suite software.⁷²

3.5. Study of Superconducting RF-Cavities

Since an inherently important factor of the sensitivity of a cavity resonator, the unloaded Q-factor, is dependent on the conductivity (or more precisely the surface resistance) of the material the cavity is made of, it is not surprising that the idea of producing such a cavity from superconducting materials was brought up very quickly after the proof of concept of the SIKIVIE type detector. In the following section, after briefly introducing the concept of superconductivity, the vortex state - a typical issue of cavities in external magnetic fields - and the current state of the art solutions of SRF cavities are discussed.

3.5.1. Superconductors Type I and II

Superconductivity is a physical concept, observable in specific materials, so-called superconductors, where electrical resistance vanishes completely. While the resistance of ordinary conductors like copper decreases steadily with decreasing temperatures, a superconducting material has a critical temperature T_c below which the resistance drops to 0. This critical temperature depends on the type of superconductor. Type I superconductors, as discovered in 1911 by Heike Kamerlingh Onnes,⁷³ usually have critical temperatures of $T_c \approx 0 - 10$ K while Type II superconductors, as discovered in 1935 by Rjabinin and Shubnikov,⁷⁴ can have much higher critical temperatures $T_c > 10$ K, hence being called high-temperature superconductors.




Figure 3.11.: Phase diagram of a Type II superconductor.⁷⁵ Above the critical temperature T_c and/or the critical magnetic field B_{c2} the material behaves like a regular conductor, so there is no superconductivity in region c. In region a, below T_c and B_{c1} , the material is completely superconductive and behaves like a Type I superconductor. For temperatures below T_c and external magnetic field $B_{c1} < B_{\text{ext}} < B_{c2}$, in region b, vortices can form in the conductor and some of the magnetic field can penetrate the material.

Another very important difference between Type I and Type II superconductors is that beside a critical temperature there is also a critical magnetic field B_c above which superconductivity breaks. For Type I superconductors this is usually below a field of $B_c < 1$ T while Type II superconductors can maintain superconductivity at higher fields. This clearly is of major importance for the envisioned experiment since the cavity at hand is planned to work in an external magnetic field of $B_{\text{ext}} = 14$ T. While Type I superconductors only have one critical field B_c above which superconductivity breaks and below which the Meissner effect⁷⁶ stops all magnetic flux from penetrating the material, Type II superconductors have two critical magnetic fields B_{c1} and B_{c2} as displayed in Figure 3.11. The first threshold B_{c1} is, analogously to Type I, the field below which there is only the Meissner effect. However, for an external field $B_{c1} < B_{\text{ext}} < B_{c2}$ there is a mixed state in which superconductivity does not break completely but magnetic field vortices can capture some field lines and thereby penetrate the conductor material. These vortices add resistivity and thereby losses for the RF fields inside the cavity. A detailed analysis of this vortex effect has recently been published, backing its results with experimental data of two existing cavities of different geometries.⁷⁷

3.5.2. Vortex State of SRF Cavities

Superconducting cavities are best known from accelerator applications, where niobium cavities routinely achieve extremely high quality factors. For axion haloscopes, however, the operating conditions are fundamentally different. Since the expected signal power scales with $B^2 Q_0$, one would like to combine the highest possible magnetic field with the highest possible quality factor. But this is precisely the regime in which conventional SRF technology meets its main limitation. In the Tesla-scale magnetic fields of haloscopes, a superconductor is generally no longer in the pure Meissner state but in the mixed, or vortex, state.

It is important to distinguish between the formation of vortices and their RF-driven motion. The external static magnetic field is what creates vortices once the relevant field scale exceeds the lower critical field B_{c1} . In addition, real materials are never perfectly homogeneous, so magnetic flux can also remain trapped at defects during cooldown. The RF field, whether actively injected with a VNA or simply present as thermal noise, is therefore not needed to create vortices, but rather drives the vortices that are already present. The oscillating surface currents exert a Lorentz force on the flux lines, causing them to oscillate around their equilibrium positions.

This vortex motion is dissipative and therefore increases the surface resistance. Roughly speaking, each vortex contains a core in which superconductivity is strongly suppressed, so moving a vortex through the material converts part of the stored electromagnetic energy into heat. In an SRF cavity this appears as an additional contribution to the surface resistance and therefore as a reduction of Q_0 , recall eq. (3.99). Vortex motion does not necessarily destroy superconductivity immediately, but rather lowers the quality factor by gradually increasing RF losses.

The strength of this loss depends mainly on three factors. First, a larger external magnetic field implies a larger density of vortices and therefore a larger potential for dissipation. Second, the loss is geometry dependent and becomes particularly harmful in surface regions where the external B -field penetrates the cavity walls and the local RF magnetic field is also large. This is why not only the choice of material, but also the cavity shape, matters. Third, the loss depends on frequency through the depinning frequency. At frequencies well below

3. Theory

this characteristic scale, vortices remain largely confined to their local pinning wells and the losses stay relatively small. Pinning refers to the trapping of vortices at (possibly deliberate) material defects or inhomogeneities, constraining their RF-driven motion. As the drive frequency approaches or exceeds the depinning frequency, pinning becomes progressively less effective and the response approaches the more dissipative flux-flow regime.

This frequency dependence is directly relevant for the planned operating range around 8 GHz. For Nb₃Sn, Posen et al. estimate an effective depinning frequency of about 12 GHz for their cavity model,⁷⁷ while independent microwave measurements on bulk Nb₃Sn in fields up to 12 T found depinning frequencies around 4.5 GHz at temperatures of ≈ 6 K.

Possible mitigation strategies follow directly with one option being geometric optimisation. In the previously mentioned paper Posen et al. developed a cigar-shaped Nb₃Sn cavity specifically to reduce vortex losses, which, at 6 T and a shared resonance frequency of 3.9 GHz outperformed a TESLA-like cavity by about an order of magnitude in Q_0 .

A second route is to use materials that are known, or at least expected, to perform well in strong magnetic fields. This is the most common solution in cavity haloscopes today, mostly relying on REBCO or YBCO tapes. This approach has been proven to work repeatedly, one example being CAPP showing no degradation of the quality factor at 6.93 GHz up to 8 T with a YBCO tape coated polygon shaped cavity.⁶⁷ However, tape-based solutions have also proven to be rather restricted to flat surfaces and forming a high-quality film on a concave surface is technically very difficult, as explicitly pointed out in the aforementioned paper and similar problems being reported by other haloscope groups as well with regard to REBCO tapes. While the tapes are highly promising from the point of view of RF performance in strong fields, they are less attractive when the cavity geometry is not composed of nearly exclusively flat surfaces. Chapter 5.3 discusses one experimental approach in trying to find a non-tape-based solution.

4. The SUPAX Experiment

This chapter presents the experimental details of the SUPAX experiment. Conceptual work on the experiment started in 2022, with the final setup being assembled and physics data being taken for the first time in 2024. From 2026 on it has largely moved to the University of Bonn and is being integrated into the larger GravNet collaboration. This chapter starts with the general experimental setup in Mainz and a discussion of the cavity used in the conducted experimental runs. It then details the engineering efforts taken to track and manage the factors of pressure and temperature inside the cryostat during operation.

4.1. Experimental Setup

The experimental setup is shown schematically in Fig. 4.1. The cavity is hung from the top of the cryostat (“KONTI IT Magno He” by CryoVac) at the end of the inset, which is shown in Fig. 4.2. To determine the loaded quality factor Q_L with a transmission measurement,




Figure 4.1.: Experimental setup of the cryostat and the electronics in the magnet.

the weakly coupled port at the bottom of the cavity can be connected to the VNA. This is done via a 20 dB attenuator, which is installed to eliminate the 300 K thermal noise of the

4. The SUPAX Experiment

terminator outside of the cryostat, once the system has been characterised. The (ideally) critically coupled port on the top of the cavity is connected to a circulator. The circulator enables the determination of the read-out antenna's coupling coefficient β with a reflection measurement. This input is also fitted with an attenuator due to the thermal noise. The readout uses a cryogenic low noise amplifier (LNA). In the experiment a simple high electron mobility transistor (or short HEMT) from LNF was used (LNF-LNC4 16B).





Figure 4.2.: Inset of the cryostat bearing the cavity, heat exchanger and electronics. The temperature sensors located at the cavity and the Pre-Amp are not visible in the picture, only their copper-coloured cables.

A first data acquisition inside a LHe dewar (measurement in November 2022) without circulator revealed some important insights regarding the data acquisition chain. First of all, the simple cavity characterisation with a VNA does not work when a Pre-Amp is installed. The signal fed by the VNA into the weakly coupled port is attenuated but still delivers data for the reflection on the weakly coupled antenna and the transmission through the cavity, so Q_L can be determined. However, the signal fed through the strongly coupled port never reaches the cavity. It gets stopped and reflected by the Pre-Amp, which only passes signals directionally from an input to an output port, not the other way around.

After the characterisation, when trying to take data, a larger problem was discovered. Cable reflections during data taking with the RSA cause an oddly shaped peak. This does not stand out while the cool-down is monitored with the VNA since the power coupled into the cavity by the VNA overpowers any thermal noise. But once this pumping effect is gone and only the cavity's thermal noise is left, one can see that reflections of thermal noise somewhere above the strongly coupled port are of the same order of magnitude as the thermal noise of the cavity itself, as can be seen in Fig. 4.3. Therefore, the circulator was added between cavity and the Pre-Amp. Since it was unclear whether a switch would work at 4 K temperatures and was clearly not operable in the magnetic field it was decided to find a way to characterise and take data in one single setup without the need for a switch.

In this setup including the circulator one cannot determine all S-parameters in one single measurement. Instead, one needs to construct the network of S_{ij} with two measurements and combine them in the following way. One 'transmission' measurement in which the power of the VNA is fed into the weakly coupled port - this provides all S-parameter but S_{11} . We obtain this missing S-parameter by feeding the power of the VNA through the third port of the circulator. The signal gets reflected at the readout antenna, so we can effectively use the S_{12} measurement of this run - named 'reflection' measurement - as S_{11} data for the mixed network. The analysis and results of this mixed network are discussed in chapter 5.2.



- (a) VNA reflection measurement. The signal of the VNA fed into the strongly coupled port (here port 2) does not reach the antenna but only the Pre-Amp, so no reflection measurement on the antenna and subsequent characterisation of it is possible.
- (b) Non-Lorentzian peak shape, even after removing the RSA gain curve. This has most likely been caused by cable reflections and thereby a superposition of reflection on the readout antenna and actual thermal noise of the cavity.

Figure 4.3.: Data taken in a liquid helium bath in November 2022 without a circulator installed. Subsequent measurements taken with a circulator show no such peak shape.

4.2. Data Acquisition

The data acquisition chain is shown schematically in Fig. 4.4. The cavity characterisation is detailed in chapter 5. Once the cavity and antenna are characterised and data acquisition can begin the bottom setup of Fig. 4.4 is assembled. The RSA is accessed via a custom built script. It streams IQ value pairs in real time from the RSA to the PC, which then applies a Fast Fourier Transform to a freely chosen data acquisition time - in all data sets presented here this integration time was set to 1 s - and then saves the resulting spectra in five-minute batches. After each batch a phase alignment is performed. The spectra are stored locally as `.root` files and analysed offline.

4.3. The Cavity

The centrepiece and what comes closest to being truthfully called *the* detector is the copper cavity. The design was chosen based on the need to meet several criteria while maximizing the available volume. The cryostat has an inner diameter of 52 mm, which means that the cavity including its holding structure, as well as a cryogenic SMA cable for the lower antenna had to be no wider than that at any point. Additionally, the dimension should yield a geometry in which the next higher resonant mode is at least 50 MHz higher, such that one can easily distinguish the modes, avoid mode crossings and generally stay clear of any possible confusion. The resonance frequency should be between 8 – 9 GHz as this high mass range is still relatively untouched by other experiments. These constraints led to the decision of constructing the cavity with inner dimensions 150 mm × 30 mm × 22.8 mm, from a solid copper block of 160 mm × 40 mm × 26.8 mm, resulting in an expected resonance at 8.26 GHz with 60 MHz distance to the next higher mode and a higher cross section, and hence volume, than a classic cylindrical cavity with the same resonance frequency (6.84 cm²

4. The SUPAX Experiment




Figure 4.4.: Data acquisition chain in all three combinations. Top and middle configuration are for characterising the antenna coupling and loaded quality factor respectively. The relevant parts are the cavity (A), circulator (B), cryogenic amplifier (C), -20 dB attenuators (D) and terminators for unused ports (E).

rectangular vs. 6.06 cm^2 circular). The cavity is milled in two halves, held together by a set of M2.5 screws, with rounded edges and corners of radius 9 mm to reduce boundary effects. It is hung from the cryostat's inset by an M6 threaded rod, allowing smoothly adjustable mounting height of the cavity via the holding structure shown in Fig. 4.5.

As the cavity will operate in a strong magnetic field all screws and holding structure materials had to be non-magnetic, mainly aluminium and brass. To allow for more efficient cooling, excess weight of the cavity was trimmed, reducing the final weight to roughly 400 g versus 700 g before trimming, with an extra ≈ 60 g for the holding structure and screws.

4.4. Tracking of Environmental Conditions during Data Taking

It has become apparent very quickly that no real data taking is possible without meticulously controlling the environmental conditions, specifically the pressure inside the cryostat (and thereby the dielectric constant inside the cavity) and the temperatures of all involved components: the cavity itself, the helium gas inlet and at least one temperature sensor at the electronics, better one each at the circulator and the Pre-Amp. This is achieved by recording all these data at a fixed time interval into a time series database (InfluxDB), allowing for an efficient slow control by monitoring the InfluxDB in a Grafana display (see Fig. 4.6).




Figure 4.5.: 3D render of the cavity before weight trimming with its custom holding structure and physical cavity after weight trimming.

The pressures inside the internal Helium dewar and the sample chamber are tracked by self constructed sensors, the temperatures are measured with three separate temperature sensors. Two DT-670 silicon diode and one carbon ceramic resistor are used. This has two reasons. Firstly, resistors are much cheaper than professionally constructed temperature sensors, especially for very low temperatures. However, it is the second reason that is crucial. Most cryogenic temperature sensors that are available off the shelf use diodes, which do not work reliably in magnetic fields. On the flip side, the characterisation of the resistor is not stable over long time frames. This however is no big problem. Every cooldown gives the opportunity of recording the current characterisation curve of the resistor with the precisely working temperature sensors. When the setup is cold and the diode based temperature sensors start to drift the freshly characterised resistor remains steady and gives precise temperature values independently of the magnetic field.

4.5. Temperature and Cooling Management

The first measurement with the superconducting NbN coated cavity, which later yielded the data for the first dark photon sweep, also constituted the first successful test of the circulator in the setup. The circulator suppresses reflections from the readout chain and thereby isolates the cavity from impedance mismatches in the electronics and noise sources further up the readout chain. Without the circulator the measured resonance exhibited a distorted peak shape caused by the superposition of the thermal noise in the cavity and reflections from the readout antenna. After installing the circulator this impedance mismatch was eliminated.

4. The SUPAX Experiment



Figure 4.6.: Grafana display for tracking the environmental conditions in the sample chamber and the dewar.

When adapting the inset for operation inside the external magnet, the cryogenic electronics, specifically the circulator and cryogenic preamplifier, had to be relocated further away from the cavity. These components must operate outside the magnetic field region, the circulator in particular has a magnetic shield that can stop fields up to a maximum of 150 mT. Consequently, the electronics were positioned about 20 cm above the magnet (see Fig. 4.1 left side), where the stray field is no higher than 60 mT at full power (14.1 T), presumably lower in the experimental run, as we measured at 12 T due to previous quenching problems at fields higher than that. This relocation introduced an additional challenge related to thermal gradients inside the cryostat.

During cooldown the resonance frequency of the cavity was continuously monitored using the VNA. In this configuration the VNA injects RF power into the cavity via the weakly coupled port and the characteristic Lorentzian peak with its resonance frequency are easily visible. However, when switching the measurement chain to the RSA, which measures only the thermal noise emitted by the cavity without active RF power injection, the resonance appeared inverted. This effect is illustrated in Fig. 4.7 and is independent of the magnetic field being applied and is apparent for temperatures lower than ≈ 60 K.

The inverted resonance indicates that the measured signal is dominated by reflections at the readout antenna rather than by the thermal noise inside the cavity. The spectral shape implies that the electronics are significantly warmer than the cavity due to the temperature gradient in the cryostat, which is a common problem in cavity haloscopes, however the magnitude in this case is problematic.

This scenario is plausible given the geometry of the cryostat. The helium inlet is located at the bottom of the chamber, hence the lowest temperature is near the bottom. The top of the chamber, where the spent helium is pumped out, constitutes the warmest point of the chamber, resulting in a temperature gradient along the inset where components located




Figure 4.7.: Left: typical reflection measurement obtained with a VNA, where the reflected power is minimised at resonance. Right: inverted peak observed with the RSA when the electronics are warmer than the cavity. In this case reflected noise dominates over the thermal noise emitted by the cavity itself.

further above the inlet tend to remain warmer. Since the cavity is mounted at the lowest point while the circulator and Pre-Amp are installed above it, insufficient cooling of the electronics can result in a higher noise temperature than that of the readout chain. In this case reflected noise from the electronics can dominate over the cavity’s thermal emission and produce a spectral shape such as the one seen in Fig. 4.7. Independent measurements of the system noise temperature T_{sys} later confirmed this interpretation (see Sec. 6.1).

To mitigate this effect the thermal coupling of the electronics to the colder regions of the cryostat needed to be improved. In a first step two heat reflectors were removed and a copper plate was mounted to the Pre-Amp, acting as a heat sink. The plate was thermally connected to the Pre-Amp and the circulator using copper tape and was extended 21 cm downward along the inset in order to establish thermal contact with colder regions closer to the helium inlet. This modification already reduced the inversion of the resonance peak, but did not eliminate it entirely.

Further improvement focused on reducing the thermal gradient between the cavity and the electronics. The entire electronics assembly was moved closer to the cavity by 9 cm, placing it nearer to the helium inlet while still remaining below the magnetic field threshold of the circulator. In addition, the routing of the SMA cable was modified. Instead of connecting the Pre-Amp directly to the cryostat feed-through at the top, the cable was first guided downward to the bottom of the cryostat before looping back upward. This ensures that the thermal power* conducted through the cable is dissipated in the coldest region of the cryostat before reaching the electronics. As a result, the cable temperature approaches that of the cavity, thereby reducing the thermal load on the electronics and acting as an additional way of conducting heat away from them. Despite these efforts the cooling of the electronics remained insufficient. To further increase the cooling power a dedicated heat exchanger was designed and implemented. The underlying principle is similar to heat exchangers used in dilution fridges, where the flow of helium is guided through thermally conductive structures to maximise heat transfer. Usually such exchangers consist of two rings connected by thin fins, in order to maximise the available surface area exposed to the

* “Thermal power” in this context does not refer to the signal as usually but rather to the actual thermal conductance of the cable, i.e. a *warm* cable conducting heat from the outside onto the Pre-Amp.

4. *The SUPAX Experiment*

helium flow. Due to manufacturing constraints, however, a simplified design was realised in which the heat exchanger was constructed from a solid copper disk containing drilled channels that guide the helium flow. The implemented design is marked in the final inset (Fig. 4.2).

The heat exchanger is thermally connected to the electronics and forces the helium to pass through the copper structure, thereby improving heat transfer from the electronics to the cold gas. This modification significantly enhanced the cooling of the electronics, although it did not completely eliminate the residual distortion of the peak. The remaining effect becomes evident when comparing the measurement without the magnet, where the electronics could be placed close to the cavity, to the measurement inside the magnet where the electronics are positioned further away (see Sec. 6.7, in particular Figs. 6.10 and 6.16).

5. Cavity Characterisation

This chapter introduces the software used to simulate the cavity before the experimental phase and the subsequent analysis of data obtained by simulation and experimental characterisation data taken with a VNA. It closes by presenting the first experimental results, studying the suitability of NbN as a superconducting coating for the cavity with and without external magnetic field.

5.1. Simulations in CST Studio

The expected characteristics of the cavity were investigated using *CST Studio Suite*,⁷² a commercial software widely used for numerical simulations of RF fields, among other use cases. CST employs methods such as the finite element method, finite integration technique and the transmission line matrix method to solve Maxwell's equations in complex three-dimensional geometries. It is particularly well suited for the design and optimisation of RF components, including resonant cavities and antennas. In this work, CST was used mainly to determine the appropriate antenna lengths required to achieve the desired coupling conditions of the cavity.

To characterise the cavity, two ports are required. As discussed in section 3.4.1, one antenna must be strongly undercoupled, serving as the input port for RF excitation during VNA measurements and the second antenna should be critically coupled for efficient signal extraction with minimal losses. In the first step a detailed three-dimensional CAD model of the cavity was constructed in CST, shown in Fig. 5.1.




Figure 5.1.: 3D model of the cavity and its antennas in CST Studio.

Subsequently, material properties were assigned to each component, according to the real setup. This means copper for the cavity and antenna, vacuum within the cavity volume and PTFE surrounding each antenna's conducting core.

For the relevant temperature regimes (room temperature and cryogenic temperatures) the properties (mainly conductivity, but also thermal contraction) of the conducting elements were then probed in steps to match the results of prior measurements, giving a set of values

5. Cavity Characterisation




Figure 5.2.: Simulation results for various lengths of the critically coupled antenna (left) and the undercoupled antenna (right). A diameter of 1 corresponds to critical coupling, while larger diameter shows an overcoupled and smaller diameter an undercoupled antenna.

which could be used to optimise the antenna length for the measurements searching for DPs and axions.

Once the electric parameters of the cavity and antennas were known, a parameter sweep was performed in which both antenna lengths were varied in steps of 0.1 mm. This resulted in network data like those shown in Fig. 5.2. The coupling behaviour was evaluated using those simulated scattering parameters. The left side of Fig. 5.2 shows the polar representation of the S_{11} parameter for a sweep in which the antenna length of the strongly coupled port is varied between 1 mm and 2 mm, while the second antenna is fixed. In this representation, a critically coupled port is characterised by a circle of diameter 1 through the origin. Overcoupling is indicated by a larger diameter, whereas undercoupled configurations produce a smaller diameter as shown for the deliberately undercoupled port on the right side of Fig. 5.2. The next chapter details how those data are analysed to determine the expected quality factor and coupling parameters.

5.2. Methodology of Characterisation

The analysis of the cavity characterisation data acquired with the VNA and the simulation data from CST Studio follows these steps:

1. Determine loaded quality factor Q_L via the 3 dB method.
2. Plot reflection measurement in a polar plot and apply a circle fit to the data in the peak region. Determine the maximum value of the fit. Rescale the data.
3. Plot the magnitude of the reflection on a linear scale and fit the peak region with an inverted Lorentzian curve.
4. Determine the minimum of the fit and use it to calculate the coupling factor β .
5. Calculate the unloaded quality factor Q_0 .

Step 1: Determine loaded quality factor Q_L via the 3 dB method.

As described in the theory section the loaded quality factor is defined as

$$Q_L = \frac{f_0}{\text{FWHM}} \quad (5.1)$$

where f_0 is the resonance frequency of the cavity and FWHM is the bandwidth 3 dB below the power at resonance frequency. Since the flanks of a very sharp peak (especially important for superconducting cavities) can be very steep and the number of measurement points the VNA takes is limited it is very unlikely that 3 dB below the peak hits a data point perfectly. This is problematic since the calculation of the quality factor is very sensitive to the bandwidth. For this reason, on both sides of the peak, we take the two points that are closest to 3 dB below the peak and interpolate between them. This ensures that the bandwidth is calculated very precisely. Nevertheless we conservatively assume for both the resonance frequency as well as the 3 dB points on both sides of the peak (Δf_b , Δf_a for frequency below and above the peak respectively) an uncertainty of 1 VNA bin width ν . For the error calculation of Q_L this means the following.

$$\Delta f_b = \Delta f_a = \pm \frac{\nu}{2} \quad (5.2)$$

$$\text{FWHM} = f_a - f_b \quad (5.3)$$

$$\Rightarrow \Delta \text{FWHM} = \sqrt{(\Delta f_b)^2 + (\Delta f_a)^2} = \sqrt{2 \cdot \left(\frac{\nu}{2}\right)^2} = \frac{\nu}{\sqrt{2}} \quad (5.4)$$

For the error of Q_L follows

$$\Delta Q_L = \sqrt{\left(\frac{\Delta f_0}{\text{FWHM}}\right)^2 + \left(\frac{f_0}{\text{FWHM}} \cdot \Delta \text{FWHM}\right)^2} \quad (5.5)$$

$$= \sqrt{\left(\frac{\nu}{2 \cdot \text{FWHM}}\right)^2 + \left(\frac{f_0 \nu}{\sqrt{2} \cdot \text{FWHM}^2}\right)^2}. \quad (5.6)$$

Step 2: Plot reflection measurement in a polar plot and apply a circle fit on the data in the peak region. Determine the maximum value of the fit. Rescale data.

As shown in Fig. 5.3, S_{11} (assuming a strongly/critically coupled port 1) is plotted in a polar plot and the peak region (what is visible as a circle here) is fitted with a circle. The maximum magnitude of the fit is determined ($\max(|S_{11}|) = 5.18$ in this example) and all values are scaled by this factor. This ensures that, in the peak region, the effects of the Pre-Amp and cable losses are eliminated and the data is rescaled in a way such that the maximum reflection of the antenna matches the input power. Also, in this plot one determines if the antenna is over- or undercoupled. If the circle fit of the rescaled data has a diameter of more than one the antenna is overcoupled. In this case the origin of the coordinate system is inside the circle. Consequently, if the diameter is less than one, it is undercoupled. A diameter of exactly one represents the ideal case of critical coupling.

5. Cavity Characterisation




Figure 5.3.: On the left is the raw S_{11} data plotted on a complex plane. On the right is the rescaled S_{11} data so that the maximum of the circle fit has $|S_{11}| = 1$ ensuring that in the peak region amplification (from Pre-Amplifier) and cable losses are accounted for. Figure from [78].

Step 3: Plot the magnitude of the reflection on a linear scale and fit the peak region with an inverted Lorentzian curve.

