AARHUS UNIVERSITY

Measuring the Properties of Antihydrogen

Peter Granum

June 2022

PhD Dissertation

Supervisor: Prof. Jeffrey Scott Hangst

The ALPHA Collaboration, CERN

Department of Physics and Astronomy
Aarhus University, Denmark

Abstract

This thesis describes the latest results of the on-going efforts to measure the properties of antihydrogen within the ALPHA collaboration. More specifically, it covers the construction and commissioning of the ALPHA-g experiment [1], and the plans to measure how antimatter behaves in Earth's gravitational field. A special emphasis is on the ALPHA-g magnet system used to confine and manipulate the antihydrogen atoms. Tests of methods for calculating magnetic fields relevant for simulations [2] are covered as well. Amongst the described results from the ALPHA-2 experiment is the characterisation of the fine structure of antihydrogen [3]. The combined result of the measured 1S–2P$_{1/2}$ and 1S–2P$_{3/2}$ transitions agrees with the prediction of quantum electrodynamics to 16 parts per billion. The thesis also describes the demonstration of the first laser cooling of antimatter [4], which paves the way for a measurement of the 1S–2S transition in antihydrogen with hydrogen-like precision, and a measurement of antigravity with 1% precision. Both are future goals of the ALPHA collaboration.

Resumé

Denne afhandling beskriver de seneste resultater af ALPHA-gruppens fortsatte arbejde med at måle antibrints egenskaber. Mere præcist beskriver den konstruktionen og de første test af ALPHA-g eksperimentet [1] og planerne for at måle, hvordan antibrint opfører sig i Jordens tyngdefelt. Et særligt fokus er på beskrivelsen af magnetsystemet i ALPHA-g, som bruges til at fange og manipulere antibrintatomerne. Tests af metoder til beregning af magnetiske felter relevante for simuleringer [2] er også beskrevet. Blandt de beskrevne resultater fra ALPHA-2 eksperimentet er en karakteristik af antibrints finstruktur [3]. De samlede resultater af de målte 1S–2P$_{1/2}$ og 1S–2P$_{3/2}$ overgange er i overensstemmelse med kvanteelektrodynamik indenfor 16 milliardendedele. Afhandlingen beskriver også en demonstration af det første laserkølede antistof [4], hvilket baner vejen for en måling af 1S–2S overgangen i antibrint med brint-lignende præcision og en måling af antityngdekraft med 1% præcision. Begge er fremtidige mål for ALPHA-gruppen.

Contents

Abstract **iii**

Resumé **v**

Acknowledgements **ix**

Preface **xi**

List of Abbreviations **xiii**

Publications **xv**

1 Introduction **1**

2 Theoretical Background **3**

2.1 History of Antimatter 3
2.2 CPT Violation and BSM Physics 5
2.3 The CERN Antimatter Community 7
2.4 Charged Particles in Penning Traps 8
2.5 Trapping Antihydrogen 13
2.6 The Spectrum of (Anti)Hydrogen 15

3 The AD and ELENA **21**

3.1 The AD 21
3.2 ELENA 22

4 The ALPHA-2 Apparatus **23**

4.1 Motivation 23
4.2 Setup 24

5 ALPHA-2 Results **35**

5.1 Finestructure Measurements 35
5.2 Laser Cooling of Antihydrogen 40

5.3 Sympathetic Cooling of Positrons 44

6 The ALPHA-g Apparatus 47
6.1 Motivation 47
6.2 Overall Setup 48
6.3 Magnet Layout 55
6.4 Electrodes 68
6.5 Magnet Control System 68
6.6 The External Solenoid 78
6.7 Tickle Power Supply 84
6.8 DCCT Precision 85
6.9 KEPCO and DCCT Tests 88
6.10 Magnetometry 90

7 ALPHA-g Commissioning 93
7.1 The 2018 Run 93
7.2 Babcock Cryogenics 94
7.3 Babcock Field Maps 98
7.4 BNL Magnet Commissioning 100
7.5 Magnet Inductances 105
7.6 Particle work 109

8 Magnetic Field Calculations 117
8.1 Calculating Magnetic Fields of Solenoids 118

9 Conclusion and Outlook 125

A Appendix 127
A.1 Detailed Magnet Tables 127
A.2 SC Circuit Diagram 131
A.3 Babcock Diagram 132

Bibliography 135

Acknowledgements

I would like to thank my supervisor Prof. Jeffrey Hangst for giving me the opportunity to do my PhD with the ALPHA collaboration. ALPHA is an incredible experiment, and it is run by a fantastic group of people. I cannot imagine that I will ever find a group of such skilled and hard-working people, and I am grateful to have been part of this collaboration. Thanks to present and former members of ALPHA.

A special thanks goes to Danielle Hodgkinson, Dr. William Bertsche, Dr. Daniel Maxwell, Jaspal Singh, and Dr. Emma Anderson in the Magnet Team for the great teamwork, the discussions, the guidance, all the hours of hard work, and the shared highs and lows – in particular to Danielle Hodgkinson, with whom I have been lucky to share the PhD journey. The magnet work described in this thesis is credited to them. A special thanks also goes to Dr. Joseph McKenna and Prof. Joel Fajans for many good discussions and their great help. I cannot imagine to have done my PhD without these people.

Thanks to Dr. Nathan Evetts and Dr. Emma Anderson for leading the NMR campaign, to Dr. Chukman So for sharing detailed information and figures of the ALPHA-g magnet system, to Adam Powell for leading the efforts on ECR measurements and sharing the obtained data, and to Andrew Christensen for leading the patch potential campaign and sharing his calculated ALPHA-g magnet inductances.

I would like to thank family – the most important people of all – for always supporting me. Thanks to my friends at ALPHA and around Geneva for all the good memories. Also thanks to the good people at Aarhus University, who have helped and guided me throughout my studies, and to everybody who helped me proofread this thesis.

Thanks to the Carlsberg Foundation for their generous support of the ALPHA experiment and thereby my PhD.

Preface

During my PhD studies, I have been based full-time at CERN, where the ALPHA experiment is located. The ALPHA collaboration consists of more than 50 people from about 15 different universities/institutions around the world. Some people are based full time at CERN, while others are based at their home institutions.

Working at ALPHA means that you contribute to all parts of the experiment rather than just a single project. This includes the daily running of the experiment, construction, testing, and commissioning of equipment and experimental protocols, and the general physics-discussions. The achievements of the collaboration are thus the result of a huge teamwork. I will therefore mainly use the pronoun *we* to narrate this thesis.

The results from the ALPHA-2 experiment presented in chapter 5 were obtained early in my time at ALPHA, when my main contribution was to the general operation of the experiment. When the focus of the collaboration was later changed to ALPHA-g, my efforts were focussed on the magnet system and the general commissioning and operation of the apparatus. Over the last years, the ALPHA-g Magnet Team has consisted of 3–4 people. The efforts on the ALPHA-g magnets described in this thesis are credited to the Magnet Team and the rest of the collaboration.

My time at ALPHA has included the Long Shut-down 2 (LS2) at CERN, during which the accelerator infrastructure was upgraded, and antiprotons were unavailable for the antimatter community. The shut-down began in November 2018 and ended in August 2021. In the summer prior to LS2, we focussed our efforts on obtaining the results in ALPHA-2 described in chapter 5, before the focus changed to the construction and commissioning of ALPHA-g, with the ambition of having it operational before LS2 as described in chapter 6 and 7.

COVID-19 has of course delayed the operations at ALPHA. Besides complicating working on a physics experiment within a group, is has also meant delays in deliveries from external contractors. Examples are the external and the internal magnets that arrived about a year late – just a few months before the return of antiprotons in 2021 – but it is hard to quantify the impact that COVID has had on the experiment.

A few technical notes: for the trapped low-energy particles in our experiment, temperature and energy might be used interchangeably, as they are related via Boltzmann's constant $k_b = 8.52 \cdot 10^{-5}$ eV/K. Similarly for atomic transitions, which are given in units of either frequency, energy, or wavelength, which are related via the speed of light, c, and Planck's constant, h. Typically, the plasmas in ALPHA do not fulfil the technical definitions of a plasma, but they are referred to as plasmas anyway. Unless otherwise indicated, *cooling of particles* refers to a reduction of their total energy.

Regarding atomic notation, for antihydrogen, the nL_J notation is used (for example $2\mathrm{P}_{3/2}$), where n, S, L, and J are the usual quantum numbers. For other systems with a single valence electron, the notation $(nl)^{2S+1}L_J$ is used, where the quantum numbers in the parenthesis refer to the valence electron, and the others are for the total system.

In accordance with GSNS rules, parts of this thesis were also used in the progress report for the qualifying examination.

List of Abbreviations

AD Antiproton Decelerator

ALPHA Antihydrogen Laser PHysics Apparatus

APS Applied Power Systems (company)

BNL Brookhaven National Laboratory

BSM Beyond the Standard Model

CERN Conseil Européen pour la Recherche Nucléaire

CPT Charge Parity Time

cRIO Compact Real time Input/Output system

DCCT Direct Current Current Transformer

ECR Electron Cyclotron Resonance

ELENA Extra Low ENergy Antiproton ring

EVC EVaporative Cooling

FFT Fast Fourier Transform

FPGA Field-Programmable ate Array

FWHM Full Width Half Maximum

HTS High Temperature Super conductor

I16/I32 16/32 bit Integer

IGBT Insulated Gate Bipolar Transistor

I/O Input/Output

IRR Infinite Impulse Response

LS2 Long Shutdown 2

LTS Low Temperature Superconductor

MCP MicroChannel Plate

MCS Magnet Control System

MOSFET Metal–Oxide–Semiconductor Field-Effect Transistor

NI National Instruments (company)

NMR Nuclear Magnetic Resonance

PMT PhotoMultiplier Tube

PS Power Supply

QD Quench Detection

QED Quantum Electro Dynamics

QPS Quench Protection System

rTPC radial Time Projection Chamber

RW Rotating Wall

SC Signal Conditioner

SCL SuperConducting Lead

SCR Silicon-Controlled Rectifier

SDR Strong Drive Regime

SiPM Silicon PhotoMultiplier

SM Standard Model

SME Standard Model Extension

TOF Time Of Flight

VCL Vapour Cooled Lead

Publications

Peer reviewed published articles:

- M. Ahmadi et al. Investigation of the fine structure of antihydrogen. *Nature*, 578(7795), February 2020
- C. J. Baker et al. Laser cooling of antihydrogen atoms. *Nature*, 592(7852):35–42, mar 2021.
- C. J. Baker et al. Sympathetic cooling of positrons to cryogenic temperatures for antihydrogen production. *Nature Communications*, 12(1), oct 2021
- Peter Granum, Magnus Linnet Madsen, Joseph Tiarnan Kerr McKenna, Danielle Louise Hodgkinson, and Joel Fajans. Efficient calculations of magnetic fields of solenoids for simulations. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 1034:166706, jul 2022.

Articles pending publication:

- C. J. Baker et al. Precision spectroscopy of the 1S-2S transition in antihydrogen: hyperfine structure and CPT invariance. *Pending publication.*
- C. J. Baker et al. Design and Performance of a Novel Low Energy Multi-Species Beamline for the ALPHA Antihydrogen Experiment. *Pending publication.*
- C. J. Baker et al. The Effects of Patch Potentials in Penning-Malmberg Traps. *Pending publication.*

Outreach articles:

- Peter Granum. Jagten på antityngdekraft og forklaringen på vores eksistens. *KVANT*, October 2020

Chapter 1

Introduction

Antimatter, which can be described as the mirror image of matter, is almost completely absent from our universe. Since its discovery in 1932 [5], scientists have only laid their hands on tiny amounts of antimatter and only the simplest particles and systems. In itself, the scarcity of antimatter and the lack of measurements makes it extremely interesting to measure its properties, but in addition, such measurements test our understanding of the fundamental principles of Nature.

The antimatter facility at CERN (Conseil Européen pour la Recherche Nucléaire) is the only place in the world, where antiatoms can be created and trapped, to allow measurements of their properties. The ALPHA experiment is the only experiment that has achieved to create and trap antihydrogen [6] – the simplest antiatom possible – and they have made numerous measurements of its properties [7, 8, 9, 10].

This thesis reports on the efforts to measure antigravity – that is, how antimatter behaves in the gravitational field of the Earth – with the ALPHA-g experiment [1], as well as the latest results from the ALPHA-2 experiment: an improved measurement of the fine structure of antihydrogen [3], and the first laser cooling of antimatter [4].

The outline of this thesis is as follows: chapter 2 provides a summary of the theory required to understand the work described in the rest of the thesis. Chapter 3 covers the production of 5 keV antiprotons at CERN, while chapter 4 and 5 describe the ALPHA-2 apparatus and the recent results. Chapter 6 introduces the ALPHA-g apparatus and the planned antigravity measurements, while chapter 7 describes the current status and the results obtained during the commissioning of ALPHA-g. Finally, chapter 8 deals with methods for calculating magnetic fields with an emphasis on simulating the behaviour of antihydrogen in ALPHA's magnetic minimum traps.

Chapter 2

Theoretical Background

This chapter briefly covers some of the theory needed to understand the motivation for studying antihydrogen, the results obtained, and the current work towards future measurements. It is by no means intended to be exhaustive or to introduce concepts beyond what is generally known by people within the field.

2.1 History of Antimatter

In 1928, P.A.M. Dirac published his famous relativistic wave equation for spin-1/2 particles – the Dirac equation [11] – which in natural units is

$$(i\gamma^{\mu}\partial_{\mu} - m)\,\psi = 0 \tag{2.1}$$

where γ^{μ} are the Dirac matrices, ∂_{μ} is the 4-derivative, m is the mass of the particle in question, and ψ is the Dirac spinor field. The Dirac equation incorporates both quantum mechanics and special relativity. The Dirac equation leads to solutions with negative energy, which is unphysical, but by quantising the Dirac field and thereby turning it into a quantum field theory, these negative energy solutions are avoided. Dirac quantum field theory is treated in many places; for example [12, 13], from which a few central equations are repeated here. The Lagrangian density of the Dirac field is

$$\mathcal{L}_{Dirac} = \bar{\Psi}\,(i\gamma^{\mu}\delta_{\mu} - m)\,\Psi \tag{2.2}$$

where the spinor field Ψ and the adjoint spinor field $\bar{\Psi}$ can be written in terms of two pairs of creation and annihilation operators $a_s^{\dagger}(\vec{p})$, $a_s(\vec{p})$, $b_s^{\dagger}(\vec{p})$, and $b_s(\vec{p})$ for two types of particles (a and b), the integral over all momentum states $\vec{p}$, the sum over all spin states $s = 1, 2$, and the eigenfunctions of the Dirac Hamiltonian $u^{(s)}(\vec{p})e^{i\vec{p}\cdot\vec{x}}$ and $v^{(s)}(\vec{p})e^{-i\vec{p}\cdot\vec{x}}$ with eigenvalues E_p

$$\Psi = \frac{1}{(2\pi)^3} \int \frac{d^3\vec{p}}{\sqrt{2E_p}} \sum_{s=1,2} \left(a_s(\vec{p}) u^{(s)}(\vec{p}) e^{-i\vec{p}\cdot\vec{x}} + b_s^\dagger(\vec{p}) v^{(s)}(\vec{p}) e^{i\vec{p}\cdot\vec{x}} \right) \tag{2.3}$$

$$\bar{\Psi} = \frac{1}{(2\pi)^3} \int \frac{d^3\vec{p}}{\sqrt{2E_p}} \sum_{s=1,2} \left(a_s^\dagger(\vec{p}) \bar{u}^{(s)}(\vec{p}) e^{i\vec{p}\cdot\vec{x}} + b_s(\vec{p}) \bar{v}^{(s)}(\vec{p}) e^{-i\vec{p}\cdot\vec{x}} \right) \tag{2.4}$$

The total momentum and the total charge are given as

$$\vec{P} = \frac{1}{(2\pi)^3} \int d^3\vec{p} \sum_{s=1,2} \vec{p} \left(a_s^\dagger(\vec{p}) a_s(\vec{p}) + b_s^\dagger(\vec{p}) b_s(\vec{p}) \right) \tag{2.5}$$

$$Q = \frac{1}{(2\pi)^3} \int d^3\vec{p} \sum_{s=1,2} \left(a_s^\dagger(\vec{p}) a_s(\vec{p}) - b_s^\dagger(\vec{p}) b_s(\vec{p}) \right) \tag{2.6}$$

Since the particle number operators for the two particle types, a and b, are given as

$$N_s^a(\vec{p}) = a_s^\dagger(\vec{p}) a_s(\vec{p}) \tag{2.7}$$

$$N_s^b(\vec{p}) = b_s^\dagger(\vec{p}) b_s(\vec{p}) \tag{2.8}$$

it is clear that the two particles species contribute equally to the momentum, so they must have the same mass, and equally but oppositely to the total charge, so the particles must have opposite charges. At this point, it seems natural to predict the existence of antimatter.

In 1932, Carl D. Anderson studied particle tracks caused by cosmic rays in a cloud chamber immersed in a magnetic field. He observed tracks with a curvature which corresponded to the electron charge to mass ratio, but with the opposite charge. Thereby, he discovered the positron and confirmed the existence antimatter [5].

The antiproton was observed for the first time in 1955 [14], but it was not until 1995 that a positron and an antiproton were combined to form the first antiatom at the LEAR facility at CERN [15], but the antihydrogen was too energetic to allow studies of its properties. The efforts of antimatter studies at CERN then concentrated on synthesising and trapping antihydrogen at the Antiproton Decelerator facility. In 2010, the ALPHA collaboration succeeded in trapping antihydrogen for the first time in history [6], which enabled the studies conducted on antihydrogen to date [7, 8, 10, 3].

2.2 CPT Violation and BSM Physics

The Standard Model (SM) incorporates Lorentz and CPT (Charge, Parity, Time) invariance. As a result, the SM predicts that antimatter is the exact "mirror image" of matter. Since matter and antimatter are always observed to be created in equal amounts, the universe must have consisted of half matter and half antimatter at the Big Bang. Today however, the universe is observed to consist almost exclusively of matter, which indicates some asymmetry between matter and antimatter. In 1967 A. Sakharov suggested three criteria [16], known as the Sakharov criteria, that must be fulfilled to cohere equal amounts of matter and antimatter at the Big Bang with the imbalance observed today:

- Baryon number violation
- C and CP violation
- Interactions out of thermal equilibrium

Interactions that violate baryon number are predicted by the SM, but they have so far not been observed in experiments. P violation (and indirectly C violation) was confirmed in Co-60 beta decays in 1956 [17], and CP violation was confirmed in kaon decays in 1964 [18]. The Sakharov criteria assume CPT symmetry, and as a consequence, the third criterion can be replaced by CPT violation [19], which has not yet been observed.

The search for Beyond the Standard Model (BSM) physics, such as CPT violating fifth forces, has been going on for decades. As a result, many of the proposed Standard Model Extensions (SME) are already highly constrained by a huge diversity of experiments. Pure antimatter systems offer a unique environment that can contribute to this search [20, 21]. The SM predicts that antimatter has exactly the same or opposite properties of matter, so by measuring its properties, the SM is tested.

Antihydrogen is the only pure atom-like antimatter system that has been trapped so far. Since trapping it in 2010, the ALPHA collaboration has pioneered measurements of the properties of antihydrogen. The most precisely measured property is the 1S–2S transition, which has been measured with a relative precision of $2 \cdot 10^{-12}$ [10]. The spectroscopy measurements of antihydrogen can be compared with those undertaken by the hydrogen spectroscopy group at the Max

Planck Institute, who have measured the 1S–2S transition with a relative precision of $4 \cdot 10^{-15}$ [22, 23].

Besides being a good candidate for matter-antimatter comparisons, the 1S–2S transition is also a good probe of fundamental constants as discussed in section 2.6. For the spectroscopy measurement of hydrogen, the values used for the fine structure constant and the proton to electron mass ratio are taken from other experiments, and two hydrogen transitions are measured to determine the Rydberg constant and the proton charge radius [24]. Measurements of the 1S–2S transition in antihydrogen have the potential to improve these determinations.

SMEs that affect the properties of hydrogen and/or antihydrogen are reviewed in [21]. Examples of such theories include a fifth force interacting with the baryon minus lepton (B-L) charge, or a *gravivector* boson, which appears in some supersymmetry (SUSY) theories. The gravivector would act in addition to the normal *graviton*, but as opposed to the graviton, the gravivector would not couple to gluons. As quarks only contribute with about 1% to the proton mass, the difference in acceleration in a gravitational field between the proton and antiproton due to the gravivector would be on the 1% level. Efforts at CERN to do gravity measurements of antimatter are summarised in section 2.3. Such measurement test the free fall weak equivalence principle (ffWEP) and probe for BSM physics.

2.2.1 Antigravity

As a consequence of CPT symmetry, two antimatter bodies must experience the same gravitational force as the force experienced by the two equivalent matter bodies, but CPT symmetry does not predict the force between a matter and an antimatter body. One might imagine that they repel each other, so the acceleration experienced by the antimatter body in a field created by a matter body, $\bar{g}$, is opposite to the acceleration between two matter bodies, g; that is, $\bar{g} = -g$. Various arguments against such "anti-gravity" are reviewed by Nieto and Goldman [25]. The perhaps most famous of these arguments, namely Morrison's *gedankenexperiment* from 1958 [26], is repeated here.

Imagine an electron-positron pair that are raised from a height h_1 to h_2 in a gravitational field generated by a matter body. Assuming $\bar{g} = -g$, this will not cost any energy. At h_2, they annihilate to produce two photons emitted in opposite directions. By using heavy mirrors, these photons could be made to meet up at h_1 and recreate the electron-positron pair. As the photons travel downwards through the gravitational field, they get blue-shifted. Hence, the recreated electron-positron pair would have more energy than the original pair. Thus to conserve energy, $\bar{g}$ must be greater than $-g$, so it will cost energy to move the electron-positron pair upwards.

2.3 The CERN Antimatter Community

The antimatter community at CERN today consists of the AEgIS, ALPHA, ASACUSA, BASE, GBAR, and PUMA collaborations, which are at different stages of construction/operation. In short, the declared goals of the experiments are the following:

- AEgIS: gravity measurement of antihydrogen [27, 28].
- ALPHA: spectroscopy, gravity and other measurements of antihydrogen [3, 4, 1, 10].
- ASACUSA: spectroscopy of antiprotonic helium, pionic helium and antihydrogen [29, 30, 31].
- BASE: comparing properties of protons and antiprotons [32, 33].
- GBAR: gravity measurement of antihydrogen [34, 35].
- PUMA: transport antiprotons in a portable trap to the ISOLDE facility for studies of nuclei [36].

Although some of the goals are shared between different collaborations, the approaches differ. For example, while spectroscopy of antihydrogen in ALPHA is done on antihydrogen in a magnetic trap, the ASACUSA experiment plans to do spectroscopy on an antihydrogen beam in a field-free region. The methods of antihydrogen production differ as well. In ALPHA and ASACUSA, the antihydrogen is made by mixing antiproton and positron plasmas, while AEgIS and GBAR intend to create antihydrogen my mixing antiprotons and positronium.

Regarding gravity measurements, GBAR plan to produce $\bar{\mathrm{H}}^+$ ions, which could be sympathetically cooled with laser cooled Be^+ ions to obtain antihydrogen ions colder than 10 μK, and then photo-detach the extra positron with a laser to do a free fall measurement of antihydrogen [37, 38]. AEgIS intend to apply an electric field gradient across the created antihydrogen and thereby use Stark acceleration to create an antihydrogen beam directed towards a Moiré deflectometer, which can be used to measure the predicted 10–20 μm free fall of antihydrogen during the time of flight. Finally, in ALPHA we plan to measure gravity by letting cold antihydrogen escape from a magnetic trap as described in chapter 6.

2.4 Charged Particles in Penning Traps

Penning traps are a crucial tool to low energy particle physics and a key part of the ALPHA experiment. The behaviour of charged particles in Penning traps is described thoroughly in the literature – for example [39]. The equations in this section describe the motion and behaviour of charged particles in Penning traps to the extend that is relevant for this thesis.

A classical penning trap consists of two endcap electrodes, which are electrically connected, a ring electrode, and an external solenoid, which immerses the electrodes in an external magnetic field, B (see figure 2.1). A voltage, V_0, between the ring and the endcap electrodes, gives the electrostatic potential

$$\phi(x, y, z) = \frac{V_0}{R^2}(2z^2 - x^2 - y^2) \tag{2.9}$$

where $R^2 = \rho_0^2 + 2z_0^2$ is determined by the geometry of the trap (see figure 2.1). The motion of a particle with charge q and mass m in this potential consists of a radial and an axial motion. The axial motion is a simple harmonic oscillator with axial frequency[1]

$$\omega_z = \sqrt{\frac{4qV_0}{mR^2}} \tag{2.10}$$

The radial motion is a superposition of two circular motions. Defining the cyclotron frequency as

$$\omega_c = \frac{qB}{m} \tag{2.11}$$

[1]Also referred to as the bounce frequency.

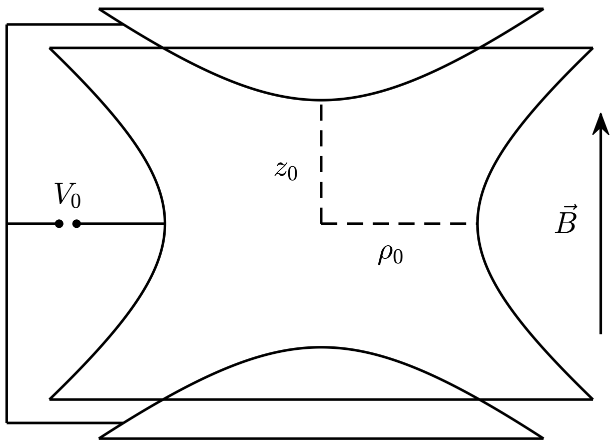

Figure 2.1: A classical Penning trap. A voltage between the ring and end cap electrodes provides axial confinement of charged particles, while a solenoid (not shown) produces an axial magnetic field to provide radial confinement.

the frequencies of the two circular motions can be written as

$$\omega_c' = \frac{\omega_c}{2} + \sqrt{\frac{\omega_c^2}{4} - \frac{\omega_z^2}{2}} \tag{2.12}$$

$$\omega_m = \frac{\omega_c}{2} - \sqrt{\frac{\omega_c^2}{4} - \frac{\omega_z^2}{2}} \tag{2.13}$$

called the modified cyclotron motion and the magnetron motion respectively. This also yields the useful relation

$$\omega_c = \omega_m + \omega_c' \tag{2.14}$$

Note that equations 2.10–2.14 only apply to a *perfect* Penning trap with a potential given by equation 2.9. However, such Penning traps are inconvenient for experimental purposes, as access for particles, lasers, etc. is limited, so cylindrical Penning traps[2] are more common. For cylindrical electrodes, the equations 2.10–2.12 are only approximately true, and equation 2.14 is modified to give the Brown-Gabrielse theorem

[2] Cylindrial Penning traps consist of a minimum of three concentric cylindrical electrodes at different axial positions.

Frequency	**Electron**	**Proton**
f'_c [GHz]	28	0.015
f_z [MHz]	42	0.99
f_m [kHz]	32	32

Table 2.1: Frequencies of an electron and a proton in a typical Penning trap ($B = 1$ T, $V_0/R^2 = 10^5$ V/m^2).

[40].

$$\omega_c^2 = \omega_m^2 + \omega_c'^2 + \omega_z^2 \tag{2.15}$$

Values for the frequencies for an electron and a positron in a typical Penning trap ($B = 1$ T, $V_0/R^2 = 10^5$ V/m^2) are given in table 2.1.

ECR Magnetometry

From equation 2.11 it is clear that the magnetic field strength in the Penning trap can be measured by measuring the cyclotron frequency of trapped particles. The technique is known as Electron Cyclotron Resonance (ECR), as it is often done with electrons, although the principle holds for any charged particle. One way to measure the cyclotron resonance is by using microwaves [41]. An electron plasma will be heated by the microwaves, if the wave frequency is $\omega = \omega_c' + n\omega_z$, where $n \in \mathbb{Z}$. Under the assumption that $\omega_c' = \omega_c$, the magnetic field strength can be derived by determining the frequency of the $n = 0$ resonance.

By varying the trapping potential and thereby ω_z, the axial frequency sideband peaks will move, and one can identify the stationary $n = 0$ resonance. The central peak has a substructure with rotational sidebands, but the substructure of the central peak is often not relevant for the desired precision of the field strength[3]. For an electron in a typical Penning trap (like the one in table 2.1), the error introduced by assuming $\omega_c' = \omega_c$ is around 32 kHz, which translates to a error of 1.1 μT.

[3]The rotational substructure is on the order of 10–100 kHz [41], which translates to a ΔB of about 0.4–4.0 μT at 1 T.

Cyclotron Cooling

A non-relativistic charged particle that is accelerated will radiate energy as described by the Larmor formula [42]

$$P = \frac{q^2 a^2}{6\pi\epsilon_0 c^3} \tag{2.16}$$

where P is the power radiated, q is the particle charge, and a is the acceleration. As a charged particle rotates around magnetic field lines, it is constantly accelerated and will thus lose energy. This type of cooling is referred to as cyclotron cooling [43, 44]. For a charged particle following a circular motion with the cyclotron frequency, the acceleration is

$$a = \omega^2 r = \frac{qB}{m} v_\perp \tag{2.17}$$

Considering a particle in a non-neutral plasma with a high collision rate[4], so the energy follows the equipartition theorem, the perpendicular energy can be written as

$$E_\perp = \frac{1}{2} m v_\perp^2 = \frac{2}{3} E \tag{2.18}$$

and one obtains

$$P = \frac{q^2}{6\pi\epsilon_0 c^3} \left(\frac{qB}{m}\right)^2 \frac{2E_\perp}{m} \Rightarrow \tag{2.19}$$

$$\frac{dE}{dt} = \frac{3}{2}\frac{dE_\perp}{dt} = -P = -\frac{q^4 B^2}{3\pi\epsilon_0 c^3 m^3} E_\perp \Rightarrow \tag{2.20}$$

$$E_\perp(t) = E_\perp(0) e^{-t/\tau_c} \tag{2.21}$$

where the cooling time is given as

$$\tau_c = \frac{3}{2}\frac{3\pi\epsilon_0 c^3 m^3}{q^4 B^2} \tag{2.22}$$

The parallel component of the energy only cools through mixing between the energy components, so weak energy mixing can limit the plasmas temperature cooling rate. Equation 2.22 shows that the cooling rate strongly depends on the charge, the mass, and the magnetic field strength. As a consequence, cyclotron cooling is effective for electrons and positrons (τ_c = 3.88 s at 1 T), but not feasible for

[4] If we also assume that the plasma follows a Maxwell-Boltzmann distribution, we can use energy and temperature interchangeably.

antiprotons ($\tau_c = 2 \cdot 10^{10}$ s at 1 T). However, by mixing antiprotons with electrons, they will thermally equilibrate (sympathetic cooling), while the electrons are cooled via cyclotron cooling.

Rotating Wall Compression

In a real-life Penning trap, background gas and imperfect fields will cause a rotating plasma to slowly expand and eventually hit the trap walls. To keep the charged particle plasmas compressed, we use the so-called rotating wall (RW) technique [45, 46], which applies a time-dependent transverse electric field to drive the rotation of the plasma. As the rotation accelerates, the plasma will compress due to conservation of angular momentum. The electric field, oscillating with frequency f_{RW}, is generated by an electrode azimuthally divided into N segments, where the n'th segment has a potential given as

$$V_n(t) = A \cdot \cos\left(m\left(2\pi f_{RW} t \pm \frac{2\pi n}{N}\right)\right) \tag{2.23}$$

where A is the voltage amplitude, and m is the drive mode ($m = 1$ is the dipole mode, $m = 2$ is the quadropole, etc.). The amplitude of the RW has proven to define two different regimes of plasma behaviour [47]: at low amplitude, the compression happens at discrete frequencies, whereas the compression follows a more linear response for higher amplitudes. The latter is referred to as the strong drive regime (SDR), and together with evaporative cooling (EVC) [48], it allows consistent production of plasmas in the ALPHA experiment [49, 50].

Sympathetic Cooling

If two plasmas with different temperatures are mixed, they will equilibrate via either Coulomb (elastic) collisions if the plasmas are charged, or elastic collisions if the plasmas are neutral. This process is called sympathetic cooling. The technique is used to achieve lower temperatures than would otherwise be possible for a given plasma by mixing it with a plasma that can be cooled even further. In an elastic collision between two electrically neutral spheres with diameter d, the velocity of particle 1 after the collision, $\vec{v}_1'$ is given as

$$\vec{v}_1' = \vec{v}_1 - \frac{2m_2}{m_1 + m_2} \frac{(\vec{v}_1 - \vec{v}_2) \cdot \vec{d}}{||\vec{d}||^2} \vec{d} \tag{2.24}$$

where m_1 and m_2 are the masses of the two particles, $\vec{v}_1$ and $\vec{v}_2$ are the velocities before the collision, and $\vec{d}$ is the distance between the centres at the time of the collision. The cooling is seen to be most effective, when $m_1 = m_2$, which is often not achievable, so m_2 is chosen to get as close to m_1 as possible. Be^+ ions can be used to cool light, positive particles, as Be^+ can be laser cooled, and electrons are often used to cool light, negative particles, as electrons can be cyclotron cooled (see chapter 4 and 5).

2.5 Trapping Antihydrogen

As antihydrogen is electrically neutral, it cannot be confined in Penning traps. Instead, it can be confined via its magnetic dipole moment, $\vec{\mu}$, generated by its angular momentum, which interacts with magnetic fields. The potential seen by a magnetic dipole in a magnetic field and the resulting force is given as

$$U = -\vec{\mu} \cdot \vec{B} \Rightarrow \tag{2.25}$$

$$\vec{F} = \nabla(\vec{\mu} \cdot \vec{B}) = \pm\mu\nabla B \tag{2.26}$$

where we have used that the quantised magnetic dipole moment aligns itself either parallel or antiparallel to the magnetic field. From equation 2.25 it is seen that only atoms/antiatoms with a magnetic moment pointing in the opposite direction of the magnet field will be confined in a magnetic minimum configuration. These atoms are said to be in a low-field-seeking state as opposed to untrappable atoms in a high-field-seeking state.

Ignoring the contribution to the magnetic moment from the proton[5], the total magnetic moment of (anti)hydrogen is given as [51]

$$\vec{\mu} = \vec{\mu}_s + \vec{\mu}_l = \mp\frac{\mu_B}{\hbar}\left(g_s\vec{S} + g_l\vec{L}\right) \tag{2.27}$$

where $g_l = 1$ and $g_s \approx 2$ are the electron g-factors, $\mu_B = e\hbar/2m_e$ is the Bohr magneton, and the $\mp$ is for matter/antimatter. We cover the equations for antimatter, so we will keep the + sign. In high magnetic fields, where $\vec{S}$ and $\vec{L}$ are decoupled, and the wavefunction[6] is given

[5] The magnetic moment of the proton is $\mu_p = 2.79\mu_N$, where $\mu_N/\mu_B = m_e/m_p$, so its contribution is suppressed by more than three orders of magnitude.

[6] The notation $\vec{S}\,|s,l,m_s,m_l\rangle = m_s\hbar\,|s,l,m_s,m_l\rangle$, $\vec{L}\,|s,l,m_s,m_l\rangle = m_l\hbar\,|s,l,m_s,m_l\rangle$, and $\vec{J}\,|s,l,j,m_j\rangle = m_j\hbar\,|s,l,j,m_j\rangle$ is used.

by $|\psi\rangle = |s, l, m_s, m_l\rangle$, the magnetic potential energy is given by

$$E = -\frac{\mu_B}{\hbar} \langle\psi| g_s\vec{S} + g_l\vec{L} |\psi\rangle B \tag{2.28}$$

$$= -\mu_B (g_s m_s + g_l m_l) B \tag{2.29}$$

In a low magnetic field, where $\vec{L}$ and $\vec{S}$ couple to the total angular momentum, $\vec{J}$, and the wavefunction is $|\psi\rangle = |s, l, j, m_j\rangle$, the energy is given by

$$E = -\frac{\mu_B}{\hbar} \langle\psi| g_s\vec{S} + g_l\vec{L} |\psi\rangle B \tag{2.30}$$

$$= -\frac{\mu_B B}{\hbar} \left(g_s \left\langle \psi \left| \frac{\vec{S}\cdot\vec{J}}{J^2}\vec{J} \right| \psi \right\rangle + g_l \left\langle \psi \left| \frac{\vec{L}\cdot\vec{J}}{J^2}\vec{J} \right| \psi \right\rangle \right) \tag{2.31}$$

$$= -\mu_B g_j m_j B \tag{2.32}$$

where g_j is the Landé g-factor.

$$g_j = g_l \frac{j(j+1) + l(l+1) - s(s+1)}{2j(j+1)} + g_s \frac{j(j+1) + s(s+1) - l(l+1)}{2j(j+1)} \tag{2.33}$$

For an antihydrogen atom in the ground state, equations 2.29 and 2.32 are equal, and they give a trap depth of 58 $\mu eV/T$ for low-field-seeking atoms. The low-field-seeking states for antihydrogen are the ones with positron spin antiparallel to $\vec{B}$ (spin down), while for ground state hydrogen, it is the states with electron spin parallel to $\vec{B}$ (spin up). Using Boltzmann's constant to convert the trap depth to the commonly used unit of K/T, one gets 0.67 K/T. In ALPHA, a magnetic well depth of 1 T is typically used, so the confined antihydrogen have velocities of less than 110 m/s.

A magnetic minimum trap (a Ioffe-Pritchard trap [52]) can be made using an octupole and two coaxial solenoids at different axial positions. The octupole provides radial confinement, as it generates a field, whose size is $|\vec{B}(\rho)| \propto \rho^3$ [53], and the solenoids provide the axial confinement of the particles. Antihydrogen trap design will be elaborated in chapter 4, and calculations of the field of solenoids is elaborated in chapter 8.

Trapped antihydrogen can in principle be *adiabatically* cooled by slowly expanding the trapping volume, which will cause the antihydrogen to lose energy in collisions with the slowly retracting potential barriers. Simulations of adiabatic expansion suggest that it can reduce the mean energy by at least 38% [54]. It is expected that adiabatic expansion will be used increasingly in both future gravity and spectroscopy experiments at ALPHA.

Simulations show that trapped antihydrogen atoms can be divided qualitatively into two categories [55, 54]: “mixing” antihydrogen, where energy is exchanged between the parallel and the perpendicular components[7] of the kinetic energy, and “no-mixing” antihydrogen, where no energy is exchanged. The mixing is mainly due to the azimuthal field asymmetries introduced by the octupole. It is possible for antihydrogen atoms to change from one category to the other. The rate with which energy is exchanged [55] is crucial for the effect of axial adiabatic expansion or axial laser cooling, as both techniques only decrease the parallel energy directly. Energy mixing is also highly relevant for gravity measurements (see section 6.2.1).

2.6 The Spectrum of (Anti)Hydrogen

The purpose of this section is to outline the spectrum of hydrogen, which is assumed to be equal to the spectrum of antihydrogen in the SM. The sensitivities to physical constants are highlighted, as spectroscopy measurements serve as tests of the QED’s (Quantum ElectroDynamics) predictions of these constants. Derivations of the hydrogen energy levels are usually done either using perturbation theory [51, 56] or from the Dirac equation. A detailed tabulation of energy contributions to the hydrogen spectrum can be found in [57].

Bohr Energy Levels

To first order, the electronic energy level in an atom with nuclear charge Z is given as

$$E_n = -\frac{m_e}{2\hbar^2}\left(\frac{Ze^2}{4\pi\epsilon_0}\right)^2\frac{1}{n^2} \tag{2.34}$$

$$= -\frac{m_e c^2}{2}\frac{Z^2}{n^2}\alpha^2 \tag{2.35}$$

$$= -R_\infty\left(hc\frac{Z^2}{n^2}\right) \tag{2.36}$$

[7] Parallel and perpendicular relative to the axis of symmetry.

where n is the principal quantum number, and α is the fine structure constant

$$\alpha = \frac{e^2}{4\pi\epsilon_0\hbar c} \approx \frac{1}{137} \tag{2.37}$$

and R_∞ is the Rydberg constant

$$R_\infty = \frac{\alpha^2 m_e c}{4\pi\hbar} \tag{2.38}$$

Higher order corrections to the energy levels are often referred to with their dependence of the fine structure constant relative to the principal energy levels. Since $E_n \propto \alpha^2$, an energy correction term that is proportional to α^n is therefore said to be sensitive to α^{n-2}.

Fine Structure

The first of the corrections to the Bohr energy levels is due to the interaction between the angular momenta of the electron. The effect can be understood intuitively in a semi-classical model, where the electron orbits the nucleus. Seen from the electron's rest frame, the charged nucleus orbits the electron and thereby generates a magnetic field, which interacts with the spin magnetic moment of the electron, and gives an addition to the Hamiltonian of the system, $H_\mu = -\vec{\mu}\cdot\vec{B}$. Hence, the spin-orbit coupling causes an energy shift. Using perturbation theory, one obtains [51]

$$E_{so} = -(g_s - 1)Z^2\frac{E_n\alpha^2}{2n}\left(\frac{j(j+1) - l(l+1) - s(s+1)}{l(l+1/2)(l+1)}\right) \tag{2.39}$$

where g_s is the gyromagnetic ratio of the electron[8]. E_{so} is seen to depend on α^2 relative to the principal energy. For hydrogen, spin-orbit coupling causes the 2S and 2P states to be split by 11 GHz at no external magnetic field, but the shift is cancelled by the other fine structure effects.

In the derivation of the principal energy levels, it is assumed that the electron is non-relativistic. A classical calculation for an electron in the ground state of hydrogen gives the velocity of the electron to

[8] g_s is often put equal to 2 (as predicted by Dirac theory) in derivations of fine structure energy corrections, as its actual value of approximately 2.002 gives a relatively small correction. I choose to keep it as g_s to make the dependence transparent.

be less than $0.01c$, so the relativistic correction is small, but it is significant compared to experimental sensitivity. The expression for the relativistic kinetic energy

$$E = \sqrt{p^2c^2 + m^2c^4} - mc^2 \tag{2.40}$$

$$= mc^2 \left(\sqrt{\frac{p^2}{m^2c^2} + 1} - 1 \right) \tag{2.41}$$

can be Taylor expanded, and the second order term can be treated as a perturbation to the Hamiltonian, which yields the following energy correction

$$E_{rel} = -\frac{E_n^2}{2mc^2} \left(\frac{4n}{Z^2(l + 1/2)} - 3 \right) \tag{2.42}$$

Furthermore, rapid fluctuations of the electron's position – *zitterbewegung* [58] – alters the electrostatic interaction between the electron and the nucleus, since the electron charge can be considered to be distributed over a larger volume. This leads to the Darwin term, which contributes to the total energy

$$E_D = 2n \frac{E_n^2}{m_e c^2} \delta(l = 0) \tag{2.43}$$

The effect only applies to electrons in the s-orbital due to the overlap of their wavefunctions with the nucleus. Both the relativistic correction and the Darwin term are of order α^2, like the spin orbit correction. Combining the above corrections yields the same result as the Dirac equation to order α^4.

According to the predictions of the Dirac equation, the $2S_{1/2}$ and $2P_{1/2}$ states should have the same energy, but in 1947 W. Lamb and R. Retherford measured a 1 GHz difference between the two energy levels [59], which would later be known as the Lamb shift. Quantum field theory is needed to describe this effect, which originates from the interaction between the electron's electrical charge and the energy fluctuations of the electromagnetic vacuum field. This only affects states with orbital angular momentum equal to zero, and the effect can be shown to depend on the fine structure constant to one order higher than the fine structure effects [60, 61]. The Lamb shift is therefore referred to as an α^3-effect. From the Lamb shift, one can derive the proton charge radius, so Lamb shift measurements [62] are a key part of the proton radius puzzle [63].

Hyperfine Structure

The magnetic moment of the nucleus affects the electronic energy levels as well. In the following, we will just consider a hydrogen atom. The proton magnetic moment is given as

$$\vec{\mu}_p = g_p \frac{\mu_N}{\hbar} \vec{I} \tag{2.44}$$

where $g_p \approx 5.59$, μ_N is the nuclear magneton, and $\vec{I}$ is the nuclear spin operator. The proton magnetic moment is orders of magnitude smaller than the electron magnetic moment, as it is suppressed by the electron/proton mass ratio ($\mu_N/\mu_B = m_e/m_p$). The proton magnetic moment generates a magnetic field, which interacts with the electron and contributes to the Hamiltonian. The resulting energy change depends on the relative orientation of the proton spin and electron angular momentum. For $l = 0$ states, the energy change for the triplet state, where the electron and proton spins add to 1, and for the singlet state, where the spins add to 0, are given as [51]

$$E_{hf} = -\frac{8}{3} g_p \frac{m_e}{m_p} E_n n^2 \alpha^2 \begin{cases} \frac{1}{4} & \text{triplet state} \\ -\frac{3}{4} & \text{singlet state} \end{cases} \tag{2.45}$$

Although the hyperfine energy shift depends on α^2 like the fine structure terms, it is suppressed by the electron to proton mass ratio, and the energy shift is therefore about 1.4 GHz for the ground state, which is similar in size to the Lamb shift.

All the above expressions for the hydrogen energy levels apply to a zero external magnetic field. An external magnetic field will interact with the magnetic moment of the quantum states, as described in section 2.5, so the energy levels will change[9]. Energy diagrams for the relevant energy levels in antihydrogen are shown in chapter 5.

Transitions

To this date, the measured transitions in antihydrogen have been between the $1S_{1/2}$, $2S_{1/2}$, $2P_{1/2}$, and $2P_{3/2}$ states. The transition between the $1S_{1/2}$ and $2S_{1/2}$ state is forbidden, so the $2S_{1/2}$ state is metastable with a lifetime of about 1/8 s. In comparison, the lifetime of the 2P state is about 1.6 ns. To excite the forbidden $1S_{1/2}$–$2S_{1/2}$

[9]The Zeeman effect for weak external fields, and the Paschen-Back effect for strong external fields.

transition, a two-photon transition is used. Since antihydrogen travels at a finite speed in the magnetic traps of ALPHA, it experiences an electric field, if it travels perpendicularly to the magnetic field. This causes the $2S_{1/2}$ and the $2P_{1/2}$ state to mix, which allows the $2S_{1/2}$ state to decay to the groundstate via the $2P_{1/2}$ state [64].