With the data rescaled in a way that eliminates the effect of the Pre-Amp and cable losses, we can now determine the minimum of the reflection S-parameter S_{11} . For this we fit the linear values of $|S_{11}|$ with an inverted Lorentzian, including an offset and a linear term. The fit function is

$$S_{\text{fit}}(f) = \frac{-A\gamma^2}{(f - f_0)^2 + \gamma^2} + o + l(f - f_0) \quad (5.7)$$

where A is a scaling factor, γ corresponds to the peak width, o is the offset and l the linear factor. The minimum of the reflection can be found exactly at the resonance frequency $f = f_0$.

$$S_{\text{min}} = S_{\text{fit}}(f_0) \quad (5.8)$$

$$= -A + o \quad (5.9)$$

$$\Rightarrow \Delta S_{\text{min}} = \sqrt{(\Delta A)^2 + (\Delta o)^2} \quad (5.10)$$

The uncertainties for A and o are obtained from the covariance matrix of the fit. This method of determining ΔS_{min} is only one of two methods. The other is investigating whether widening (narrowing) the fitting window, so increasing (lowering) the amount of points fitted, by 10 % has a noticeable effect on S_{min} . If increasing or narrowing the fitting window has a larger impact on S_{min} than is captured in ΔS_{min} as described in eq. (5.10) ΔS_{min} is

5.3. Performance of a NbN Coated Cavity in a Strong Magnetic Field

determined as

$$\Delta S_{\min} = \frac{S_{\min}^{10\% \text{more}} - S_{\min}^{10\% \text{less}}}{2}. \quad (5.11)$$

In conclusion, $\Delta S_{\min}$ is the larger of eq. (5.10) and eq. (5.11).

Step 4: Determine the minimum of the fit and use it to calculate the coupling factor β .

The coupling factor β can be determined from here with

$$\beta = \frac{1 + S_{\min}}{1 - S_{\min}} \quad (5.12)$$

$$\Rightarrow \Delta\beta = \frac{2\Delta S_{\min}}{(1 - S_{\min})} \quad (5.13)$$

if the antenna is overcoupled and

$$\beta = \frac{1 - S_{\min}}{1 + S_{\min}} \quad (5.14)$$

$$\Rightarrow \Delta\beta = \frac{2\Delta S_{\min}}{(1 + S_{\min})} \quad (5.15)$$

if the antenna is undercoupled. Note here that if the antenna is critically coupled, so $S_{\min} = 0$, both equations yield $\beta = 1 \pm 2\Delta S_{\min}$.

Step 5: Calculate the unloaded quality factor Q_0 .

Finally, calculate the unloaded quality factor and its error. As described in the theory section the unloaded quality factor is given as

$$Q_0 = (1 + \beta)Q_L \quad (5.16)$$

and now we can calculate the error as

$$\Delta Q_0 = \sqrt{((1 + \beta) \cdot \Delta Q_L)^2 + (Q_L \cdot \Delta\beta)^2}. \quad (5.17)$$

5.3. Performance of a NbN Coated Cavity in a Strong Magnetic Field

Section 3.5.2 mentioned that, to this day, the usual approach for superconducting coatings in the context of RF cavities is the use of layered HTS tapes such as YBCO. These tapes bring some inherent disadvantages. One problem is that one relies heavily on industrial manufacturers, some of which have presented themselves as rather reluctant regarding custom orders already manufactured specifically to the needs of a certain cavity. This, albeit annoying, can still be solved by finding a way to precisely cut the tapes into the desired shapes in-house without damaging the structure of the tape. A more difficult problem to solve is that these tapes, due to their layered structure, tend to be very brittle if one tries to apply them to a

5. Cavity Characterisation

rounded surface. The trend at the moment seems to be directed toward problem avoidance rather than problem solving. One example of this is recent discussions about only coating the side walls and not the caps of the cavity, simply avoiding the rounded surfaces. However, recent theoretical studies suggest that spherical cavities might be advantageous for the search of gravitational waves induced by primordial black hole mergers.^{79,80} Attempting to find flat surfaces in a sphere to apply a superconducting tape is literally trying to square the circle, so this approach of problem avoidance is unsuitable here.

Therefore, we have decided to try several superconductors with high critical magnetic fields which can simply be sputtered onto a surface, thereby not relying on a flat surface but enabling the application of the superconductor to rounded surfaces as well. As a first candidate NbN (niobium nitride) was chosen. To test the suitability of NbN as a superconductor in this context, a cavity, identical to those used in previous experimental runs, was produced and shipped to the Instituto Nazionale di Fisica Nucleare di Legnaro (INFN) for electropolishing. After receiving the cavity back, we quickly tested for an improvement of the quality factor. The testing was only done at room temperature but the extrapolated increase in Q_0 is of the order of 8 – 10 %, so, although not large, it is still clearly measurable, even in a non-cryogenic setting. The cavity was then shipped to the University of Siegen where a group around Aleksandr Zubtsovskii deposited a $2\,\mu\text{m}$ thick layer of NbN on the copper cavity via sputtering.




Figure 5.4.: Images of the cavity under test. Left: the electropolished copper cavity. Right: after depositing the NbN superconductor.

Two tests have been conducted. The first was without external magnetic field to find the critical temperature at which the cavity becomes superconducting. The second test was conducted inside the magnet to see how well the superconductivity could withstand the external magnetic field. The critical temperature of the substrate has been determined to be $T_c = 10.14\,\text{K} \pm 0.38\,\text{K}$ using a sigmoidal function to fit the data below 20 K, as shown in Fig. 5.5. The highest reached quality factor is $Q_0 = 342,607 \pm 1,868$. Remarkably, at a relatively constant temperature of 7.4 K, several measurements of Q_0 range from below $300k$ up to roughly $330k$, a difference of $\sim 10\%$. This gives two main implications. Firstly, determining the quality and coupling factor in a physics run must happen in very stable conditions regarding temperature and pressure. Secondly, the superconducting coating clearly works. With the quality factor improving from roughly $40k$ to $340k$ it is very clear that, without the application of an external magnetic field, the NbN coating instantly improves the sensitivity (by increasing the unloaded quality factor) by a factor of at least 8. For probing dark photons it is therefore considered appropriate. The results of one such physics run are discussed in

5.3. Performance of a NbN Coated Cavity in a Strong Magnetic Field




Figure 5.5.: Left: Loaded quality factors for temperatures up to 200 K without the application of a magnetic field. A steep incline at temperatures below 20K is visible. On the right this steep incline is fitted with a sigmoid curve, described to find the critical temperature of superconductors, determined here to be $T_c = 10.14 \text{ K} \pm 0.38 \text{ K}$.

section 6.7.1. The looming question is whether this holds when an external magnetic field is applied. This would imply that it is also suitable for axion detection. To test this, the cavity was placed in the magnet and Q_0 determined repeatedly at different external magnetic fields. The process here was as follows. Cool down as far as possible, take data with the VNA, warm up to above 20 K and change the magnetic field, repeat. The warming up and cooling down again was done to make sure that the magnetic field can actually penetrate into the cavity itself. If the superconductivity holds and the external magnetic field is changed it might freeze-in inside of the cavity and therefore yield false results. Finally, the results can be seen in Fig. 5.6 and Table 5.1.




Figure 5.6.: The unloaded quality factor of the NbN coated cavity, the copper reference cavity (same geometry but without coating) and a similar Nb₃Sn coated cavity⁶⁵ of the RADES collaboration are shown against the applied external magnetic field.

5. Cavity Characterisation

From eq. (3.99) we know that the relation of the unloaded quality factor and the conductivity of the cavity should approximately take the form

$$Q_0 \approx k\sqrt{\sigma} \quad (5.18)$$

and fitting simulations to several measurements yields $k \approx 2.28\sqrt{\frac{\text{m}}{\text{S}}}$ for our cavity. This approximation is accurate to 1%.

B -Field in [T]	Q_L	β	Q_0	conductivity σ in $[\frac{\text{S}}{\text{m}}]$
0.0	322647 ± 1712	0.073 ± 0.001	346103 ± 1854	$2.2 \cdot 10^{10}$
0.25	14297 ± 171	5.690 ± 0.123	95644 ± 2097	$1.7 \cdot 10^9$
0.55	14316 ± 171	4.232 ± 0.068	74904 ± 1317	$1.0 \cdot 10^9$
1.0	13306 ± 148	3.787 ± 0.086	63690 ± 1348	$7.7 \cdot 10^8$
1.5	12745 ± 136	3.126 ± 0.023	52590 ± 629	$5.1 \cdot 10^8$
2.0	12183 ± 124	3.325 ± 0.081	52695 ± 1117	$5.1 \cdot 10^8$

Table 5.1.: Results of the characterisation of the NbN coated cavity inside the magnet. There was no data taken beyond $B = 2$ T since the trend was clear. Note how Q_0 adheres to the previously established dependence $Q_0 \propto \sqrt{\sigma}$.

After measuring several iterations as just described, ramping the magnetic field up to 2.0 T, we concluded that the NbN, just as the previously⁶⁵ tested Nb₃Sn, does not withstand the magnetic field in a way that is suitable for probing axions or gravitational waves in which you want to maximise the magnetic field. Still, NbN as a material is not to be ruled out categorically. It has been shown that at chip scale it can withstand magnetic fields very well, given a thickness of the NbN film at the nanometre scale. However, for such large surfaces as our cavities there is no trivial way of applying nanometre-scale thin films to the inner cavity walls. With sufficient technical development, however, there is room for reevaluating NbN as a possible candidate material, if it can be applied with chip-level precision and thickness to larger surfaces.

6. Axion and Dark Photon Analysis

In this chapter the main results of the SUPAX experiment are presented. First, the analysis of characterisation data regarding the system noise and signal attenuation is shown. Then the main and alternative analysis approaches are discussed in detail, followed by the results of two separate data acquisitions: one with and one without external magnetic field. The results also discuss the choice of analysis approach used to set the claimed limits on DP and axion coupling, their response to a signal injection and how the results match with the expected sensitivity.

6.1. System Noise Temperature

To calibrate the system noise temperature a heated load was placed as a source in the readout chain and the power at a frequency beside any resonance was read out. The load was heated to several temperatures, constantly monitoring the temperature with a sensor, and the resulting spectra were recorded with the RSA. The resulting data is displayed in Table 6.1.

T [K]	V_{rms} [V]
40.00	$(1.92095 \pm 0.00004) \times 10^{-7}$
30.00	$(1.70594 \pm 0.00004) \times 10^{-7}$
24.00	$(1.56089 \pm 0.00005) \times 10^{-7}$
20.00	$(1.46175 \pm 0.00003) \times 10^{-7}$
16.00	$(1.35470 \pm 0.00003) \times 10^{-7}$
10.00	$(1.17428 \pm 0.00003) \times 10^{-7}$
8.00	$(1.10426 \pm 0.00004) \times 10^{-7}$
4.00	$(9.48064 \pm 0.00004) \times 10^{-8}$
3.00	$(9.14664 \pm 0.00003) \times 10^{-8}$
2.32	$(8.96782 \pm 0.00003) \times 10^{-8}$

Table 6.1.: Measured V_{rms} values with uncertainties at different ambient temperatures T_{amb} .

At ambient temperatures reached in the cryostat used the system noise is simply the sum of the ambient temperature of the cavity and the gain of the amplifier.

$$T_{\text{sys}} = T_{\text{amb}} + T_{\text{gain}} \quad (6.1)$$

This approximation holds well for ambient temperatures of > 0.5 K, while for ambient temperatures $T_{\text{amb}} < 0.5$ K the quantum noise becomes the dominating factor. The noise power in a bin of width $\Delta\nu$, including an amplifier with gain G , can therefore be written as

$$P_n = Gk_B\Delta\nu T_{\text{sys}} = Gk_B\Delta\nu(T_{\text{amb}} + T_{\text{gain}}) \quad (6.2)$$

6. Axion and Dark Photon Analysis

$$P_n = \frac{V_{\text{RMS}}^2}{R} \quad (6.3)$$

where $R = 50 \Omega$ is the impedance of the system, and the system temperature T_{sys} can be expressed as the sum of ambient temperature T_{amb} and T_{gain} , where T_{gain} corresponds to the noise temperature of the electronics including the amplifier. Equating them gives

$$\frac{V_{\text{RMS}}^2}{R} = Gk_B\Delta\nu T_{\text{amb}} + Gk_B\Delta\nu T_{\text{gain}} \quad (6.4)$$

where we linearly fit $\frac{V_{\text{RMS}}^2}{R}$ against T_{amb} to get the slope m and intercept n .

$$m = Gk_B\Delta\nu \quad (6.5)$$

$$n = Gk_B\Delta\nu T_{\text{gain}} = mT_{\text{gain}} \quad (6.6)$$

$$\Rightarrow T_{\text{gain}} = \frac{n}{m} \quad (6.7)$$




Figure 6.1.: Noise power of a heated load at specific temperatures and resulting system temperature fit.

The errors of the fit parameters m and n propagate to T_{gain} .

$$\Delta T_{\text{gain}} = \sqrt{\left(\frac{\Delta n}{m}\right)^2 + \left(\frac{n\Delta m}{m^2}\right)^2} \quad (6.8)$$

Inserting the data gathered with the heated load (see Table 6.1) we get

$$m = (1.5309 \pm 0.0082) \cdot 10^{-17} \frac{\text{W}}{\text{K}} \quad (6.9)$$

$$n = (1.2235 \pm 0.0106) \cdot 10^{-16} \text{ W} \quad (6.10)$$

$$\Rightarrow T_{\text{gain}} = (7.99 \pm 0.08) \text{ K}. \quad (6.11)$$

As usual in RF systems we write the gain

$$G = \frac{m}{k_B \Delta \nu} \pm \frac{\Delta m}{k_B \Delta \nu} = 1108.79 \pm 5.92 \quad (6.12)$$

in dB as follows.

$$G_{\text{dB}} = 10 \cdot \log_{10}(G) \pm \frac{10 \cdot \Delta G}{G \cdot \ln(10)} = (30.45 \pm 0.02) \text{ dB} \quad (6.13)$$

Consulting the datasheet of the LNA the amplifier noise is expected to be $T_{\text{gain}} < 4 \text{ K}$ with gain $28 \text{ dB} < G_{\text{dB}} < 36 \text{ dB}$ depending on the power at which the LNA is operated at. From this we can deduce that the gain calculated in eq. (6.13) is within the specs of the device and therefore realistic, although it has apparently not been operated at the ideal power input. The electronic noise T_{gain} is, however, significantly higher than expected. Possible reasons for this include the LNA being warmer than expected during operation, the circulator generating noise, noise from outside the cryostat injecting power into the cavity or simply the RSA producing a lot of noise. Most of these possibilities hint towards a loose SMA connection, although it is not possible to diagnose the true culprit in hindsight. Due to constrained access to the experimental facilities there was no possibility of repeating the measurement to find out the true reason.

6.2. Signal attenuation

The signal attenuation factor η is determined via a fit as it can be approximated by the cable losses up to the Pre-Amp. The attenuation is given by the loss per distance, usually in $\frac{\text{dB}}{\text{m}}$. The loss of a coaxial cable

$$\alpha = \alpha_c + \alpha_d \quad (6.14)$$

can be written as the sum of the loss in the conductor α_c and the dielectric material α_d . Consulting Pozar's *Microwave Engineering*,⁸¹ we know* that the dielectric loss scales with f and $R \propto \sqrt{f}$ with $\alpha_c \propto R$, hence the conductor loss α_c scales with $\sqrt{f}$. We can therefore write eq. (6.14) as

$$\alpha(f) = c\sqrt{f} + df \quad (6.15)$$

where c and d are constants. Inside the cryostat two types of cables have been used, the Qaxial® S02S02-30 and S02S02-40 assembly series in various lengths, which use the Qaxial® HF086-50 and HF141-50 cables respectively. These display the loss of the cable per meter for several frequencies so we can fit eq. (6.15). The fit result for the total loss is plausible, given that at $f = 8 \text{ GHz}$ the datasheets state $\alpha = 1.30 \text{ dB/m}$ (fit 1.321 dB/m) for the S02S02-40 assembly and $\alpha = 2.15 \text{ dB/m}$ (fit 2.149 dB/m) for the S02S02-30 assembly. For the transition from 300 K to 4 K there are no official data from Qaxial. However, the Japanese producer of similar semi-rigid coaxial cables Coax Co., Ltd. provides the loss for both temperatures at

6. Axion and Dark Photon Analysis

	HF141-50	HF086-50
c	0.350 dB/(m√GHz)	0.621 dB/(m√GHz)
d	0.042 dB/(m GHz)	0.049 dB/(m GHz)
α_c at 8.4 GHz	1.014 dB/m	1.798 dB/m
α_d at 8.4 GHz	0.348 dB/m	0.413 dB/m
α at 8.4 GHz	1.362 dB/m	2.212 dB/m

Table 6.2.: Fit of the cable loss parameters for both cable types used at room temperature.




Figure 6.2.: Fit of the cable loss parameters for both cable types used at room temperature.

10 GHz which we can use to find a figure of merit for the cables used in the experiment.⁸² The datasheet of the used cables states that the inner conductor is silver-plated copper, the dielectric is PTFE, and the outer conductor is tin-plated copper. While no tin-plated versions of copper are reported, they specify beryllium copper (copper alloy with 0.5 – 3% beryllium) and cupronickel (copper alloy with 10 – 40% nickel). Table 6.3 shows that for the Cupronickel cables the attenuation at 4 K is typically smaller by a factor of 2.1 while for the Beryllium copper based cables it is smaller by a factor of 2.9 across all offered outer diameters (3.58 mm and 2.10 mm in case of the cables used in the experiment). Since the exact copper (alloy) composition and quality is unknown we stay conservative and scale down the loss parameters, determined in Table 6.2 at 300 K, by a factor of 2.0 to 4 K.

$$\alpha_{4\text{K}}(8.4\text{ GHz}) = \begin{cases} 0.681\text{ dB/m} & \text{HF141} - 50 \\ 1.106\text{ dB/m} & \text{HF086} - 50 \end{cases} \quad (6.16)$$

*from equations (3.28)-(3.30) and (1.98), (2.85a) in [81]

Cable ID	Outer diameter	Conductor	Attenuation [dB/m]		
			300 K	4 K	ratio
SC-086/50-SCN-CN	0.86 mm	Cupronickel	9.5	4.6	2.1
SC-119/50-SCN-CN	1.19 mm	Cupronickel	6.7	3.3	2.0
SC-219/50-SCN-CN	2.19 mm	Cupronickel	3.8	1.8	2.1
SC-358/50-SCN-CN	3.58 mm	Cupronickel	2.1	1.0	2.1
SC-086/50-SB-B	0.86 mm	Beryllium copper	6.3	2.2	2.9
SC-119/50-SB-B	1.19 mm	Beryllium copper	4.4	1.5	2.9
SC-219/50-SB-B	2.19 mm	Beryllium copper	2.5	0.9	2.8

Table 6.3.: Cable specifications of Coax Co., Ltd. coaxial cables in which PTFE is the dielectric and the inner conductor is the silver-plated variant of the outer conductor. Attenuation is specified at $f = 10$ GHz.⁸²

We can now calculate the signal attenuation factor η

$$\eta = \frac{P_{\text{out}}}{P_{\text{in}}} \quad (6.17)$$

via the total loss L along a cable of length l with loss (hence the negative sign) $\alpha_{4\text{K}}$

$$L = -\alpha_{4\text{K}} \cdot l = 10 \cdot \log_{10} \left(\frac{P_{\text{out}}}{P_{\text{in}}} \right) \quad (6.18)$$

$$\Rightarrow \eta = 10^{\frac{L}{10}}. \quad (6.19)$$

For the DP run the setup specifies that just one 20 cm HF141-50 cable was between the cavity and the Pre-Amp. The attenuation therefore is

$$L = 0.681 \text{ dB/m} \cdot 20 \text{ cm} = -0.1362 \text{ dB} \quad (6.20)$$

$$\eta'_{\text{nomagnet}} = 0.97 \quad (6.21)$$

while for the axion run in the magnet the same 20 cm HF141-50 plus an additional 40 cm and 10 cm HF086-50 cable was used.

$$L = 0.681 \text{ dB/m} \cdot 20 \text{ cm} + 1.106 \text{ dB/m} \cdot (40 \text{ cm} + 10 \text{ cm}) = -0.6892 \text{ dB} \quad (6.22)$$

$$\eta'_{\text{magnet}} = 0.85 \quad (6.23)$$

Assuming a temperature scaling factor of 2.9 instead of the rather conservative 2.0, the attenuation factors would be slightly higher at $\eta'_{\text{nomagnet}} = 0.98$ and $\eta'_{\text{magnet}} = 0.90$, i.e. lower loss of signal power. We'll use this to estimate the error $\Delta\eta$ and assume this error to be symmetrical. We therefore arrive at the final values and their errors.

$$\eta'_{\text{nomagnet}} = 0.97 \pm 0.01 \quad (6.24)$$


$$\eta'_{\text{magnet}} = 0.85 \pm 0.05 \quad (6.25)$$

6.3. Analysis Overview


Having discussed all necessary system characterisations, the following sections will concentrate on the actual analysis and subsequent limit setting of the acquired data. In this section all analysis steps are outlined, to get an overview of the data processing stages. The analysis follows these steps:

1. Study all acquired 1 s integrated spectra and filter out those with unusually high peak width as this indicates a non-stable peak, for example during rapidly changing pressure or temperature of the helium in the cryostat.
2. Each raw spectrum (Fig. 6.3a) is truncated around the resonance peak, which is also fitted.
3. Based on the centre frequency the peak each raw spectrum is sorted into a bin such that spectra of (near) equal cavity resonance can be integrated (Fig. 6.3b) and a histogram (Fig. 6.3c) shows how much data was acquired at each frequency.
4. To remove the cavity resonance structure an SG filter is applied to each integrated spectrum.
5. The SG-filtered spectra are weighted by the Lorentzian cavity structure and combined into one Grand Unified Spectrum (GUS) weighted by their inverse variance (Fig. 6.3d).
6. The sensitive area of the GUS is identified and truncated accordingly (Fig. 6.3e).
7. The limit setting can be done via Bayesian statistics or, simpler and more conservative by treating the GUS based on its weakest point. The former will be called "Main analysis" here, the latter "Alternative analysis" (Fig. 6.3f).


6.3. Analysis Overview




(a) Representative raw (1s integration time) spectrum, including truncation limits in red dashed lines.




(b) Integrated spectra with their respective Lorentzian fit in dashed lines, sorted by bins.




(c) Distribution and binning of raw spectra, based on the frequency of the resonance in each raw spectrum.



(d) Full GUS, marking an excluded bin and boundaries deemed as sensitive in red.



(e) Sensitive region of the GUS with 1, 2 and 3σ highlighted in dashed red lines.



(f) Upper limit of kinetic mixing for randomly polarised DPs via the main and alternative analysis with experimental uncertainty bands.

Figure 6.3.: Overview of the data-processing and analysis chain, from raw spectra through combination into the Grand Unified Spectrum to the two limit-setting approaches.

6.4. Constructing the Grand Unified Spectrum

The data acquisition software writes spectra in V_{RMS}/Hz , integrated over a freely chosen time. We decided on 1 second of integration time. One such raw spectrum can be seen in Fig. 6.3a. The first step of the analysis is truncating the spectra on both sides, only using 6 MHz around the resonance frequency, due to the bandpass filter of the RSA. Then, each spectrum is fitted with a Lorentzian function and grouped by its resonance frequency. Fig. 6.4 shows how the resonance shifts over time due to the (intentional) pressure changes of the helium. Spectra which fall into time intervals in which the change in resonance is very quick are eliminated from further analysis, as are spectra which have an unusually high FWHM, or a distorted peak due to non-stationary resonance. The spectra are now sorted by their




Figure 6.4.: Cavity resonance frequency f_0 over the time of data acquisition after eliminating spectra with high FWHM due to an unstable peak for the dedicated dark photon run (superconducting cavity, no B -field).

resonance frequency, the bin width of the distribution is chosen to be the mean FWHM of the peaks. This way we can construct 5 integrated spectra by simply averaging the spectra which belong to one bin in units of power. The resulting spectra and the associated histogram can be seen in Fig. 6.3b and 6.3c.

Each bin in these integrated spectra (which from now on shall be referred to as simply ‘spectra’) is called δ_{ik}^r where i denotes the index of the spectrum and k denotes the index of the bin, referring to the frequency axis. The upper index, r for *raw* to start with, denotes the processing stage of the analysis. The resonance peak of the cavity is still visible but also some electronic oscillations of the RSA become apparent. To remove these a Savitzky-Golay (SG) filter is used (see sec. 6.5). The process of choosing the right parameters is discussed in the following section. However, to know how to optimise the Grand Unified Spectrum (GUS) we first need to understand how to construct it.

6.4. Constructing the Grand Unified Spectrum

The SG filter is applied to δ_{ik}^r to construct a ‘fit’ (which is really several fits) δ_{ik}^{SG} and then used to filter out the identified large structures in the following way to construct the *filtered* (f) spectra.

$$\delta_{ik}^f = \frac{\delta_{ik}^r}{\delta_{ik}^{SG}} - 1 \quad (6.26)$$

The filtered spectra are then rescaled by dividing them by the corresponding cavity Lorentzian, which is normalised to be 1 at the resonance frequency f_0 . Note that every spectrum i has its own resonance frequency f_0^i .

$$L_{ik}(f) = \frac{1}{1 + 4Q_L^2(f/f_0^i - 1)^2} \quad (6.27)$$

$$\delta_{ik}^s = \frac{\delta_{ik}^f}{L_{ik}} \quad (6.28)$$

For these *rescaled* (s) spectra δ_{ik}^s , the noise of every bin σ_{ik}^s is scaled analogously to eq. (6.28) where σ_i^f is simply the standard deviation of the i th spectrum.

$$\sigma_{ik}^s = \frac{\sigma_i^f}{L_{ik}} \quad (6.29)$$

Finally, the construction of the GUS is achieved by adding all spectra weighted by their inverse variance.

$$\delta_k^g = \sum_i \delta_{ik}^s \cdot \omega_{ik} \quad (6.30)$$

$$\omega_{ik} = \frac{(\sigma_{ik}^s)^{-2}}{\gamma} \quad (6.31)$$

$$\gamma = \sum_i (\sigma_{ik}^s)^{-2} \quad (6.32)$$

The noise of the resulting spectrum σ_k^g is found similarly.

$$\sigma_k^g = \sqrt{\sum_i (\sigma_{ik}^s)^2 \cdot \omega_{ik}^2} \quad (6.33)$$

The resulting GUS over the whole frequency range and zoomed into the sensitive region is shown in Fig. 6.12. As the SG-filter exerts a certain degree of correlation onto the different bins of the spectrum, it is necessary to quantify the bin-to-bin correlation of neighbouring bins, by fitting

$$\delta_{k+1}^g = \rho \cdot \delta_k^g \quad (6.34)$$

in the sensitive region of the GUS. To test if the resulting ρ is statistically significant we simply compare it to the relevant test statistic z for a given confidence level CL, and if

$$|\rho| > \frac{z}{\sqrt{N}} \stackrel{\text{CL}=95\%}{\approx} \frac{1.96}{\sqrt{N}} \quad (6.35)$$

where N is the number of bins in the GUS, one considers ρ to be a significant bin-to-bin correlation, which will be relevant for the limit-setting process.

6.5. Savitzky-Golay Filter Optimisation

To eliminate large structures such as the cavity resonance peak or remainders of the RSA gain curve while leaving small structures such as a mass peak of an axion or DP untouched, the integrated spectra are divided by a Savitzky-Golay (SG) filtered version of themselves. An SG filter is a smoothing filter used to reduce noise in data while preserving important features like peaks. It works by fitting a low-degree polynomial to a moving window of data points using least squares, then replacing the centre point with the fitted value. The goal is to choose a configuration of the SG filter which minimises the required polynomial and maximises the fitting window (the number of bins taken into every fitting step) while ensuring that the resulting SG-normalised data has no remaining structures and its distribution is Gaussian. To ensure that the SG-normalised data behaves Gaussian several configurations of the SG-filter are tested independently, constructing the GUS with each configuration and fitting a Gaussian.




Figure 6.5.: Left: Reduced χ^2 (first row) and Coefficient of determination R^2 of Gaussian distribution fitted to resulting SG-normalised GUS after applying SG filter with different order and window size. Right: ($n=10$) moving average of l.h.s. to identify trends more easily.

The fit quality of the Gaussian for these configurations can be seen in Fig. 6.5, showing two common fit measures: reduced χ^2 and the coefficient of determination R^2 .

$$\chi_{\text{red}}^2 = \chi^2/df = \frac{1}{df} \sum_i \frac{(y_i - f_i)^2}{\sigma_i^2} \quad (6.36)$$

$$R^2 = 1 - \frac{\sum_i (y_i - f_i)^2}{\sum_i (y_i - \bar{y})^2} \quad (6.37)$$

A $\chi_{\text{red}}^2 \approx 1$ generally indicates a good fit whereas $\chi_{\text{red}}^2 \gg 1$ suggests a poor fit and $\chi_{\text{red}}^2 \ll 1$ suggests overfitting. R^2 measures how much of the data variation is explained by the fit in %. Several observations can be made from Fig. 6.5. As there is no immediately clear trend on the left side of Fig. 6.5, one can inspect the fit quality measures in a running average, as can be seen on the right side. Here it becomes clear that the results of using SG-filters of orders 2 & 3 and orders 4 & 5 do not significantly differ from each other. This hints at the conclusion that the uneven polynomials do not contribute significantly to the SG fit. Furthermore one can see the (expected) general trend that with higher polynomial orders and lower window sizes the SG filter is more effective in filtering out unwanted structures (lower χ^2 , better fit of the Gaussian distribution), a sanity check. This might lead to the conclusion that we have several candidate configurations which should be equally good at around the 200 bin mark. For this we shall conduct a first visual inspection of the GUS and see if we can identify any remaining structures in it. The resulting GUS with SG filters of order 3 and 6 have been plotted for illustration in Fig. 6.6 (a) and (b) respectively.

While both of them have a distribution that is well described by a Gaussian, it is apparent that the GUS which has been constructed with an SG filter of order 3 still carries a clearly defined fine structure, an oscillation which is likely an electronic residual of the RSA. This observation shows that it is not enough to simply look at the distribution of the SG-filtered data but one needs to quantify how much of this structure is still present in a constructed GUS, even if its values might seem Gaussian distributed. Since the remaining structure is of oscillating nature it is intuitive to test the power spectral density (PSD). The PSD shows how the power of a signal is distributed across different frequencies. It helps identify dominant oscillatory components by highlighting which frequencies contain the most energy. It is being used here to reveal whether there is still a remaining oscillation in the SG-normalised data. If there is an oscillation, as can be seen with the naked eye in Fig. 6.6 (a) there should be a clear peak in the PSD.

The PSD of the GUS δ_k^g is estimated using Welch's method,⁸³ which divides the spectrum into overlapping segments, applies a window function, computes a discrete Fourier transform of each segment, calculates the squared magnitude of each transform, and averages these spectra to obtain the PSD.

$$PSD(f) = \frac{1}{M} \sum_{m=1}^M \frac{1}{NU} \left| \sum_{\ell=0}^{N-1} w[\ell] \delta_m^g[\ell] e^{-2\pi i f \ell} \right|^2 \quad (6.38)$$

Here, $\delta_m^g[\ell]$ denotes the ℓ -th entry of the m -th segment of the grand unified spectrum δ_k^g , $w[\ell]$ is the window function, N the segment length, M the number of segments, and $U = \frac{1}{N} \sum_{\ell=0}^{N-1} w^2[\ell]$ a normalisation constant. In this work, the PSD is computed using the implementation provided by `scipy.signal.welch` with a segment length of $n = 256$. The resulting PSDs for exemplary grand unified spectra with SG filters of order 3 and 6 can be seen in Fig. 6.6 (e) and (f). It shows remarkably how the remaining oscillation is caught by the power spectral density. Therefore, we shall run this analysis for each configuration of the SG filter.

6. Axion and Dark Photon Analysis




Figure 6.6.: Comparison of GUS, distribution of GUS, and PSD constructed with SG filters of order 3 (left column) and order 6 (right column), window size 201 for both. Note the y-axis scaling in the PSD.

Then we investigate whether any frequency differs greatly from the mean (the presence or absence of such can be seen in Fig. 6.6 (e) or (f) respectively). This ‘Flatness’ of the power spectral density

$$\text{PSDFlatness} = \frac{\sigma(\text{PSD})}{\mu(\text{PSD})} \quad (6.39)$$

is the relative standard deviation of the PSD and is shown in Fig. 6.7.




Figure 6.7.: PSD Flatness for all tested configurations of the SG filter.

This captures the SG filter’s capability of filtering out all systematic structures much better than the initial Gaussian fit analysis. Once more it is visible that orders 2 & 3 and 4 & 5 behave virtually identical, so the hypothesis that odd-order polynomials do not play a significant role in the SG filter is confirmed here. Furthermore it is visible that order 2 & 3 SG filters fail to filter the oscillations already at window sizes around 110 while SG filters of order 4 & 5 manage to filter the RSA remainders up to a window size of 205 very reliably. SG filters of order 6 and higher are efficient at significantly larger window sizes but start to filter out small structures (like an axion signal) as well, so they will not be considered and are only included for illustrative purposes.

The SG filter configuration using order 4 and window size 190 is used for the DP run and window size 150 for the axion run, giving the best trade-off between large window size and low polynomial order. The final resulting GUS of the DP run is shown in Fig. 6.12, while Fig. 6.8 shows the Gaussian distribution and flat PSD (again, note the scale of the y-axis) for this configuration.




Figure 6.8.: Left: Distribution of GUS in green and Gaussian fit in red. Right: PSD of the GUS. The final configuration of the SG filter results in a GUS that is Gaussian distributed and shows no residual oscillations, indicated by the lack of a peak in the PSD.

6.6. Limit Setting

We can search for an axion or DP signal in the Grand Unified Spectrum constructed using the steps in section 6.4. For this we must make at least some assumptions about the lineshape of the axion or DP, since it spans over more than one bin in the spectrum. This is done in two different ways, detailed in sections 6.6.1 and 6.6.2, the results of which are subsequently compared in sec. 6.7.

The main analysis approach is, in its core, a maximum likelihood calculation using the expected lineshape of a dark matter particle. It considers the effects of the analysis steps performed to obtain the GUS, including how the SG filter distorts the mass peak. This main approach is also used to set final limits. The alternative approach uses the GUS to construct a signal-to-noise ratio (SNR) and identifies the power excess in every bin as a Gaussian variable. It then sets the limit with a certain confidence level in a frequency range if there are no breaches of a defined threshold in the GUS. If the threshold is passed at a specific frequency there is no limit set before a rescan is performed at this specific frequency. The results of this analysis are used to confirm the validity of the results obtained from the main analysis.

6.6.1. Main Analysis Approach

For our analysis we assume a spherical, isothermal and fully virialised DM halo. We can therefore assume a velocity distribution of the particles to follow a Maxwell-Boltzmann distribution in the galactic rest frame. However, the lab frame we are in is moving with Earth, and more importantly the Sun, relative to the galactic rest frame with approximately the velocity of the sun around the centre of our galaxy $v_s = 220$ km/s. The resulting distribution is calculated in [84, 85] to be

$$f(\nu) = \frac{2}{\sqrt{\pi}} \left(\sqrt{\frac{3}{2}} \frac{1}{r} \frac{1}{\nu_a \langle \beta^2 \rangle} \right) \sinh \left(3r \sqrt{\frac{2(\nu - \nu_a)}{\nu_a \langle \beta^2 \rangle}} \right) \exp \left(-\frac{3(\nu - \nu_a)}{\nu_a \langle \beta^2 \rangle} - \frac{3r^2}{2} \right) \quad (6.40)$$

where ν_a is the the frequency corresponding to the axion rest mass and $\langle \beta^2 \rangle = \langle v^2 \rangle / c^2$ is the second moment of the Maxwell-Boltzmann distribution defined as $\langle v^2 \rangle = \frac{3}{2} v_c^2$ so with the previously mentioned speed of rotation we get $\langle v^2 \rangle = (270 \text{ km/s})^2$ and for the radius (distance from the galactic centre) we get $r = v_s^2 / \langle v^2 \rangle \approx \sqrt{2/3}$. This lineshape needs to be discretised, since we cannot measure a continuous function but discrete values at a defined bin width $\Delta\nu$. This discretised lineshape D_q can be calculated as follows where q relates to the bin number in the discretised lineshape. Note that both continuous and discretised lineshapes are normalised by design, proof of that can be found in Appendix A.

$$D_q = \int_{\nu_a + (q-1)\Delta\nu}^{\nu_a + q\Delta\nu} f(\nu) d\nu \quad (6.41)$$

Here, however, we need to take into account that if we apply an SG filter to our data we also inevitably apply it to any potential signal of a DM candidate and thereby alter its signal shape. Therefore it must be studied how the SG filter distorts this mass peak. The result can be seen in Fig. 6.9. The entire information of the spectrum is captured within 200 bins




Figure 6.9.: Effect of the SG filter on a virialised mass peak for a particle with rest mass corresponding to f_0 .

(≈ 64 kHz). We can now search for the SG-filtered lineshape

$$D_q^S = \text{SGfilter}(D_q) \quad (6.42)$$

in the GUS with the following algorithm. Fig. 6.9 shows clearly that the lineshape holds significantly more power before the SG filter is applied than after.

$$1 \equiv \sum_q D_q > \sum_q D_q^S \quad (6.43)$$

By design this now considers the inevitable attenuation introduced by using the SG filter on the spectra. To quantify this attenuation we can fit a signal efficiency parameter η' in the form of

$$D_q^S = \eta'_{\text{fit}} \cdot D_q \quad (6.44)$$

which mainly depends on the choice of the SG filter parameters. For the dedicated DP measurement we get $\eta'_{\text{fit}} = 0.407$, meaning that almost 60% of the signal power is lost due to the SG filter treatment of the spectra. For the axion measurement the ≈ 2.4 times larger SG-filter window helps reduce this loss to less than 30% attenuation ($\eta'_{\text{fit}} = 0.704$). One other option of quantifying the signal attenuation is by simply comparing the maximum values of the peak.

$$\eta'_{\text{peak}} = \frac{\max(D_q^S)}{\max(D_q)} \quad (6.45)$$

6. Axion and Dark Photon Analysis

In this case the signal attenuation in the DP run is $\eta'_{\text{peak}} = 0.571$ and for the axion run $\eta'_{\text{peak}} = 0.789$. For the final steps we define the (bin-wise) function

$$y(q, A) = A \cdot D_q^S \quad (6.46)$$

where A is the parameter which in essence describes how large of a signal we can exclude from being present in the GUS with a given confidence level. To find this upper limit on A we can use Bayesian statistics.⁸⁶ From Bayes' theorem we know that the posterior function $p(A|\delta_k^g)$ is defined as

$$p(A|\delta_k^g) = \frac{L(\delta_k^g|A) \pi(A)}{\int L(\delta_k^g|A') \pi(A') dA'} \quad (6.47)$$

where $L(\delta_k^g|A)$ is the likelihood function and $\pi(A)$ is the prior probability distribution function (p.d.f.) for A . Since a negative signal power ($A < 0$) is unphysical we can choose the prior p.d.f. to be

$$\pi(A) = \begin{cases} 1 & A \geq 0 \\ 0 & A < 0. \end{cases} \quad (6.48)$$

From here we can determine the upper limit A_{UL} via

$$1 - \alpha = \int_{-\infty}^{A_{\text{UL}}} p(A|\delta_k^g) dA = \frac{\int_{-\infty}^{A_{\text{UL}}} L(\delta_k^g|A) \pi(A) dA}{\int_{-\infty}^{\infty} L(\delta_k^g|A) \pi(A) dA} \quad (6.49)$$

where a confidence level (CL) of 95% corresponds to $\alpha = 5\%$. Inserting the previously defined prior $\pi(A)$ the lower integration boundaries become 0. The only remaining puzzle piece is the likelihood function. Since we assume δ_k^g Gaussian distributed the likelihood function of each bin should converge to a χ^2 distribution of δ_k^g and $y(k, A)$. The lineshape is N bins wide and we assume these bins to be uncorrelated and Gaussian distributed. We can therefore consult [86] once more and find that we can interpret these N bins as a N -dimensional random vector and therefore the joint p.d.f. is the product of N Gaussians.

$$g = \prod_{k=1}^N \frac{1}{\sqrt{2\pi\sigma_k^g}} \exp\left(-\frac{1}{2} \frac{(\delta_k^g - y(k, A))^2}{(\sigma_k^g)^2}\right) \quad (6.50)$$

Dropping the normalisation terms, which cancel out in eq. (6.49), we can write the likelihood-function as simply the product of the exponential functions and since the product of exponential functions can be written as the exponential of a sum we can write the likelihood function L as

$$L = e^{-\frac{1}{2} S(\delta_k^g, A)} \quad (6.51)$$

$$S(\delta_k^g, A) = \sum_k \frac{(\delta_k^g - y(k, A))^2}{(\sigma_k^g)^2} \quad (6.52)$$

where S can be interpreted as a χ^2 distribution. Here it is important to note that eq. (6.50) only holds under the assumption that there is no significant bin-to-bin correlation in the

GUS, i.e. that the fluctuations in different bins are statistically independent. Under this assumption, the covariance matrix

$$C = \text{diag}((\sigma_1^g)^2, (\sigma_2^g)^2, \dots, (\sigma_N^g)^2) \quad (6.53)$$

of the N -bin GUS window δ^g is diagonal, so that the joint p.d.f. factorises into the product of N one-dimensional Gaussian distributions. If there is significant bin-to-bin correlation in the GUS, the covariance matrix is no longer diagonal but takes the form

$$C = \begin{pmatrix} \sigma_1^2 & \rho \sigma_1 \sigma_2 & 0 & \dots \\ \rho \sigma_1 \sigma_2 & \sigma_2^2 & \rho \sigma_2 \sigma_3 & \dots \\ 0 & \rho \sigma_2 \sigma_3 & \sigma_3^2 & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix} \quad (6.54)$$

with all other off-diagonal elements zero, as we only consider next-bin correlations. The correlation ρ refers to the bin-to-bin correlation of the respective window, the size of which is determined by the SG filtered axion lineshape and the bin width. With this we must generalise the χ^2 distribution to

$$S = (\delta^g - \mathbf{y})^T C^{-1} (\delta^g - \mathbf{y}). \quad (6.55)$$

Inserting this generalised form into eq. (6.51) we no longer treat neighbouring bins as fully independent but respect their correlation. The effect of this correction is small. By design, if $\rho < 0$ this treatment improves the final limit on average. After defining the normalisation factor

$$Z = \int_0^\infty e^{-\frac{1}{2}S(\delta_k^g, A)} dA \quad (6.56)$$

we can bring eq. (6.49) into its final form*

$$1 - \alpha = \frac{1}{Z} \int_0^{A_{\text{UL}}} e^{-\frac{1}{2}S(\delta_k^g, A)} dA. \quad (6.57)$$

The dimensionless quantity A_{UL} can now be used to determine the maximum signal power P_S we can exclude to be present in our data with a confidence level of 95% for every bin k .

$$P_S = A_{\text{UL}} \cdot P_N \quad (6.58)$$

$$P_N = k_B T_{\text{sys}} \Delta\nu \quad (6.59)$$

The expected (Gaussian) noise power per bin P_N is determined by the system noise temperature T_{sys} and the bin width $\Delta\nu$. For every bin the value of A_{UL} is determined numerically such that the integral in eq. (6.57), falls within the interval $[1 - \alpha - \epsilon, 1 - \alpha + \epsilon]$ where $\epsilon = 10^{-4}$ is the tolerance chosen for the confidence level, technically giving the confidence level an uncertainty of $\text{CL} = (95 \pm 0.01)\%$. The algorithm proceeds in two stages. At first, starting from an initial guess, A_{UL} is adjusted exponentially. If the integral is below the

*A mathematically equivalent but computationally much leaner, and hence faster, way of determining A_{UL} is shown in appendix B.

6. Axion and Dark Photon Analysis

lower bound of the interval, A_{UL} is repeatedly doubled until the integral enters or surpasses the target range. Analogously, if the integral is above the upper bound of the interval A_{UL} is halved. Phase two is a classic binary search. Once an upper and a lower bound of A_{UL} is identified (the last value of A_{UL} from phase one, which overshoot the tolerance interval, and the second to last value), a binary search is used to identify the value of A_{UL} with which eq. (6.57) falls within the tolerance interval of the confidence level.

The limit setting process for axions and dark photons is identical until this point. The signal power of an axion coupling to a Standard Model photon is different from a Dark Photon coupling to a SM photon. As has been described numerous times before axions require an external magnetic field while dark photons do not, except for hypothesised scenarios with ludicrously high quality factors of $Q > 10^{21}$ for μeV axions⁸⁷ (16 orders of magnitude higher than our cavity) in which axion detection might be possible without the use of an external magnetic field. Therefore every axion data run (with magnetic field) can be considered a dark photon data run but not vice versa. Another consequence is that should one actually detect an axion signal with sufficient confidence it is, until here, indistinguishable from the signal a randomly polarised dark photon field would show*. To exclude the possibility that, instead of axions, dark photons are responsible for the power excess one would simply switch off the magnet, re-scan the exact frequency at which the signal was observed and if the signal vanishes the only conclusion can be that the signal must have come from coupling to the magnetic field, hence an axion signal. These rescans should either be done over a time frame of at least a few hours or multiple times spread across different times of day to avoid the coincidental alignment of a fixed polarised DP field orthogonal to the B -field (showing a signal) in the primary data taking and subsequent parallel alignment to the B -field (showing no or very weak signal) in the second data taking.[†] For axions the signal power is

$$P_S = P_a = g_{a\gamma\gamma}^2 \rho_{\text{DM}} \frac{\beta}{\beta + 1} \frac{1}{m_a} B^2 V_{\text{eff}} Q_L \eta \delta_a \quad (6.60)$$

and for dark photons it is

$$P_S = P_{DP} = \chi^2 \rho_{\text{DM}} \frac{\beta}{\beta + 1} m_{DP} V_{\text{eff}} Q_L \eta \delta_{DP}. \quad (6.61)$$

These resemble the equations for the signal power from theory sections, namely eq. (3.66) for axions and (3.70) for dark photons, with the addition of the factor η due to cable signal attenuation, discussed in sec. 6.2 and the factors

$$\delta_a = \hbar^2 c^5 \varepsilon_0 \quad (6.62)$$

$$\delta_{DP} = \begin{cases} \hbar^{-1} \langle \cos^2 \theta \rangle_T & \text{random polarisation} \\ \hbar^{-1} \langle \cos^2 \theta \rangle_T^{\text{excl.}} & \text{fixed polarisation} \end{cases} \quad (6.63)$$

*If by all means one does not want to switch off the magnetic field a fixed polarisation DP could be tested for by conducting the same measurement at a later time of day, due to the rotation of the Earth but simply switching off the magnet is much more straightforward.

[†]Realistically, an axion discovery would most certainly prompt several other experiments to scan and analyse the same frequency to verify or falsify the results. However, should a sudden wave of trust strike the physics community for the very first time in history and every other experimentalist say “Well, looks like our 20 years of research were simply in the wrong frequency range, congratulations to you.” one should at least verify the type of particle one claims to have discovered.

accounting for the vector particle nature of the dark photon and the fact that the theory chapter is written in the usual particle physics convention of natural units $\hbar = c = \varepsilon_0 = 1$.

We can derive the couplings $g_{a\gamma\gamma}$ for axions and χ for dark photons with their respective experimental uncertainties by equating (6.58) with (6.60) for axions and (6.61) for DPs. This way we obtain

$$g_{a\gamma\gamma} = \sqrt{A_{\text{UL}} \frac{k_B T_{\text{sys}} \Delta\nu m_a}{\rho_{\text{DM}} B^2 V_{\text{eff}} Q_L \eta \delta_a} \cdot \frac{\beta + 1}{\beta}} \quad (6.64)$$

$$\Delta g_{a\gamma\gamma} = g_{a\gamma\gamma} \cdot \sqrt{\left(\frac{\Delta T_{\text{sys}}}{2T_{\text{sys}}}\right)^2 + \left(\frac{\Delta\beta}{2(\beta^2 + \beta)}\right)^2 + \left(\frac{\Delta B}{2B}\right)^2 + \left(\frac{\Delta Q_L}{2Q_L}\right)^2 + \left(\frac{\Delta\eta}{2\eta}\right)^2 + \left(\frac{\Delta V_{\text{eff}}}{2V_{\text{eff}}}\right)^2} \quad (6.65)$$

for axions, while for dark photons we get the following.

$$\chi = \sqrt{A_{\text{UL}} \frac{k_B T_{\text{sys}} \Delta\nu}{\rho_{\text{DM}} m_{\text{DP}} V_{\text{eff}} Q_L \eta \delta_{\text{DP}}} \cdot \frac{\beta + 1}{\beta}} \quad (6.66)$$

$$\Delta\chi = \chi \cdot \sqrt{\left(\frac{\Delta T_{\text{sys}}}{2T_{\text{sys}}}\right)^2 + \left(\frac{\Delta\beta}{2(\beta^2 + \beta)}\right)^2 + \left(\frac{\Delta V_{\text{eff}}}{2V_{\text{eff}}}\right)^2 + \left(\frac{\Delta Q_L}{2Q_L}\right)^2 + \left(\frac{\Delta\eta}{2\eta}\right)^2} \quad (6.67)$$

To recast the axion limit at frequency f into a dark photon limit, one can equate (6.60) and (6.61) to show

$$\chi = g_{a\gamma\gamma} \frac{B\sqrt{\varepsilon_0 \hbar^3 c^5}}{m_{\text{DM}} |\cos \theta|} \quad (6.68)$$

where $m_{\text{DM}} = m_{\text{DP}} = m_a = hf$.

6.6.2. Alternative Analysis Approach

An alternative approach to limit setting, as opposed to the maximum likelihood fit, is in essence a comparison of the expected (Gaussian) noise to the signal-to-noise ratio. For this we need to alter the way in which we calculate the GUS. The scaling factor $L_{ik}(f)$ given in (6.27) is altered in such a way that the normalised spectra δ_{ik}^f are transformed into a normalised power excess. This is done by using the alternative scaling factor

$$\tilde{L}_{ik}(f) = L_{ik}(f) \cdot \frac{P_{ik}(f)}{k_B T_{\text{sys}} \Delta\nu} \quad (6.69)$$

where

$$P_{ik}(f) = (g_{a\gamma\gamma}^R)^2 \rho_{\text{DM}} \frac{\beta}{\beta + 1} B^2 V_{\text{eff}} Q_L \eta \delta_a \frac{1}{hf} \quad (6.70)$$

is the power an axion would deposit at a given frequency in a single bin, given a hypothetical axion coupling $g_{a\gamma\gamma}^R$, which is temporarily chosen here freely. Now, we can write

$$\tilde{\delta}_{ik}^s = \frac{\tilde{\delta}_{ik}^f}{\tilde{L}_{ik}} \quad \text{and} \quad \tilde{\sigma}_{ik}^s = \frac{\tilde{\sigma}_i^f}{\tilde{L}_{ik}} \quad (6.71)$$

6. Axion and Dark Photon Analysis

where, in the absence of a signal, we can identify every bin k of every spectrum i as a Gaussian variable with expected mean $\mu_{ik} = 0$ and standard deviation $\tilde{\sigma}_{ik}^s$, but more importantly we can write

$$R_{ik}^s = (\tilde{\sigma}_{ik}^s)^{-1} \quad (6.72)$$

as the signal-to-noise ratio (SNR) an axion with coupling $g_{a\gamma\gamma}^R$ would leave in that bin. We construct the GUS just as before

$$\tilde{\delta}_k^g = \sum_i \tilde{\delta}_{ik}^s \cdot \tilde{\omega}_{ik} \quad (6.73)$$

$$\tilde{\omega}_{ik} = \frac{(\tilde{\sigma}_{ik}^s)^{-2}}{\tilde{\gamma}} \quad (6.74)$$

$$\tilde{\gamma} = \sum_i (\tilde{\sigma}_{ik}^s)^{-2} \quad (6.75)$$

$$\tilde{\sigma}_k^g = \sqrt{\sum_i (\tilde{\sigma}_{ik}^s)^2 \cdot \tilde{\omega}_{ik}^2} \quad (6.76)$$

and can now identify

$$R_k^g = \sqrt{\sum_i (R_{ik}^s)^2} = (\tilde{\sigma}_k^g)^{-1} \quad (6.77)$$

as the SNR resulting from an axion signal with coupling $g_{a\gamma\gamma}^R$. We can also identify $\tilde{\delta}_k^g / \tilde{\sigma}_k^g$ as the normalised power excess of an axion signal in every bin k , given the coupling $g_{a\gamma\gamma}^R$. This, however, assumes that the axion deposits all its energy in one single bin. Since the lineshape of a virialised axion (see eq. (6.40)) is much broader than the bin width of our spectra this needs to be addressed. We change the notation, such that

$$D_k^g = \frac{\tilde{\delta}_k^g}{(\tilde{\sigma}_k^g)^2} \quad (6.78)$$

is now a Gaussian variable with mean 0 and standard deviation R_k^g . Now we apply a weighted moving average, where the weights

$$w_{k,q} = \frac{A_{k,q}^S}{\sum_r A_{k,r}^S} \quad (6.79)$$

are given by the discretised, SG-filtered lineshape A_q^S (see eq. (6.42)* and Fig. 6.9). We can write this as

$$D_k^l = \sum_q w_{k,q} D_{k+q-1}^g \quad (6.80)$$

$$R_k^l = \sqrt{\sum_q w_{k,q}^2 (R_{k+q-1}^g)^2} \quad (6.81)$$

*altered notation for the lineshape from $D_{k,q}^S$ in (6.42) to $A_{k,q}^S$ to avoid confusion with the previously defined D_k^g

6.6. Limit Setting

where we define the new index l for *lineshape weighted*. Analogously to eq. (6.77) and (6.78) we obtain the normalised power excess after lineshape weighting δ_k^l/σ_k^l via

$$\delta_k^l = \frac{D_k^l}{(R_k^l)^2} \quad (6.82)$$

$$\sigma_k^l = (R_k^l)^{-1} \quad (6.83)$$

where in the absence of an axion signal we still obtain mean $\mu(\delta_k^l/\sigma_k^l) = 0$ and standard deviation $\sigma(\delta_k^l/\sigma_k^l) = 1$.