Chapter 3

The AD and ELENA

At CERN, antiprotons are produced by making bunches of high energy protons collide with a solid irridium target [65, 66]. Some of the created antiprotons are directed into a beam, which is steered towards the decelerators and the experiments in the Antimatter Factory at CERN. The decelerators available are the Antiproton Decelerator (AD), and the Extra Low ENergy Antiproton ring (ELENA).

To get the protons to the required energy, a linear accelerator first accelerates the them, before they are sent to the Booster and then to the Proton Synchrotron, where the protons reach an energy of 26 GeV. Colliding the proton beam with an irridium target results in inclusive production of antiprotons with an energy of a few GeV [66]. The antiprotons directed into a beam are decelerated by the AD to 5.3 MeV. Prior to LS2, the antiprotons would then be delivered to the experiments, but they are now directed to ELENA, which decelerates them to 100 keV before delivering them to the experiments.

3.1 The AD

The AD [65, 67] came into operation in 2000. It is a circular decelerator, which cools[1] the antiprotons and decelerates them to 5.3 MeV. The deceleration is done with a radio frequency cavity, while the cooling is done by bunch rotation, stochastic cooling, and electron cooling. Typically, more than $5 \cdot 10^7$ antiprotons are received by the AD, but as some are lost during the deceleration and cooling process, only $3 \cdot 10^7$ are left after reaching 5.3 MeV [66, 68]. The cooling cycle takes around

[1]Cooling is this context is used to described a decrease in the energy spread and transverse emittance of the antiproton bunch.

100 s. Prior to LS2, when the AD delivered antiprotons directly to the experiments, it was only able to direct antiprotons to one experiment at a time. The antiprotons were distributed over three eight-hour shifts per day.

3.2 ELENA

As 5 MeV antiprotons are too energetic to be trapped directly by the experiments at the Antimatter Facility, the experiments are forced to further reduce the antiproton energy, before they can be caught (see chapter 4). In ALPHA, the related antiproton loss was as high as 99.9%. To reduce the energy gap between the delivered and the trappable antiproton energy, ELENA was constructed to decelerate the antiprotons from the AD [69], before sending them to the experiments. ELENA is about 30 m in circumference and placed within the AD hall. It uses radio frequency and electron cooling to decelerate and cool the antiproton bunches from 5.3 MeV to 100 keV over a period of about 20 s. About 40% of the antiprotons are lost during the process. The remaining 18 million antiprotons are split in four bunches of around 4.5 million, which can be directed to four different experiments at a time [68].

Chapter 4

The ALPHA-2 Apparatus

The ALPHA-2 atom trap and the associated apparatus such as the catching trap and the positron accumulator are described in this chapter. The ALPHA-2 apparatus has been described in great detail in the literature [70, 71, 9], and this chapter is only meant to serve as an overview of the major components and the techniques used. The purpose of the chapter is to lay the foundation for explaining the results we obtained in ALPHA-2 in chapter 5.

As the results obtained in ALPHA-2 were largely from before the construction of ALPHA-g, and as ALPHA-g and its associated beamline are of little relevance to the results, a description of ALPHA-g and the beamline will be postponed to chapter 6. For the present chapter, the experimental setup will largely be described as it was in 2018.

4.1 Motivation

When founded in 2005, the purpose of the ALPHA collaboration was to create and trap antihydrogen to allow measurements of its properties and thereby test CPT symmetry. Being largely a continuation of the ATHENA collaboration [72], ALPHA succeeded in trapping antihydrogen for the first time in history in 2010 [6]. This was done in the apparatus referred to as ALPHA-1 [70]. Two years later, the first transitions (hyperfine transitions) ever detected in antihydrogen were measured [73], and the collaboration later set a limit to the electric charge of antihydrogen [74, 7] and developed a technique to measure the gravitational mass of antihydrogen [75].

ALPHA-1 was upgraded to ALPHA-2 in 2012 with the aim of measuring laser-induced transitions between principal energy levels. This of course required the addition of lasers and laser access to the trapping region. The first transition measured was the 1S–2S transition [71, 10], and later the 1S–2P [9]. It is also in ALPHA-2 that techniques for accumulating large samples of antihydrogen atoms have been demonstrated [76]. To this day, there has not been any disagreement between the properties measured in antihydrogen and the ones measured in hydrogen.

Although the transitions measured in antihydrogen so far agree with the corresponding transitions in hydrogen, a difference in the energy levels could show in other transitions, or the difference could be smaller than the current level of uncertainty [21], so there is good reason to continue the spectroscopy measurement campaign. Furthermore, it is feasible that the accuracy of the measurements of antihydrogen could exceed the ones of hydrogen. This could improve the QED test currently based on hydrogen, as well as improving the direct comparisons of matter and antimatter systems and thereby test CPT symmetry.

4.2 Setup

As of 2018, the ALPHA experiment consisted of three main components: the catching trap, the ALPHA-2 atom trap, and the positron accumulator. A sketch of the apparatus can be seen in figure 4.1. The catching trap to the left is where the antiprotons from the AD (later ELENA) are trapped and cooled, before they get transferred to the atom trap. To the right is a Na-22 source, which emits positrons that are trapped and accumulated in the positron accumulator. When enough positrons have been accumulated, they are transferred to the atom trap, where the positron and antiproton plasmas are mixed, and antihydrogen atoms form. Microwaves and lasers can be sent into the atom trap to manipulate the particles and perform spectroscopy.

4.2.1 Positrons

Via β^+ decays, the Na-22 source provides positrons with energies in the 100 keV range[1]. The positrons emitted in the direction of the accumulator will pass through a solid neon moderator, which slows the

[1] Data obtained from BetaShape [77] via IAEA.

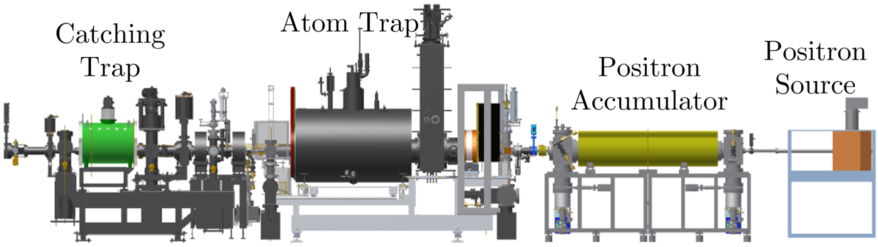

Figure 4.1: Sketch of the ALPHA apparatus prior to the installation of ALPHA-g and the beamline. To the left is the catching trap, which traps the antiprotons received from the AD. To the right is the Na-22 positron source and the positron accumulator, and in the middle is the atom trap, where antihydrogen is created.

positrons via collisions. Most of them will annihilate on the moderator, but the about 1% that survive have energies around 80 eV. These positrons are guided into the accumulator, which is a three stage Penning trap called a Surko trap [78].

The axial trapping potential of the positron accumulator, Φ, which is seen in figure 4.2, is generated by 7 electrodes. An external solenoid provides the radial confinement. The trapping volume is filled with nitrogen gas via an inlet near E2, and the different diameters of the electrodes, together with a cryopump at each end of the trap, ensures a pressure gradient across the electrode region. The nitrogen causes the positrons to slow down via collisions and make them accumulate at the left end of the trap. The fourth electrode is split into six segments, and is used to apply RW compression of the plasma, as described in section 2.4. When enough positrons have accumulated, the axial blocking potential is lowered, and the trapping potential raised, so the positrons can be transferred to the atom trap. Normally, a few million positrons are transferred successfully.

4.2.2 The Catching Trap

The purpose of the catching trap is to trap the antiproton bunches received from the AD. It is a Penning-Malmberg trap with 20 electrodes and a 3 T external solenoid. The bunches are received at 5.3 MeV, but to be trappable, their energy must be less than the 5 keV potential generated by the electrodes. To decelerate the antiprotons, the bunch is shot through a 0.2 mm beryllium and aluminium degrading foil. About 0.1% of the antiprotons survive the interaction and become

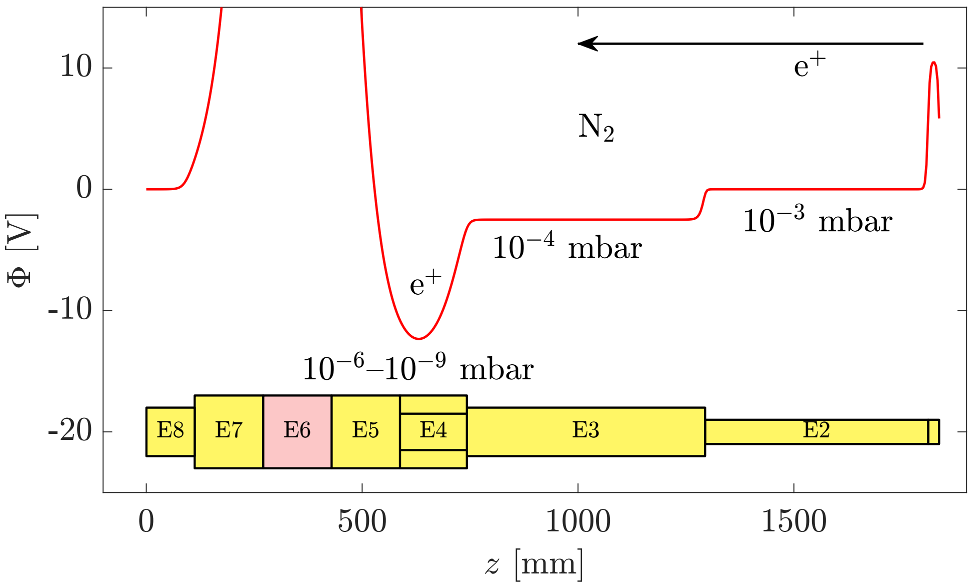

Figure 4.2: The axial trapping potential, Φ, and the electrode stack in the positrons accumulator. Positrons enter from the Na-22 source on the right. Due to collisions with the nitrogen buffer gas, they lose energy and accumulate at the bottom of the trap. The electrode stack is illustrated at the bottom of the figure. The radial dimensions are not to scale. E4 is segmented to allow RW compression, and E6 is a high voltage electrode that supplies the blocking potential.

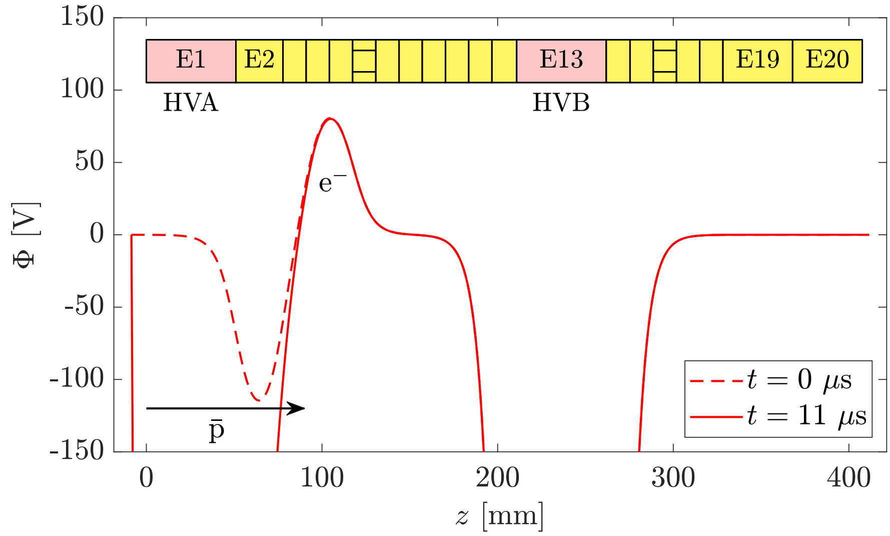

Figure 4.3: The axial trapping potential, Φ, and the electrode stack in the catching trap. An antiproton bunch enters from the left. Once the bunch is inside the trap, the high voltage electrode E1/HVA ramp up to trap the antiprotons, which sympathetically cool with the preloaded electrons. Note that the positive antiprotons are attracted by higher potentials.

trappable. When the bunch is inside the trap, the axial high voltage potential on the entry side is raised to confine the particles, as sketched in figure 4.3.

The antiprotons are too heavy to be cyclotron cooled on any relevant timescale (see section 2.4), so they are mixed with electrons, with which they sympathetically cool. The electrons can be injected into the trap from a barium-oxide filament from the right. About 85 million electrons are normally loaded in the catching trap prior to receiving the antiprotons [79]. Segmented electrodes are installed to compress or expand the plasmas using RW compression.

After cooling the antiprotons, the electrons are extracted by lowering the axial potential barrier for about 100 ns, before the potential is restored. This allows the light electrons to escape, while the heavier and slower antiprotons remain confined. The process is repeated a few times with intermediate RW compression, and the duration and amplitude of each "kick" is carefully tuned to minimise heating of the antiprotons from the rapid potential changes. After extracting

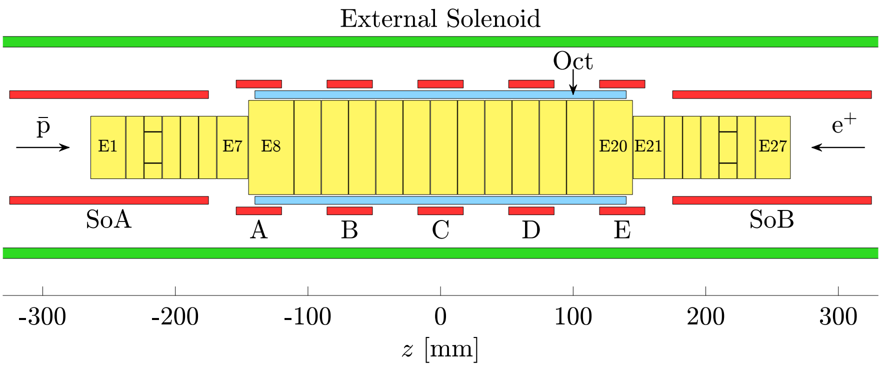

Figure 4.4: Sketch of the electrode and magnet layout of the ALPHA-2 atom trap. Antiprotons and positrons are trapped at each end, before they are mixed in the central region, where the octupole and mirrors A to E can trap the created antihydrogen. The external solenoid extends outside the figure. Electrodes 3 and 25 are segmented to allow RW compression.

the electrons, the antiprotons are transferred to the atom trap by lowering the end potential of the catching trap. As of 2018 up to 10^5 antiprotons would be successfully transferred to the atom trap and used for antihydrogen production [4].

4.2.3 The Atom Trap

The atom trap consists of a Penning-Malmberg trap for charged particles embedded into a Ioffe-Pritchard trap for neutral particles [52]. A sketch of the magnet and electrode layout can be seen in figure 4.4. The Penning-Malmberg trap is generated by 34 electrodes and a 1 T external solenoid. Its purpose is to trap and manipulate the charged particle plasmas. Two internal solenoids (SoA and SoB) can add to the magnetic field to increase cyclotron cooling. The Ioffe-Pritchard trap is generated by an octupole and five short solenoids/mirror coils. It is able to trap the electrically neutral antihydrogen by interacting with its dipole moment (see section 2.5), and it generates a bathtub-shaped potential as seen in figure 4.5.

To make antihydrogen, the positrons and the antiprotons are mixed in the central region between electrode E8 to E20, where the octupole and the mirror coils can confine the created antihydrogen. As the positrons and antiprotons have opposite charges, they are confined in

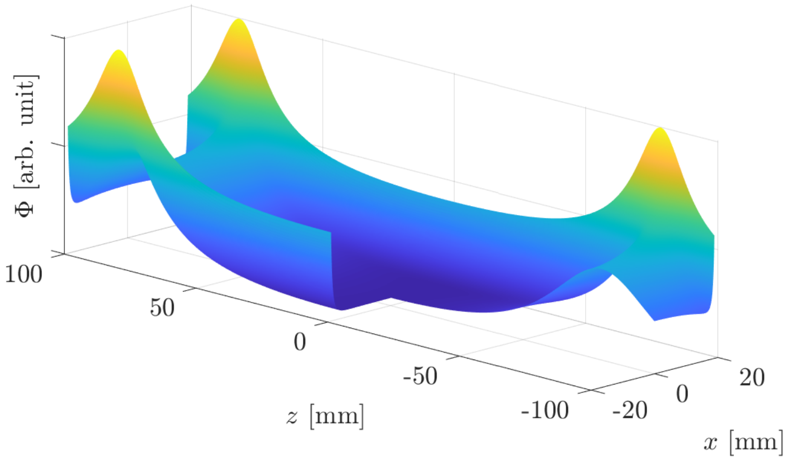

Figure 4.5: Illustration of the potential, Φ, seen by low-field-seeking antihydrogen atoms. The radial confinement is provided by an octupole, and the axial confinement is provided by two mirror coils. A section of the potential has been cut out to improve readability.

a nested potential like the one seen in figure 4.6. To bring the particles together, the potentials are slowly made shallower, so the plasmas will overlap. In this way, the most energetic antiprotons will overlap with the most energetic positrons. We typically allow the process to go on for about 1 second.

The formation of antihydrogen is typically a three-body process, where one positron binds to the antiproton, and one positron takes away the excess energy: $\bar{p} + e^{+} + e^{+} \rightarrow \bar{H} + e^{+}$. Only antihydrogen created with an energy smaller than the trap depth, which is about 50 μeV or 0.5 K, are trapped. By mixing about 10^6 antiprotons with a few million positrons, up to 50000 antihydrogen atoms are created, of which up to 30 are trapped [4]. Previously, ALPHA has demonstrated how "stacking" antihydrogen over a few mixing cycles can be used to trap 10s of antihydrogen atoms [76]. The number can be increased to more than 1000 by accumulating over more cycles (see section 5.2). The antihydrogen lifetime in ALPHA-2 has been demonstrated to be at least 66 hours [80], so effectively it does not pose a problem to antihydrogen accumulation.

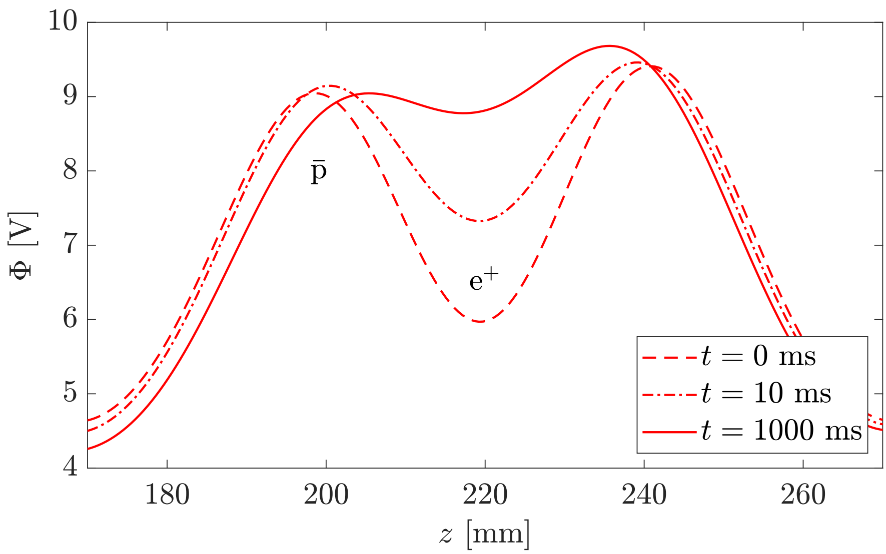

Figure 4.6: The axial potential, Φ, used to confine the positive positrons and the oppositely charged antiprotons in the atom trap. The positron trapping potential is raised to slowly merge the positrons into the antiprotons. The self-potential of the positrons is not included in the figure, but it adds to the total field and flattens the potential region around the particles.

4.2.4 The Detector

The ALPHA-2 atom trap is surrounded by a three-layer silicon vertex detector. The detector consists of multiple panels of silicon strips, which cross and overlap to create a fine grid. When an antihydrogen atom annihilates on the trap wall, the positron will produce two gamma rays, and the antiproton will typically create three to five pions (both neutral and charged pions) [72, 81]. When a charged pion passes through each of the silicon layers, it deposits energy that excites electron-hole pairs, which is detected as an electric signal. The signals allow for reconstruction of the trajectory of the pion in the magnetic field and reconstruction of the annihilation vertex [82].

As the detector has low sensitivity to gamma rays, we cannot rely on the detector to distinguish between annihilating antiprotons and antihydrogen atoms. However, this is not necessary, as all antiprotons can be ejected from the trap by ramping down the electric field after mixing and applying an electrical *sweep*. ALPHA has previously proved that quasi trapped antiprotons are not an issue [6]. The antihydrogen will remain trapped and can be ejected later.

Cosmic rays can mimic annihilation events in the detector and are the main source of background signal. The cosmic rays can be filtered out, as their temporal and spatial signatures are very different from those of antiproton annihilations. The filtering is done by a background rejection algorithm, which is created using supervised machine learning based on boosted and bagged decision trees [4, 83]. The amount of training data varies between different measurements, but in general, more than 100000 annihilation events containing about 1% percent background, and more than 100000 pure background events are used. The rate of false positives is tuned between experiments, but it can often be kept below the 0.1% level, with a rate of $1.0 \cdot 10^{-3}$ s^{-1} [10].

4.2.5 The Beryllium Source

To the left of the ALPHA-2 atom trap (see figure 4.1) is a beryllium source, which makes positive beryllium ions available in the atom trap. Positive beryllium ions are interesting, as they can be laser cooled and used to sympathetically cool positrons [84, 85]. The beryllium source is a small sheet of beryllium metal, from which Be^+ ions are ablated

Purpose	**Wavelength/frequency**
$\bar{H}$ 1S–2P transition	121 nm
$\bar{H}$ 1S–2S two-photon transition	243 nm
Beryllium laser cooling	313 nm
Beryllium ablation	355 nm
$\bar{H}$ hyperfine transitions, ECR	15–31 GHz

Table 4.1: A list of the microwave frequencies and laser wavelengths available in the ALPHA-2 experiment and their purpose.

by shining a 355 nm pulsed laser[2] on it. The ions ejected towards the atom trap can be caught by a trapping potential, where the end barrier is quickly lowered to allow the ions to enter, and then raised to trap them. A single bunch typically contain 10^6 ions. In the atom trap, a 313 nm laser is cooling the ions.

4.2.6 Lasers and Microwaves

It is possible to irradiate the trapped antihydrogen with both lasers and microwaves in ALPHA-2. Laser access is limited by the geometry of the trap, so the lasers and microwaves are injected from far outside the trapping region as sketched in figure 4.7. The lasers available for spectroscopy are the 121 and 243 nm lasers, and the lasers related to the beryllium cooling and production are 313 and 355 nm. Microwaves are used to induce hyperfine transitions and ECR measurements. A list of the frequencies available is seen in table 4.1.

The 243 nm laser used for the 1S–2S transition is generated by frequency doubling a 972 nm laser twice [71]. The 121 nm laser is a pulsed laser, which is generated by frequency doubling a 730 nm pulsed laser, and then generating third harmonics in a Kr-Ar gas cell [9]. The 243 nm laser is referenced to a frequency comb to determine the frequency. The frequency comb is referenced to a quartz oscillator, which is steered by a GPS reference. Post LS2, the laser system has received various upgrades, which will be described in chapter 9.

4.2.7 Plasma Diagnostics

Throughout the experiment, various diagnostic tools are available. Some are mounted on "sticks", which can be translated transversely

[2] Pulse duration is 6.3 ns, pulse energy is about 75 μJ, and fluence is about 3 J/cm^2.

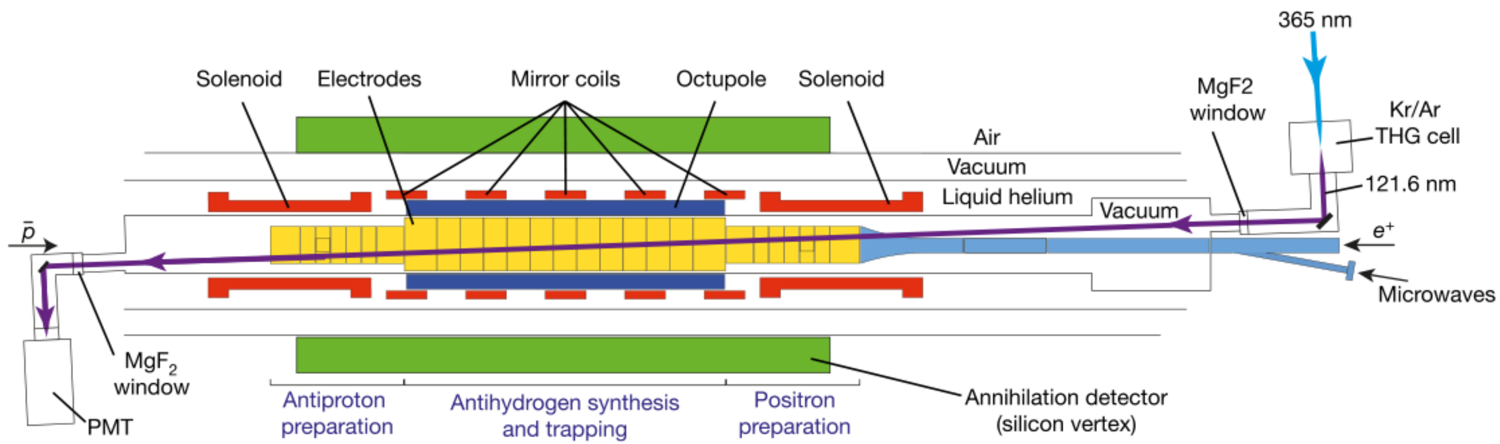

Figure 4.7: A sketch of the ALPHA-2 atom trap. Antiprotons enter from the left, and positrons from the right. Indicated on the sketch are the electrodes, mirror coils, octupole, solenoids, and the detector. The 121 nm laser used for the 1S–2P transitions is sketched coming in at a small angle compared to the axis of symmetry. The microwave injection is seen on the right as well. From [3].

to the beam path, so different tools can be inserted in the beam. In this way, Faraday cups and MicroChannel Plate (MCP)/phosphor detectors are available to diagnose the plasmas (destructively). A Faraday cup is a piece of conducting metal, which will get a small charge, when charged particles/ions hit it. This charge will discharge and cause a small current, which is proportional to the number of charges hitting the metal. The Faraday cups in ALPHA are sensitive to particle numbers greater than a few 10^5s.

The MCP/phosphor detectors consists of an MCP, a phosphor screen and a CCD camera [86]. When a charged particle hits one of the channels in the MCP, the channel works as a electron multiplier. The electrons exit the back of the MCP and strike a phosphor screen, which will emit photons that are imaged by the CCD camera. The MCP detectors are used to determine the radial profile of the plasmas. The Faraday cup can be used to calibrate the MCP, so the MCP can be used to determine the number of charged particles as well. In ALPHA, the MCPs have single particle sensitivity [87]. Furthermore, to detect antiproton annihilations, we use external scintillator panels, which are panels of scintillating material paired with photo multiplier tubes (PMTs) or Silicon PhotoMultipliers (SiPMs).

Another application of the MCP/phosphor detector is to determine the temperature of a plasma [87]. This is done by slowly lowering the potential confining the plasma to allow the most energetic particles to escape and hit the MCP. By counting the number of particles detected as a function of the trap depth, the energy distribution of the plasma is determined. An alternative to imaging the photons from the phosphor screen with a CCD camera, is to use a SiPM [49]. As the SiPM is less sensitive to noise, it can diagnose smaller and colder plasmas. Temperature measurements of plasmas as cold as 10 K and as small as 300 electrons have been demonstrated [87].

Chapter 5

ALPHA-2 Results

This chapter describes the results obtained with the ALPHA-2 apparatus just before and during LS2. These represent an improvement of the existing measurement of the fine structure in antihydrogen, a demonstration of the first ever laser cooling of antimatter, and sympathetic cooling of positrons with laser cooled beryllium. Both laser cooled antihydrogen and colder positrons should lead to colder antihydrogen and more precise measurements of its properties.

5.1 Finestructure Measurements

In 2020, we published measurements of the 1S-2P transitions, also known as the Lyman-α transitions, in antihydrogen [3]. More specifically, we measured the frequencies of the $1S_{1/2}$–$2P_{1/2}$ and $1S_{1/2}$–$2P_{3/2}$ transitions. Both were in agreement with the corresponding transitions in hydrogen, and by averaging the deviations of the measured transitions from the predicted values, they test CPT invariance to the 16 parts per billion level (1 standard deviation). By combining the $1S_{1/2}$–$2P_{1/2}$ measurement with the previous measurement of the 1S–2S transition [10], we infer a zero-field Lamb shift ($2S_{1/2}$–$2P_{1/2}$) that agrees with the theoretical value for hydrogen within 11%.

Placing an antihydrogen atom in a magnetic field lifts the degeneracy of the fine structure states as seen in figure 5.1, which shows the energy levels as a function of the magnetic field strength. For the 1S–2P measurements, the field minimum was 1.0329 ± 0.0004 T. The $2P_{3/2}$ state split into four states ($2P_{a,b,c,d}$), and the $2P_{1/2}$, $2S_{1/2}$, and $1S_{1/2}$ split into two states each ($2P_{e,f}$, $2S_{ab,cd}$, and $1S_{ab,cd}$). Taking the antiproton spin into account, each substate splits into two hyperfine states. The S states split into $S_{a,b,c,d}$, and the P states each split into two states indicated with a $\pm$ subscript. The resolution in the experiment presented is not high enough to resolve the $n = 2$ hyperfine states.

As only the low-field-seeking states are trappable in the ALPHA-2 atom trap, the positron spin is naturally polarised in the trapped antihydrogen sample. By injecting microwaves into the trap, the $1S_c$ state can be driven into the untrappable $1S_b$ state, thereby creating a double spin-polarised sample, as the $1S_c$ states are eliminated. The transition is indicated by a solid black arrow in figure 5.1.

The fine structure measurement scheme involves measurements of the single polarised sample of the $1S_{cd}$–$2P_{f\pm}$ and $1S_{cd}$–$2P_{c\pm}$ transitions, and measurements of the double polarised sample of the $1S_d$–$2P_{f-}$ and $1S_d$–$2P_{c-}$ transitions as indicated with solid blue and red lines in figure 5.1. When the exited states decay (indicated with dashed lines) after a few nanoseconds, there is a chance that they will decay to a non-trappable state. We refer to this as the antihydrogen having "spin-flipped". The following annihilation will give a signal in the detector, which indicates that the laser is on resonance. The laser used to probe the transitions is the 121 nm linearly polarised pulsed laser with a pulse energy of about 0.5 nJ.

A measurement series was done for each of the chosen transitions, and each series consisted of two to four runs. For each run about 500 antihydrogen atoms were accumulated, and the sample was irradiated for two hours. The frequency of the laser was set to scan around the resonance frequency expected for hydrogen. A total of 12 different frequencies within an about GHz window were chosen, and the laser was set to each frequency for a period of 20 seconds before changing to another frequency.

After the two hours, the trap was ramped down and the remaining antiatoms detected. In this way we could determine that 40–60% of the antiatoms had experienced a laser induced spin flip. The resulting resonance curves can be seen in figure 5.2. The transition probability is determined by taking the number of detections at each point relative to the total number of antiatoms and then normalising to the laser power.

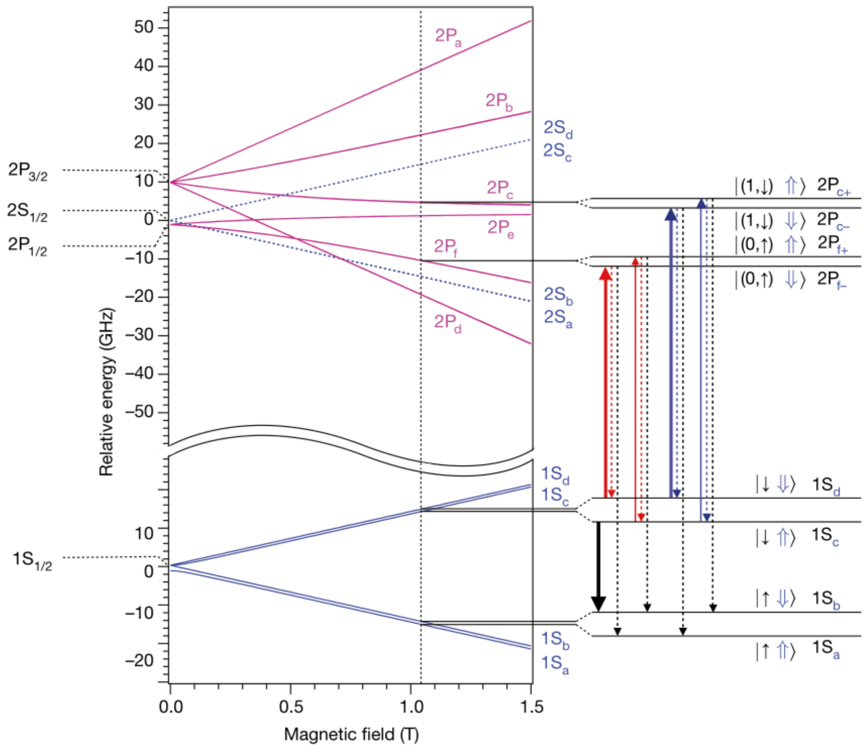

Figure 5.1: The energy levels in antihydrogen as a function of magnetic field. The 1S state is seen to split into two states ($2S_{ab,cd}$), and the $2P_{1/2}$ and $2P_{3/2}$ states split into the $2P_{e,f}$ and $2P_{a,b,c,d}$ states respectively. Each state further split into two hyperfine states, indicated with a $\pm$ for the 2P states. Microwaves can drive the $1S_c$–$1S_b$ transition indicated with a black arrow to produce a double spin polarised sample. The transitions driven with the 121 nm laser are indicated by red and blue arrows. The decays of the excited 2P states to trappable states are shown with dashed blue and red arrows, while the decays to untrappable states are indicated with dashed black arrows. From [3].

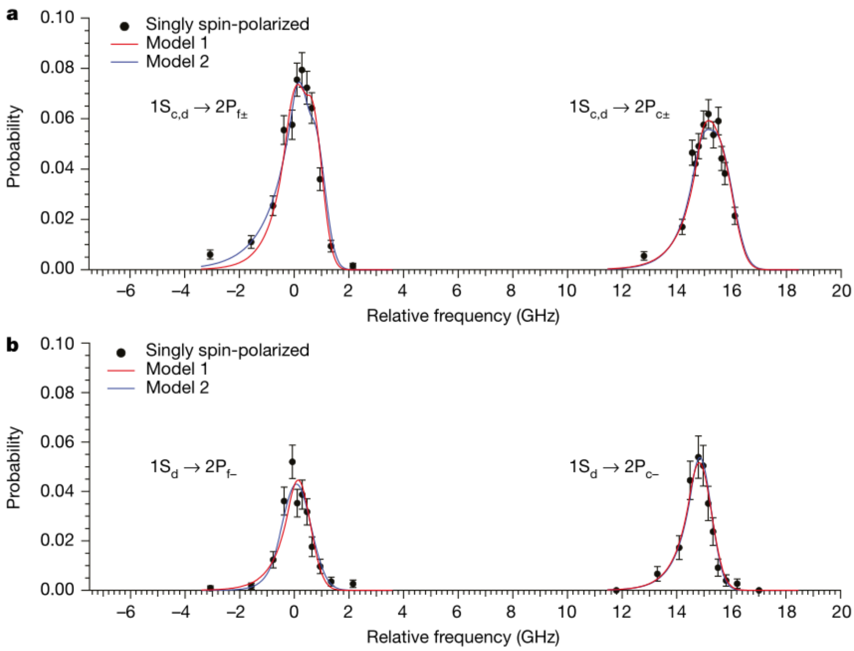

Figure 5.2: The measured line shapes for the single polarised (a) and double polarised (b) samples. The transition probability is determined as the number of antiatoms escaped due to resonance transitions relative to the total number of trapped antihydrogen atoms and normalised to the laser power. The data are fitted to the two models described in the text. From [3].

Transition	$\bar{H}$ **f(exp) [MHz]**	**Deviation from H theory [MHz]**
$1S_{cd}$–$2P_{c\pm}$	2,466,051,659(62)	34
$1S_{cd}$–$2P_{f\pm}$	2,466,036,611(88)	-31
$1S_{d}$–$2P_{c-}$	2,466,051,189(76)	-81
$1S_{d}$–$2P_{f-}$	2,466,036,395(81)	108

Table 5.1: The measured frequencies and their uncertainties (1 standard deviation) for the four different types of measurements described in the text. The deviation from the theoretical result for hydrogen in a field of 1.0329 T is listed as well. From [3].

The inhomogeneity of the magnetic trap causes the asymmetry of the lineshape seen in figure 5.2, as stronger magnetic field causes a lower resonance frequency. Extensive simulations have been done to understand this feature, and they show that the lineshape data can be described by a Gaussian that splices into an exponential in its tails. The efficiency of the microwave sweep also affects the lineshape. The sweep efficiency is assumed to be 95%, which is consistent with earlier experiments [8] and the recorded data, and this is incorporated into the simulation. Two different fits are shown in figure 5.1. In model 1, the simulation results are used to constrain the parameters of the fit, whereas no constraints are used in model 2. The models are seen to generally agree with each other.

The absolute values for the measured frequencies and their uncertainties are listed in table 5.1 together with the predicted values for hydrogen. Besides the significant statistical uncertainties, major uncertainty contributions come from the calibration and drift of the wavemeter, frequency shifts of the laser, the imperfect microwave sweep, and the fitting model. By taking the average of the uncertainties listed in table 5.1, the measurements combine to give a relative uncertainty of $16 \cdot 10^{-9}$.

The uncertainty of the 1S–2P transition measured in antihydrogen is greater than for the 1S–2S transition ($2 \cdot 10^{-12}$ [10]), but although the 1S–2P transition is not the transition most precisely determined, it is a probe for CPT breaking effects that might not appear in the 1S–2S transition [21, 20] as discussed in section 2.2. Combining the 1S–2S transition with a more precise measurement of the Lamb shift would allow a determination of the antiproton charge radius purely based on antimatter results [88], and thereby be independent of the proton charge radius puzzle [89, 90]. Finally, the 1S–2P transition has the potential to be used to laser cool antihydrogen.

5.2 Laser Cooling of Antihydrogen

The 1S–2P transition allows laser cooling of antihydrogen, which we demonstrated [4] shortly after the characterisation of the transition [3]. The effect of the laser cooling was seen in the lineshapes of the 1S–2P and 1S–2S transitions as well as the antihydrogen's time of flight from ejection to annihilation. Reconstruction of the transverse energies show a reduction in the median energy by more than a factor of 10 to around 1 μeV.

The preperation of the antihydrogen sample is similar to the one described in section 5.1. We accumulated about 1000 antihydrogen atoms in the ALPHA-2 atom trap with a trap depth of about 50 μeV. The energy levels in the 1 T field split as shown in figure 5.3, where the relevant energy levels and transitions are shown in the right-hand side. The notation of the energy levels is described in section 5.1.

In the experiment, the antihydrogen sample was made double spin-polarised by using microwaves to drive the $1S_c$ state into the untrappable $1S_b$ state. The $1S_d$ state is excited to the $2P_{a-}$ state by a 121 nm laser, which can be detuned to cool or heat the antihydrogen as indicated by the red and blue arrow respectively. In a magnetic field, where the $2P_{3/2}$ degeneracy is lifted, the $1S_d$–$2P_{a-}$ transition is closed, and thus suitable for laser cooling. To probe the antihydrogen sample, the $1S_d$–$2P_{c-}$ transition is excited (purple arrow), after which the antihydrogen can decay to a trappable or untrappable state. Finally, the $1S_d$–$2S_d$ two-photon transition is indicated as well. The $2S_d$ state can decay to both trappable and untrappable states (see section 2.6), or get photo ionised and thereby escape the trap.

After accumulation followed a 2 to 4 hour period dedicated to potential laser cooling. We did four types of experiments: one with the red detuned light (laser cooling), one with blue detuned light (laser heating), one without a laser, and one with laser cooling during the accumulation as well as during the dedicated cooling time. This was followed by a 2 to 4 hour probing period, where we probed the 1S–2P transition.

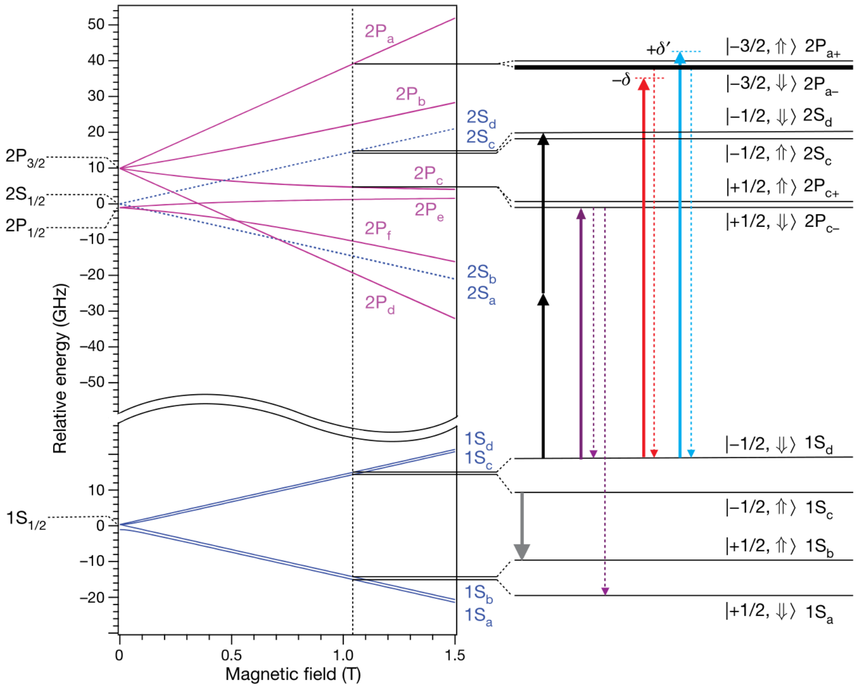

Figure 5.3: The (assumed) energy levels in antihydrogen as a function of magnetic field. The 1S state is seen to split into two states ($2S_{ab,cd}$), and the $2P_{1/2}$ and $2P_{3/2}$ states split into the $2P_{e,f}$ and $2P_{a,b,c,d}$ states respectively. Each state further split into two hyperfine states, indicated with a $\pm$ for the 2P states. Microwaves can drive the $1S_c$–$1S_b$ transition indicated with a black arrow to produce a double spin polarised sample. A 121 nm laser is used to probe the $1S_d$–$2P_{c-}$ transition, and it is detuned to cool/heat the antihydrogen by driving the $1S_d$–$2P_{a-}$ transition (red/blue arrow). Finally, the two-photon 1S–2S transition is indicated as well. From [4].

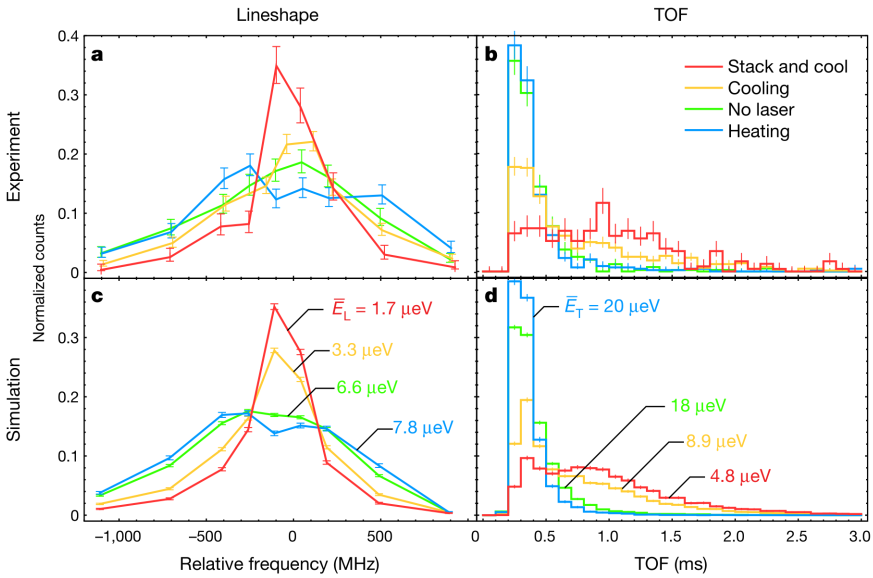

Figure 5.4: **A:** Measured lineshapes of the $1S_d$–$2P_c$ transition of the trapped antihydrogen sample for the different laser settings: cooling during and after stacking/accumulation (red), cooling after stacking (orange), no laser (green), and heating (blue). **B:** Histogram of the time-of-flight (TOF) from when laser sent a pulse to an annihilation was detected. **C:** Simulation of the expected lineshapes. **D:** Simulation of the expected times of flight. From [4].

Probing the 1S–2P transition does two things: it measures the lineshape, whose linewidth is influenced by the temperature via Doppler broadening, and it allows us to measure the time between when the laser sends a pulse to when any spin-flipped antihydrogen atoms annihilate on the electrode wall. The obtained lineshapes and the time-of-flight (TOF) spectra are seen in figure 5.4.

The measured lineshapes in figure 5.4a show that the linewidth is narrower (less Doppler broadened) for the samples subject to the cooling laser, indicating that the antiatoms have been cooled longitudinally. Similarly, the TOF spectra in figure 5.4b show that the samples illuminated with the red detuned laser have a wider TOF spectrum than the other samples, indicating that the antiatoms are slower/colder transversely. Both datasets agree with the simulated results in figure 5.4c and d.