Similarly to the main analysis approach we need to remind ourselves that eq. (6.81) only holds if D_k^g has no significant bin-to-bin correlation ρ . As the axion measurement has significant bin-to-bin correlation we gain an additional term and calculate R_k^l via

$$R_k^l = (\hat{\sigma}_k^l)^{-1} \quad (6.84)$$

$$(\hat{\sigma}_k^l)^2 = (\sigma_k^l)^2 + 2(\sigma_k^l)^4 \sum_q w_q w_{q+1} \frac{\text{Cov}(\delta_{k+q}, \delta_{k+q+1})}{(\sigma_{k+q}^g \sigma_{k+q+1}^g)^2} \quad (6.85)$$

as shown in [56].

In this normalised power excess we choose a threshold Θ where for bins k in which $\delta_k^l/\sigma_k^l > \Theta$ we consider these bins as candidate signals and flag them for a rescan. By varying Θ one can control the trade-off between the achievable limit with the available data and the amount of rescan candidates. No rescans were conducted for both datasets, hence

$$\Theta = \max(\delta_k^l/\sigma_k^l). \quad (6.86)$$

If we assume an axion signal with $\text{SNR} = R_T$ in any bin k , then δ_k^l/σ_k^l is Gaussian with mean $\mu(\delta_k^l/\sigma_k^l) = R_T$ and standard deviation $\sigma(\delta_k^l/\sigma_k^l) = 1$. This means, that for R_T the choice of $R_T = \Theta$ is poor, since the probability that $\delta_k^l/\sigma_k^l > \Theta$ is only 50%. The appropriate choice to exclude $\delta_k^l/\sigma_k^l > \Theta$ at $\text{CL} = 95\%$ is

$$R_T = \Theta + \Phi^{-1}(\text{CL}). \quad (6.87)$$

Remember, R_k^l is the measured SNR, given an axion coupling $g_{a\gamma\gamma}^R$, so we can now relate the SNR R_T with the SNR at every bin k . We define a helper variable

$$G_k = \sqrt{\frac{R_T}{R_k^l}} \quad (6.88)$$

and can now, for every bin k , relate the axion coupling $g_{a\gamma\gamma}$ we can exclude at a confidence level of $\text{CL} = 95\%$ to the previously chosen $g_{a\gamma\gamma}^R$ as follows.

$$g_{a\gamma\gamma} = G_k \cdot g_{a\gamma\gamma}^R \quad (6.89)$$

We can now use (6.68) to recast this limit for dark photons in the random and fixed polarisation scenarios. Equations (6.67) and (6.65) still apply for the calculation of the uncertainties $\Delta\chi$ and $\Delta g_{a\gamma\gamma}$.

6.7. Results

In this section the results of a dedicated dark photon measurement, performed on April 4th 2024 in the absence of an external magnetic field, and of an axion searching measurement, performed on October 10th 2024, are presented. Limits are set for the dark photon for both random and fixed polarisation scenarios in both measurements and for the axion in the second measurement. The experimental parameters and uncertainties are given in Table 6.4.

Parameter	Dark Photon measurement	Axion measurement
ρ_{DM}	0.45 GeV/cm ³	
V_{eff}	(56280 $\pm$ 170) mm ³	
$\langle \cos^2 \theta \rangle_T$	1/3	
$\langle \cos^2 \theta \rangle_T^{\text{excl.}}$	0.0414	0.0479
Q_L	119647 $\pm$ 11961	27618 $\pm$ 653
β	1.87 $\pm$ 0.04	0.87 $\pm$ 0.06
η	0.97 $\pm$ 0.01	0.85 $\pm$ 0.05
T_{sys}	(11.99 $\pm$ 0.26) K	(9.99 $\pm$ 0.50) K
$\Delta\nu$	320 Hz	960 Hz
τ	8438 s	10 318 s
B	–	(12.0 $\pm$ 0.1) T

Table 6.4.: Summary of experimental parameters and uncertainties for the dedicated dark photon measurement and the axion measurement.

We claim limits that result from the main analysis approach and validate them by comparing the results to those obtained with the alternative analysis approach. The reasoning behind this choice is contained in the last section of this chapter.

6.7.1. Dedicated Dark Photon Measurement

On April 4th 2024 the first data taking took place, in which a total of 8438 s $\approx$ 140 min of data was taken, after excluding spectra which were obtained during times when the resonance frequency of the cavity was unstable due to temperature or pressure fluctuations. Since the resonance peak has a width of $f_0/Q_L \approx 70$ kHz and each spectrum records 6 MHz of data, the raw spectra are truncated at red dashed lines for the remainder of the analysis, as indicated in Fig. 6.10. Then, the individual raw spectra are binned, based on the frequency of the resonance peak within them and then integrated. One such integrated spectrum is seen on the right side of Fig. 6.10. The number of bins is defined by two criteria. No integrated spectrum should have less than 10 min of integration time, and the bins should have no larger distance than the FWHM of the resonance. The results of binning and integrating in this way are presented in Fig. 6.11. The spectra shown on the right side of this plot can be identified with δ_{ik}^r in eq. (6.26).




Figure 6.10.: Left side: A raw (1 s integration time) spectrum, including truncation limits in red dashed lines. Right side: Example of an integrated spectrum.




Figure 6.11.: Distribution and binning of raw spectra, based on the frequency of the resonance in each raw spectrum. On the right are the resulting integrated spectra in light colour with their respective Lorentzian fit in dashed lines.

To the left of the peaks seen in this plot, at a frequency just below 8.4710 GHz we can observe a single bin, which far exceeds the noise in the rest of the spectrum. This single bin power excess is even more visible once we follow through the construction of the grand unified spectrum as described in sec. 6.4 and arrive at Fig. 6.12.




Figure 6.12.: Grand unified spectrum of the dedicated Dark Photon run. On the left the full GUS, marking the excluded bin and boundaries deemed as sensitive in red. The sensitive region is shown on the right side with 1, 2 and 3 σ highlighted in dashed red lines. Its properties are shown in Fig. 6.8.

6. Axion and Dark Photon Analysis

Since a dark matter mass peak is far wider than a single bin - more than 10 kHz (see Fig. 6.9) vs a bin width of 0.32 kHz in this data set - we would expect a dark photon to leave a trace in roughly 30 adjacent bins. Therefore, we can safely exclude a discovery in this instance but this anomaly should still be investigated in the future. The outlook section contains a paragraph on how to confirm that this is almost certainly a relic of the readout electronics and will therefore be excluded. The GUS is now considered in the sensitive area of the experiment (see right side of Fig. 6.12) and the main analysis approach can now directly set the limit with the calculations demonstrated in sec. 6.6.1. The results can be seen in blue in Fig. 6.15.




Figure 6.13.: Comparison of normalised SNR before (first row) and after (second row) lineshape weighting.

For the alternative analysis approach we calculate the normalised signal to noise ratios before (δ_k^g / σ_k^g) and after (δ_k^l / σ_k^l) weighting it with the lineshape. Both, the weighted and unweighted SNR, are described well by a Gaussian distribution with mean 0 and standard deviation 1, although the lineshape weighted SNR has slightly higher standard deviation than expected. The choice of the threshold Θ in the distribution of δ_k^l / σ_k^l determines how candidate signals need to be rescanned. This trade-off is demonstrated visually in Fig. 6.14. As stated in equation (6.86), Θ has been set to the maximum of the distribution of δ_k^l / σ_k^l and therefore no rescan has been conducted. With the values R_T and frequency dependent R_k^l derived from δ_k^l / σ_k^l , we can set the limits with the alternative method and compare it to the results obtained with the main analysis method. No bin-to-bin correlation correction is applied, as $\rho \approx -0.017$ is not statistically significant in the GUS. If it were applied it would improve the limit by 1.6% on average with the greatest improvement of a single bin of 11.5%

and the greatest degradation of a single bin being 5.8%. The final limits are shown in Fig. 6.15 and compared to the results of other experiments in Fig. 6.22.




Figure 6.14.: Connection between Θ , R_T , confidence interval and rescan candidates. Since we set Θ to be the maximum of the normalised SNR, the shaded area describing the rescan candidates is empty in this case.




Figure 6.15.: Upper limit of kinetic mixing for random and fixed polarisation scenario via the main and alternative analysis approach against rest mass of the dark photon in GHz. The shaded area is the experimental uncertainty.

The limits obtained from both methods are compatible. By design the limit obtained via the alternative method is more conservative, utilizing only the least sensitive point in the GUS. The best limit for every frequency f can only be obtained via the main analysis. Section 6.7.3 further discusses how both limits compare.

6.7.2. Axion Search

On October 10th 2024 the measurement inside the magnet was conducted, in which a total of ≈ 172 min* of data was taken. As described in section 4.1 the setup in the magnet required the cryo Pre-Amp to be mounted further up the cryostat to be outside the stray field of the magnet. This likely caused the Pre-Amp to be slightly warmer than the cavity. This had the effect of introducing a tiny amount of thermal noise into the cavity. Since the antenna is optimised to have ideal coupling at the cavity resonance frequency, this is visible in the spectra by means that the readout noise power at the resonance is smallest, since there is more noise coupled into the cavity than out. The effect on the sensitivity is minimal but spotting the resonance frequency must now be done by fitting a Lorentzian with negative amplitude, as can be seen on the left side of Fig. 6.16.




Figure 6.16.: Left side: A raw (1 s integration time) spectrum. Right side: Integrated ≈ 5 min spectrum.

Just as previously, the individual raw spectra are binned, based on their resonance frequencies and then integrated. The number of bins is, once more, defined by two criteria. Firstly, the bin centre frequency is still required to be no further apart than the FWHM of the resonance peak. The FWHM is calculated here via the loaded quality factor.

$$\text{FWHM} = \frac{f_0}{Q_L} \approx 300 \text{ kHz} \quad (6.90)$$

The bins are ≈ 200 kHz apart. This time, due to the larger overall frequency range scanned, we fixed the minimum integration time of an integrated spectrum to 5 min, obtaining the eight bins shown in Fig. 6.17. The spectra shown on the right side of this plot can once more be identified with δ_{ik}^r in eq. (6.26), so we can now construct the grand unified spectrum shown in Fig. 6.18.

*again excluding spectra obtained in unstable temperature or pressure conditions with consequently unstable resonance frequency




Figure 6.17.: Distribution and binning of raw spectra, based on the frequency of the resonance in each raw spectrum. On the right are the resulting integrated spectra in light colour with their respective Lorentzian fit in dashed lines.




Figure 6.18.: Top: Grand unified spectrum of the axion measurement with 1, 2 and 3 σ highlighted in dashed red lines. Bottom: Properties of the GUS. It is Gaussian distributed and shows no residual oscillations, indicated by the lack of a peak in the PSD.

6. Axion and Dark Photon Analysis

For the alternative analysis approach we will again calculate the normalised signal to noise ratios before ($\tilde{\delta}_k^g/\tilde{\sigma}_k^g$) and after (δ_k^l/σ_k^l) weighting them with the lineshape.




Figure 6.19.: Comparison of normalised SNR before (first row) and after (second row) line-shape weighting.

Here, again both are Gaussian with $\mu \approx 0, \sigma \approx 1$, as can be seen in Fig. 6.19. The deviation of μ and σ from their theoretical values is significantly smaller than in the measurement without magnet, so one might assume slightly better statistics, possibly due to the longer total integration time. Again, Θ is set to the $\max(\delta_k^l/\sigma_k^l)$, so no rescan has been conducted. Having obtained R_T and R_k^l from δ_k^l/σ_k^l , we can set the limits and compare them to the main analysis results. In the GUS of this measurement there is significant bin-to-bin correlation $\rho \approx 0.099$, hence we apply the appropriate corrections to both analyses. Compared to the limit one obtains without corrections, the final limit is weakened by 1.1% on average. The effect ranges from a maximum improvement of a single bin of 10.9% to a maximum degradation of 14.6% in the main analysis. In the alternative analysis the average limit is weakened by 1.3% with a maximum improvement of 8.7% and maximum degradation of 16.2%.

The final limits are shown in Fig. 6.20 and compared to the results of other experiments in Fig. 6.21 for the DP kinetic mixing and 6.22 for the axion coupling.




Figure 6.20.: The upper plot shows the limits for the kinetic mixing χ for the random and fixed polarisation scenarios. The lower plot shows the limit for the axion coupling $g_{a\gamma\gamma}$, obtained by the main (blue) and alternative (red) analysis method. Shaded bands are experimental uncertainty.

6. Axion and Dark Photon Analysis




Figure 6.21.: Comparison of the measured DP limit of the axion run (lower mass) and the DP run (higher mass) in the fixed polarisation scenario to other results. Code and limits to create the plot taken from [88].




Figure 6.22.: The measured axion-photon coupling limit in comparison to other results. Code and limits to create plot taken from [88].

6.7.3. Signal Injection Test

To more rigorously test the validity of the results we can inject a signal into the raw data and observe how this influences both analyses. We can use the results from this test to decide if we should claim the limits obtained from the main or the alternative analysis method. We start by calculating the signal power of a randomly polarised dark photon field with coupling equal to the limit claimed by the main approach at two different frequencies: 8.471 217 GHz and 8.471 225 GHz. These are chosen to test the two extreme cases of an “optimistic” and “pessimistic” limit within the bounds of the claimed limits. For this we calculate the total power via eq. (6.61), simulate its distribution (eq. (6.40)) in the discrete power spectrum via eq. (6.41) and then add its corresponding voltage $V = \sqrt{P_S/R}$ onto every raw spectrum. With these signal-injected raw spectra the entire data analysis is carried out.

In the case of the pessimistic limit at 8.471 225 GHz we see a $> 5\sigma$ excess in the lineshape weighted normalised SNR (δ_k^l/σ_k^l in the alternative analysis) - a very clear red flag that would be worth investigating further. Arguably even more interesting is the effect it has on the limit set by both methods when this injected signal is included in the data. For the alternative approach the limit set over the entire scanned frequency is degraded, given the experimental constraint of no possible rescans. This also is the main reason why we choose to set the limit using the main approach rather than the alternative one. To effectively utilise the alternative approach one needs to have the possibility to rescan potential candidates and hence choose the threshold Θ at a level where there is a manageable amount of candidates to rescan and hence obtain a better limit from the exiting data by obtaining a more favourable R_T at all frequencies, subject to the rescan candidates. For the presented analysis the limited available time slots for data acquisition did not allow for rescan candidates. The main approach meanwhile sets the limit for all frequencies just as before, except for the area around the frequency at which the signal was injected - here the limit of previously $\chi \approx 1.91 \cdot 10^{-14}$ degrades to $2.60 \cdot 10^{-14}$.

For the rather optimistic limit at 8.471 217 GHz we do not observe any change in the limit set by the alternative approach when we insert the signal with the power corresponding to $\chi \approx 0.88 \cdot 10^{-14}$ - the limit set with the main approach. This is no surprise since the injected signal is far smaller than the limit excluded by this approach ($2.15 \cdot 10^{-14}$). However, the result of the main approach does change, as it should. Here the claimed limit of previously $0.88 \cdot 10^{-14}$ degrades to $1.09 \cdot 10^{-14}$. It should be emphasised that, when a signal with power corresponding to the claimed coupling limit is injected into the raw data, the main method adjusts dynamically and sets the limit more conservatively than in the absence of such an artificial injection, as seen in Fig. 6.23.

Overall one can conclude that by design the alternative approach is the more conservative one and yields more homogeneous limits which is desirable when aiming for a fixed coupling such as the KSVZ or DSFZ bands for axions. However, in order to utilise the data effectively one needs to have the ability to rescan candidate frequencies at command, which we unfortunately did not have during data taking for practical reasons such as magnet time availability. The main method yields limits that vary stronger over the scanned frequency range but it utilises the available data more effectively while still maintaining certainty of not setting limits that are overly ambitious without the need to rescan candidate frequencies with high SNR.

6. Axion and Dark Photon Analysis




Figure 6.23.: Comparison of the claimed upper limit DP kinetic mixing (blue) and the resulting limits after injecting a signal (red and purple) at two different frequencies with signal power equivalent to the claimed limit. It shows that in the presence of a DP signal of given frequency and kinetic mixing the limit setting process does not exclude said candidate.

6.7.4. Expected Sensitivity

This section evaluates the expected noise fluctuations of the GUS and uses them to assess the statistical consistency of the derived exclusion limits.

In the following the dagger will indicate expected values, while the non-daggered values are the measured quantities. The fact that the expected power fluctuation

$$\sigma_k^\dagger = \frac{1}{\sqrt{\Delta\nu \cdot \tau}} \quad (6.91)$$

scales purely with the bin width $\Delta\nu$ and the total integration time τ only holds for simply averaging over a number of unprocessed spectra, assuming they consist of white noise. Since in the construction of the GUS we undergo several steps we must take these into account when calculating the expected power fluctuations. First we construct the raw spectra δ_{ik}^r by binning and integrating. Then we filter these binned raw spectra by the SG filter as described in eq. (6.26) to obtain δ_{ik}^f . Since the bins do not consist of equal amounts of spectra (number of spectra = integration time in seconds as we chose 1 s integration time per raw spectrum), and therefore have differing expected noise. We must consider this weighting going forward. The goal in choosing the SG-filter parameters is to calculate the filtered spectra δ_{ik}^f such that they can be considered white noise. However, in designing the GUS we rescale the filtered spectra and their uncertainties by L_{ik} as per eq. (6.28) and (6.29), so effectively we alter the magnitude of σ_k^g by a factor $1/L_{ik}$ versus the simple weighted average of δ_{ik}^f . Assuming

that there is no signal and therefore only white noise in the GUS the ML weights used in constructing it (eq. (6.30) and (6.31)) and a simple weighted average are equivalent, since for white noise the power fluctuation follows $\sigma \propto 1/\sqrt{t}$. In the sensitive area shown in Figures 6.12 and 6.18 we assume the GUS to be Gaussian distributed, so we can now investigate how the uncertainties σ_k^g compare to the theoretical noise of the rescaled spectra given by

$$(\sigma_k^s)^\dagger = \frac{\sigma_k^\dagger}{\bar{L}_k} \quad (6.92)$$

where, for every bin k , $\bar{L}_k$ is the average of L_{ik} weighted by the amount of seconds that went into the i -th spectrum, shown in Fig. 6.11 and 6.17 for the DP and axion run respectively. To gain sufficient statistics we temporarily approximate the noise to be constant in the sensitive region of N bins.

$$(\sigma^s)^\dagger = \frac{1}{N} \sum_{k=1}^N (\sigma_k^s)^\dagger \quad (6.93)$$

We compare these to the measured statistical uncertainties of the GUS, σ_k^g , in Fig. 6.24.




Figure 6.24.: Rescaled expected $((\sigma^s)^\dagger)$ and measured (σ_k^g) noise power fluctuations for the DP run (left side) and axion run (right side). For both data sets they agree apart from an offset that is likely to originate from the maximum likelihood weighting of the measured GUS giving an edge over the simple integration time weighting of the expected noise.

For both data sets the measured power fluctuations agree with the expected values, so that we can confirm that the SG filter effectively eliminates the structure of the cavity resonance peak. We can now generate several spectra of white noise with the magnitude of the expected power fluctuations $(\sigma^s)^\dagger$ and calculate the expected A_{UL} n times. We can numerically integrate the distribution of A_{UL} to obtain the 1σ , 2σ and 3σ bands and compare them to the calculated A_{UL} used for setting the DP and axion limits. Due to the choice of the prior (see eq. (6.48)) the resulting distribution of A_{UL} is not symmetric. The $n\sigma$ bands are constructed such that 1, 2 and 3σ entail $\approx 68\%$, 95% and 99.7% of the distribution. Lastly, we rescale with the mean and sigma bands of A_{UL} with $(\sigma_k^s)^\dagger/(\sigma^s)^\dagger$.

6. Axion and Dark Photon Analysis




Figure 6.25.: Distribution of expected A_{UL} and comparison of measured A_{UL} with the expected values for DP run (left) and axion run (right). For both data sets $n = 12000$ simulations of A_{UL} have been conducted.

7. High-Frequency Gravitational Waves

The SUPAX setup, and similar axion-searching RF-cavity haloscopes, can be modified in a minimal way and be used to search for High Frequency Gravitational Waves (HFGW). These HFGW can have several sources, one particularly interesting one is primordial black holes (PBH), which can emit gravitational waves via several different mechanisms. This section will discuss the theoretical background of the GW emission by PBHs and how to detect them in an axion searching RF-cavity haloscope. The chapter will also showcase GRAVNET - a proposed global detector network to search for HFGW.

7.1. Theoretical Background

Primordial black holes, black holes produced in the early universe through non-stellar mechanisms, have been discussed as possible dark matter candidates for a long time. A very appealing argument for PBHs making up (at least a part of) DM is the fact that, in contrast to axions or dark photons, they do not necessarily require the introduction of a new fundamental particle. This does not mean that a relevant PBH abundance arises automatically in standard cosmology, it generally requires sufficiently large small-scale density perturbations or other early universe formation mechanisms, detailed in [89]. In the conventional PBH dark matter scenario PBHs form before Big Bang Nucleosynthesis (BBN) and do not participate in nuclear reactions that determine the cosmic baryon abundance. Their abundance is therefore not constrained by the baryon-density budget inferred from BBN and the CMB in the way that ordinary baryonic matter is. This is what qualifies them for the non-baryonic cold dark matter component of the universe in the cosmological accounting shown in Fig. 1.2. This is what makes them fundamentally different from ordinary (non-primordial) black holes, which have already been observed, e.g. by LIGO in their groundbreaking 2016 discovery. Astrophysical black holes, such as stellar mass remnants, originate from matter already included in the cosmological baryon budget of $\approx 5\%$ of the energy density.

For PBH masses spanning more than 27 orders of magnitude the fraction of dark matter constituted by PBHs

$$f_{\text{PBH}} = \frac{\Omega_{\text{PBH}}}{\Omega_{\text{DM}}} \quad (7.1)$$

has been constrained by several experimental observations since the search for them began in the late 60s. These constraints are shown in Fig. 7.1 and from there it is also clear that the range around $10^{-16} - 10^{-11} M_{\odot}$, where $M_{\odot} \equiv$ solar mass, PBHs could constitute up to 100% of dark matter.⁹⁰ Such low mass PBHs emit gravitational waves with very high frequencies. This is ideal for RF-cavities working in the GHz regime and another advantage of probing for GWs in the GHz band is that for frequencies above the kHz range there are no known background sources of GWs allowing for a somewhat undisturbed search for PBH induced

7. High-Frequency Gravitational Waves




Figure 7.1.: Constraints on the fraction of DM in the form of PBHs, f_{PBH} , with mass M_{PBH} , coming from various types of experiments.⁹⁰ Plot from [91].

signals. These signals can generally be classified into one of three categories: stochastic, coherent (and persistent) or transient.

The first class is the *stochastic gravitational-wave background* (SGWB). In this case one does not observe individual events separately. Instead, the detector is exposed to the superposition of a very large number of GWs arriving from many directions and epochs. In this sense, the SGWB is conceptually similar to the CMB, both are all-sky backgrounds that carry information about the history of the universe. The important difference is that GWs decouple essentially immediately after production and can therefore probe epochs long before recombination, which are inaccessible to electromagnetic observations like the CMB. In PBH scenarios such a background can arise already at formation, for example from GWs generated by the same enhanced primordial perturbations that produce PBHs. A stochastic background can also be built up later from Hawking radiation of evaporating black holes or from a superposition of many unresolved PBH mergers. Primordial backgrounds are expected to be close to isotropic, while a background sourced by PBH binaries can in principle show stronger anisotropies because it is generated by a discrete population of compact objects.^{89,92}

The second class consists of *coherent and persistent* signals. These are individually resolvable sources whose phase evolution remains well defined over an observation time much longer than a single oscillation period, so that the signal appears as a narrow spectral line of near-constant frequency, similar to what one would expect an axion signal to look like. In contrast to the stochastic background, such an emission is not noise-like but coherent, which makes a haloscopes' typical long integration time especially powerful. One realisation of such a source is a very light PBH binary during a very early inspiral stage, where the

emitted frequency changes very slowly, such that it stays in the resonance band of a cavity for a sufficiently long time to integrate. Another example, though more model-dependent, is gravitational wave emission from boson clouds around spinning PBHs through black-hole superradiance. These signals are nearly monochromatic and can persist from days to months or even longer, making them natural targets for haloscopes.^{92,93}

The third class, and most relevant in the following, are *transient* signals. These are short-lived individually resolvable events for which amplitude, phase and frequency evolve significantly during the observation. They are the most familiar type of signals from the perspective of GW astronomy, as exemplified by the binary black hole mergers observed by LIGO. In the case of PBH binaries, observable transient signals are expected mainly from the inspiral phase, though merger and ringdown are also observable in principle. As the orbit of the binary shrinks, the emitted GWs chirp to higher frequencies and amplitudes. For PBHs of sufficiently low mass, these events occur at much higher characteristic frequencies than their stellar-mass counterparts, potentially extending into the GHz regime of cavity haloscopes. At such high frequencies the evolution can be extremely rapid, so that the signal may remain inside a narrow detector band, like a cavity, only for a very short time. For transient sources, detectability is therefore controlled not only by the strain amplitude, but also by the signal duration and by how many cycles are spent inside the detector bandwidth.^{92,94} In the following, the focus will be on these transient signals from PBH mergers.