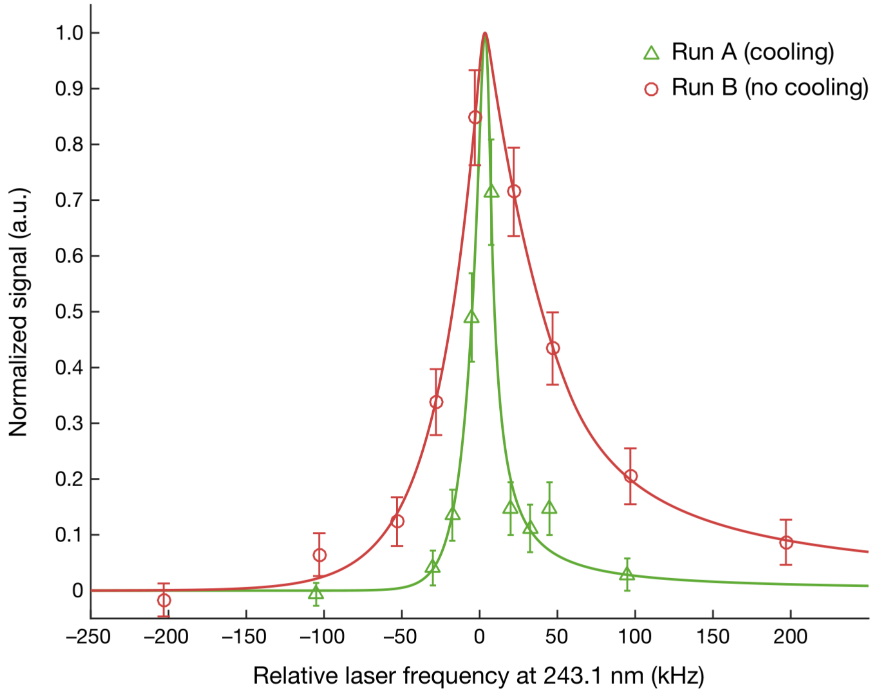

Figure 5.5: Measured 1S–2S lineshape of a laser cooled (run A) and an uncooled (run B) sample of antihydrogen. The spectra are normalised to their amplitude and the background has been subtracted to make the lineshapes line up and help the comparison. From [4].

Searching for further evidence of laser cooling, we measured the lineshape of the 1S–2S transition of a laser cooled sample. After accumulating antihydrogen for 9 hours while laser cooling it, and laser cooling it for another 6 hours, we obtained the lineshape seen in figure 5.5 by probing it with the 243 nm laser for 15 minutes. The figure shows the lineshape obtained for a non-laser cooled sample as well. The linewidth (FWHM) of the cooled sample is seen to be approximately a fourth of the uncooled sample.

The full potential of the demonstrated laser cooling is yet to be explored. The 1S–2S lineshape presented in figure 5.5 should only be seen as a demonstration of the effect of laser cooling, as the cooling procedure has not been optimised, and not enough data points were taken to characterise the lineshape. Current upgrades to the laser will increase the overall stability of the system and increase the pulse repetition rate from 10 to 50 Hz, which should increase the cooling rate. Besides offering narrower lineshapes due to reduced Doppler, Zeeman, and transit-time broadening, cold antihydrogen also offers more uniform responses to manipulations of the sample by for example magnetic fields. Colder antihydrogen is thus also desirable for gravity experiments.

5.3 Sympathetic Cooling of Positrons

The way antihydrogen is currently produced in ALPHA (see chapter 4), the antiprotons and positrons thermally equilibrate before forming antihydrogen. As there are about ten times more positrons than antiprotons, it is the temperature of the positron plasma that determines the antihydrogen temperature distribution, and thereby how many of the formed antihydrogen atoms are trappable. The lower experimental limit of the positron temperature in ALPHA-2 has been around 17 K. To cool the positrons further, we have demonstrated how positrons can be sympathetically cooled (see section 2.4) with laser cooled beryllium ions [84, 85].

Beryllium is the lightest positive ion that can be laser cooled in its ground state, which is done with a 313 nm laser. The transition used for the laser cooling is the $(2s)^2S_{1/2}$–$(2p)^2P_{3/2}$ transition[1]. In ALPHA-2, we have been able to cool a sample of pure Be^+ ions to about 200 mK.

Sympathetically cooling positrons with beryllium must be compatible with subsequently forming antihydrogen. Hence, the number of beryllium ions cannot be too high, as it will increase the risk of antiproton capture and annihilations. Figure 5.6 shows the axial positron temperature in a mix of $2.6 \cdot 10^6$ positrons and $3.8 \cdot 10^5$ Be^+ ions as a function of the laser detuning. The detuning is meaured in units of the natural linewidth of the $(2s)^2S_{1/2}$–$(2p)^2P_{3/2}$ transition, Γ. The figure shows that a positron temperature of about 7 K is achieved at a detuning of -7Γ to -2Γ. Our studies suggests that the effect of the sympathetic cooling saturated at this temperature, which might be explained by centrifugal separation of the two species. It is estimated that decreasing the positron temperature from 17 to 7 K would increase the fraction of trappable antihydrogen atoms created by up to a factor of 5.

[1] In this notation for the atomic state, the first two symbols describes the state of the valence electron, while the last three define the quantum numbers $^{2S+1}L_J$.

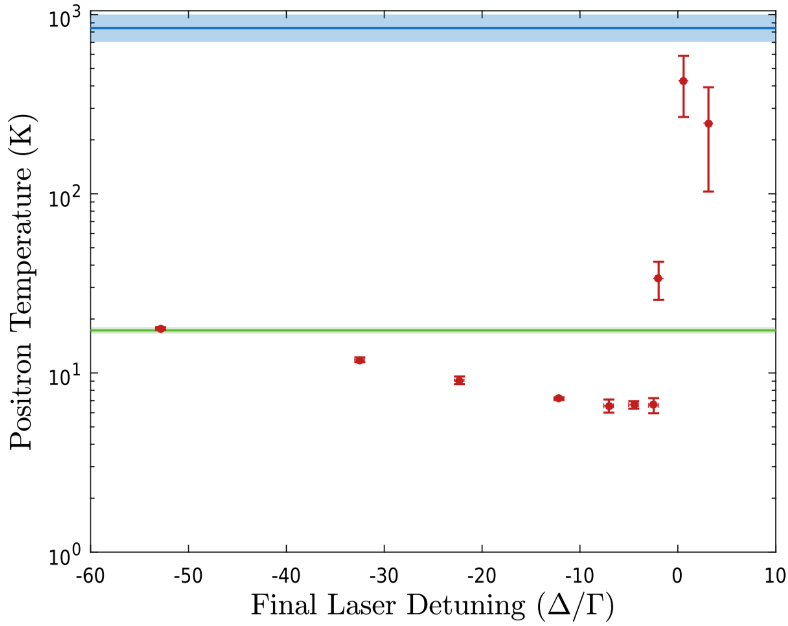

Figure 5.6: The axial positron temperature as a function of the laser detuning, Δ, measured in units of Γ. The data points are equipped with 1 standard deviation error bars. The blue line indicates the positron temperature with the Be^+ ions present but without the cooling laser, and the green line indicates the positron temperature with no Be^+ ions present. From [84].

As the octupole field in the atom trap would be present in an antihydrogen production run, the radial size of the beryllium and positron plasmas must be kept small enough to avoid interaction with the octupole field. RW compression (described in section 2.4) would be one way to compress the plasmas, but there is currently no segmented electrode in the central region of the ALPHA-2 atom trap, where the overlap with the lasers is the greatest, so the electrode stack would require modification.

Chapter 6

The ALPHA-g Apparatus

The purpose of this chapter is to describe the ALPHA-g experiment, which was constructed as an addition to the existing apparatus in 2018. Elements like the catching trap and the positron accumulator, which are shared by ALPHA-g and ALPHA-2, have already been described in chapter 4. This chapter will focus on ALPHA-g itself and especially the magnet system. The beamline, which was added alongside ALPHA-g in 2018, is described briefly as well.

6.1 Motivation

As described in section 2.2, CPT invariance does not describe how *antimatter* is accelerated in a gravitational field created by *matter*. The effect is hard to measure on charged antimatter, as the electromagnetic force, on for example an antiproton, is many orders of magnitude greater than the gravitational force[1]. A gravity measurement would thus require extreme knowledge and control of the surrounding electric fields, which has made such a measurement infeasible thus far. With neutral antimatter or neutral systems containing antimatter, the issue of stray electric fields is largely avoided[2]. However, experiments with antineutrons or positronium have been deemed infeasible [91, 34], while experiments on antihydrogen and muonium seem promising [7, 92].

In 2013, the ALPHA collaboration demonstrated how the effect of Earth's gravitational field on antihydrogen can be measured [75].

[1]For two protons, the Coulomb force between them is $1.2 \cdot 10^{36}$ times stronger than the gravitational force between them.

[2]With a non-uniform charge distribution of the system, there can be polarisation effects.

Defining $F = \bar{g}/g$, where $\bar{g}$ is the acceleration of antimatter in a Earth's gravitational field, F was measured to be within the interval $-65 < F < 110$ (2 standard deviations). The measurement was done in the horizontal ALPHA-2 atom trap, which has not been designed with such measurements in mind. ALPHA-g is a vertical atom trap dedicated to a gravity measurement of antihydrogen, and it is expected to restrict the value of F.

6.2 Overall Setup

The ALPHA-g trap is essentially a 3 meter vertical version of the ALPHA-2 trap. To incorporate ALPHA-g with the pre-existing apparatus, the positron source and accumulator were moved away from ALPHA-2, and a beamline was constructed to connect the positron accumulator to ALPHA-2. An interconnect section, which can bend the charged particles upwards to ALPHA-g [93], is placed between the accumulator and ALPHA-2. A sketch of the key parts of the full apparatus can be seen in figure 6.1. The main components of ALPHA-g are the external solenoid, the detector, the cryogenic system, the magnets, and the electrodes at the bore of the system.

The electrodes and magnets can be divided into three regions: an upper and a lower trapping/full strength region, and an analysis region in the middle (see figure 6.2). Each of the trapping regions works like the ALPHA-2 atom trap with a Penning-Malmberg embedded in a Ioffe-Pritchard trap. The control of the magnetic field in these regions is not robust enough to perform a "precision measurement" of antigravity, whereas the analysis region offers a higher degree of control.

The analysis region is not designed to trap charged particles and create antihydrogen, but antihydrogen can be transferred from one of the trapping regions to the analysis region. Having a trapping region on each side of the analysis region ensures symmetry around the analysis region, which is important due to the fields created by Eddy currents. A full strength region at each end of the detector also ensures good detector coverage above/below the region, in case antimatter falls up/down.

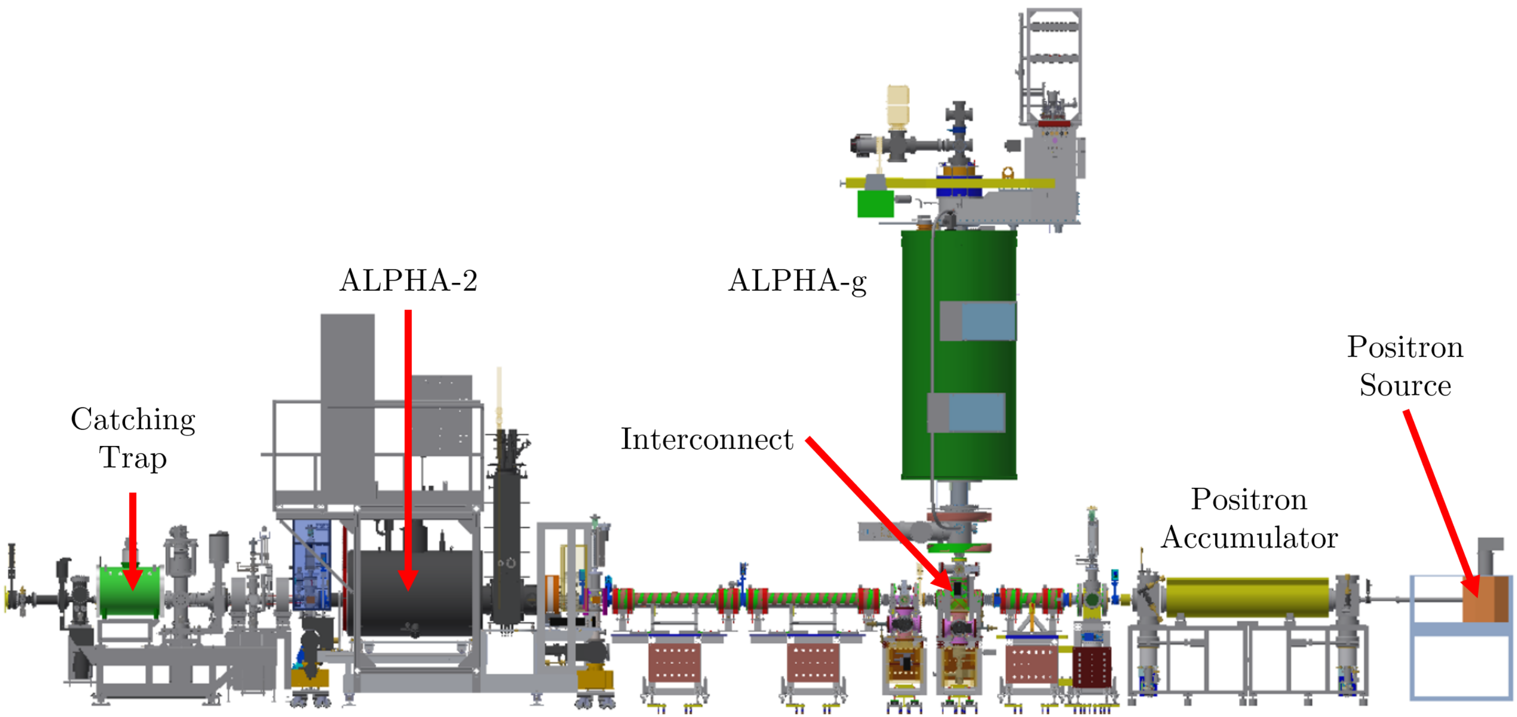

Figure 6.1: Simplified sketch of the key parts of the entire ALPHA apparatus. To the left is the CT, which traps the antiprotons from the AD/ELENA, and to the right is the Na-22 positron source and the positron accumulator. The ALPHA-2 and ALPHA-g atom traps are installed between the CT and the accumulator. The interconnect is seen below ALPHA-g.

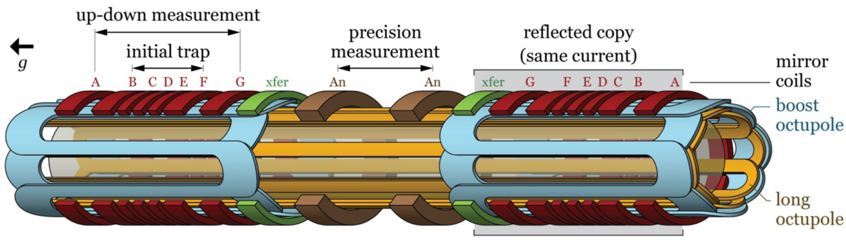

Figure 6.2: Simplified 3D model of the inner ALPHA-g magnet system. The long octupole (orange) and the two analysis coils (brown) make up the analysis region in the middle, whereas the two short/boost octupoles (blue) and the two sets of mirrors A-G (red) make up the trapping region at each end. The two transfer coils (green) are shown as well. Courtesy of Chukman So.

6.2.1 Measurement Scheme

In principle, the gravity measurement scheme is simple. When antihydrogen has been trapped in one of the vertical atom traps, the mirror coils at each end are gradually lowered, so the antihydrogen can escape axially. Let us define R_E as the ratio of antiatoms escaping upwards and downwards. A gravitational effect will translate into a difference between the number of antiatoms escaping in each direction, so $R_E \neq 1$. Rather than using R_E to derive the gravitational acceleration, we will balance the axial magnetic field to cancel the effect of the gravitational field as illustrated in figure 6.3, so the same number of antiatoms should escape downwards and upwards ($R_E = 1$). The surrounding detector can determine the axial location of the resulting annihilations.

Define ΔB_g as the difference in magnetic field it would take to cancel the effect of gravity for a given difference in height, assuming $\bar{g} = g$. For the initial measurement, we intend to make at least two types of measurements: one with $\Delta B = \Delta B_g$ and one with $\Delta B = -\Delta B_g$. The first would offset gravity if $\bar{g} = g$, and the second would offset gravity if $\bar{g} = -g$. By measuring which of the two experiments gives R_E closer to 1, we can determine the sign of $\bar{g}$. We refer to this experiment as the "up/down measurement".

The distance between the centres of the outermost mirror coils in the full strength region (A and G), where the up/down measurement is intended to be done, is 256 mm. Assuming $\bar{g} = g$, the distance

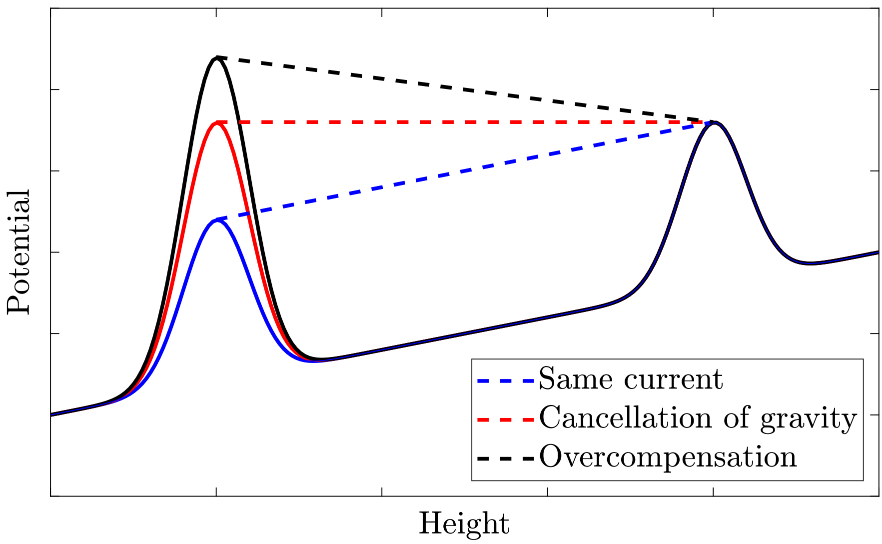

Figure 6.3: A sketch of the total on-axis potential (gravitational plus magnetic potential) seen by an antihydrogen atom. The concept of the gravity measurement is to offset the gravitational potential by increasing the current in the mirror where the gravitational potential is lowest, as shown by the red curve. The antihydrogen will then have an equal chance to escape upwards and downwards, when the fields are lowered, as described in the text.

translates into a potential energy difference of $4.20 \cdot 10^{-27}$ J or 26.3 neV for an antihydrogen atom. If the antihydrogen is in its ground state, its magnetic moment is approximately one Bohr magneton (see section 2.5), so the potential energy difference can be offset by a difference in magnetic field of $\Delta B_g = 4.53 \cdot 10^{-4}$ T. Hence, the maximal uncertainty of the field produced by each of the two mirror coils must be smaller than $\sigma_B = 3.20 \cdot 10^{-4}$ T.

The analysis region is designed to determine the value of $\bar{g}$ within $\pm g/100$. We refer to this measurement as the "1% measurement". The measurement scheme is the same as for the up/down measurement, but with measurements done at magnetic field offsets separated by $2\Delta B_g/100$ rather than $2\Delta B_g$. As the analysis region is 400 mm long, the maximal uncertainty of the field produced by each of the mirrors must be smaller than $\sigma_B = 5.01 \cdot 10^{-6}$ T.

Neutral particles in the trap that have limited mixing between the energy components or change from the no-mixing to the mixing category (see section 2.5) pose a problem to the gravity measurement scheme. If energy is suddenly transferred from the perpendicular to the parallel component, on a time-scale that is fast compared to the ramp down time, the antihydrogen atom will get enough energy to escape in a random direction. The effect sets a limit to the extent to which the escape energy of the antihydrogen can be controlled. To mitigate the problem, the axial potential barriers will be partially lowered to let the most energetic antiatoms escape, before a gravity measurement is initiated. Simulations show that the optimal initial field of the end barriers is about 0.10 – 0.15 T for the up/down measurement, but the number depends on the experimental parameters.

Assuming that the gravitational potential is exactly cancelled by the magnetic field, there should be an equal chance for the antihydrogen to escape upwards or downwards. To suppress the statistical uncertainty, simulations show that we need about 1000 trapped antihydrogen atoms initially. Such numbers have been demonstrated to be achievable in ALPHA-2 [4].

6.2.2 The Beamline

The beamline connects the positron accumulator, ALPHA-g, and ALPHA-2 as shown in figure 6.1. The details are described in [93, 94]. The beamline is designed to transport positrons and antiprotons with energies lower than 100 eV by using magnetic fields generated by a number of solenoids. Antiprotons travel from the catching trap to either ALPHA-2 or ALPHA-g, while positrons travel from the positron accumulator to ALPHA-2 or ALPHA-g. For both particle species the beam path is about 5 m. The transfer efficiency has proved to be in the 70–80 % range for both species for the bunch sizes typically used. The interconnect, which is placed directly under ALPHA-g, is where particles can be steered upwards to ALPHA-g or be shot straight through from the positron accumulator to ALPHA-2.

In principle, the polarity of the magnetic fields in the beamline does not matter, as long as it matches the regions connected to the beamline to make the field lines continuous. With the current field configuration, the polarities along parts of the beam path are therefore different, dependent on whether particles are transferred to ALPHA-g or ALPHA-2. However, by introducing a zero-field region between ALPHA-2 and the interconnect, only the beamline sections (with bipolar power supplies) will need to change polarity, when running in the ALPHA-g configuration compared to the ALPHA-2 configuration. Such a field configuration is currently under investigation, and the results are promising.

6.2.3 The Detector

When an antihydrogen atom annihilates on the trap walls, the antiproton typically creates 3–5 charged pions [81, 72], which can be detected. The detector that surrounds the internal magnets is a radial time projection chamber (rTPC) [95] with a 90-10 mix of Ar-CO_2 gas. 265 anode wires parallel to the axis of the detector run along the outer perimeter, while the inner surface is covered by 18432 cathode pads. Hence, there is a radial electric field pointing inwards, as well as the axial magnetic field generated by the outer solenoid.

When a charged particle passes through the detector, it ionises the gas along its track. The released electrons drift towards the wires and generate an electric signal. A simulation of such a process can be seen in figure 6.4. The signal in the anode wires gives information about the azimuthal position of the particle track. The cathode pads provide the axial position, and the relative timing of the signals in the pads and the wires can be used to derive the radial position. The axial resolution of the rTPC is designed to be about 1 cm.

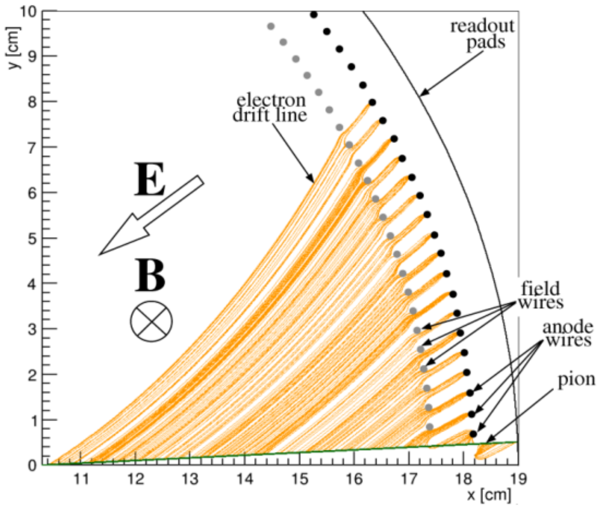

Figure 6.4: Simulation of a pion track (in green) through the rTPC with a 1 T field and an electric field pointing radially inwards. The pion ionises the Ar-CO_2 gas along its way. The released electrons (in orange) drift towards the anode wires, where they create a signal. From [1].

Due to the insensitivity to the photons from the positron annihilations, antihydrogen and antiproton annihilations are practically indistinguishable in the rTPC. However, by removing the electric potential confining the antiprotons and applying an electrical *sweep*, we are guaranteed to be left with only antihydrogen atoms in the trap [6]. The detector readout rate is approximately 200 Hz, which should be enough to avoid saturation, during antihydrogen release. In the planned measurement scheme, hundreds of antihydrogen atoms will be released over tens of seconds.

Just like in ALPHA-2, the only significant background comes from cosmic radiation. The rTPC does not have sufficiently good time resolution to distinguish traversing cosmic rays from particles originating from the bore, so it it surrounded by a Barrel Veto [96]. The Barrel Veto consists of 64 plastic scintillators, and has a time resolution of hundreds of picoseconds, which is fast enough to distinguish a particle travelling through the detector from an event creating multiple tracks inside the detector. A machine learning-based algorithm similar to the one described for ALPHA-2 (see chapter 4) is used to distinguish background from signal.

6.2.4 Plasma Diagnostics

The ALPHA-g apparatus is equipped with diagnostics in a similar fashion to the rest of the experiment as described in section 4.2.7. Key diagnostic tools such as Faraday cups and MCPs are placed on sticks above and below the ALPHA-g magnet region, which allows analysis of particles dumped to one of these diagnostic stations. Multiple scintillator panels are mounted outside the external solenoid to detect antiproton annihilations.

6.3 Magnet Layout

This section will focus on the details of the internal magnets [97] relevant for the measurement scheme, for designing the magnet control system, and operating the magnets. The external solenoid, which provides a 1 T background field, will be described in a separate section (section 6.6).

The internal magnets were manufactured at Brookhaven National Laboratories (BNL). Initially, in 2018, the magnet system was delivered with only the magnets relevant for the lower full strength region installed, as only the lower region is required to do an up/down measurement, which was the goal of the 2018 run. During LS2, the remaining magnets were installed, and the full system was delivered to us in spring 2021.

A 2D sketch of the ALPHA-g magnet system can be seen in figure 6.5. A list of all magnets and their abbreviations and axial field strengths is given by table 6.1, and a more detailed table can be found in appendix A.1. The two full strength regions consist of the blue and orange octupoles and mirror coils A to G in red. The analysis trap for the precision measurement is formed by the yellow long octupole and the brown analysis coils. The transfer and background coils in green can help transfer charged particles or antihydrogen between regions. Finally, the two outermost solenoids in red do the initial trapping and cooling of the charged particles. The magnet system has been designed to produce fields similar to the ones in ALPHA-2.

BNL carefully documented the wiring process by taking pictures of each wire layer with a 24 megapixel camera, as the magnets were wound. The pictures are analysed by an image analysis algorithm, which can determine the axial position of the wires down to about

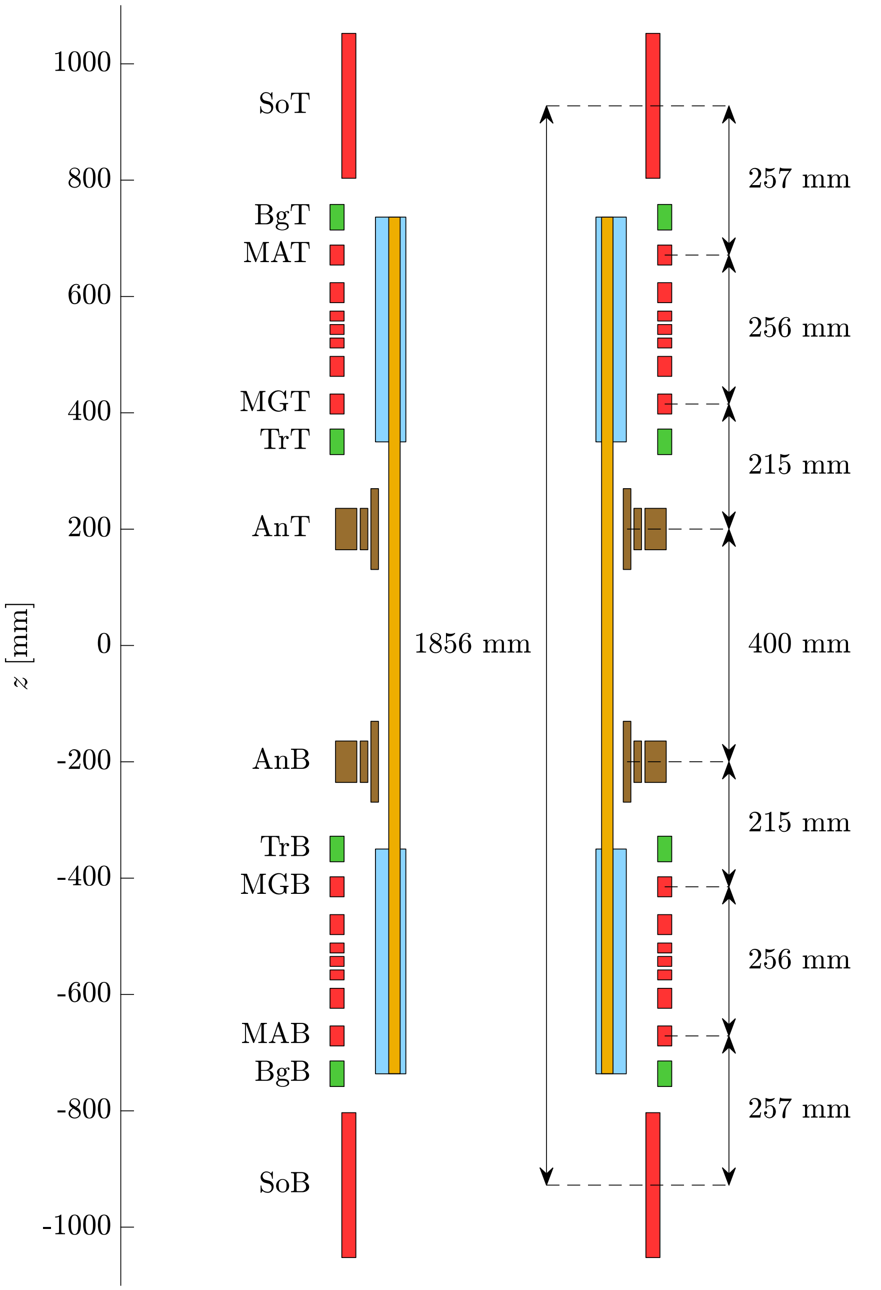

Figure 6.5: 2D sketch of the ALPHA-g magnet system. The octupoles are shown in blue and orange, and the solenoids are shown in red, green, and brown. See table 6.1 for a complete list of the magnets. The radial dimensions are not to scale, and some dimensions have been exaggerated to improve readability. The external solenoid is not shown for simplicity.

Magnet Name	**Abbreveation**	B_{Trap}/I **[mT/A]**
Bottom Magnets		
Capture Solenoid	SoB	33
Background Coil	BgB	10.6
Mirror A	MAB	10.4
Mirror B	MBB	10.5
Mirror C	MCB	5.66
Mirror D	MEB	5.67
Mirror F	MFB	5.64
Mirror E	MEB	10.6
Mirror G	MGB	10.6
Transfer Coil	TrB	12.8
Analysis Coils		
Analysis Coil Bottom	AnB	22.1
Corrector Coil Bottom	CCB	2.27
Corrector Octupole Bottom	COB	0.111
Analysis Coil Top	AnT	22.4
Corrector Coil Top	CCT	0.63
Corrector Octupole Top	COT	0.111
Top Magnets		
Transfer Coil Top	TrT	13.1
Mirror G	MGT	10.6
Mirror F	MFT	10.7
Mirror E	MET	5.56
Mirror D	MDT	5.71
Mirror C	MCT	5.68
Mirror B	MBT	10.7
Mirror A	MAT	10.6
Background Coil	BgT	10.9
Capture Solenoid	SoT	33.4
Others		
Long Octupole	LOc	0.10
Short Octupole Bottom	OcB	0.55
Short Octupole Top	OcT	0.55
External Solenoid	Babcock	5.26
External Shim Coil	Babcock Shim	0.70

Table 6.1: Table of the magnets in ALPHA-g. For each magnet, its abbreviation and the trapping field per current, B_{Trap}/I, seen by an antihydrogen atom are listed.

5 μm. By measuring the diameter of a wire layer at multiple places several times, the radial position of each wire is determined down to 25 μm. The wire position data is used as a basis for a detailed Biot-Savart model of the magnet fields as described in chapter 8.

It is crucial for a gravity measurement that the antiatoms are able to sample both end-mirrors as they orbit in the trap, as there would otherwise be a preferred escape direction independent of the gravitational potential. To avoid local minima, the gradient of the magnetic potential must be smaller than the gradient of the gravitational potential. Thus, with the gravitational energy differences calculated in section 6.2.1, assuming $\bar{g} = g$, the field gradient must be less than $1.77 \cdot 10^{-6}$ T/mm for the up/down measurement and less than $1.77 \cdot 10^{-8}$ T/mm for the 1% measurement. If the field gradients are higher, antiatoms below a certain energy cannot sample both ends, so they must thus be ignored in the gravity measurement.

Persistent currents[3] make it hard to fulfil the strict requirements for the fields in the analysis region [97]. The problem is helped by reducing the amount of superconducting material used to make the magnets in the analysis region, by increasing the distance between the analysis coils and the full strength regions, and by making the magnet system symmetric around the analysis region. As a consequence, the analysis trap is too shallow to trap enough antihydrogen from mixing positrons and antiprotons, so antihydrogen must be created (and accumulated) in a full strength region and transferred to the analysis region.

Other measures taken to make the field in the analysis region as smooth as possible, are to have the end-turns of the long octupole as far away from the analysis region as possible, and to wind the octupoles with as thin a wire as possible, to reduce the amount of material that can have persistent currents. By displacing the long octupole end-turns from the analysis region, they overlap with the full strength region. Hence, the short octupoles can be made with fewer layers, at the cost of using both octupoles to create the radial confining field of the full strength trap.

Most of the bottom region magnets are connected in series with their upper region counterpart to maximise field symmetry for the analysis

[3] Currents that flow without being driven by a power source. For example in superconducting magnets.

region. The magnets operated in series are mirrors B to F, the short octupoles, and the background and transfer coils. Mirrors A and G, the analysis coils, the corrector coils and octupoles, and the capture solenoids are all run independently.

The power supplies used to energise the magnets are the Delta Elektronika SM 15-100 model for the capture solenoids and the background and transfer coils, the Sorensen SGA 10/1200 for the octupoles, and the CAENels Fast-PS-1k5 100-15 for the mirrors. The external solenoid is run by two CAENels in parallel, and the external shim coil is run by a CAENels Fast-PS 1020-200 model.

6.3.1 Capture solenoids

To capture the charged particles, as they are directed towards ALPHA-g from the interconnect, a capture solenoid is placed at each end of the magnet region. Together with a series of electrodes, each capture solenoid forms a Penning-Malmberg trap. The capture solenoids are designed to produce a field of 3 T in addition to the external solenoid field. They are run independently, as it would not be possible to extract their stored energy fast enough in case of a quench (see section 6.5.3), if they were run in series.

The capture solenoids sit in the fringe region of the external solenoid, and the wiring pattern and density are designed to compensate for the magnetic field gradient of the external solenoid. As shown in figure 6.6, a capture solenoid consists of nine layers of uniform wire density, one layer of non-uniform wire density, and 8 layers of axially short shim windings at each end. The shim windings serve to flatten the field of the capture solenoid in its fringe region.

6.3.2 The Full Strength Regions

The full strength regions consist of the long octupole and the top and bottom set of mirrors A through G and the short octupole. The two regions are symmetric around the analysis region. A 2D sketch of the bottom full strength region is shown in figure 6.7. Mirrors A and G will be used to provide the axial confinement for the up/down measurement, as these mirrors have the greatest gravitational potential energy difference between them.

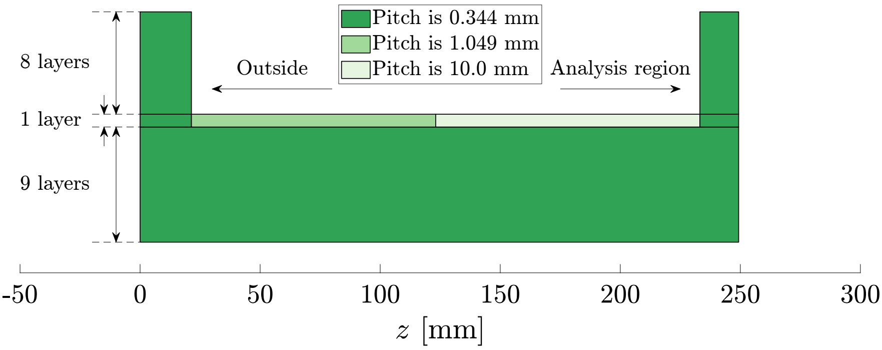

Figure 6.6: Cross section of the capture solenoid that illustrates the wiring pattern. There are more turns on the end facing the outside of the ALPHA-g magnet region, as the capture solenoid is located in the fringe region of the external solenoid.

In more advanced measurement schemes, antihydrogen is first confined between mirrors B and F, and the trapping region is then expanded to A and G to adiabatically cool the antihydrogen. The central mirrors can be used to deepen the trapping potential between B and F, and make the central potential flatter and the edge potentials steeper[4]. Using three magnets to reshape the potential between B and F allows the region to be axially shorter than if only one magnet was used to achieve the same effect.

The long octupole extends over the full strength regions to get the end-turns as far away from the analysis region as possible. However, the long octupole cannot provide a sufficiently strong radial confinement for the full strength regions, so it is supplemented by the short octupoles. All octupoles are made of bi-layers, which are two layers of wire that have oppositely wound end-turns to minimise the solenoid-like field of the end-turns. The short octupoles have one bi-layer on the inside of the long octupole and two on the outside. The wiring pattern for each layer is approximately the same for the long octupole, so the field produced by the short octupole is roughly three times stronger than the field of the long octupole.

[4]This makes the potential more *square-like*, which increases the chance that the antihydrogen is formed with a potential energy close to the potential minimum.

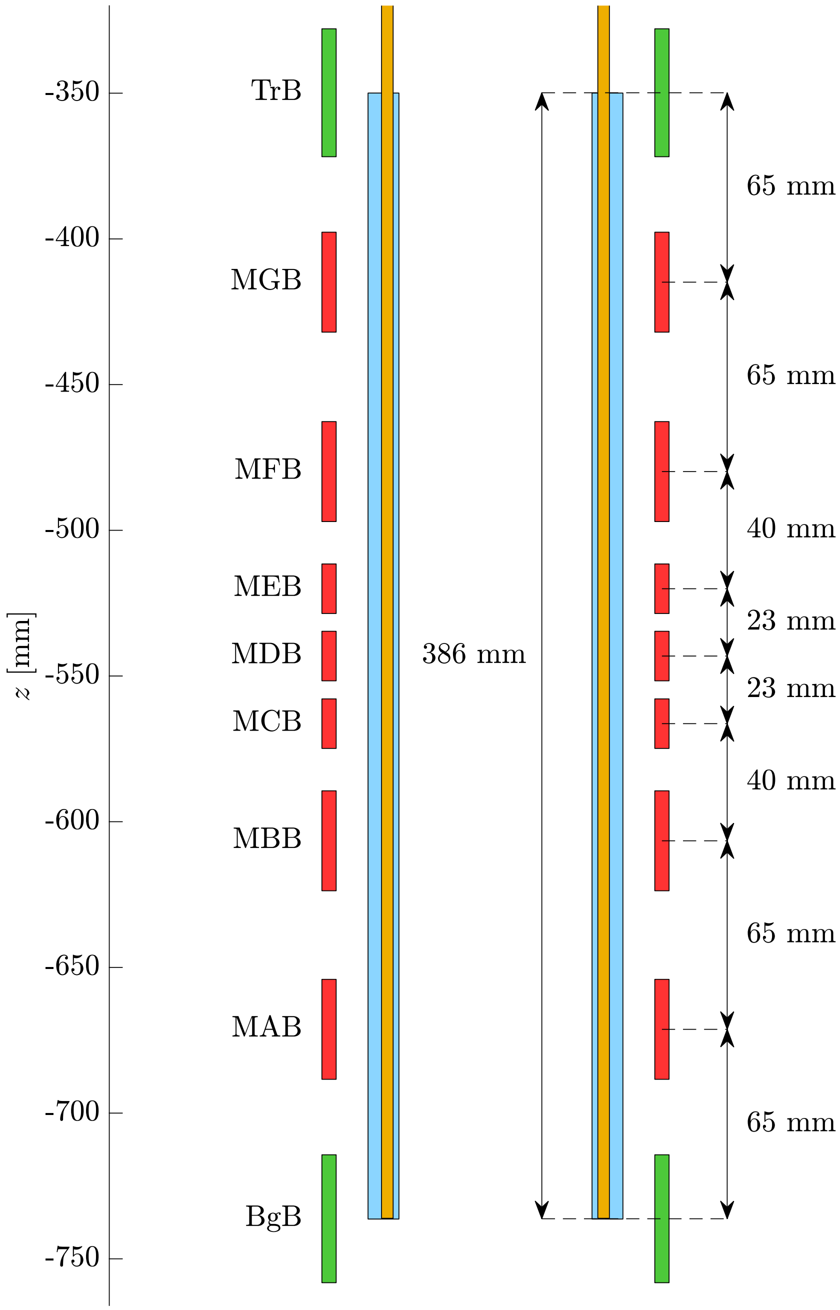

Figure 6.7: Sketch of the magnets in the bottom full strength region. The octupoles are shown in blue and orange, and the mirrors are shown in red and green. Radial dimensions are not to scale, and some dimensions have been exaggerated to improve readability.

6.3.3 The Analysis Region

The main components of the analysis region are the two analysis coils and the long octupole. These magnets are designed to be used for the precision measurement as described in section 6.2.1. Given that the gravitational potential energy is $U_g > 0$ at the axial centre of the upper coil and 0 at the lower coil, one chooses the current in the coils to be $I_{lower} > I_{upper}$ (assuming $\bar{g} = g$), so the magnetic potentials fulfil $U_{lower} = U_{upper} + U_g$. However, if this is true on the axis at $\rho = 0$, it will not be true for larger $\rho > 0$, so $U_{lower}(\rho > 0) > U_{upper}(\rho > 0) + U_g$.

For a given set of currents, I_{lower} and I_{upper}, $|\vec{B}(\rho)|$ at the axial centre could be made the same at each magnet by having the coil with the stronger current be axially longer than the other[5]. However, for a given set of magnet dimensions, the field curvatures will only match for a specific current. A solution to this problem is to add a long low strength "corrector" solenoid at the same axial position as each of the analysis coils, which will allow the effective length of the analysis coils to be adjusted [97].

For higher radii, the situation becomes further complicated due to the contribution from the octupole to the total field. The octupole field, $B_\perp$, is perpendicular to the axis, and it adds to the parallel field at the axial solenoid centre, $B_{||}$, of the solenoids. Since the sum of the vectors is smaller than the sum of the absolute values, the size of the total magnetic field is smaller at the upper mirror, compared to the lower given $I_{lower} > I_{upper}$. To counter this effect, a short low strength corrector octupole is added between the analysis and corrector coils. A sketch of the magnet configuration around the analysis coil can be seen in figure 6.8.

6.3.4 Transfer and Background Coils

To distance the short octupole end-turns from the analysis region, there is an approximately 120 mm gap between the short octupoles and the analysis coils. These gaps are too big to perform a lossless transfer of antihydrogen from the trapping to the analysis region with only the 1 T background field. Hence, a transfer coil has been added in each gap. The coils introduce an asymmetry around the full strength region, so identical coils, referred to as background coils, are added outside mirrors A (see figure 6.5).

[5] Assuming the same number of turns and same current.

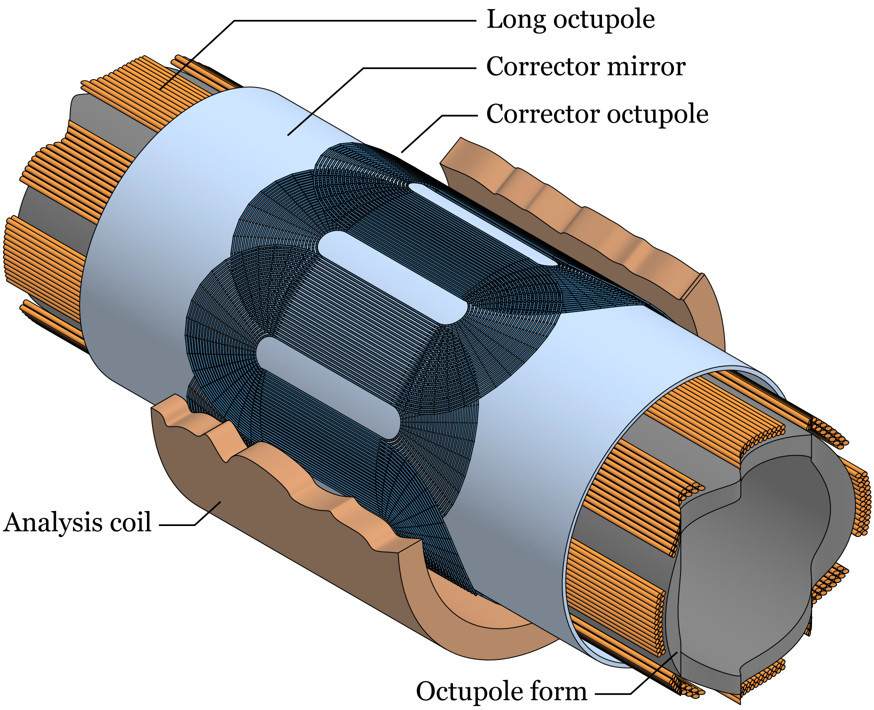

Figure 6.8: Detailed sketch of the region around an analysis coil. To make the magnetic field match the field of the other analysis coil for all ρ, a corrector coil/mirror and a corrector octupole are added to each of the main analysis coils. Courtesy of Chukman So.

B-field [T]	Critical Current [A]
0.5	383.6
1.0	298.0
2.0	213.1
3.0	171.4
4.0	143.5
5.0	120.7
6.0	97.8

Table 6.2: The critical current as a function of magnetic field strength for a straight single-strand sample of the superconducting wire at 4.208 K.

In case the external solenoid is not operational, the field from the background coil can be used to help transfer charged particles from the capture solenoid to the full strength region. Together with the transfer coil and the rest of the solenoids in the full strength region, it can also be used to create a background field across the full strength region as a replacement of the external solenoid field, although the field would be less uniform. Finally, since the background solenoid overlaps the ends of the short and long octupoles, it can be used to cancel the axial field component generated by the octupole end-turns.

6.3.5 Superconducting Wire

The ALPHA-g magnets are wound with a mix of superconducting niobium-titanium (NbTi) and normalconducting copper wires. The superconducting wire is the VSF-678 (Very Small Filament, 678 filaments per wire) model from Supercon Inc, which has NbTi filaments sitting in a copper matrix. This wire is designed to be one of the best high field superconducting wires available. The critical current[6] of a straight single-strand wire sample was measured by BNL at 4.2 K as a function of the magnetic field. The results can be seen in table 6.2.

Of the internal magnets, the top and bottom capture solenoid generate the strongest field and are the ones closest to the critical current. When the capture solenoids are at the maximum intended current (100 A), it is immersed in a 4.3 T field (including the 1 T field from the external solenoid), which should still be comfortably below the critical current specified by table 6.2.

[6]The current at which the wire quenches.