7.2. Gravnet

Having outlined the theoretical motivation for PBH induced HFGWs, we will now discuss how GRAVNET, the *Global detector network to search for high-frequency gravitational waves*, plans to search for such signals experimentally. Its central idea is to extend the experimental toolbox developed for axion haloscopes to search for GWs in the MHz – GHz regime. This section is largely based on the conceptual design report of GRAVNET.⁹³

In an axion haloscope a static magnetic field enables the conversion of axions into photons inside a resonant cavity. However, a passing GW can also induce an electromagnetic response in the presence of a strong background field, such that a cavity can again amplify an otherwise extremely weak signal. Although the underlying interaction is different, see Fig. 7.2, the experimental approach is similar. Both combine a high- Q RF cavity, a strong magnetic field and cryogenic low-noise readout in order to search for a small excess of RF power. In this sense, SUPAX provides much of the technological foundation on which the search for HFGW can build.

While there exist plenty of other experimental methods for the detection of GWs, such as Weber bars and similar mechanical resonators or laser interferometers as used by LIGO, GRAVNET focuses on cavity based setups as this is currently the most promising approach for near-term progress in the high frequency regime. Also, instead of developing a completely new detector technology from scratch, the project reuses methods, and in part hardware, that are already well established in axion searches. This makes the project quite scalable and at the same time opens a new observational window, since the MHz – GHz GW range is still largely unexplored.

An important difference from axion searches is that PBH binary induced HFGWs do not appear as stationary narrow lines in the frequency spectrum and instead evolve over

7. High-Frequency Gravitational Waves




Figure 7.2.: Comparison of stimulated axion-to-photon conversion via inverse Primakoff effect (left) and GW-to-photon conversion via inverse Gertsenshtein effect⁹⁵ (right).

time as the binary inspirals, as discussed earlier. This means that the cavity does not need to be tuned in the same way as in a conventional axion scan. If the signal sweeps through the resonance band, the detector can still respond resonantly during the time the GW frequency remains close enough to the cavity’s resonance frequency. For GRAVNET, this is a major practical advantage, as it allows the use of simple cavities without elaborate tuning mechanisms while still accessing a broad and astrophysically interesting range of possible signal sources.

The first generation of GRAVNET combines a large, low-frequency cavity with several smaller, high-frequency cavities. On the low-frequency side, a large ($\approx 4.8 \text{ m}^3$) cavity is planned to operate in the FINUDA magnet at the Laboratori Nazionali di Frascati, with a resonance around 150 MHz. On the high-frequency side, smaller cavities ($\approx 2 \text{ L}$) operating in strong commercial magnets (9 – 12 T) at 10 mK are planned for the GHz regime.

The most distinctive feature of GRAVNET, however, is already contained in its name: it works as a *network*, not as an isolated experiment. The reasoning for this is that in a single cavity a small excess, especially a temporary or non-persistent one, can always be caused by a local disturbance of any sort. Electronic noise, thermal fluctuations, spurious RF pickup, or site-specific backgrounds (some of which are discussed in sec. 8.2) are just a few of many suspects. In contrast, a real GW signal should produce correlated responses in detectors at spatially separated locations. These responses need not be identical, since they depend on several factors such as the relative detector orientation, the GW polarisation, and the arrival direction, but they should occur with propagation-time delays consistent with a signal travelling at the speed of light. This makes coincidence and correlation analyses a central element of GRAVNET’s analysis strategy. Another idea might be that instead of having several spatially distant experimental setups with identical resonance frequency, one could use several cavities of different, but similar, resonance frequency. In case of a PBH merger induced GW signal these cavities should register a signal consecutively as the GW frequency sweeps through each of their resonance bands.

This network approach of spatially distant setups provides several advantages. It reduces the likelihood of false positives, since mimicking a correlated event across distant laboratories purely by local noise would be a statistical miracle in itself. With identical reasoning, combining information from several detectors increases the statistical significance of a candidate signal. Furthermore, the spatial separation of the sites makes it possible, at least in principle, to infer information about the direction of the source from relative arrival times.

This is described in detail in sec. 7.3.3. In this respect, the logic of GRAVNET is analogous to established GW observatories such as the LIGO–Virgo–KAGRA network, where multiple detectors are not only included for sensitivity, but also for validation and source reconstruction.

In practice, every detector station will perform measurements and signal searches locally and in (or near) real-time. When a potential signal is identified, this triggers the storage of the corresponding time window at all detector stations, taking the expected delays into account, such that potential signals can be investigated offline in more detail. This makes the experiment scalable, as new nodes can always be added to the network over time. Even different detector technologies could, in principle, become part of the network. In the first stage, GRAVNET will build a network with setups in Bonn, Mainz, and Frascati, with the longer-term goal of expanding toward a more global network.

7.3. Detecting HFGW in RF-Cavities

When a resonant cavity is driven at its resonant frequency, the stored energy $E(t)$ rises towards its steady state value E_{ss}

$$E(t) = E_{\text{ss}} \left(1 - e^{-t/\tau}\right)^2 \quad (7.2)$$

where $\tau = 2Q/\omega_0$ is the ring-up time of the cavity. To assume resonant enhancement of the cavity with its quality factor Q the number of cycles $\mathcal{N}_{\text{cyc}}$ to consider the cavity rung-up is relevant. To reach a chosen energy fraction $\eta = E(t)/E_{\text{ss}}$, the required time is

$$\frac{t}{\tau} = -\ln(1 - \sqrt{\eta}) . \quad (7.3)$$

Using $T = 1/f_0 = 2\pi/\omega_0$, the corresponding number of RF cycles is

$$\mathcal{N}_{\text{cyc}} = \frac{t}{T} = -\frac{Q}{\pi} \ln(1 - \sqrt{\eta}) \quad (7.4)$$

Now it is just a matter of definition at which fraction of stored to steady state η one considers the cavity sufficiently rung-up to assume resonant enhancement. Table 7.1 gives the approximate number of cycles a cavity with quality factor Q_L needs to be rung up to η .

Energy fraction η $\approx \mathcal{N}_{\text{cyc}} [Q_L]$	
90%	0.95
91.5%	1
99.0%	1.69
99.9%	2.42

Table 7.1.: Approximate number of cycles to ring up a cavity for different choices of at which η a cavity is considered "fully" rung-up.

From this we can deduce that to ring up such a cavity one needs at least $\mathcal{N}_{\text{cyc}} \gtrsim Q_L$. For the inspiral of a binary system, consisting of two PBHs of equal mass M_b , the number of

7. High-Frequency Gravitational Waves

cycles spent in the detector bandwidth $f \in [f_{\min}, f_{\max}]$ where

$$f_{\max}, f_{\min} = f_0 \pm \frac{f_0}{Q_L} = f_0 \left(1 \pm \frac{1}{Q_L} \right) \quad (7.5)$$

is given by⁹⁶

$$\mathcal{N}_{\text{cyc}} = \frac{1}{32\pi^{8/3}} \left(\frac{GM_c}{c^3} \right)^{-5/3} \left(f_{\min}^{-5/3} - f_{\max}^{-5/3} \right) \quad (7.6)$$

where the chirp mass is defined as

$$M_c = \frac{(m_1 m_2)^{3/5}}{(m_1 + m_2)^{1/5}} \quad (7.7)$$

and the total mass of the merger is simply $m = m_1 + m_2$. For a binary of two equal mass PBHs ($m_1 = m_2 = M_b$) this results in

$$M_c = \frac{M_b^{6/5}}{(2M_b)^{1/5}} = \frac{M_b}{\sqrt[5]{2}} \quad (7.8)$$

so that one can set an upper limit on M_b by demanding that the merger spends enough cycles near the cavity's resonance frequency to ring it up. For example, demanding $\mathcal{N}_{\text{cyc}} \gtrsim 10^5$ for a cavity with $f_0 = 10$ GHz and $Q_L = 10^5$ yields $M_b \lesssim 10^{-12} M_\odot$. One might notice that, for a fixed $\mathcal{N}_{\text{cyc}}$, the upper limit on the merger mass scales with $M_b \propto 1/f_0$, so for a cavity with the same quality factor but $f_0 = 1$ GHz we get $M_b \lesssim 10^{-11} M_\odot$. For merger events of larger mass than this Gatti et al.⁹⁷ define the effective quality factor

$$Q_{\text{eff}} = \min(Q_L, \mathcal{N}_{\text{cyc}}) \quad (7.9)$$

to account for the not fully rung up cavity and hence not full amplification by Q_L .

As shown in [96] we can express the rate of change of the GW frequency as function of the frequency itself.

$$\dot{f} = \frac{96}{5} \pi^{8/3} \left(\frac{GM_c}{c^3} \right)^{5/3} f^{11/3} \quad (7.10)$$

This allows us to compute the frequency of a GW emitted by a merger of chirp mass M_c at a given time τ where we assume the GW has frequency $f_0 \equiv f(t=0)$. This way we can simulate a time series starting at a frequency of our choice and observe how the frequency evolves in time. Solving eq. (7.10) reveals that $f(t)$ diverges at a finite time t_m corresponding to the time of the merger itself. However, for times $t < t_m$ before the merger, or more precisely $t \leq t_{\text{ISCO}}$, we can solve the differential equation.

$$\int_{f(t=0)}^{f(t=\tau)} f^{-11/3} df = \frac{96}{5} \pi^{8/3} \left(\frac{GM_c}{c^3} \right)^{5/3} \int_0^\tau dt \quad (7.11)$$

$$-\frac{3}{8} f^{-8/3} \Big|_{f(t=0)}^{f(t=\tau)} = \frac{96}{5} \pi^{8/3} \left(\frac{GM_c}{c^3} \right)^{5/3} \tau \quad (7.12)$$

$$f(\tau)^{-8/3} - f_0^{-8/3} = -\frac{256}{5} \pi^{8/3} \left(\frac{GM_c}{c^3} \right)^{5/3} \tau \quad (7.13)$$

$$f(t) = \left[f_0^{-8/3} - \frac{256}{5} \pi^{8/3} \left(\frac{GM_c}{c^3} \right)^{5/3} t \right]^{-3/8} \quad \text{for } t \leq t_{\text{ISCO}} \quad (7.14)$$

where in the last step τ has been renamed $\tau \rightarrow t$. The frequency at the time when the merger objects have reached their Innermost Stable Circular Orbit (ISCO) is

$$f(t_{\text{ISCO}}) = f_{\text{ISCO}} = \frac{1}{12\pi\sqrt{6}} \frac{c^3}{Gm} \approx 4.4 \text{ kHz} \left(\frac{M_\odot}{m} \right) \quad (7.15)$$

in the test mass limit ($m_1 \gg m_2$). At this point, it is useful to examine the (often implicit) assumptions underlying this derivation of the ISCO. First, the spin of the PBHs is neglected. In addition, we are assuming a circular orbit at the time of observation. For an elliptical orbit, the emitted signal would not exhibit a nearly constant amplitude, but would instead be modulated, with peaks occurring near periastron passage. While this assumption is not strictly necessary, it is well motivated, as the emission of GWs reduces orbital eccentricity over time, driving the system toward a quasi-circular orbit. Finally, there is an inconsistency in the mass assumptions. In eq. (7.8), we assumed that both PBHs have equal mass $m_1 = m_2 = M_b$, whereas eq. (7.15) assumes $m_1 \gg m_2$. The ISCO depends, among other factors, on the symmetric mass ratio

$$\nu = \frac{m_1 m_2}{(m_1 + m_2)^2}. \quad (7.16)$$

In experimental practice one usually considers near-equal-mass systems ($\nu = 1/4$), rather than $\nu = 0$ as implied in eq. (7.15). For equal masses, $m_1 = m_2 = M_b$, one obtains⁹⁸

$$f_{\text{ISCO}} \approx 4.7 \text{ kHz} \left(\frac{M_\odot}{m} \right) = 2.35 \text{ kHz} \left(\frac{M_\odot}{M_b} \right). \quad (7.17)$$

To know the maximum simulation time t_{ISCO} for a given start frequency f_0 one can just rearrange eq. (7.13).

$$t_{\text{ISCO}} = \frac{5}{256} \pi^{-8/3} \left(\frac{GM_c}{c^3} \right)^{-5/3} \left(f_0^{-8/3} - f_{\text{ISCO}}^{-8/3} \right) \quad (7.18)$$

For the simulation we now just need the rate of frequency change at time t

$$\dot{f}(t) = \frac{96}{5} \pi^{8/3} \left(\frac{GM_c}{c^3} \right)^{5/3} f(t)^{11/3} \quad (7.19)$$

and can hence write the amplitude of the wave as

$$\Psi(t) = \sin(2\pi \cdot \phi(t)) \quad (7.20)$$

where $\phi(t)$ is the phase of a GW at time t , assuming that at time $t_0 = 0$, corresponding to the start of the observation, or simulation, it has phase $\phi(t_0) = 0$.

$$\phi(t) = \int_{t_0}^t f(\tau) d\tau = \int_{f_0}^{f(t)} \frac{f}{\dot{f}} df \quad (7.21)$$

7. High-Frequency Gravitational Waves

For now $\Psi(t)$ is bound between $[-1, 1]$, the amplitude will be discussed later on. First we find an explicit term for $\phi(t)$ by expressing $\dot{f}$ in terms of f using eq. (7.19).

$$\phi(t) = \frac{5}{96} \pi^{-8/3} \left(\frac{GM_c}{c^3} \right)^{-5/3} \int_{f_0}^{f(t)} \frac{f}{f^{11/3}} df \quad (7.22)$$

$$= \frac{5}{96} \pi^{-8/3} \left(\frac{GM_c}{c^3} \right)^{-5/3} \left(-\frac{3}{5} f^{-5/3} \right) \Big|_{f_0}^{f(t)} \quad (7.23)$$

$$= \frac{1}{32} \pi^{-8/3} \left(\frac{GM_c}{c^3} \right)^{-5/3} \left(f_0^{-5/3} - f(t)^{-5/3} \right) \quad (7.24)$$

In practice one does not know the phase at the beginning of the observation, so we cannot assume $\phi(t_0) = 0$. Rather one can define a random misalignment $\delta \in [0, 2\pi)$ and write

$$\Psi(t) = \sin(2\pi \cdot \phi(t) + \delta) \quad (7.25)$$

where for the goal of using several cavities, it will be the task of the analysis to align the phases of all separate setups to constructively combine the acquired data of several experimental setups.

It is to be mentioned here, that eq. (7.26), and all that follows from it, are expressions in the *proper detector* (PD) frame. Recently, the *transverse-traceless* (TT) gauge has been argued to be better suited for the experimental case of cavity-based HFGW searches with the main argument “that if the GW wavelength and the size of the experiment are comparable, the detector appears to be moving in the proper detector frame while being at rest in transverse-traceless (TT) gauge”.⁷⁹ This discussion, however, is beyond the scope of this thesis and can be found in [79] and [33]. While assuming TT gauge, rather than PD, changes eq. (7.26) slightly, the derived findings of this section still hold qualitatively.

The power generated by a GW in a cavity of volume V_{cav} and quality factor Q inside an external magnetic field B_0 has been derived in [33] as

$$P_{\text{sig}} = \frac{1}{2} Q \omega_{gw}^3 V_{\text{cav}}^{5/3} (\eta_n h_0 B_0)^2 \quad (7.26)$$

in which Q is the loaded quality factor, the antenna coupling is assumed as $\beta = 1$, and the GW frequency is exactly aligned with the resonant frequency of the cavity $\omega_{gw} = (2\pi)f_0$. Here, f_0 corresponds to the cavity’s resonant frequency of mode n , and η_n is the overlap of this mode and the electric field induced by the interaction of the GW with the static B -field. More generally one can write

$$P_{\text{sig}} = \frac{\beta}{2(\beta + 1)} f_{gw}^3 (2\pi)^3 V_{\text{cav}}^{5/3} (\eta_n h_0 B_0)^2 L(f_{gw}, f_0, Q_{\text{eff}}) \quad (7.27)$$

where the Lorentzian

$$L(f_{gw}, f_0, Q_{\text{eff}}) = \left| \frac{-i\omega_{gw}^2}{\omega_0^2 - \omega_{gw}^2 - i\frac{\omega_{gw}\omega_0}{Q_{\text{eff}}}} \right| \quad (7.28)$$

$$= \left| \frac{-if_{gw}^2}{f_0^2 - f_{gw}^2 - i\frac{f_{gw}f_0}{Q_{\text{eff}}}} \right| \quad (7.29)$$

accounts for the partial enhancement by the cavity for frequencies away from the resonance $f_{gw} \neq f_0$ while for $f_{gw} = f_0$ it simply gives $L = Q_{\text{eff}}$ as in eq. (7.26).^{*} Here we can insert the dependency of the strain h_0 on the mass, distance and GW frequency, explicitly mentioned in [33].

$$h_0 \sim 10^{-29} \times \left(\frac{1\text{pc}}{D} \right) \left(\frac{M_b}{10^{-11}M_\odot} \right)^{5/3} \left(\frac{2\pi \cdot f_{gw}}{1\text{GHz}} \right)^{2/3} \quad (7.30)$$

We isolate the signal power's dependency on the cavity (f_0 & Q_{eff}) and f_{gw}

$$P_{\text{sig}} = \mathcal{A} \times f_{gw}^3 \cdot f_{gw}^{4/3} \cdot L(f_{gw}, f_0) \quad (7.31)$$

$$= \mathcal{A} \times f_{gw}^{13/3} \cdot \left| \frac{-if_{gw}^2}{f_0^2 - f_{gw}^2 - i\frac{f_{gw}f_0}{Q_{\text{eff}}}} \right| \quad (7.32)$$

where

$$\mathcal{A} = \underbrace{\left(10^{-29} \times \left(\frac{1\text{pc}}{D} \right) \left(\frac{M_b}{10^{-11}M_\odot} \right)^{5/3} \left(\frac{2\pi}{1\text{GHz}} \right)^{2/3} \right)^2}_{h_0/f_{gw}^{2/3}} \frac{\beta}{\beta + 1} (2\pi)^3 V_{\text{cav}}^{5/3} (\eta_n B_0)^2 \quad (7.33)$$

is a composite of all parameters that are not dependent on either of the three. The signal's high order of the f_{gw} -dependence suggests that it might be sensible to measure at frequencies that are not inside the cavities resonant enhancement but high enough for the $\propto f_{gw}^{13/3}$ dependency to overpower the cavity's absent enhancement. Fig. 7.3 shows the signal power of eq. (7.32) against f_{gw} for several cavity parameters. It shows that indeed the signal power does get larger than at the resonance peak for large f_{gw} , however we can already see that this is not the case until f_{gw} is approximately one order of magnitude larger[†] than f_0 and thereby hardly feasible. We can find this characteristic frequency f_c for which $P_{\text{sig}}(f_{gw} > f_c) > P_{\text{sig}}(f_{gw} = f_0)$ analytically. For this we need to find the local minimum to the right of the peak by solving

$$\frac{dP_{\text{sig}}}{df_{gw}} \stackrel{!}{=} 0 \quad (7.34)$$

resulting in a value of

$$f_{\text{min}} = \frac{-8if_0 + \sqrt{247f_0^2Q_{\text{eff}}^2 - 64f_0^2}}{13Q_{\text{eff}}}. \quad (7.35)$$

For $Q_{\text{eff}} \gg 1$ this can be approximated as

$$f_{\text{min}} \approx \frac{\sqrt{247}}{13} f_0 \approx 1.21 f_0. \quad (7.36)$$

^{*}In light of the discussions on broadband readout one thing to remember here is that this describes the power in the single mode n , hence the factor η_n describing the overlap the GW-induced E -field and the E -field of mode n . If the observed frequency range of f_{gw} is near the resonance f_0 of that mode this is the dominating contribution to the total energy in the cavity. However, if the observed frequency strays further from the resonance frequency of the observed mode one has to consider the influence of other modes.

[†]This result is even more extreme, if the calculation is done in the TT rather than PD.

7. High-Frequency Gravitational Waves




Figure 7.3.: Signal power in units of $[A \cdot \text{Hz}^{13/3}]$ against signal frequency for cavities with different resonant frequencies and quality factors.

For $f_{gw} > f_{\min}$ the signal power grows monotonically, hence there exists an inverse function to eq. (7.32), i.e. $f_{gw}(P_{\text{sig}})$, for which we can compute

$$f_c = f_{gw}(P_{\text{sig}}(f_{gw} = f_0)) = f_{gw}(A f_0^{13/3} Q_{\text{eff}}). \quad (7.37)$$

While it might be possible to find $f_{gw}(P_{\text{sig}})$ analytically, in order to have a simple rule of thumb relation of f_c and the cavity parameters one can simply calculate eq. (7.32) and find a fit in the relevant ranges of f_0 and Q to find f_c numerically. For frequencies in the range of $f_0 \sim 1 \text{ MHz} - 100 \text{ GHz}$ and $Q \sim 10^2 - 10^7$ a decent fit is

$$f_c \approx f_0 \cdot Q^{0.23} \quad (7.38)$$

which, for the cases illustrated in Fig. 7.3, deviates less than 1% from the exact values highlighted in the dashed lines (see Table 7.2).

Q_{eff}	f_0	f_c	f_c^{fit}	$f_c/f_c^{\text{fit}} - 1$
10k	1 GHz	8.35 GHz	8.31 GHz	0.5%
10k	10 GHz	83.49 GHz	83.18 GHz	0.4%
100k	1 GHz	14.23 GHz	14.13 GHz	0.7%
100k	10 GHz	142.21 GHz	141.25 GHz	0.7%

Table 7.2.: Approximation of the critical frequency f_c for which the signal power P_{sig} at frequencies $f > f_c$ is larger than P_{sig} in the resonance frequency f_0 .

7.3.1. Simulating Time-Series Data of a PBH Merger

For given parameters such as ϕ_0 , the merger mass and distance, etc. we can write the signal power of the GW, as shown in eq. (7.26). This signal power is the result of a sinusoidal voltage $U(t)$ with $P_{\text{sig}} = U_{\text{RMS}}^2/Z$ for a given impedance Z of our read out device, usually $50\ \Omega$. As $U_{\text{RMS}} = U_0/\sqrt{2}$ for a sinusoidal voltage with amplitude U_0 , we can write the signal voltage

$$U_{\text{gw}}(t) = \sqrt{2P_{\text{sig}}Z} \cdot \Psi(t) \quad (7.39)$$

with the signal shape $\Psi(t)$ from eq. (7.25).

In the hardware, e.g. the Tektronix RSA518 used in the axion and dark photon measurements, the input voltage $U_{\text{gw}}(t)$ passes through the $50\ \Omega$ input, attenuators and pre-selection circuitry. It is then mixed with an internal local oscillator (LO) at the selected centre frequency f_c to translate the signal to an intermediate frequency (IF). Ideally, this mixing operation can be written as

$$U_{\text{IF}}(t) = U_{\text{gw}}(t) \cdot \cos(2\pi f_c \cdot t), \quad (7.40)$$

followed by a band-pass filter of the defined bandwidth around the IF. The IF signal is subsequently digitised by a high-speed ADC. In the simulation, we do not model the detailed RF path and IF stages. Instead, we work directly with a discretised version of the RF signal

$$v_{\text{rf}} = U_{\text{gw}} \left(\frac{n}{f_{\text{cx}}} \right) \quad n = 0, 1, \dots, N-1 \quad (7.41)$$

at the analyser input in $N = f_{\text{cx}} \cdot t_{\text{sim}}$ discrete time steps, where f_{cx} is the complex sample rate and t_{sim} the total time simulated. For example, the RSA518 has a complex sample rate of 56 MS/s at full bandwidth of 40 MHz.

Modern RSAs internally generate a digital LO and perform digital down conversion to obtain a complex baseband representation of the signal. Conceptually, after digitisation an in-phase and quadrature LO, $\cos(2\pi f_c \cdot t)$ and $-\sin(2\pi f_c \cdot t)$, are applied and low-pass filtered to yield an analytic (complex) signal.

In the simulation, we model this step by directly multiplying the discrete RF signal with a complex exponential at the centre frequency f_c evaluated at the sample times $t_n = n/f_{\text{cx}}$:

$$\ell_I[n] = \cos(2\pi f_c \cdot t_n) \quad (7.42)$$

$$\ell_Q[n] = -\sin(2\pi f_c \cdot t_n) \quad (7.43)$$

$$s_{\text{bb}}[n] = v_{\text{rf}}[n] (\ell_I[n] + i\ell_Q[n]) = v_{\text{rf}}[n] \cdot e^{-i2\pi f_c t_n}. \quad (7.44)$$

This corresponds to an ideal complex digital down-conversion to baseband and is the discrete-time counterpart of the IQ demodulation performed in the RSA.

In a physical RSA, the user selects an analysis bandwidth (BW). This is implemented as a digital channel filter acting on the complex baseband data. The filter specifies which frequency components around DC (corresponding to frequencies around f_c at RF) are passed to the output I/Q stream and with what amplitude.

7. High-Frequency Gravitational Waves

In the simulation, this behaviour is modelled in the frequency domain. First a discrete Fourier Transform (or Fast Fourier Transform, FFT) of the complex baseband signal

$$S[k] = \sum_{n=0}^{N-1} s_{\text{bb}}[n] \cdot e^{-i2\pi nk/N}, \quad k = 0, 1, \dots, N-1 \quad (7.45)$$

$$f_k = \frac{k}{N} f_{\text{cx}} \quad (7.46)$$

is computed, where f_k represent the discrete frequency bins (with the usual interpretation of positive and negative frequencies after reordering).