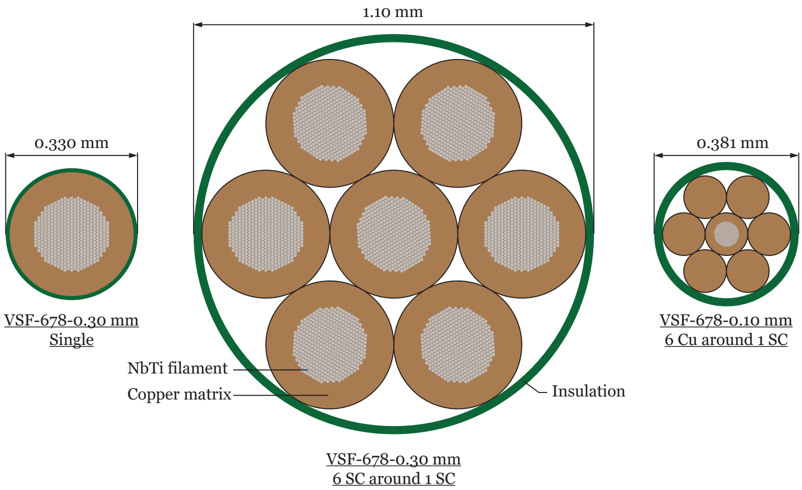

Figure 6.9: The wire configuration used for: **a)** All solenoids except the analysis coils, **b)** the octupoles, **c)** the analysis and corrector magnets. Courtesy of Chukman So.

Different wire configurations are used to wind different magnets. All solenoids (except the analysis coils) are wound using a single superconducting wire strand (figure 6.9a). A cable with 7 superconducting wire strands is used for the long and short octupoles (6.9b), and the analysis and corrector magnets are wound using a thin superconducting wire strand surrounded by six normalconducting copper strands for mechanical support (6.9c).

The superconducting wire was delivered in spools of more than 1800 m, which is long enough that each magnet could be wound from a single spool. Voltage taps are added at several places along the magnets (these will be discussed further in section 6.5). The taps are connected by removing the wire insulation and soldering/splicing a normal conducting wire to the superconducting wire.

6.3.6 Inductances

The self-inductance of a long solenoid with one layer of wires is given as

$$L_{\text{inf}} = \frac{\mu_0 N^2 \pi R^2}{l} \tag{6.1}$$

where R is the radius of the solenoid, l is the length, and N is the number of turns. For finite solenoids, equation 6.1 gives an inductance that is too large, as the field strength declines near the ends of the coil. By treating the solenoid as a thin current sheet, the inductance can be corrected by multiplying with the Nagaoka coefficient [98], $k_L \in [0, 1]$, thus obtaining Lorenz's formula for the inductance of a thin-sheet-solenoid [99]

$$L = \frac{\mu_0 N^2 \pi R^2}{l} k_L \tag{6.2}$$

with

$$k_L = \frac{4}{3\pi} \frac{1}{k'} \left(\frac{k'^2}{k^2} \left(K(k) - E(k)\right) + E(k) - k \right) \tag{6.3}$$

$$= \frac{4}{3\pi} \frac{D}{l} \left(\frac{2k^2 - 1}{k^3} E(k) + \frac{1 - k^2}{k^3} K(k) - 1 \right) \tag{6.4}$$

and

$$D = 2R \tag{6.5}$$

$$k^2 = \frac{D^2}{D^2 + l^2} \tag{6.6}$$

$$k'^2 = \frac{l^2}{D^2 + l^2} \tag{6.7}$$

$K(k)$ and $E(k)$ are the complete elliptic integrals of first and second kind. Using the average radius of the internal magnets as R, this method yields the inductances given in table 6.3. Equation 6.2 must be expected to give a poor result for magnets with a non-square cross section in the (ρ, z)-plane like the capture solenoids.

The mutual inductance, M_{ij}, relates the flux through the area enclosed by magnet i, Φ to the current in magnet j, I_j, which will induce a voltage across magnet i

$$V_i = \frac{d\Phi}{dt} = M_{ij} \frac{dI_j}{dt} \tag{6.8}$$

where $M_{ii} = L$ is the self-inductance , and $M_{ij} = M_{ji}$. As ALPHA-g contains many high inductance magnets close to each other, induced currents are expected to be significant.

Magnet	Estimated self-inductance [mH]
A, B, F, G	56
C, D, E	18
Tr, Bg	83
Capture Solenoid	1800
Analysis coil	310
Corrector coil	3.2
Babcock main	28600

Table 6.3: Self-inductances as calculated by equation 6.2. The value for the external solenoid is given by the manufacturer but included for completeness.

	Bab	**SoB**	**MAB**	**MBB**	**MGB**	**TrB**	**AnaB**
Bab	28600	116	24.8	24.8	24.8	31.7	64.6
SoB		768	2.30	0.96	0.19	0.17	0.16
MAB			51.8	7.02	0.23	0.15	0.10
MBB				51.8	0.53	0.30	0.16
MGB					51.8	9.19	1.03
TrB		Unit: mH				75.1	3.67
AnB							249

Table 6.4: Self- and mutual inductances for selected magnets. The values are derived by applying equation 6.9 to a numerical model of the magnets.

To determine the self-inductances, an alternative to equation 6.2 is to exploit the relation between the energy stored in the field produced by the magnet and the self-inductance of the magnet.

$$E = \frac{1}{2\mu_0} \int B^2 dr^3 = \frac{1}{2} L I^2 \tag{6.9}$$

where the integral is over all space. By programming a detailed magnet model (see for example chapter 8) based on the dimensions of the magnet and calculating the integral numerically, both self- and mutual inductances can be derived. Based on a magnet model, where each solenoid wire layer is treated as a finite length thin solenoid, the values in table 6.4 are obtained. The table only shows the values for a selection of magnets. The self-inductance calculated for mirrors C, D, E is 17.0 mH. For the mirrors, the values agree with the ones in table 6.3. The self- and mutual inductances in table 6.4 can be compared to the experimentally determined values in section 7.5.

6.4 Electrodes

Electrodes are needed in ALPHA-g to manipulate charged particles and form Penning traps. A sketch of the electrode stack in the bottom full strength region can be seen in figure 6.10. The electrodes cover the region from SoB to MGB. Each electrode has either a 75 or a 150 V bipolar amplifier, which sets the voltage on the electrode. Electrodes 3 and 12 are split into six segments, so they can be used for RW compression as described in section 2.4.

6.5 Magnet Control System

Besides controlling the magnets and doing general monitoring, the most important tasks of the magnet control system (MCS) are to provide quench[7] detection and protection. A simplified sketch of the MCS can be seen in figure 6.11. The elements in the figure will be described in the following subsections.

If a quench happens in a small region of the magnet, the current will heat up the region due to the non-zero resistance, which will make the surrounding region cross the temperature threshold for superconductivity. The timescale for this process is on order of ms, which means that the entire magnet can go into a normal conducting state practically instantly compared to human reaction time, so an automatic quench response system is required.

6.5.1 General Control and Monitoring

The low response time required for the monitoring system is achieved by programming the control software on the FPGAs of cRIOs (Compact Real time Input/Output systems) from NI. The model chosen for the MCS is cRIO-9039, which has a 1.91 GHz Quad-core CPU and a Kintex-7 325T FPGA. The cRIO has 8 slots for NI C-modules. The NI-9403 is used as a digital I/O module, and the NI-9205 module is used for analogue input. The minimum update time for these modules is 7 and 4 μs respectively[8]. These are the majority of the modules used in the MCS, but other modules such as 9239 (for high resolution analogue BNC input), and 9870 and 9871 (for serial communication) are used as

[7] A *quench* is when a superconducting magnet "suddenly" goes normalconducting during operation.

[8] Specifications for NI hardware are available on National Instruments' website www.ni.com

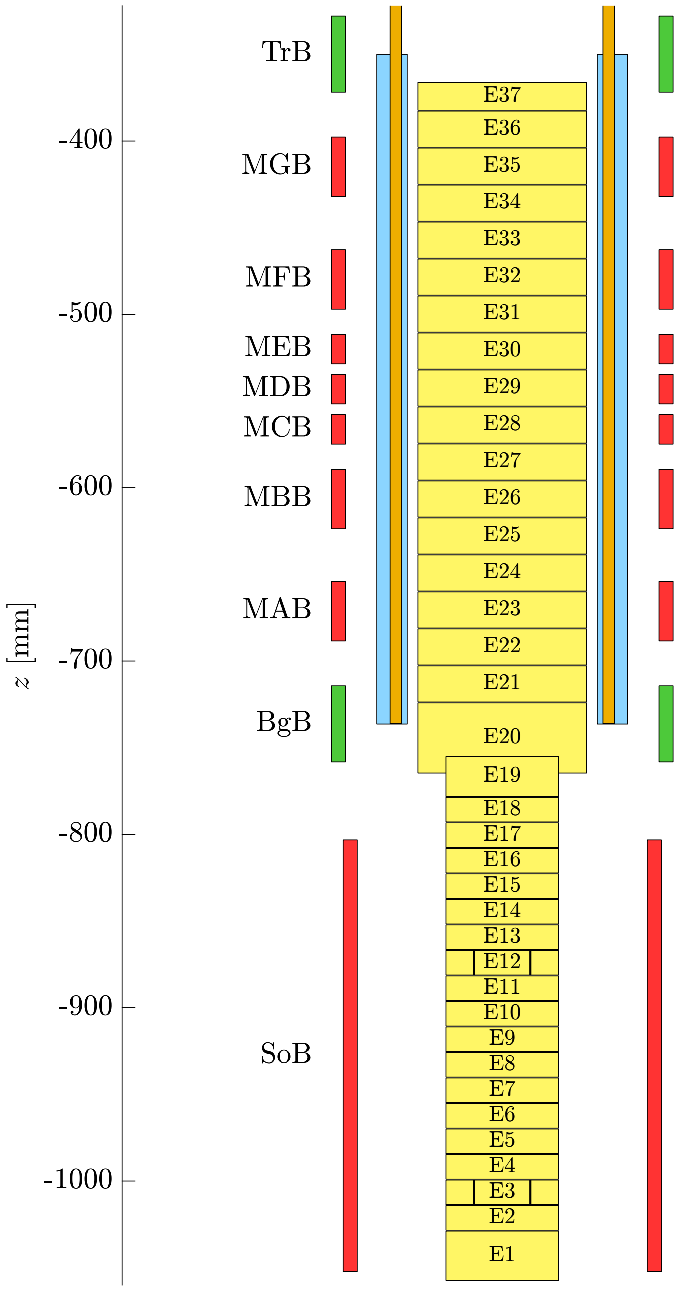

Figure 6.10: Sketch of the bottom electrode region. E20 overlaps E19, and E3 and E12 are segmented. Radial dimensions are not to scale, and some dimensions have been exaggerated to improve readability.

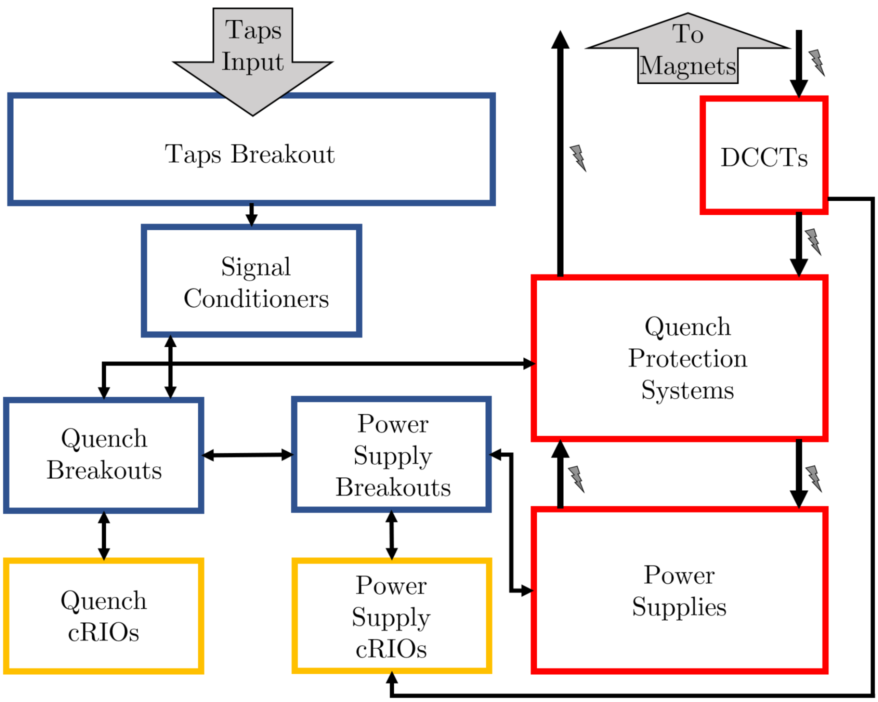

Figure 6.11: Simplified schematic of the ALPHA-g MCS. The signals from voltage taps are via the taps breakout sent into the signal conditioners. The amplified signals are via the quench breakouts sent to the QD cRIOs, which determines whether a quench has happened. The PS cRIOs control the power supplies via the PS breakout. The current generated by the PSs goes through the QPSs and DCCTs to the magnets. The two types of cRIOs can communicate quickly via their breakouts.

well. There are two different types of cRIO controllers for the MCS: a Power Supply cRIO, which mainly deals with controlling the magnets, and a Quench Detection cRIO, which mainly handles the quench detection. Once programmed, the code runs autonomously on the cRIOs.

To get a precise measurement of the current sent through the magnets, DCCTs (Direct Current Current Transducer) sit on all returning current leads. They can measure the current more precisely than the power supplies themselves, and besides being used to monitor the current, the DCCTs can be used in a feedback loop to control the current through the magnet. The details and performance of the DCCTs in ALPHA-g will be described in section 6.8.

To allow us to use the DCCT reading in a feedback loop, we need to know whether the DCCT is in a good state and outputs a sensible reading. Each DCCT outputs status signals, which we feed into an DFRduino Nano. The Arduino converts the inputs to a single digital signal, which tells whether the DCCT is in a good state or not. This signal is fed into the cRIO to be used as an interlock. The system also interlocks on the helium level in the cryostat, various temperatures, the operating status of various parts of the system, etc.

6.5.2 Quench Detection

To detect quenches, the voltage is measured at multiple points along the magnet and its current leads. An example of the locations of the voltage taps can be seen in figure 6.12, where each voltage tap is indicated with an arrow. For every magnet, the current leads consist of a few sections made with different wire types: the vapour cooled lead (VCL), the high temperature superconductor (HTS), and two low temperature superconducting leads (LTS AMI and LTS ALPHA). On each side of a joint between two wire types, there is a tap. The magnet itself has a voltage tap at each end and one or more voltage taps around the centre.

Nominally, only a few of the voltage taps are used for quench detection. The voltage drops across the magnets and their leads are divided into three categories. These are referred to as the VCL, SCL, and LTS drops:

- VCL drop: The normalconducting section of the leads. In figure 6.12, it corresponds to the VCL.

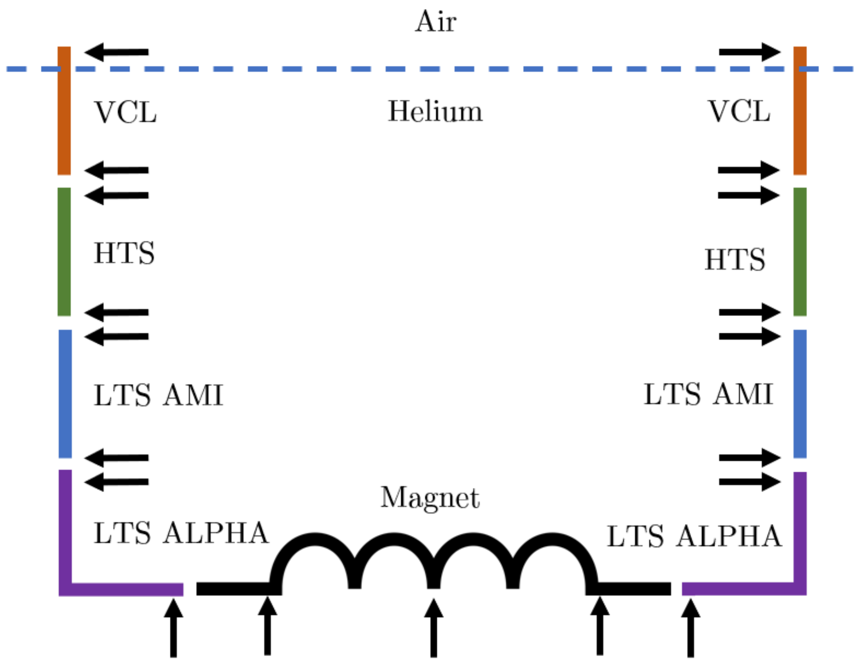

Figure 6.12: Sketch of current leads and the voltage taps for a mirror coil. Different lead sections are indicated in different colours, and the magnet is represented by an inductor. The voltage taps, which are indicated with arrows, are placed along the current leads and the magnet to detect quenches. The dashed blue line indicates what is interior and exterior to the cryostat.

- SCL drop: The superconducting section of the leads. In figure 6.12, it corresponds to the total voltage across HTS, LTS AMI, and LTS ALPHA.
- LTS drop: The inductive part of the system. In figure 6.12, it corresponds to the magnet. For historical reasons, LTS is used to describe the inductive part, although it is unrelated to the LTS sections of the leads.

Often, it is more practical to describe the magnet and its leads with these terms, than the ones in figure 6.12. When the magnet is superconducting, the voltage drops across these sections should be

$$V_{VCL} = IR \tag{6.10}$$

$$V_{SCL} = 0 \tag{6.11}$$

$$V_{LTS} = L\frac{dI}{dt} \tag{6.12}$$

By monitoring these voltages, it is easy to detect a quench on the SCL drop, as one just needs to set a voltage threshold comfortably above the noise level. If the lead goes normalconducting, the resulting resistive voltage drop will quickly exceed the threshold. Similarly for the VCL drop, a threshold that is comfortably above the voltage drop caused by the maximal current is set. Violation of these thresholds will rarely be the primary trigger for a quench response, but if violated, the system is certainly in a bad state.

As the voltage induced on the LTS drop by changing currents can easily exceed the voltage caused by a small section of the magnet going normalconducting, a simple voltage threshold will not work as a method for detecting quenches on the magnet itself. Instead, we measure the voltage drops across the two halves of the magnet and do a weighted subtraction

$$V_{sum} = c_1 V_{LTS,1} - c_2 V_{LTS,2} \tag{6.13}$$

$$= (c_1 L_1 - c_2 L_2) \frac{dI}{dt} \tag{6.14}$$

where c_1 and c_2 are positive constants that can be tuned to make $V_{sum} = 0$ during normal operation. Note that the inductances of the two halves are not necessarily equal. If one of the two halves quenches and becomes resistive, the two terms will no longer cancel, V_{sum} will cross a threshold, and the quench is detected.

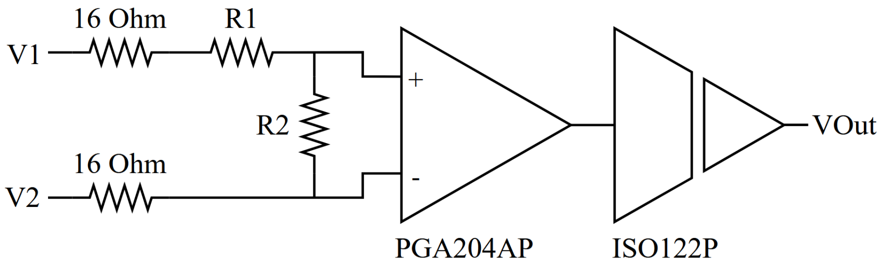

Figure 6.13: A simplified circuit diagram for a signal conditioner channel. The input voltage, V_1 and V_2, are amplified by the programmable amplifier (PGA204AP), and the signal is sent through an isolated amplifier (ISO122P).

The voltage taps are wired out from the cryostat, through a breakout board, and into a signal conditioner (SC). The SCs are custom built electronic circuits, which make a differential measurement between two voltage taps and amplify and filter the signals. The SCs amplify the signal by a gain of 1, 10, 100, or 1000 by a programmable amplifier, before the signal is sent through an isolated amplifier to reference the signal to a common ground. For some channels, the gain is set manually using switches and is therefore practically fixed, whereas for other channels the gain is varied programmatically.

A simplified diagram of an SC channel is shown in figure 6.13. A more detailed circuit diagram can be found appendix A.2. Each channel has the option to install a voltage divider by choosing values for R1 and R2. By default, $R_1 = 0$ and $R_2 = \infty$. A voltage divider will essentially divide the voltage drop by a given factor, which can be useful for two reasons; firstly, the isolated amplifier saturates at about 10 V, so the voltages presented to the amplifier during normal operation should not exceed this limit. The second reason has to do with the data treatment on the cRIO FPGA and will be elaborated later in this section.

The amplified signals are sent from the SC to the cRIO via the QD breakout. The cRIO C-module, which has an input range of ± 10 V, converts the voltage to an 16-bit integer (I16), which is read by the FPGA. The software sums N voltage readings[9], $V_{raw,i}$, to make an average, although the sum is not divided by N, as division is a computationally expensive operation on the FPGA, and it is thus

[9] N is typically between 1 and 5.

avoided. The sample time[10] with the current code is around 140 μs. The average is then multiplied by $1000/g$, where g is the gain applied in the SC to the given measurement. In summary

$$V_{avg} = \frac{1000}{g}\sum_{i=0}^{N} V_{raw,i} \tag{6.15}$$

V_{avg} is stored as an I16, but for large voltage drops, V_{avg} can exceed the size of its data type. Installing a voltage divider in the SC can help this issue as described above. Another future solution is to extend the data type to I32. The conversion from V_{avg} to a voltage drop across the magnet or lead, V_{actual}, is

$$V_{actual} = \frac{10}{2^{(16-1)}}\frac{V_{avg}}{N}\frac{1}{1000} \tag{6.16}$$

6.5.3 Quench Protection

In ALPHA-g, we operate our magnets at tens to hundreds of amps. The energy dissipated by such currents would melt the magnet wire, if it was normalconducting, so in case of a quench, the current must be ramped down quickly. Quick ramp downs represent a problem for the power supplies connected to the magnets, as the high inductances of the coils mean that a lot of energy would be deposited in the power supply over a short period of time. Quick ramp downs also generate large inductive voltage drops, which can potentially create current arcs.

A quench protection system (QPS), which sits between the power supply and the magnet, addresses both the above issues. An example of a QPS can be seen in figure 6.14. When a quench is detected, the QPS places a short across the power supply within 100 μs, so most of the current from the power supply will not go through the magnet, and the energy stored in the magnet will not be deposited in the power supply. The power supply is disabled as well. The QPS also breaks the circuit that connects the power supply to the magnet. The QPS installed for the solenoids are commercial QPS units from Applied Power Systems (APS), while the high current QPS for the octupoles is designed and constructed at ALPHA, based on the design of the ALPHA-2 QPS.

The core components of the QPS are the IGBT (Insulated Gate Bipolar Transistor) and the SCR (Silicon-Controlled Rectifier). A simplified

[10]The time is takes to read and process a value of $V_{raw,i}$.