The RSA analysis bandwidth is represented by a smooth, cosine-tapered low-pass mask $H(f)$ centred at DC, chosen to approximate the effective channel filter of a Tektronix RSA518. For a user-selected two-sided bandwidth B , we define

$$B_{\text{pass}} = \frac{B}{2} \quad (7.47)$$

$$B_{\text{roll}} = 0.2 B \quad (7.48)$$

$$B_{\text{stop}} = B_{\text{pass}} + B_{\text{roll}} \quad (7.49)$$

$$(7.50)$$

and set the discrete filter values at the frequencies f_k as

$$H(f_k) = \begin{cases} 1 & |f_k| \leq B_{\text{pass}} \\ \frac{1}{2} \left[1 + \cos \left(\pi \frac{|f_k| - B_{\text{pass}}}{B_{\text{roll}}} \right) \right] & B_{\text{pass}} < |f_k| < B_{\text{stop}} \\ 0 & |f_k| \geq B_{\text{stop}} \end{cases} \quad (7.51)$$

This idealised mask captures the essential passband and roll-off behaviour of the instrument's digital channel filter without modelling its detailed implementation.

Applying this filter in the frequency domain gives

$$\tilde{S}[k] = S[k] \cdot H(f_k) \quad (7.52)$$

and the filtered complex baseband signal is obtained by the inverse FFT.

$$\tilde{s}_{\text{bb}}[n] = \frac{1}{N} \sum_{k=0}^{N-1} \tilde{S}[k] e^{+i2\pi nk/N}. \quad (7.53)$$

In the real instrument, the corresponding operation is carried out by the internal digital IF and decimation filters.

The RSA finally outputs a complex time record at an effective sample rate f_{cx} , which is interpreted as in-phase and quadrature components of the filtered baseband signal. In the simulation, we identify

$$I[n] = \text{Re}\{\tilde{s}_{\text{bb}}[n]\} \quad (7.54)$$

$$Q[n] = \text{Im}\{\tilde{s}_{\text{bb}}[n]\} \quad (7.55)$$

together with the time stamps t_n . The triplet $\{t_n, I[n], Q[n]\}$ thus represents the idealised I/Q data that a Tektronix RSA-type analyser would deliver for the gravitational-wave-induced voltage $U_{\text{gw}}(t)$, under the approximations of perfect mixing, ideal sampling and a channel filter described by eq. (7.51).

7.3.2. Energy Deposition

If the readout is done with non-quantised methods, such as these discussed so far, the current consensus* is, that the total amount of energy deposited in the cavity can be arbitrarily small. However, as GRAVNET also envisions quantised methods such as single-photon readout, in this case the deposited energy needs to be large enough to convert into (at least) one photon of energy $E_\gamma = hf_0$ for such readout techniques to work.

To calculate the total number of photons in the cavity, if any, we calculate the total energy deposited in the cavity, approximated by the integral of the signal power over the time the signal remains in the bandwidth, and divide by the energy of a single photon at the resonance frequency. With the signal power from eq. (7.32), the deposited energy is

$$E = \int_{\Delta t} P_{\text{sig}} dt = \int_{f_0 - \frac{\Delta f}{2}}^{f_0 + \frac{\Delta f}{2}} \frac{P_{\text{sig}}}{f_{gw}} df_{gw}. \quad (7.56)$$

Introducing the constant

$$\alpha \equiv \frac{96}{5} \pi^{8/3} \left(\frac{GM_c}{c^3} \right)^{5/3} \quad (7.57)$$

and inserting eq. (7.10) into (7.56) one can write

$$E = \int_{f_0 - \frac{\Delta f}{2}}^{f_0 + \frac{\Delta f}{2}} \frac{\mathcal{A} \cdot f_{gw}^{13/3}}{\alpha \cdot f^{11/3}} \cdot \left| \frac{-if_{gw}^2}{f_0^2 - f_{gw}^2 - i \frac{f_{gw}f_0}{Q_{\text{eff}}}} \right| df \quad (7.58)$$

$$= \frac{\mathcal{A}}{\alpha} \int_{f_0 - \frac{\Delta f}{2}}^{f_0 + \frac{\Delta f}{2}} f^{2/3} \cdot \left| \frac{-if_{gw}^2}{f_0^2 - f_{gw}^2 - i \frac{f_{gw}f_0}{Q_{\text{eff}}}} \right| df \quad (7.59)$$

where the integral is calculated numerically. This deposited energy can then simply be divided by the energy of one photon at the cavity's resonance frequency $E_\gamma = hf_0$ and by rounding to the next lowest integer one obtains the maximum number of converted photons. Especially, if $E < E_\gamma$ there is not enough energy to convert the GW signal into a photon, resulting in a clearly defined sensitivity barrier for single-photon readout.

*to whatever degree something as a consensus exists within the theoretical physics community

7. High-Frequency Gravitational Waves

7.3.3. Triangulating the Source Origin

If 3 geographically distant setups record a correlated signal, the GW propagation direction, and hence its source, can be reconstructed up to a two-fold mirror degeneracy from the relative arrival times at the different setups.

To demonstrate this concept, let's choose three axion haloscope experiments which are very practical regarding their positioning on the globe: ORGAN in Perth (Australia), ADMX in Livermore (California) and, of course, our very own SUPAX in Mainz.

Experiment	Location	Latitude	Longitude
ORGAN	Perth, Australia	31°57' S	115°51' E
ADMX	Livermore (CA), USA	37°41' N	121°41' W
SUPAX	Mainz, Germany	50°0' N	8°16' E

Table 7.3.: The three chosen setups to demonstrate the procedure of finding the direction of a GW source, should multiple setups record a correlated event and their location on the globe.

What qualifies this combination is that they all have the hardware suitable for this kind of search and they are spread very far across the globe. The distance between the setups is calculated by transforming the geographical coordinates into common spherical coordinates (r, ϕ, θ) . We can write

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} r \cos \theta \cos \phi \\ r \cos \theta \sin \phi \\ r \sin \theta \end{pmatrix} \quad (7.60)$$

with the inverse transformation

$$r = \sqrt{x^2 + y^2 + z^2} \quad (7.61)$$

$$\phi = \text{atan2}(y, x) \quad \phi \in (-\pi, \pi] \quad (7.62)$$

$$\theta = \arcsin\left(\frac{z}{\sqrt{x^2 + y^2 + z^2}}\right) \quad \theta \in [-\pi/2, \pi/2] \quad (7.63)$$

where ϕ corresponds to the longitude and θ to the latitude. Positive signs correspond to East and North while negative signs correspond to West and South respectively. The baseline between two points j and i is then

$$\vec{d}_{ij} = \vec{x}_j - \vec{x}_i \quad (7.64)$$

$$|\vec{d}_{ij}| = \sqrt{(x_j - x_i)^2 + (y_j - y_i)^2 + (z_j - z_i)^2} \quad (7.65)$$

for all three of them. The approximate distances of the three setups are shown in Fig. 7.4, where we have assumed a constant earth radius, i.e. no elevation differences. The time that passes between measuring the signal in two setups i and j is defined as

$$\Delta t_{ij} = t_j - t_i \quad \text{and} \quad \Delta t_{ij} = \frac{\vec{d}_{ij} \cdot \hat{k}}{c} \quad (7.66)$$

where $\vec{d}_{ij} = \vec{x}_j - \vec{x}_i$ and $\hat{k}$ is the propagation direction of the GW with $|\hat{k}| \equiv 1$.

A single time-delay measurement between two setups constrains the GW propagation direction $\hat{k}$ to a cone of possible solutions around the baseline vector $\vec{d}_{ij}$. Only in the edge case

$$\Delta t = \frac{|\vec{d}_{ij}|}{c} \quad (7.67)$$

does the cone collapse to a single direction, corresponding to the GW propagating parallel to the baseline. Since $\vec{d}_{ij}$ is the maximum projection of the baseline onto the propagation direction, this relation also implies a maximum allowed time delay. For the longest baseline of our system, between ORGAN and ADMX, eq. (7.67) yields

$$\Delta t_{\max} \approx 39 \text{ ms}, \quad (7.68)$$

indicating that sub-millisecond precision is required to meaningfully resolve the source direction of the GW.*

To advance from the cone of possible wave-front planes we now include the third detector. This reduces the continuum of solutions to two discrete candidates. These two solutions are related by a mirror reflection with respect to the plane spanned by the detector baselines and therefore cannot be distinguished using timing information alone in a three-detector network. The propagation direction $\hat{k}$ is obtained by solving

$$\hat{k} \cdot \vec{d}_{ij} = c \cdot \Delta t_{ij} \quad (7.69)$$

$$\hat{k} \cdot \vec{d}_{il} = c \cdot \Delta t_{il} \quad (7.70)$$

$$\|\hat{k}\| = 1 \quad (7.71)$$

where one might want to remind themselves that

$$\hat{k} \cdot \vec{d}_{lj} = c \cdot \Delta t_{lj} \quad (7.72)$$

is not a third independent equation but a linear combination of eq. (7.69) and (7.70).

$$\vec{d}_{lj} = \vec{x}_j - \vec{x}_l = (\vec{x}_j - \vec{x}_i) - (\vec{x}_l - \vec{x}_i) = \vec{d}_{ij} - \vec{d}_{il} \quad (7.73)$$

$$\Delta t_{lj} = t_j - t_l = (t_j - t_i) - (t_l - t_i) = \Delta t_{ij} - \Delta t_{il} \quad (7.74)$$

Equations (7.69) and (7.70) describe two planes in $\mathbb{R}^3$ with normal vectors $\vec{d}_{ij}$ and $\vec{d}_{il}$. The intersection of these planes is a line that is orthogonal to both normals and therefore given by

$$\vec{n} = \vec{d}_{ij} \times \vec{d}_{il}. \quad (7.75)$$

*We neglect Earth's rotation of $\lesssim 10^{-4}$ degrees during Δt in the following calculation.




Figure 7.4.: Approximate distances of three possible experimental setups participating in GRAVNET.

7. High-Frequency Gravitational Waves

To obtain a solution of the two plane equations $\hat{k}_0$ we impose the condition that this solution has no component in this direction.

$$\hat{k}_0 \cdot \vec{n} = 0 \quad (7.76)$$

This gives a 3×3 linear system

$$\begin{pmatrix} \vec{d}_{ij}^\top \\ \vec{d}_{il}^\top \\ \vec{n}^\top \end{pmatrix} \hat{k}_0 = \begin{pmatrix} c \cdot \Delta t_{ij} \\ c \cdot \Delta t_{il} \\ 0 \end{pmatrix} \quad (7.77)$$

with a unique solution for $\hat{k}_0$ as $\vec{d}_{ij}$, $\vec{d}_{il}$ and $\vec{n}$ are linearly independent. The general solution of the two plane equations can be written as

$$\hat{k}(\alpha) = \hat{k}_0 + \alpha \vec{n}, \quad \alpha \in \mathbb{R}. \quad (7.78)$$

Applying eq. (7.71) to this we get

$$\|\hat{k}_0 + \alpha \vec{n}\|^2 = \|\hat{k}_0\|^2 + 2\alpha \underbrace{(\hat{k}_0 \cdot \vec{n})}_{=0} + \alpha^2 \|\vec{n}\|^2 = 1 \quad (7.79)$$

where the mixing term is eliminated by eq. (7.76) and therefore

$$\alpha^2 = \frac{1 - \|\hat{k}_0\|^2}{\|\vec{n}\|^2}. \quad (7.80)$$

The two candidate solutions of eq. (7.78), shown in Fig. 7.5, are hence

$$\hat{k}_\pm = \hat{k}_0 \pm \sqrt{\frac{1 - \|\hat{k}_0\|^2}{\|\vec{n}\|^2}} \vec{n}. \quad (7.81)$$

The normalisation condition yields two mirror solutions, $\hat{k}_\pm$, which correspond to two possible propagation directions located on opposite sides of the plane spanned by $\vec{d}_{ij}$ and $\vec{d}_{il}$. Both candidates satisfy the measured time delays by construction. Identifying a unique propagation direction requires additional information beyond the relative arrival times, e.g. a fourth detector site. In this case we get a set of three linear equations

$$\hat{k} \cdot \vec{d}_{ij} = c \cdot \Delta t_{ij} \quad (7.82)$$

$$\hat{k} \cdot \vec{d}_{il} = c \cdot \Delta t_{il} \quad (7.83)$$

$$\hat{k} \cdot \vec{d}_{im} = c \cdot \Delta t_{im} \quad (7.84)$$

similar to earlier, with the difference that the third non-trivial timing differential Δt_{im} results in a unique solution $\hat{k}$ straight away. As noise and limited timing resolution in the real-world experiment will likely yield a solution $\|\hat{k}\| \neq 1$ the practical implementation will likely be solving for least squares

$$k_0 = \arg \min \|Dk - b\|^2 \quad (7.85)$$

$$D = \begin{pmatrix} \vec{d}_{ij}^\top \\ \vec{d}_{il}^\top \\ \vec{d}_{im}^\top \end{pmatrix}, \quad b = \begin{pmatrix} c \cdot \Delta t_{ij} \\ c \cdot \Delta t_{il} \\ c \cdot \Delta t_{im} \end{pmatrix} \quad (7.86)$$

and then setting

$$\hat{k} = \frac{k_0}{\|k_0\|} \quad (7.87)$$

where one could report $\|D\hat{k} - b\|$ as fit quality.




Figure 7.5.: Left: Three detectors define two cones, which overlap in two directions. Right: To determine which one of $\hat{k}_\pm$ is the true source, a fourth setup is needed. With the resulting three cones, $\hat{k}_+$ is identified as the true source in this example.

For visualisation purposes, $\hat{k}$ in spherical coordinates can then be converted back into geographical coordinates. The resulting geographical coordinates can then be interpreted in the way that a vector from the Earth's core to this position is the propagation direction of the wave. To find the source one would need to look at the antipodal coordinates

$$\theta_{\text{antipodal}} = -\theta \quad (7.88)$$

$$\phi_{\text{antipodal}} = 180^\circ - \phi \quad (7.89)$$

as this is where the GW is originating from.

8. Future Prospects

8.1. Further SRF Cavity Development

Beyond our initial NbN approach, which has proven to be unsuitable for use within the magnetic field, and thereby for axion searches, two possible development directions appear promising in the quest to become less dependent on the current dependence on RE-BCO/YBCO tapes for superconducting cavities used in haloscopes. The first is defect engineering. Since vortex losses arise from vortex motion, deliberately introduced defects can be beneficial when they act as pinning centres and restrict that motion. This idea is well established and in YBCO-based materials artificial pinning centres⁹⁹ and irradiation induced defects¹⁰⁰ have been shown to enhance pinning and improve performance. However, whether a comparable strategy is equally effective in other candidate materials is unclear. The second direction is to investigate other superconductors whose vortex dynamics may be more favourable in the high-frequency, high-field range relevant to haloscopes. Iron-based superconductors appear to be promising in this context.¹⁰¹

8.2. That Single-Bin Power Excess in the Dedicated Dark Photon Measurement

The abnormal bin in the dedicated DP measurement should be investigated further. The most likely scenario is that it simply is an electronic relic. A series of measurements needs to be conducted on this. Firstly, without any cavity, a considerable amount of noise should be recorded with the RSA at this specific frequency, varying the centre frequency of the RSA and checking if a) the excess is still there and, if so, b) it varies with the centre frequency of the RSA. If that is the case one can conclude that it must be a relic of the RSA's conversion from time domain to frequency domain. If one observes a power excess at the same frequency, independent of the RSA's centre frequency it should be determined if this is an actual RF signal. For this a cavity, with resonance frequency near this suspicious frequency, can be utilised. Data is collected for two cases: cavity resonance tuned to be aligned with the suspicious frequency and tuned to be deliberately not aligned with it. If the signal is of RF nature, it should be enhanced with the cavity's quality factor when the frequencies align vs when they do not align. Should this yield the result that it really is an RF signal and not a relic of the readout electronics we must consider an outside signal. Some sources known to operate in this frequency region of 8 – 10 GHz (commonly called X-band in IEEE or I-band in NATO jargon) are ATC and radar systems. The nearby Frankfurt Airport installed a new GBAS* in 2022 as addition to the widely used ILS* to help landings in bad weather conditions. Furthermore, with Ramstein Air Base nearby there is certainly no shortage of

*GBAS = Ground Based Augmentation System (a form of Precision Approach Radar)
ILS = instrument landing system

8.3. Achievable Axion Limits with One Year of Data Acquisition

radar systems in the area. However, all of these civil and military radar systems typically operate relatively broadband with signals over several MHz and are therefore very unlikely to be the source of this one-bin (< 1 kHz) power excess. Narrowband emitters in this frequency range include satellites and protected radio astronomy observations around 8.4–8.5 GHz. Satellite carriers can indeed be only a few hundred hertz wide, but would generally not be expected to remain within this narrow frequency bin over several hours of data acquisition. With these possibilities excluded, the most plausible explanation is that the feature arises from within the measurement chain itself, for example as a local oscillator leakage which naturally produces very narrow and highly stable signals in the gigahertz range. This can only be avoided by using a different RSA or shielding the experiment from the RSA.

If this excess had been measured in an experimental run with an external magnetic field the signal might have actually been a possible candidate. If one assumes that dark matter is not distributed entirely homogeneously and one considers the possibility of axion tidal streams, there might be signal signatures as narrow as $10^{-6} - 10^{-2}$ Hz in the probed mass range, falling well within one single bin.¹⁰² However, for dark photons, no such mechanisms are known yet.

8.3. Achievable Axion Limits with One Year of Data Acquisition

The frequency scanning speed $\frac{d\nu}{dt}$ at which a certain axion-photon limit can be reached can be derived directly from the signal and system noise, as shown in [103]. The signal power is known from eq. (3.66) and the system noise

$$P_{\text{noise}} = k_B T_{\text{sys}} \sqrt{\frac{\Delta\nu}{\Delta t}} \quad (8.1)$$

is described as fluctuations in the noise power P_N within an integration time Δt over the bandwidth $\Delta\nu$. If we aim to reach the QCD-axion limits, we must restate the coupling factor $g_{a\gamma\gamma}$ into the model-dependent, dimensionless factor

$$g_\gamma = \frac{\pi\sqrt{\chi}}{m_a\alpha} g_{a\gamma\gamma} \quad (8.2)$$

as seen earlier in eq. (3.22). This factor is $g_\gamma^{\text{KSVZ}} = -0.97$ for the KSVZ and $g_\gamma^{\text{DFSZ}} = 0.36$ for the DFSZ model. The resulting scanning rate has been derived in [103] and using eq. (3.22) becomes

$$\frac{d\nu}{dt} = \left[g_\gamma^2 \frac{\alpha^2}{\pi^2 \chi} \frac{m_a \hbar^2 c^5 \varepsilon_0}{k_B T_{\text{sys}}} \frac{V_{\text{eff}} \rho_{\text{DM}} B_0^2}{\text{SNR}} \frac{\beta}{\beta + 1} \right]^2 Q_a \min(Q_L, Q_a) \quad (8.3)$$

where we will assume the usual $\text{SNR} = 5$ for a 95% confidence level, as we did in the analysis (sec. 6.7.2). Table 8.1 shows the scanning rate and possible scanned frequency range when aiming for exclusion of both main theoretical models. The existing setup assumes the experimental characteristics of the analysis presented earlier (right side of Tab. 6.4). The future setup assumes a cylindrical cavity with round cross section where the effective volume is given by

$$V_{\text{eff}} = \pi r^2 L \cdot C_{010} = \pi (4.5 \text{ cm})^2 14 \text{ cm} \cdot 0.69 = 610 \text{ mL} \quad (8.4)$$

8. Future Prospects

and the external magnetic field is $B_0 = 12$ T. The form factor C_{010} here has a different convention as in sec. 3.4.2 where the three indices stood for x, y, z . Here, as we have a cavity with a circular cross section, calling for indices θ, r, z corresponding to cylindrical coordinates. It is the equivalent of the TM_{110} mode used in the rectangular cavity. The form factor and resonant frequency have been derived in [59] as

$$\nu_{0n0} = \frac{x_n}{2\pi} \frac{c}{r} \stackrel{n=1}{=} 2.55 \text{ GHz} \quad (8.5)$$

$$C_{0n0} = \frac{4}{x_n^2} \stackrel{n=1}{=} 0.69 \quad (8.6)$$

where $x_n \stackrel{n=1}{=} 2.405$ is the n -th (first) root of the zeroth-order Bessel function $J_0(x)$. The readout antenna is coupled to $\beta = 2$ as this is optimal for maximizing the scanning rate. This scan rate optimised coupling assumes $Q_L < Q_a$ and is obtained as a simple optimizing problem of eq. (8.3) where one has to recall $Q_L = Q_0/(\beta + 1)$.

$$0 \stackrel{!}{=} \frac{d}{d\beta} \left(\frac{d\nu}{dt} \right) \quad (8.7)$$

$$= \frac{d}{d\beta} \left(\frac{\beta^2}{(\beta + 1)^3} \right) \quad (8.8)$$

$$= -\frac{(x - 2)x}{(x + 1)^4} \quad (8.9)$$

$$\Rightarrow \beta_{\max} = 2 \quad (8.10)$$

Furthermore, the planned setup aims at an ambient temperature of 20 mK and a readout chain at the Standard Quantum Limit (SQL). The details of this are discussed extensively elsewhere, as it is the gold standard of cavity haloscopes. Here it shall suffice to say that with readout at the SQL the system noise temperature is not simply the sum of the ambient temperature and added noise of the read out but rather

$$T_{\text{sys}} = \frac{h\nu}{k_B} \left(\frac{1}{e^{h\nu/k_B T_{\text{amb}}} - 1} + 1 \right) \quad (8.11)$$

which, for $T_{\text{amb}} = 20$ mK and $\nu = 5.1$ GHz, gives $T_{\text{sys}} \approx 245$ mK.

	KSVZ		DFSZ	
	scanning rate	scanned freq. range (1 year)	scanning rate	scanned freq. range (1 year)
Existing setup	31.6 Hz/day	8.8 kHz	0.6 Hz/day	168 Hz
Future setup ($Q_L = 40k$)	6.4 MHz/day	(1.8 GHz)	0.12 MHz/day	(34 MHz)
Future setup ($Q_L = 400k$)	64 MHz/day	(18 GHz)	1.2 MHz/day	(341 MHz)

Table 8.1.: Expected scanning rate, and expected scanned frequency range if the experiment takes data for 40 weeks (1 experimental year) in total at the respective rate, assuming KSVZ and DFSZ limit sensitivity for the existing setup and the planned setup at the University of Bonn.

8.3. Achievable Axion Limits with One Year of Data Acquisition

Looking at Table 8.1, there is no point in setting KSVZ or DFSZ-level limits at a reasonable frequency range with the current setup. The planned setup holds a significant improvement in the system noise, due to the implementation of a dilution fridge rather than the LHe cryostat, and possibly the unloaded quality factor, given that an applicable superconductor is found that can bring Q_0 into the range of $10^5 - 10^6$ in the magnetic field. It also holds the advantage of being at a much lower frequency, the effects of which are discussed later in this section. The 1-year scanned frequency range shown in the table however, is also not a very sensible estimate here as eq. (8.3) assumes frequencies near the resonance frequency. While “near” is a fuzzy term, it is clear that, for $\nu_0 = 2.55$ GHz, a range of 1.8 GHz (let alone 18 GHz) cannot be considered “near”. Additionally, scanning at such speeds relies on a very precise tuning mechanism. The currently employed pressure tuning is ruled out for several reasons. Firstly, the range of tuning by pressure is limited to a few MHz at best and secondly, the new setup is placed in a dilution fridge rather than a LHe cryostat. Therefore, a classic tuning rod will be installed, which modifies the form factor (typical values for cylindrical cavities with a tuning rod are around $C_{010} \sim 0.5$), the quality factor and the total volume, as the volume taken up by the tuning rod is not accessible as detector volume. Additionally, it is planned that besides the $r = 4.5$ cm cavity several more will be operated to access higher frequencies. In our Conceptual Design Report we estimate to cover 1.64 GHz at $0.9 \times \text{KSVZ}$ ($Q_0 = 400k$) or at $1.45 \times \text{KSVZ}$ ($Q_0 = 40k$) in one experimental year using these strategies and more realistic assumptions.¹⁰⁴

Coming back to the remark about the frequency dependency of the scanning speed, it is questionable if experiments in the 8 – 10 GHz regime can ever get close to the scanning rates of setups in the ~ 1 GHz regime. That is because the low frequency cavities can be much larger, often with an effective volume in the order of 20 – 40 litres. Isolating just the dependence on the searched frequency range and effective volume we can define a figure of merit $\mathcal{F}$ as follows

$$\frac{d\nu}{dt} \propto V_{\text{eff}}^2 m_a^2 \equiv \mathcal{F} \quad (8.12)$$

to compare the size difference due to geometry, assuming that reaching a similar level of noise, magnetic field and quality factor are problems that can be solved with sufficient engineering effort. For the SUPAX cavity we get

$$\mathcal{F}_{\text{SUPAX}} = [56.28 \text{ mL} \times 8.4 \text{ GHz}]^2 \approx 0.22 [\text{L} \cdot \text{GHz}]^2 \quad (8.13)$$

while for recent results of ADMX¹⁰⁵ and CAPP,¹⁰⁶ both in the range of 1.0 – 1.3 GHz, we get

$$\mathcal{F}_{\text{ADMX}} = [0.4 \cdot 107 \text{ L} \times 1.1 \text{ GHz}]^2 \approx 2200 [\text{L} \cdot \text{GHz}]^2 \quad (8.14)$$

$$\mathcal{F}_{\text{CAPP}} = [0.6 \cdot 37 \text{ L} \times 1.1 \text{ GHz}]^2 \approx 600 [\text{L} \cdot \text{GHz}]^2. \quad (8.15)$$

This gives an advantage of roughly 10,000 (ADMX) or 2,700 (CAPP) over the SUPAX experiment by geometry alone. This way they are able to reach scanning rates of a few MHz per day, which is hardly possible with the geometry required for cavities searching for axions in the 8 – 10 GHz regime.

One final observation is worth noting. The relation $d\nu/dt \propto V_{\text{eff}}^2$ might incite the thought that maximizing volume by using a higher mode of a larger cavity can gain scanning speed.

8. Future Prospects

However, from equations (8.5) and (8.6) we learn that the *effective* volume of a circular cross section cavity

$$V_{\text{eff}} = \pi r^2 L \cdot C_{0n0} = \frac{c^2 L}{\pi \nu_n^2} \quad (8.16)$$

for a given goal frequency ν_n does not depend on whether that frequency is reached via TM_{010} of a smaller or TM_{020} of a larger (radius) cavity. The scanning rate $d\nu/dt$ also depends on Q_0 of the used mode, which is always highest for the base mode. Therefore, there is no point in using any modes other than the base mode.

9. Conclusion

The axion is a highly compelling answer to two of the most interesting open questions in modern physics. Introduced through the Peccei-Quinn mechanism as a solution to the strong CP problem, it also emerges as a viable cold dark matter candidate. A specific experimental technique used in axion searches, the cavity haloscope, is also sensitive to other interesting dark matter candidates. This includes dark photons and, with some modifications, high-frequency gravitational waves which can be emitted by primordial black holes through various mechanisms.

In this thesis, the first physics results of the SUPAX experiment have been presented. SUPAX was designed as a superconducting axion haloscope experiment operating around 8 GHz, corresponding to particle masses in the μeV regime. At these frequencies, cavity haloscopes face the challenge that the cavity volume decreases as the resonant frequency increases, while the expected signal power remains very small. The use of superconducting materials is therefore a necessity for the effort of maintaining scanning rates similar to those in the low or sub-GHz regime. In this context, the NbN-coated cavity studied in SUPAX provides a result of another material in order to answer the question whether superconducting coatings, besides the commercially available and often troublesome REBCO or YBCO tapes, can be integrated into axion haloscope operation at high frequencies and, crucially, in the presence of an external magnetic field.

The experimental work presented here includes the characterisation of the relevant cavities and the data acquisition chain. Besides the cavity, which was studied with a vector network analyser and compared to the expected properties obtained by simulation, the pressure and temperature conditions inside the cryostat were monitored and managed during experimental runs. The system noise and attenuation were characterised and efforts were undertaken to ensure efficient cooling of all active parts during cryogenic operation, a challenging task especially inside the magnet.

The first experimental run in March 2024 tested the feasibility of NbN as a superconducting coating both with and without an external magnetic field. The coating improved the unloaded quality factor by a factor of approximately 8 in the absence of a magnetic field, but the superconductivity broke down rapidly, yielding no significant improvement of Q_0 over the uncoated cavity in the presence of the magnetic field. However, as it improved the quality factor significantly without the external magnetic field, it was used in the second experimental run in April 2024, in the first test for operation as a haloscope. The obtained data was used to set an exclusion limit for the kinetic mixing of dark photons down to $\chi > 8.78 \cdot 10^{-15}$ around 8.4712 GHz ($m_{\text{DP}} \approx 35 \mu\text{eV}$), assuming randomly polarised dark photons.

After some engineering efforts in order to effectively cool the setup despite the modifications needed for the operation in a magnetic field, the third experimental run was conducted in October 2024 inside the magnet. This run sets an exclusion limit in a frequency range of 1.44 MHz around 8.265 48 GHz, corresponding to an axion mass of approximately 34.18 μeV ,

9. Conclusion

excluding axion-photon couplings down to $g_{a\gamma\gamma} > 1.67 \cdot 10^{-13} \text{ GeV}^{-1}$. This limit was then recast into an additional dark photon limit, yielding an exclusion down to $\chi > 1.97 \cdot 10^{-14}$. These limits demonstrate that the SUPAX experiment is capable of producing physics results from cryogenic microwave-cavity data at a comparable sensitivity to similar experiments.

The presented experimental setup is to be understood as a base from which research into different directions can be undertaken. It is the first experiment to use helium pressure for tuning, it tested NbN as a possible alternative to tape-based superconductors and intends to test other superconductors in the future, and it also serves as a prototype for the GRAVNET experiment.

GRAVNET is a proposed network of haloscopes, which aims to extend the use of existing axion haloscope infrastructure to more exotic searches such as high-frequency gravitational waves. By adapting the GHz cavities employed in these experiments it targets a frequency range not accessible to conventional GW detectors. Primordial black holes provide one possible source of such signals, and their connection to dark matter gives ample motivation for this search. The simulation package developed in this thesis allows the output of the present GRAVNET, or in fact SUPAX, data acquisition chain to be modelled for HFGW signals from PBH mergers. This provides a way to test analysis strategies before the full experimental setup is completed and characterised.

Summarizing, the work presented in this thesis contributes to the experimental search for dark matter in three ways. It demonstrates SUPAX's ability to set exclusion limits competitive with other axion haloscopes, it probes new materials in order to tackle the dependence of haloscope experiments on niche commercial superconducting tapes and lastly, it sets the stage for a new type of gravitational wave search utilizing haloscope cavities arranged in a global network by providing a first prototype of signal simulation, allowing the early development of analysis techniques, while the experimental setups are still under construction.

Bibliography

- [1] https://en.wikipedia.org/wiki/Standard_Model,
Accessed: 2023-09-11.
- [2] F. Zwicky, “Die rotverschiebung von extragalaktischen nebeln,”
Helvetica Physica Acta, vol. 6, pp. 110–127, Jan. 1933.
- [3] https://www.esa.int/Science_Exploration/Space_Science/Planck/Planck_reveals_an_almost_perfect_Universe,
Accessed: 2023-09-11.
- [4] H. Poincare, “The Milky Way and the Theory of Gases,”
Popular Astronomy, vol. 14, pp. 475–488, Oct. 1906.
- [5] J. H. Oort, “The force exerted by the stellar system in the direction perpendicular to the galactic plane and some related problems,”
Bulletin of the Astronomical Institutes of the Netherlands, vol. 6, p. 249, Aug. 1932.
- [6] F. Zwicky, “On the Masses of Nebulae and of Clusters of Nebulae,”
Astrophys. J., vol. 86, pp. 217–246, 1937.
DOI: 10.1086/143864.
- [7] A. A. Ponomareva, M. A. W. Verheijen, and A. Bosma, “Detailed h i kinematics of tully–fisher calibrator galaxies,”
Monthly Notices of the Royal Astronomical Society, vol. 463, no. 4, Dec. 2016.
DOI: 10.1093/mnras/stw2213.
- [8] E. Corbelli, S. Lorenzoni, R. Walterbos, R. Braun, and D. Thilker, “A wide-field h i mosaic of messier 31: Ii. the disk warp, rotation, and the dark matter halo,”
Astronomy and Astrophysics, vol. 511, A89, Feb. 2010,
ISSN: 1432-0746.
DOI: 10.1051/0004-6361/200913297.
- [9] A. Afruni, G. Pezzulli, and F. Fraternali, “Inflow of low-metallicity cool gas in the halo of the andromeda galaxy,”
Monthly Notices of the Royal Astronomical Society, vol. 509, no. 4, pp. 4849–4864, Nov. 2021,
ISSN: 1365-2966.
DOI: 10.1093/mnras/stab3237.
- [10] V. C. Rubin and W. K. Ford, Jr., “Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions,”
Astrophysical Journal, vol. 159, p. 379, Feb. 1970.
DOI: 10.1086/150317.
- [11] D. Clowe *et al.*, “A direct empirical proof of the existence of dark matter,”
The Astrophysical Journal, vol. 648, no. 2, pp. L109–L113, Aug. 2006,

Bibliography

- ISSN: 1538-4357.
DOI: 10.1086/508162.
- [12] A. A. Penzias and R. W. Wilson, “A Measurement of Excess Antenna Temperature at 4080 Mc/s.,”
The Astrophysical Journal, vol. 142, pp. 419–421, Jul. 1965.
DOI: 10.1086/148307.
- [13] N. Aghanim *et al.*, “Planck2018 results: Vi. cosmological parameters,”
Astronomy & Astrophysics, vol. 641, A6, Sep. 2020,
ISSN: 1432-0746.
DOI: 10.1051/0004-6361/201833910.
- [14] A. G. Riess *et al.*, “Cosmic distances calibrated to 1% precision with gaia edr3 parallaxes and hubble space telescope photometry of 75 milky way cepheids confirm tension with Λ CDM,”
The Astrophysical Journal Letters, vol. 908, no. 1, p. L6, Feb. 2021.
DOI: 10.3847/2041-8213/abdbaf.
- [15] T. Clifton, P. G. Ferreira, A. Padilla, and C. Skordis, “Modified gravity and cosmology,”
Physics Reports, vol. 513, no. 1–3, pp. 1–189, Mar. 2012,
ISSN: 0370-1573.
DOI: 10.1016/j.physrep.2012.01.001.
- [16] R. H. Wechsler and J. L. Tinker, “The connection between galaxies and their dark matter halos,”
Annual Review of Astronomy and Astrophysics, vol. 56, no. 1, pp. 435–487, Sep. 2018,
ISSN: 1545-4282.
DOI: 10.1146/annurev-astro-081817-051756.
- [17] R. S. Somerville and R. Davé, “Physical models of galaxy formation in a cosmological framework,”
Annual Review of Astronomy and Astrophysics, vol. 53, no. 1, pp. 51–113, Aug. 2015,
ISSN: 1545-4282.
DOI: 10.1146/annurev-astro-082812-140951.
- [18] A. S. Chou *et al.*,
Snowmass cosmic frontier report,
2022.
arXiv: 2211.09978 [hep-ex].
- [19] R. Adhikari *et al.*, “A white paper on kev sterile neutrino dark matter,”
Journal of Cosmology and Astroparticle Physics, vol. 2017, no. 01, pp. 025–025, Jan. 2017,
ISSN: 1475-7516.
DOI: 10.1088/1475-7516/2017/01/025.
- [20] T. Asaka and M. Shaposhnikov, “The vmsm, dark matter and baryon asymmetry of the universe,”
Physics Letters B, vol. 620, no. 1–2, pp. 17–26, Jul. 2005,
ISSN: 0370-2693.

- DOI: 10.1016/j.physletb.2005.06.020.
- [21] K. N. Abazajian *et al.*,
Light sterile neutrinos: A white paper,
2012.
arXiv: 1204.5379.
 - [22] C. Alcock *et al.*, “The MACHO project: Microlensing results from 5.7 years of large magellanic cloud observations,”
The Astrophysical Journal, vol. 542, no. 1, pp. 281–307, Oct. 2000.
DOI: 10.1086/309512.
 - [23] P. Tisserand *et al.*, “Limits on the macho content of the galactic halo from the EROS-2 survey of the magellanic clouds,”
Astronomy & Astrophysics, vol. 469, no. 2, pp. 387–404, Apr. 2007.
DOI: 10.1051/0004-6361:20066017.
 - [24] *The xenon experiment*,
<https://xenonexperiment.org>,
Accessed: 2023-09-11.
 - [25] E. Aprile *et al.*, “Observation of two-neutrino double electron capture in ^{124}Xe with XENON1T,”
Nature, vol. 568, no. 7753, pp. 532–535, 2019.
DOI: 10.1038/s41586-019-1124-4. arXiv: 1904.11002 [nucl-ex].
 - [26] M. Schumann, “Direct detection of WIMP dark matter: Concepts and status,”
Journal of Physics G: Nuclear and Particle Physics, vol. 46, no. 10, p. 103 003, Aug. 2019.
DOI: 10.1088/1361-6471/ab2ea5.
 - [27] G. F. Giudice, M. McCullough, and A. Urbano, “Hunting for dark particles with gravitational waves,”
Journal of Cosmology and Astroparticle Physics, vol. 2016, no. 10, pp. 001–001, Oct. 2016.
DOI: 10.1088/1475-7516/2016/10/001.
 - [28] C. Palenzuela, L. Lehner, and S. L. Liebling, “Orbital dynamics of binary boson star systems,”
Physical Review D, vol. 77, no. 4, Feb. 2008.
DOI: 10.1103/physrevd.77.044036.
 - [29] C. Palenzuela, P. Pani, M. Bezares, V. Cardoso, L. Lehner, and S. Liebling, “Gravitational wave signatures of highly compact boson star binaries,”
Physical Review D, vol. 96, no. 10, Nov. 2017.
DOI: 10.1103/physrevd.96.104058.
 - [30] T. Helfer, E. A. Lim, M. A. Garcia, and M. A. Amin, “Gravitational wave emission from collisions of compact scalar solitons,”
Physical Review D, vol. 99, no. 4, Feb. 2019.
DOI: 10.1103/physrevd.99.044046.
 - [31] G. Narain, J. Schaffner-Bielich, and I. N. Mishustin, “Compact stars made of fermionic dark matter,”

Bibliography

- Physical Review D*, vol. 74, no. 6, Sep. 2006.
DOI: 10.1103/physrevd.74.063003.
- [32] M. D. Diamond, D. E. Kaplan, and S. Rajendran, “Binary collisions of dark matter blobs,”
Journal of High Energy Physics, vol. 2023, no. 1, Jan. 2023.
DOI: 10.1007/jhep01(2023)136.
- [33] A. Berlin *et al.*, “Detecting high-frequency gravitational waves with microwave cavities,”
Physical Review D, vol. 105, no. 11, Jun. 2022.
DOI: 10.1103/physrevd.105.116011.
- [34] F. Wilczek, “Problem of strong P and T invariance in the presence of instantons,”
Phys. Rev. Lett., vol. 40, pp. 279–282, 5 Jan. 1978.
DOI: 10.1103/PhysRevLett.40.279.
- [35] F. Wilczek, “Nobel lecture: Asymptotic freedom: From paradox to paradigm,”
Rev. Mod. Phys., vol. 77, pp. 857–870, 3 Sep. 2005.
DOI: 10.1103/RevModPhys.77.857.
- [36] H. Primakoff, “Photo-production of neutral mesons in nuclear electric fields and the mean life of the neutral meson,”
Phys. Rev., vol. 81, pp. 899–899, 5 Mar. 1951.
DOI: 10.1103/PhysRev.81.899.
- [37] *Current axion limit database*,
<https://github.com/cajohare/AxionLimits>,
Accessed: 2023-09-12.
- [38] R. Bähre *et al.*, “Any light particle search II —Technical Design Report,”
JINST, vol. 8, T09001, 2013.
DOI: 10.1088/1748-0221/8/09/T09001. arXiv: 1302.5647 [physics.ins-det].
- [39] J. F. Navarro, C. S. Frenk, and S. D. M. White, “A universal density profile from hierarchical clustering,”
The Astrophysical Journal, vol. 490, no. 2, pp. 493–508, Dec. 1997,
ISSN: 1538-4357.
DOI: 10.1086/304888.
- [40] J. I. Read, “The local dark matter density,”
Journal of Physics G: Nuclear and Particle Physics, vol. 41, no. 6, p. 063101, May 2014,
ISSN: 1361-6471.
DOI: 10.1088/0954-3899/41/6/063101.
- [41] C. A. J. O’Hare, G. Pierobon, and J. Redondo, “Axion minicluster streams in the solar neighborhood,”
Physical Review Letters, vol. 133, no. 8, Aug. 2024,
ISSN: 1079-7114.
DOI: 10.1103/physrevlett.133.081001.
- [42] E. Noether, “Invariant variation problems,”

- Transport Theory and Statistical Physics*, vol. 1, no. 3, pp. 186–207, Jan. 1971,
ISSN: 1532-2424.
DOI: 10.1080/00411457108231446.
- [43] T. D. Lee and C. N. Yang, “Question of parity conservation in weak interactions,”
Phys. Rev., vol. 104, pp. 254–258, 1 Oct. 1956.
DOI: 10.1103/PhysRev.104.254.
 - [44] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson, “Experimental test of parity conservation in beta decay,”
Phys. Rev., vol. 105, pp. 1413–1415, 4 Feb. 1957.
DOI: 10.1103/PhysRev.105.1413.
 - [45] M. Goldhaber, L. Grodzins, and A. W. Sunyar, “Helicity of neutrinos,”
Phys. Rev., vol. 109, pp. 1015–1017, 3 Feb. 1958.
DOI: 10.1103/PhysRev.109.1015.
 - [46] J. H. Christenson, J. W. Cronin, V. L. Fitch, and R. Turlay, “Evidence for the 2π decay of the K_2^0 meson,”
Phys. Rev. Lett., vol. 13, pp. 138–140, 4 Jul. 1964.
DOI: 10.1103/PhysRevLett.13.138.
 - [47] S. Dar,
The neutron edm in the sm : A review,
2000.
arXiv: hep-ph/0008248 [hep-ph].
 - [48] V. Baluni, “CP-nonconserving effects in quantum chromodynamics,”
Phys. Rev. D, vol. 19, pp. 2227–2230, 7 Apr. 1979.
DOI: 10.1103/PhysRevD.19.2227.
 - [49] R. Crewther, P. Di Vecchia, G. Veneziano, and E. Witten, “Chiral estimate of the electric dipole moment of the neutron in quantum chromodynamics,”
Physics Letters B, vol. 88, no. 1, pp. 123–127, 1979,
ISSN: 0370-2693.
DOI: [https://doi.org/10.1016/0370-2693\(79\)90128-X](https://doi.org/10.1016/0370-2693(79)90128-X).
 - [50] M. D. Schwartz,
Quantum Field Theory and the Standard Model.
New York, NY: Cambridge University Press, 2014.
 - [51] S. Navas *et al.*, “Review of particle physics,”
Phys. Rev. D, vol. 110, p. 030001, 3 Aug. 2024.
DOI: 10.1103/PhysRevD.110.030001.
 - [52] R. D. Peccei and H. R. Quinn, “CP Conservation in the presence of pseudoparticles,”
Phys. Rev. Lett., vol. 38, pp. 1440–1443, 25 Jun. 1977.
DOI: 10.1103/PhysRevLett.38.1440.
 - [53] R. D. Peccei and H. R. Quinn, “Constraints imposed by CP conservation in the presence of pseudoparticles,”
Phys. Rev. D, vol. 16, pp. 1791–1797, 6 Sep. 1977.
DOI: 10.1103/PhysRevD.16.1791.

Bibliography

- [54] L. Álvarez-Gaumé and J. Ellis, “Eyes on a prize particle,” *Nature Phys.*, vol. 7, no. 1, pp. 2–3, 2010.
DOI: 10.1038/nphys1874.
- [55] A. J. Millar, G. G. Raffelt, J. Redondo, and F. D. Steffen, “Dielectric haloscopes to search for axion dark matter: Theoretical foundations,” *Journal of Cosmology and Astroparticle Physics*, vol. 2017, no. 01, pp. 061–061, Jan. 2017.
DOI: 10.1088/1475-7516/2017/01/061.
- [56] B. M. Brubaker,
First results from the haystac axion search,
2018.
arXiv: 1801.00835 [astro-ph.CO].
- [57] A. Á. Melcón *et al.*, “Axion searches with microwave filters: The RADES project,” *Journal of Cosmology and Astroparticle Physics*, vol. 2018, no. 05, pp. 040–040, May 2018.
DOI: 10.1088/1475-7516/2018/05/040.
- [58] S. Arguedas Cuendis,
“Dark matter axion search using novel RF resonant cavity geometries in the CAST experiment,”
Ph.D. dissertation, Vienna U., 2021.
DOI: 10.25365/thesis.70391.
- [59] P. Sikivie, “Detection rates for “invisible”-axion searches,” *Phys. Rev. D*, vol. 32, pp. 2988–2991, 11 Dec. 1985.
DOI: 10.1103/PhysRevD.32.2988.
- [60] A. Caputo, A. J. Millar, C. A. J. O’Hare, and E. Vitagliano, “Dark photon limits: A handbook,” *Physical Review D*, vol. 104, no. 9, Nov. 2021,
ISSN: 2470-0029.
DOI: 10.1103/physrevd.104.095029.
- [61] J. Golm, J. M. García-Barcelo, S. Arguedas Cuendis, S. Calatroni, W. Wuensch, and B. Dobrich, “Split-cavity tuning of a rectangular axion haloscope operating around 8.4 ghz,” *Frontiers in Physics*, vol. 12, 2024,
ISSN: 2296-424X.
DOI: 10.3389/fphy.2024.1372846.
- [62] S. J. Asztalos *et al.*, “Squid-based microwave cavity search for dark-matter axions,” *Phys. Rev. Lett.*, vol. 104, p. 041 301, 4 Jan. 2010.
DOI: 10.1103/PhysRevLett.104.041301.
- [63] S. J. Asztalos *et al.*, “Design and performance of the ADMX SQUID-based microwave receiver,” *Nucl. Instrum. Meth. A*, vol. 656, pp. 39–44, 2011.
DOI: 10.1016/j.nima.2011.07.019. arXiv: 1105.4203 [physics.ins-det].

- [64] K. M. Backes *et al.*, “A quantum-enhanced search for dark matter axions,” *Nature*, vol. 590, no. 7845, pp. 238–242, 2021.
DOI: 10.1038/s41586-021-03226-7. arXiv: 2008.01853 [quant-ph].
- [65] J. Golm *et al.*, “Thin Film (High Temperature) Superconducting Radiofrequency Cavities for the Search of Axion Dark Matter,” *IEEE Trans. Appl. Supercond.*, vol. 32, no. 4, p. 1 500 605, 2022.
DOI: 10.1109/TASC.2022.3147741. arXiv: 2110.01296 [hep-ex].
- [66] D. Ahn *et al.*, “Maintaining high Q-factor of superconducting $\text{YBa}_2\text{Cu}_3\text{O}_{7-x}$ microwave cavity in a high magnetic field,” Apr. 2019.
arXiv: 1904.05111 [physics.ins-det].
- [67] D. Ahn *et al.*, “Superconducting cavity in a high magnetic field,” Feb. 2020.
arXiv: 2002.08769 [physics.app-ph].
- [68] E. Berkowitz, M. I. Buchoff, and E. Rinaldi, “Lattice qcd input for axion cosmology,” *Physical Review D*, vol. 92, no. 3, Aug. 2015,
ISSN: 1550-2368.
DOI: 10.1103/physrevd.92.034507.
- [69] L. Fleury and G. D. Moore, “Axion dark matter: Strings and their cores,” *Journal of Cosmology and Astroparticle Physics*, vol. 2016, no. 01, pp. 004–004, Jan. 2016,
ISSN: 1475-7516.
DOI: 10.1088/1475-7516/2016/01/004.
- [70] M. Dine, P. Draper, L. Stephenson-Haskins, and D. Xu, “Axions, instantons, and the lattice,” *Physical Review D*, vol. 96, no. 9, Nov. 2017,
ISSN: 2470-0029.
DOI: 10.1103/physrevd.96.095001.
- [71] M. Buschmann, J. W. Foster, and B. R. Safdi, “Early-universe simulations of the cosmological axion,” *Physical Review Letters*, vol. 124, no. 16, Apr. 2020,
ISSN: 1079-7114.
DOI: 10.1103/physrevlett.124.161103.
- [72] *Cst studio suite*,
<https://www.3ds.com/de/produkte-und-services/simulia/produkte/cst-studio-suite/>,
Accessed: 2023-09-27.
- [73] H. Kamerlingh Onnes, “Further experiments with liquid helium. C. On the change of electric resistance of pure metals at very low temperatures etc. IV. The resistance of pure mercury at helium temperatures,” *Koninklijke Nederlandse Akademie van Wetenschappen Proceedings Series B Physical Sciences*, vol. 13, pp. 1274–1276, Jan. 1910.

Bibliography

- [74] J. N. Rjabinin and L. Schubnikow, “Magnetic properties and critical currents of superconducting alloys,”
Physikalische Zeitschrift der Sowjetunion, vol. 7, pp. 122–125, 1 1935.
- [75] *Type-II superconductor*,
https://en.wikipedia.org/wiki/Type-II_superconductor,
Accessed: 2023-09-14.
- [76] R. O. W. Meissner, “Ein neuer effekt bei eintritt der supraleitfähigkeit,”
Naturwissenschaften, vol. 21, pp. 787–788, 1933.
- [77] S. Posen, M. Checchin, O. Melnychuk, T. Ring, I. Gonin, and T. Khabiboulline, “High-quality-factor superconducting cavities in tesla-scale magnetic fields for dark-matter searches,”
Phys. Rev. Appl., vol. 20, p. 034004, 3 Sep. 2023.
DOI: 10.1103/PhysRevApplied.20.034004.
- [78] K. Schmieden *et al.*,
Study of nbn as superconducting material for the usage in superconducting radio frequency cavities,
2024.
arXiv: 2412.14958 [cond-mat.supr-con].
- [79] W. Ratzinger, S. Schenk, and P. Schwaller,
A coordinate-independent formalism for detecting high-frequency gravitational waves,
2024.
arXiv: 2404.08572 [gr-qc].
- [80] S. Schenk, K. Schmieden, and P. Schwaller,
Signatures of high-frequency gravitational waves in electromagnetic cavities,
2025.
arXiv: 2512.20592 [hep-ph].
- [81] D. M. Pozar,
Microwave engineering /,
4th ed.
New Delhi : Wiley, 2012,
Includes index.
- [82] https://www.coax.co.jp/en/wcaxp/wp-content/themes/coax/pdf/cryogenic_cable_catalogue.pdf,
Accessed: 28-08-2025.
- [83] P. Welch, “The use of fast fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms,”
IEEE Transactions on Audio and Electroacoustics, vol. 15, no. 2, pp. 70–73, 1967.
DOI: 10.1109/TAU.1967.1161901.
- [84] B. Brubaker, L. Zhong, S. Lamoreaux, K. Lehnert, and K. van Bibber, “Haystac axion search analysis procedure,”
Physical Review D, vol. 96, no. 12, Dec. 2017,
ISSN: 2470-0029.
DOI: 10.1103/physrevd.96.123008.

- [85] M. S. Turner, “Periodic signatures for the detection of cosmic axions,”
Phys. Rev. D, vol. 42, pp. 3572–3575, 10 Nov. 1990.
DOI: 10.1103/PhysRevD.42.3572.
- [86] G. Cowan,
Statistical data analysis.
Oxford University Press, USA, 1998.
- [87] K. C. Ruffin and G. Rybka,
Parametric resonance in rf axion haloscopes,
2024.
arXiv: 2411.17609 [astro-ph.IM].
- [88] C. O’Hare,
Cajohare/axionlimits: Axionlimits,
<https://cajohare.github.io/AxionLimits/>,
Version v1.0,
Jul. 2020.
DOI: 10.5281/zenodo.3932430.
- [89] A. Escrivà, F. Kühnel, and Y. Tada,
“Primordial black holes,”
In *Black Holes in the Era of Gravitational-Wave Astronomy*.
Elsevier, 2024,
Pp. 261–377,
ISBN: 9780323956369.
DOI: 10.1016/b978-0-32-395636-9.00012-8.
- [90] A. M. Green and B. J. Kavanagh, “Primordial black holes as a dark matter candidate,”
Journal of Physics G: Nuclear and Particle Physics, vol. 48, no. 4, p. 043 001, Feb.
2021,
ISSN: 1361-6471.
DOI: 10.1088/1361-6471/abc534.
- [91] *Current dm fraction of pbh limits*,
<https://github.com/bradkav/PBHbounds/>,
Accessed: 2026-03-30.
DOI: 10.5281/zenodo.3538999.
- [92] N. Aggarwal *et al.*, “Challenges and opportunities of gravitational-wave searches
above 10 khz,”
Living Reviews in Relativity, vol. 28, no. 1, Nov. 2025,
ISSN: 1433-8351.
DOI: 10.1007/s41114-025-00060-5.
- [93] D. Amaral *et al.*,
*Global detector network to search for high-frequency gravitational waves (gravnet):
Conceptual design*,
2026.
arXiv: 2603.24645 [astro-ph.IM].

Bibliography

- [94] G. Franciolini, A. Maharana, and F. Muia, “Hunt for light primordial black hole dark matter with ultrahigh-frequency gravitational waves,”
Physical Review D, vol. 106, no. 10, Nov. 2022,
ISSN: 2470-0029.
DOI: 10.1103/physrevd.106.103520.
- [95] M. E. Gertsenshtein and V. I. Pustovoit, “On the Detection of Low Frequency Gravitational Waves,”
Sov. Phys. JETP, vol. 16, p. 433, 1962.
- [96] M. Maggiore,
Gravitational Waves. Vol. 1: Theory and Experiments.
Oxford University Press, 2007,
ISBN: 978-0-19-857074-5.
DOI: 10.1093/acprof:oso/9780198570745.001.0001.
- [97] C. Gatti, L. Visinelli, and M. Zantedeschi, “Cavity detection of gravitational waves: Where do we stand?”
Physical Review D, vol. 110, no. 2, Jul. 2024,
ISSN: 2470-0029.
DOI: 10.1103/physrevd.110.023018.
- [98] M. Maggiore,
Gravitational Waves: Volume 2: Astrophysics and Cosmology.
Oxford University Press, 2018,
ISBN: 978-0-19-857089-9.
DOI: 10.1093/oso/9780198570899.001.0001.
- [99] N. Pompeo, R. Rogai, A. Augieri, V. Galluzzi, G. Celentano, and E. Silva, “Reduction in the field-dependent microwave surface resistance in $\text{YBa}_2\text{Cu}_3\text{O}_{7-\delta}$ with submicrometric BaZrO_3 inclusions as a function of BaZrO_3 concentration,”
Journal of Applied Physics, vol. 105, no. 1, Jan. 2009,
ISSN: 1089-7550.
DOI: 10.1063/1.3056179.
- [100] G. Ghigo *et al.*, “Heavy ion irradiation effects on the high-frequency properties of YBCO and Nb_3Sn thin films,”
Superconductivity, vol. 13, p. 100 149, 2025,
ISSN: 2772-8307.
DOI: <https://doi.org/10.1016/j.supcon.2024.100149>.
- [101] N. Pompeo, K. Torokhtii, A. Alimenti, G. Sylva, V. Braccini, and E. Silva, “Pinning properties of FeSeTe thin film through multifrequency measurements of the surface impedance,”
Superconductor Science and Technology, vol. 33, no. 11, p. 114 006, Sep. 2020,
ISSN: 1361-6668.
DOI: 10.1088/1361-6668/abb35c.
- [102] L. Visinelli and M. Naydenov,
Transient axion streams from disrupted miniclusters,
2026.

- arXiv: 2605.28005 [astro-ph.GA].
- [103] D. Kim, J. Jeong, S. Youn, Y. Kim, and Y. K. Semertzidis, “Revisiting the detection rate for axion haloscopes,”
Journal of Cosmology and Astroparticle Physics, vol. 2020, no. 03, pp. 066–066, Mar. 2020,
ISSN: 1475-7516.
DOI: 10.1088/1475-7516/2020/03/066.
 - [104] D. Chouhan, T. Schneemann, K. Schmieden, and M. Schott,
Conceptual design report of the supax experiment,
2025.
arXiv: 2505.07541 [hep-ex].
 - [105] G. Carosi *et al.*, “Search for axion dark matter from 1.1 to 1.3 ghz with admx,”
Physical Review Letters, vol. 135, no. 19, Nov. 2025,
ISSN: 1079-7114.
DOI: 10.1103/d7mg-6sqq.
 - [106] S. Ahn *et al.*, “Extensive search for axion dark matter over 1 ghz with capp’s main axion experiment,”
Phys. Rev. X, vol. 14, p. 031 023, 3 Aug. 2024.
DOI: 10.1103/PhysRevX.14.031023.

Appendix

A. Normalisation of Continuous and Discretised Lineshape

To show that by design

$$\int_0^\infty f(\nu) d\nu = 1 = \sum_q D_q \quad (\text{A.1})$$

let us consider first the definition of the virialised dark matter lineshape $f(\nu)$ as given in eq. (6.40).

$$f(\nu) = \frac{2}{\sqrt{\pi}} \left(\sqrt{\frac{3}{2}} \frac{1}{r \nu_a \langle \beta^2 \rangle} \right) \sinh \left(3r \sqrt{\frac{2(\nu - \nu_a)}{\nu_a \langle \beta^2 \rangle}} \right) \exp \left(-\frac{3(\nu - \nu_a)}{\nu_a \langle \beta^2 \rangle} - \frac{3r^2}{2} \right)$$

The first point to make is that obviously there is no power for frequencies below the frequency corresponding to the rest mass of the DM particle ν_a (index a for axion to stay consistent in the notation, but this is generally valid for all kinds of particles considered). So $f(\nu) = 0$ for $\nu < \nu_a$ and therefore

$$\int_0^\infty f(\nu) d\nu = \int_{\nu_a}^\infty f(\nu) d\nu. \quad (\text{A.2})$$

We introduce a substitution

$$t = \frac{2(\nu - \nu_a)}{\nu_a \langle \beta^2 \rangle} \quad (t \in [0, \infty)), \quad d\nu = \frac{\nu_a \langle \beta^2 \rangle}{2} dt. \quad (\text{A.3})$$

and simplify by factoring the constants (with respect to t) out of the integral.

$$\int_{\nu_a}^\infty f(\nu) d\nu = \left(\frac{2}{\sqrt{\pi}} \cdot \frac{\sqrt{3/2}}{r \nu_a \langle \beta^2 \rangle} \right) \cdot \left(\frac{\nu_a \langle \beta^2 \rangle}{2} \right) \int_0^\infty \sinh(3r\sqrt{t}) e^{-\frac{3}{2}t - \frac{3}{2}r^2} dt \quad (\text{A.4})$$

$$= \frac{\sqrt{3/2}}{r\sqrt{\pi}} e^{-\frac{3}{2}r^2} \int_0^\infty \sinh(3r\sqrt{t}) e^{-\frac{3}{2}t} dt \quad (\text{A.5})$$

Putting constants to the side for now, for the remaining integral we substitute

$$u = \sqrt{t}, \quad dt = 2u du \quad (\text{A.6})$$

$$\int_0^\infty \sinh(3r\sqrt{t}) e^{-\frac{3}{2}t} dt = \int_0^\infty 2u \sinh(3ru) e^{-\frac{3}{2}u^2} du. \quad (\text{A.7})$$

We can use the standard Gaussian integral

$$\int_0^\infty x e^{-ax^2} \sinh(bx) dx = \frac{b\sqrt{\pi}}{4a\sqrt{a}} e^{\frac{b^2}{4a}} \quad \text{for } a > 0 \quad (\text{A.8})$$

A. Normalisation of Continuous and Discretised Lineshape

found in various integral tables and identify the integrals (A.7) and (A.8) with $a = 3/2$ and $b = 3r$ to obtain

$$\int_0^\infty 2u \sinh(3ru) e^{-\frac{3}{2}u^2} du = \frac{3r\sqrt{\pi}}{2 \cdot (3/2) \cdot \sqrt{3/2}} e^{\frac{(3r)^2}{4 \cdot (3/2)}} = \frac{r\sqrt{\pi}}{\sqrt{3/2}} e^{\frac{3}{2}r^2}. \quad (\text{A.9})$$

Now we simply insert this result into eq. (A.5)

$$\int_{\nu_a}^\infty f(\nu) d\nu = \frac{1}{\sqrt{\pi}} \frac{\sqrt{3/2}}{r} e^{-\frac{3}{2}r^2} \int_0^\infty \sinh(3r\sqrt{t}) e^{-\frac{3}{2}t} dt \quad (\text{A.10})$$

$$= \frac{\sqrt{3/2}}{r\sqrt{\pi}} e^{-\frac{3}{2}r^2} \cdot \frac{r\sqrt{\pi}}{\sqrt{3/2}} e^{\frac{3}{2}r^2} = 1 \quad (\text{A.11})$$

and have shown that, independently of the rest mass ν_a and the local rotational speed determining $\langle \beta^2 \rangle$, the continuous lineshape $f(\nu)$ can be interpreted as a probability density function. It follows that for any partition of $[\nu_a, \infty)$ into disjoint intervals $I_q = [\nu_a + (q - 1)\Delta\nu, \nu_a + q\Delta\nu)$ the quantities

$$D_q = \int_{\nu_a + (q-1)\Delta\nu}^{\nu_a + q\Delta\nu} f(\nu) d\nu = \int_{I_q} f(\nu) d\nu \geq 0 \quad (\text{A.12})$$

define the probability mass on the interval I_q , or bin q , which satisfy

$$\sum_{q=1}^\infty D_q = \int_{\bigcup_q I_q} f(\nu) d\nu = \int_{\nu_a}^\infty f(\nu) d\nu = 1. \quad (\text{A.13})$$

B. Computationally Efficient Limit Setting

The following is a reformulation of the limit setting process shown in the main part. It is, in my opinion, less instructive than its representation in section 6.6.1, but for the two data sets that have been analysed in this work it is faster by a factor of approximately 10,000, which is not negligible, so it is presented here for future reference.

Since the signal model of eq. (6.46) depends only on a single parameter A , the likelihood function can be reduced to an effective one-dimensional problem. Starting from

$$S(\delta_k^g, A) = \sum_k \frac{(\delta_k^g - A D_k^S)^2}{(\sigma_k^g)^2} \quad (\text{B.1})$$

we can expand the square and get

$$S(\delta_k^g, A) = \underbrace{\sum_k \frac{(\delta_k^g)^2}{(\sigma_k^g)^2}}_F - 2A \underbrace{\sum_k \frac{\delta_k^g D_k^S}{(\sigma_k^g)^2}}_B + A^2 \underbrace{\sum_k \frac{(D_k^S)^2}{(\sigma_k^g)^2}}_C \quad (\text{B.2})$$

$$= F - 2AB + A^2 C. \quad (\text{B.3})$$

Completing the square in A gives

$$S(\delta_k^g, A) = C \left(A^2 - 2A \frac{B}{C} \right) + F \quad (\text{B.4})$$

$$= C \left[\left(A - \frac{B}{C} \right)^2 - \left(\frac{B}{C} \right)^2 \right] + F \quad (\text{B.5})$$

$$= C \left(A - \frac{B}{C} \right)^2 + F - \frac{B^2}{C} \quad (\text{B.6})$$

where the term $F - B^2/C$ is clearly independent of A . The resulting factor $\exp(F - B^2/C)$ can therefore be factored out of the integrals of eq. (6.49) and cancel out. We can therefore identify

$$\tilde{\delta}^g = \frac{B}{C} = \frac{\sum_k \frac{\delta_k^g D_k^S}{(\sigma_k^g)^2}}{\sum_k \frac{(D_k^S)^2}{(\sigma_k^g)^2}} \quad (\text{B.7})$$

$$\tilde{\sigma}^g = C^{-1/2} = \left[\sum_k \frac{(D_k^S)^2}{(\sigma_k^g)^2} \right]^{-1/2} \quad (\text{B.8})$$

to write $S(\delta_k^g, A)$ in a familiar form.

$$S(\delta_k^g, A) = \frac{(A - \tilde{\delta}^g)^2}{(\tilde{\sigma}^g)^2} + \text{const.} \quad (\text{B.9})$$

Crucially, this is just a number here as the lineshape weighting, the N -entry sum for an N -bin long signal, is done just once in equations (B.7) and (B.8) instead of in every iteration of the integral. This is what yields the large efficiency boost.

B. Computationally Efficient Limit Setting

As before, we can generalise the previous result, which assumes independent bins, to include bin-to-bin correlations. Using the signal model

$$\mathbf{y} = A \mathbf{D}^S, \quad (\text{B.10})$$

the likelihood can again be reduced to an equivalent one-dimensional problem by defining

$$\tilde{\delta}^g = \frac{(\mathbf{D}^S)^T C^{-1} \delta^g}{(\mathbf{D}^S)^T C^{-1} \mathbf{D}^S} \quad (\text{B.11})$$

$$\tilde{\sigma}^g = [(\mathbf{D}^S)^T C^{-1} \mathbf{D}^S]^{-1/2} \quad (\text{B.12})$$

One might note here that the quantity $\tilde{\delta}^g$ is similar to the lineshape weighted SNR of the alternative analysis approach (sec. 6.6.2). In fact, in terms of the uncertainty it is exactly the same: $\tilde{\delta}^g/\tilde{\sigma}^g \equiv \delta^l/\sigma^l = D^l/R^l$. Its shape appears to be imprinted with a certain structure and so does the limit set with the main method. We can try to explain this to a certain extent here. Even if the original fluctuations δ_k^g are perfectly Gaussian distributed and statistically independent between bins, the lineshape weighting, one might call it filtering, induces correlations between neighbouring values of $\tilde{\delta}^g$. To illustrate this, consider a discrete white-noise process $x[n]$ with

$$\mathbb{E}[x[n]] = 0 \quad (\text{B.13})$$

$$\mathbb{E}[x[n] x[m]] = \begin{cases} \sigma^2, & n = m \\ 0, & n \neq m. \end{cases} \quad (\text{B.14})$$

Filtering this process with a signal shape $h[k]$ yields

$$y[n] = \sum_k h[k] x[n - k] \quad (\text{B.15})$$

and the autocorrelation of τ bins of this process is then

$$R_y[\tau] = \mathbb{E}[y[n] y[n + \tau]] \quad (\text{B.16})$$

$$= \sum_k \sum_m h[k] h[m] \mathbb{E}[x[n - k] x[n + \tau - m]]. \quad (\text{B.17})$$

Using eq. (B.14), the expectation value is non-zero only for

$$n - k = n + \tau - m \quad (\text{B.18})$$

$$\Rightarrow m = k + \tau \quad (\text{B.19})$$

and therefore the double sum collapses to

$$R_y[\tau] = \sigma^2 \sum_k h[k] h[k + \tau] \quad (\text{B.20})$$

showing that the autocorrelation of the filtered process is equal to the autocorrelation of the signal shape itself. This is shown in Fig. 1 for generated white noise, the data of both runs presented in this thesis and the signal template. In our specific case that means that

Appendix

the correlations, and hence the characteristic structure, of $\tilde{\delta}^g$ (and the limit set by it, as well as D_k^l in the alternative approach) are determined by the autocorrelation properties of D_k^S . Since the SG filtered signal template (see Fig. 6.9) contains alternating positive and negative side lobes beside the main peak, its autocorrelation is also oscillating to a certain degree. Neighbouring values of $\tilde{\delta}^g$ are therefore expected to show correlated oscillatory behaviour even in the absence of any physical signal, purely induced by the lineshape weighting procedure. Since the likelihood can be written as a Gaussian in A centred at $\tilde{\delta}^g$, the upper limit A_{UL} is directly determined by $\tilde{\delta}^g$ and $\tilde{\sigma}^g$. Consequently, any oscillation-like structure in $\tilde{\delta}^g$ induced by the lineshape weighting is expected to be reflected in A_{UL} , and hence the limits χ and $g_{a\gamma\gamma}$, as well. However, A_{UL} is not simply proportional to $\tilde{\delta}^g$. The restriction to physical signal amplitudes $A \geq 0$ of eq. (6.48) maps negative and positive fluctuations of $\tilde{\delta}^g$ asymmetrically. As a result, the limit is expected to exhibit a structure similar but not identical to that of $\tilde{\delta}^g/\tilde{\sigma}^g$, which for our signal template and choice of SG-filter appears to be somewhat periodic.




Figure 1.: Autocorrelation (AC) R_y for lags $\tau = 1, \dots, 100$ for the axion (top) and DP run (bottom). In both cases the general structure of the data's AC follows that of the signal template, explaining the oscillating structure of $\tilde{\delta}^g$ and the limit.

C. Characterisation Data for the NbN Superconducting S-Curve

Temperature [K]	Q_L	β	Q_0
200	26665 ± 29	-	-
180	26942 ± 30	-	-
170	31991 ± 42	-	-
150	35210 ± 51	-	-
139	36170 ± 54	-	-
130	37712 ± 58	-	-
120	39686 ± 65	-	-
110	41857 ± 72	-	-
100	38645 ± 61	-	-
90	42419 ± 74	-	-
80	48949 ± 98	-	-
70	52213 ± 112	-	-
60	55576 ± 127	-	-
50	68714 ± 116	-	-
45	68786 ± 116	-	-
20	87436 ± 188	-	-
18.4	81096 ± 54	0.0017 ± 0.0002	81235 ± 57
14.5	84147 ± 58	0.0018 ± 0.0002	84295 ± 60
13.1	90046 ± 67	0.0033 ± 0.0001	90346 ± 68
12.1	118276 ± 57	0.0064 ± 0.0006	119037 ± 90
12	109635 ± 296	0.0072 ± 0.0001	110423 ± 298
11.6	121290 ± 60	0.0080 ± 0.0063	122258 ± 772
11.2	205887 ± 174	0.0104 ± 0.0592	208024 ± 12183
9.9	234392 ± 1352	-	-
9.7	226650 ± 211	0.0177 ± 0.0038	230658 ± 880
9.1	240845 ± 476	0.0218 ± 0.0002	246086 ± 490
7.8	288142 ± 1363	0.0285 ± 0.0002	296359 ± 1403
7.5	310879 ± 1586	0.0304 ± 0.0003	320328 ± 1637
7.4 #1	313130 ± 1609	0.0331 ± 0.0003	323509 ± 1665
7.4 #2	316923 ± 1649	0.0326 ± 0.0003	327253 ± 1705
7.4 #3	331765 ± 1807	0.0327 ± 0.0003	342607 ± 1868
7.4 #4	312193 ± 1600	0.0329 ± 0.0003	322464 ± 1655
7.38	288213 ± 341	0.0323 ± 0.0016	297512 ± 588
3.8	322647 ± 1712	0.0727 ± 0.0008	346103 ± 1854

Table 1.: Measurements of superconductivity curve of NbN coated cavity without application of external magnetic field.

D. Cryostat Operation

For a typical measurement procedure several steps are to be conducted in a certain sequence to allow for smooth operation while taking data. These are:

1. Pumping
2. Leak Testing
3. Flushing the system with helium
4. Filling the thermal nitrogen shield
5. Filling the internal helium dewar
6. Cooldown

Step 1: Pumping

After the inset is assembled to meet the measurement-specific requirements (all cables fixed, no loose SMA connections, electronic height adjusted to use with or without magnet, etc.) it is inserted and all valves are to be closed. Now the main pump is switched on, ideally pumping a vacuum overnight or at least several hours. The cavity only has a very small groove for the air inside to be sucked out, so evacuating the probe volume with the cavity inside takes some time. Simultaneously, the vacuum shield should be pumped at least overnight if its pressure is not still below 10^{-5} mbar from a previous experimental run.

Step 2: Leak Testing

Once the shield, the internal dewar and the probe volume are sufficiently pumped, a helium leak tester should be used to check if all gas system connections are tight. While a leak rate of 10^{-10} mbar l/s and lower needs to be achieved to qualify the system as *technically tight* a leak rate of $< 10^{-7}$ mbar l/s is sufficient for our use since the helium liquefier can handle some minor impurities in the returned helium. The unit mbar l/s stands for millibar litres per second. The leak tester has a small gas pistol which should be used to emit small amounts of helium in the vicinity of all connections which have been opened. If there are leaking connections the cryostat is to be vented, opened and the sealing ring should be inspected. Counterintuitively, brand new sealing rings are often more problematic than used ones. This is because many manufacturers deliver the rings with barely visible aluminium shavings from the centring ring. The best way of cleaning these rings is by simply running them between your fingers until there are no more visible impurities. ‘Cleaning cloths’ or even chemical cleaning agents usually make the problem worse either by attacking the rings’ material (in case of chemicals) or leaving new micro fibres (in the case of cloths). Once all leaky connections are fixed, repeat from Step One.

Step 3: Flushing the system

Having made sure there are no leaks in the entire system (pump, internal dewar and probe volume) and the inside is pumped to $\approx 10^{-2}$ mbar it is now time to flush out the residual air. For this the valve between cryostat and pump is closed and helium is flushed into the system

from a gas bottle (usually leak testers come with a helium bottle anyway) until the pressure is at 1 bar. Now the system is pumped again into the double digit mbar range. This process should be repeated at least 3 times. At the end the cryostat should be filled with helium at a pressure slightly above ambient pressure.

Step 4: Filling the thermal nitrogen shield

The cryostat is equipped with a thermal shield of liquid nitrogen. This compartment has two openings, one KF16 and one simple exhaust pipe. To fill the shield the KF16 blind flange is removed and the transfer line of a liquid nitrogen dewar is inserted. The exhaust pipe should be open since in the process of filling a lot of nitrogen evaporates (the shield itself has to cool down as well). The level of liquid nitrogen in the system can be checked using an acrylic glass rod (see Fig. 2), where ambient humidity will freeze only where the rod was submerged in the liquid. Once the shield is full* close the KF16 blind flange. The shield will continue to boil a bit, but settles down after a few minutes. The exhaust pipe should then be closed with a Bunsen valve. This kind of valve can be easily crafted yourself, consisting only of a hose which is closed at one end (by letting some two-component glue run into the end for example) and has a 2 – 3 cm long slit in it. This allows gas to escape when there is overpressure inside but will not suck air inside when the pressure inside falls below ambient pressure. This way we can make sure there is only nitrogen and no moisture (from air) inside but there is also no pressure building up in the shield.



Figure 2.: Stick made from acrylic glass used to determine the filling level of the liquid nitrogen shield. The stick will condense the moisture of the air and fog up where it has touched the liquid nitrogen.

Step 5: Filling the internal helium dewar

With helium at slight overpressure inside the probe volume and internal dewar, the nitrogen shield filled and the insulating vacuum pumped now finally liquid helium can be introduced into the system. For this you first disconnect the helium return line from the dewar and build slight overpressure in the dewar with the balloon, to prevent polluting the helium gas during insertion of the transfer line. The transfer line is slowly inserted into the dewar, staying above the liquid in the gaseous area. This now flushes the transfer line with helium. Since there is only very light overpressure in the dewar, this may take up to 10 minutes, depending on how safe you want to be not to introduce air into your system

*Unless you are incredibly lucky and test the LHe level with the stick 1 cm before it's full you will notice the shield being full by some light splashing of liquid nitrogen. Obviously filling should immediately be stopped then, however there is no need for increased concern due to the splashing. Splashing liquid nitrogen will not harm skin due to the Leidenfrost effect, with one exception: the eyes. As the eye holds a lot of water which can freeze very quickly, leading to permanent eye damage. That's why it is highly recommended to always wear goggles or a face shield when handling any cryogenic liquids.

Appendix

after you have spent so much time making sure there is only helium inside - patience is advised. Once the transfer line is flushed you can open the internal dewar and swiftly insert the transfer line. This will likely introduce a small amount of air (and with it moisture) into the system. For this reason the system should be flushed with gas through the transfer line for some time. For this it is important that neither the internal dewar nor the external dewar carrying the liquid helium is not connected to the return line. The helium should flow from the external dewar through the transfer line, the internal dewar, the needle valve connecting the internal dewar and the probe volume* and then into the return line. Once you feel comfortable that you flushed out any impurities and there is only gaseous helium in your system you can now finally insert the transfer line into the dewar holding the liquid helium. It is advised to do this very slowly and having a constant eye on the pressure since the transfer line comes from room temperature and will certainly boil the helium until it is cooled to ≈ 4.2 K. The level sensor in the internal dewar will show an increase in resistance once liquid helium is filling the internal dewar.

Step 6: Cooldown

If you see the internal dewar filling up you can adjust the needle valve and the pump to a configuration where the pressure is around 100 – 200 mbar and the needle valve is open to an extent where lots of gaseous helium can flow but no liquid is filling the probe volume. This, to some extent, is based on experience. You will be able to see the temperature drop in the Grafana display. A slow cooldown is approximately 2 K per minute, a fast cooldown will be at up to 10 K per minute, which is pushing it a little bit and risking filling the volume with liquid. Staying at a cooldown rate around 5 K per minute is what we are usually aiming at. With this we can cool down from room temperature to ≈ 4 K in just under an hour. The key for a swift cooldown is staying at low pressure (50 – 100 mbar) in the probe volume. This requires a fully open pump and a just barely opened needle valve.

Step 7: Starting data acquisition

Wait until the cavity is cooled to the target temperature. The temperature in Grafana might indicate the cavity is cold but since the sensor is on the outside the entire copper might not be in thermal equilibrium. As a safe indicator that the cavity is cooled entirely is if the resonance frequency is stationary, this should be monitored with the VNA. Once the cavity's resonance is at its target frequency (adjustable by pressure) and stationary you can detach the SMA cable for the input and change the output from feeding into the VNA to now feeding into the RSA. The now open RF port should be closed with a $50\ \Omega$ terminator. The port is attenuated and is only reaching the cavity's weakly coupled port so not much power should reach the cavity anyway. Still, an open RF port is effectively an antenna feeding 300 K thermal noise power into the cavity. This is effectively avoided using the terminator. Finally, the data acquisition process can be started.

*This is crucial. There has been an instance in which the flushing was not done through the needle valve. As a consequence, the ambient humidity of the air, introduced during inserting the transfer line, was cooled until it froze the needle valve shut. No helium flow into the probe volume possible meant stopping the entire measurement run, warming everything up and starting again. You do not want to waste an entire day like that because you were too impatient to wait 2 minutes.