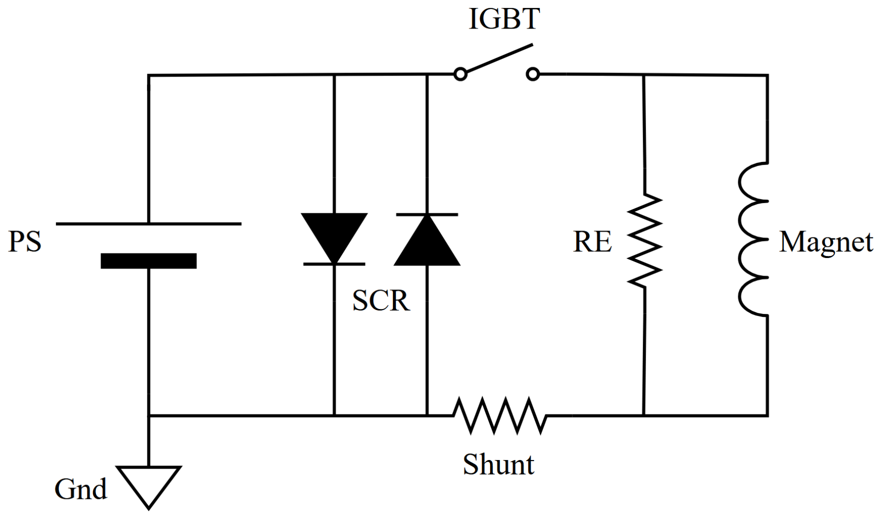

Figure 6.14: Simplified circuit diagram for a bipolar QPS. When the magnet is operating, the IGBT is closed, and the SCRs are open. In case of a quench, the SCRs close and short out the magnet, and the IGBT opens and forces the current to go through the energy extraction resistor.

circuit diagram of a QPS circuit can be seen in figure 6.14. The main difference between the octupole QPSs and the APS QPSs is that the APS QPS is bipolar, like the one in figure 6.14, whereas the octupole QPS is unipolar. Both the SCR and the IGBT work as switches controlled by an external trigger voltage. The normal state (when the magnet is operating) for the IGBT is closed, and for the SCR the normal state is open.

When the QPS triggers, the SCR closes, and the power supply is thereby shorted out of the circuit. For the APS QPSs, the response time of the SCR has been measured to be on the order of tens of μs. Hundreds of μs later, the IGBT is opened, which causes the current in the circuit to flow through the energy extraction resistor, RE. The IGBT trigger is optically isolated from the rest of the system.

In case of a quench, current flows through the magnet and the energy extraction resistor in a closed circuit. Let R_M denote the resistance of the magnet itself, and R_E be the resistance of the rest of the circuit, which is practically the resistance of the energy extraction resistor.

The current in the circuit is then given as

$$I(t) = I_{max} e^{-\frac{R_E + R_M}{L} t} \tag{6.17}$$

where I_{max} is the current that runs in the magnet before the QPS is triggered. Hence, the voltage drop across the magnet is

$$\begin{aligned} V_M(t) &= I(t) R_M + L \frac{dI(t)}{dt} && (6.18) \\ &= I(t) R_M - (R_E + R_M) I(t) && (6.19) \\ &= -R_E I_{max} e^{-\frac{R_E + R_M}{L} t} && (6.20) \end{aligned}$$

with the maximum voltage at any time being given as $R_E I_{max}$. For each magnet, R_E is chosen such that the voltage will not exceed 400 V, which is the breakdown voltage of the electric insulation, in case of a quench. The lower limit for the value of R_E is set by the need to extract energy from the magnet as quickly as possible, so it does not deposit as heat in the magnet itself. For the internal solenoids, R_E is chosen to be around 2–8 Ω, and for the octupoles it is 0.3 Ω.

In case of a quench, there is a risk that the temperature of the magnet wire will exceed its melting point. The manufacturers specify a 300 and 160 K temperature limit, for the external and the internal magnets respectively. To distribute the deposited energy over a larger area, the quench must spread as quickly as possible. The speed of the quench spread depends on the thermal properties of the magnet wire and the magnet form.

For most of the ALPHA-g magnets, the quench spreads fast enough on its own, but the capture solenoids, the external solenoid and the analysis coils have installed a quench heater to help accelerate the quench. The quench heaters are capacitor banks, which will discharge to a metal strip wound around the surface of the magnet when triggered. For the solenoids and the analysis coils, the strip is made of steel or BeCu respectively. The metal will heat up as the capacitor discharges, and the thermal contact between the strip and the magnet will heat up the magnet and help spread the quench.

6.6 The External Solenoid

The external solenoid submerges the entire ALPHA-g experiment in a 1 T field at 191 A. It was manufactured by Babcock Noell GmbH[11] and is therefore also referred to as the Babcock magnet. Besides the main coil, there is an independent shim coil, which was added after the construction of the main coil. The shim coil is intended to smooth out the field produced by the imperfect main solenoid. Its position was determined based on measurements of the field of the main coil.

The superconducting magnet is contained in a cylindrical shaped vacuum vessel, which also houses the cryocoolers, and the thermal insulation (see figure 6.15). The bore of the vessel is empty to allow for space for the detector, and the internal magnets and their cryostat. It is about 3 m tall, 1.5 m in diameter, and weighs about 2500 kg. The inner diameter is about 0.5 m.

The main coil consists of multiple sections wired in series. These sections are illustrated in figure 6.16, where the voltage taps used for quench detection are shown as well. The sections in figure 6.16 are referred to as main, outer, and boost/shim coil. These will be referred to collectively as the main coil, unless otherwise specified. To smooth out the field produced by the main coil, a separate superconducting shim coil is added outside the main coil. *The shim coil* usually refers to this external shim coil (as opposed to referring to the shim section of the main coil).

The control system for the external solenoid is implemented on a single cRIO of the type 9049. It has similar specifications and uses the same modules as the internal MCS. In general, the control system has many similarities to the control system for the internal magnets. We use Labview's Configuration Editor Framework (CEF) to generate a configuration file containing information relevant for the software about all the hardware components of the system.

6.6.1 Quench Detection

As seen in figure 6.16, the Babcock has nine voltage taps distributed along the main coil. The shim coil does not need quench protection, so it is irrelevant for description of the QD and QPS. Similarly to how quench detection is done for the BNL magnets described in section 6.5,

[11] Now Bilfinger Noell GmbH.

Figure 6.15: The Babcock vacuum vessel, which houses the main and the shim coil, the cryocoolers, and the thermal insulation. The vessel has a hollow bore to allow the detector and the internal magnets to fit inside. Credit: Bilfinger Noell GmbH.

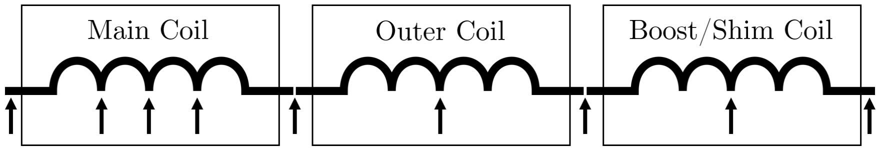

Figure 6.16: The main coil is divided in three sections in series: main, outer, and boost/shim. The voltage taps used for quench detection are indicated with arrows.

these voltage taps are used to detect quenches. The magnet was delivered with a control system designed by Babcock and its subcontractors, but we have had to do major interventions on the system, so some details are given here. The quench detection hardware specifically has been designed and built by Danfysik, and will therefore be referred to as the Danfysik QD system.

A subset of the available voltage taps are used as input for the Danfysik QD system, where they are paired up, so each half of the main, outer, and shim/boost section are monitored by the QD system. A pair of voltage drops are called a channel, so in total there are three channels in use. The circuit diagram for a single channel can be found in appendix A.2. Using V_1 and V_2 to denote the upper and lower drop respectively, a potentiometer is used to make a weighted sum of the two signals

$$V_{sum} = f \cdot V_1 + (1 - f)(-V_2) \tag{6.21}$$

where $f \in [0, 1]$. Another potentiometer and an amplifier is used to amplify V_{sum}

$$V_{amp} = 256g\,(f \cdot V_1 + (1 - f)(-V_2)) \tag{6.22}$$

where $g \in [0, 1]$. If V_{amp} is bigger than some threshold[12] for more than 50 ms, the Danfysik QD sends out a digital signal, which is used to trigger the quench protection system. Both voltage and time threshold values can be adjusted in the hardware. The simplicity of the Danfysik QD system has certain advantages, but the QD system is not advanced enough to deal with the problems with mutual inductance, which will be discussed in section 7.4.

[12]The threshold has been measured to be 0.71(2) V.

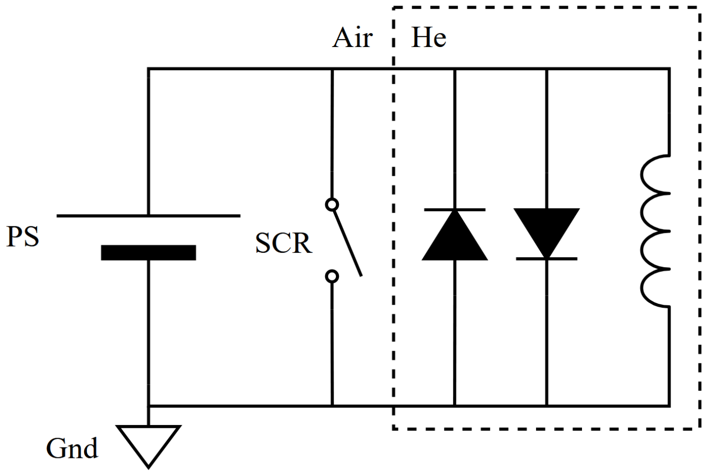

Figure 6.17: A simplified circuit diagram for the main coil, its quench protection system and its power supply. If a quench is detected, the SCR (indicated by the switch) will close. If there is a large voltage drop across the magnet, the diodes will start to conduct and current share with the magnet.

6.6.2 Quench Protection

A simplified diagram of the magnet circuit with its QPS can be seen in figure 6.17. As-delivered, the Babcock control system will disable the power supply, close the SCR, and fire a quench heater, if a quench is detected. Besides this active quench protection, the main coil is equipped with over-voltage protection diodes, which sit in parallel to the magnet and will start conducting, if the voltage across the magnet gets too high – for example as a result of a quench.

Similar to the internal magnets, the quench heater for the Babcock is a capacitor bank, which discharges into a resistive load that heats up the magnet. The heater was manufactured by Babcock subcontractor OCEM. The voltage outputted by the quench heater as a function of time is given as

$$V(t) = V_0 e^{-t/RC} \tag{6.23}$$

where $R \approx 0.4\ \Omega$ is the resistance of the external circuit, $C = 490$ mF is the capacitance of the heater, and $V_0 = 56$ V is the voltage applied to charge the capacitor. The total energy stored in the capacitors is

$$E = \frac{1}{2} C V_0^2 = 770 \text{ J} \tag{6.24}$$

According to Babcock, the critical time window for depositing energy into the magnet is 100 ms after a quench. In that time window, the energy deposited in the magnet is

$$E = -\frac{1}{R}\int_0^{100\ \mathrm{ms}} V(t)^2 dt = 491\ \mathrm{J} \tag{6.25}$$

As it unfortunately became relevant to temporarily replace the heater due to a failure in its system, we explored the option to use power supplies to deliver the required energy on the required time scale, but we found that is was not feasible due to slow ramp times of the power supplies.

When the magnet operates at full current, the energy deposited by the heater is enough to kick-start the spread of the quench, as it will make some of the magnet resistive, and the current running in the magnet will spread the quench to the rest of the magnet. However, at lower current[13], the cooling from the cryocoolers is enough to bring the magnet back to its superconducting state, but because the QPS triggered, the SCR is now closed, and the power supply is disabled. Consequently, the current is left free-running in the magnet, and without an energy extraction resistor, it will take a long time for the current to decay.

In the future, we are likely to replace the SCR with a mechanical relay. In the scenario where the magnet is quickly restored to its superconducting state after a detected quench, a relay would allow us to reintroduce the power supply to the magnet circuit, after making the power supply output match the current flowing in the magnet.

6.6.3 Cryocoolers

The Babcock is equipped with two pulse tube cryocoolers from Sumitomo Heavy Industries of the type RP-182B2S. The working principles of pulse tube cryocoolers are described elsewhere [100, 101, 102]. The system consists of a compressor (the F-100 model), a valve unit, a buffer volume, and a cold head with two temperature stages.

The compressor connects to the valve unit with 20 m helium hoses, and the valve unit connects to the cold head, which is also connected to the buffer volume as shown in figure 6.18. The pulsed tube cryocoolers have a closed helium system, so once filled, no helium needs to be

[13] At about half of the full current.

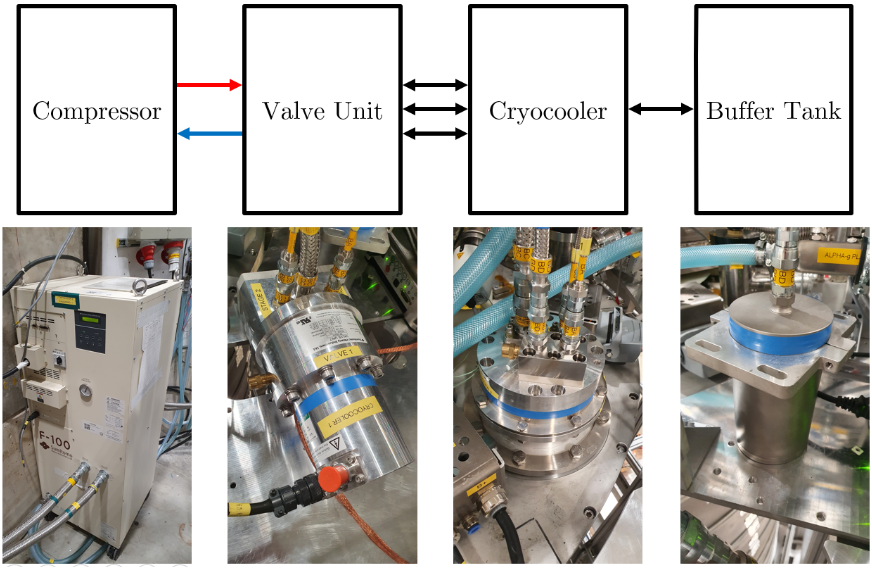

Figure 6.18: Sketch of the cryocooler system from Sumitomo. Below each part of the system is a picture of the component. The arrows illustrate the connections between the different parts.

added. The specified minimum temperature at stage 2 is less than 2.6 K, and the cooling power at stage 1 is 36 W at 48 K, and at stage 2 it is 1.5 W at 4.2 K.

According to Sumitomo, the helium supply pressure in the compressor should be 22.0–23.5 bar and it should not be changed after the cooldown is initiated. An interlock will stop the compressor, if the supply pressure falls outside 19–26 bar. We found that the decrease in pressure as a consequence of the cooldown brings the supply pressure close to 20 bar. Furthermore, we found that the cooling power increases significantly with increasing helium supply pressure, so at 4 K we typically run with a pressure of 22–25 bar. The effect of the pressure on the cooling power will be discussed further in section 7.2.

6.6.4 The Shim Coil

The main coil was delivered to ALPHA in 2018 before the installation of the shim coil. The superconducting shim coil was later installed at Babcock based on the field maps done at ALPHA. The effect of the shim coil on the total field will be discussed in section 7.3.

The shim coil consists of two identical coils with opposite polarity in series. Each coil has 4 layers with 170 turns per layer. The coils are moveable with respect to each other and with respect to the main coil. They are mounted inside the thermal insulation of the main coil and connected to the 4 K stage of the cryocoolers. The shim coil is designed to run at about 5 A and bring the field errors of the main coil from 10s of gauss to a few gauss (see section 7.3). At this level of field homogeneity, the field from magnetised surroundings is estimated to become significant.

The shim coil adds to the heat load of the system in two ways: since there is now more material connecting the 4 K stage to the room temperature environment, the static heat load is increased, and since the normalconducting leads will generate heat, when the coil is energised, the dynamic heat load is increased as well. In the first installation of the shim coil, the current leads were too thick, which meant that the static heat load was too big for the cryocoolers to handle. Exchanging the lead with a thinner version brought down the static heat load but increased the resistance and thereby the dynamic load. However, the total heat load is manageable for the cryocoolers.

6.7 Tickle Power Supply

As described in section 6.2.1, the magnetic field under mirror A and G for the up/down measurement must be known and controlled to better than $3.20 \cdot 10^{-4}$ T. With the field per ampere produced by mirrors A and G and the analysis coils (see table 6.1), the requirement translates to controlling the current to better than 30.2 mA for the up/down measurement and 0.227 mA for the 1% measurement. To achieve this performance, we could add a precise low-current power supply in parallel with the CAENels, and use the low-current power supply to stabilise the current based on feedback measurements by the DCCT. Another solution is to run mirror A and G is series and use the low-current power supply to apply the difference in current used to offset the gravitational potential. We refer to the low-current power supply as the tickle supply.

The power supply intended to be used as the tickle power supply is the KEPCO BOP 20-10 model. It is a 20 V, 10 A, four-quadrant linear power supply that can be controlled remotely. The options for remote control are analogue control and 16-bit ethernet or serial control. With the 16 bit control, the quoted accuracy is 0.3 mA for low current and 5 mA for high currents[14]. The accuracy is likely to be different when connected to a large inductive load, so it would have to be tested with the relevant magnets, before a conclusion can be made about the performance. In the case of mirrors A and G, 0.3 mA translates to 3.2 μT in the centre of the magnet. The model might be suitable as a power supply for the analysis coils as well. For the analysis coils, 0.3 mA translates to 6.7 μT.

Due to the high inductance of the mirrors and especially the analysis coils, a QPS is needed to protect the tickle supplies in case of a quench. As the current output of the tickle supply will be small, it is not a concern that the supply could damage a normal conducting magnet. The commercial APS QPSs used for the internal magnets are bipolar, but the tickle current will oscillate around 0 A, and since the IGBTs in the APS units are non-linear in the crossover region, they are unsuited for the tickle supplies.

As a QPS for the tickle supply, we tested our own circuit, which was designed at UC Berkeley. The circuit consists of a current- and a voltage-limiting board. We have tested a prototype version of the current-limiting board, whose key component is a MOSFET. The MOSFET works as a passive current limiter, and it is linear is the crossover region and thus suitable for the tickle circuit.

The test of the tickle circuit board was done with a CAENels 100 A power supply with an APS QPS connected to a 40 mH, 25 mΩ normal-conducting dipole magnet. The tickle QPS was connected across half of the dipole to simulate a scenario, where the tickle supply is used to apply the difference between mirrors A and G. We demonstrated that practically zero current was sinked into the tickle power supply. The design thus seemed promising, and we expect to use it in the final setup.

6.8 DCCT Precision

To precisely measure the current flowing through the magnets, we use DCCTs. The DCCTs used for ALPHA-g are the TOPACC zero-flux current measuring system from PM Special Measuring Systems. Zero-flux type DCCTs operate on the following principle [103] (a schematic

[14]The distinction between high and low current is in this case made at 2.5 A.

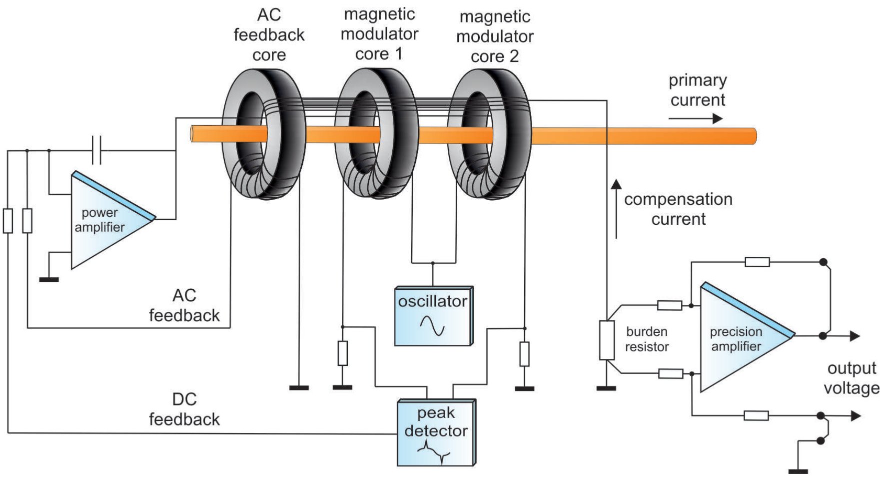

Figure 6.19: A schematic of a zero-flux DCCT system. The primary current is measured by cancelling the magnetic flux in the toroidal cores with a secondary current, which is measured to give an output signal. From [104]. © 2014 IEEE.

of the DCCT system can be seen in figure 6.19): the primary current, I_p, which is the current to be measured, runs through three magnetisable toroidal cores. Each core also has a secondary winding and an auxiliary winding. The aim is to offset the flux induced in the cores by the primary current with a current in the secondary winding, I_s.

The secondary current runs through the burden resistor, which has a well-known resistance. The generated voltage drop across the burden resistor is amplified by the precision amplifier, whose output will be proportional to I_p. The toroids are contained within a measuring head, while the rest is contained in a chassis that connects to the head via a cable.

The leftmost core in figure 6.19 is used to determine the AC component of I_p. Any change in I_p will induce a change in flux and thereby a current in the toroid, which in turn induces a current in the auxiliary winding. The auxiliary winding drives the power amplifier, which generates the secondary current.

To detect the DC component of I_p, two magnetic modulator cores are used. The auxiliary winding of both cores are connected to an oscillating voltage source, which outputs a current, I_{aux}, which magnetically saturates the toroids. As the auxiliary winding of one core is wound in the opposite direction of the other, the magnetisations of the two cores will have the opposite phase. The change in magnetisation determines how much current is re-induced in the axillary winding. If there is a DC component in the primary current, it will be seen as a constant offset to the magnetisation of the core, and I_{aux} will change the magnetisation by different amounts depending on its phase. The peak detector in figure 6.19 will therefore see a difference in the amplitude of the signals proportional to the DC component of I_p. The peak detector drives a power amplifier to add a DC component to I_s to counter I_p.

The DCCTs used to measure the currents through the magnets, must be able to match the precision required of the magnetic fields derived in section 6.2.1. The requirements translate to a current precision of better than 30.2 mA for the up/down measurement and better than 0.227 mA for the 1% measurement. The TOPACC model is specified to have a relative accuracy of less than 25 ppm and an offset of less than 2.5 ppm of its full current range. To help fulfil this, each DCCT chassis is calibrated to match a specific DCCT head from the manufacturer.

The TOPACC model can be adjusted to measure I_p in ranges from 0–100 to 0–1100 A. In the 0–100 A range, a 25 ppm relative accuracy translates to 2.5 mA accuracy, which is sufficient for the up/down measurement. For the 1% measurement, the requirement on the precision could be relaxed by looping the current lead through the DCCT head multiple times. As the analysis coils are expected to operate at a maximum of 7.8 A, one could make 12 loops through the DCCT head to better exploit the 0–100 A measuring range. 12 loops would relax the required precision to 2.72 mA, which can be fulfilled by the DCCT.

To measure the performance of the DCCTs, we approached the High Precision Measurements section at CERN. A total of 4 sets of DCCT chassis and measuring heads were tested (two matching and two mismatching sets). The DCCT relative accuracy was measured to be less than 20 ppm, and the constant offset was less than 4 ppm. As well as characterising different types of drift and noise in different frequency ranges, other conclusions from the study were:

- No recalibration of the DCCTs was required at the time of the measurements. It is recommended to characterise the performance again in two years.

- The results show a low dependence on head-chassis pairing.
- The DCCT output drifts on the 1 ppm level within the first 30 minutes of being powered. After 30 minutes, there is no further drift.
- The accuracy dependence of the range setting of the DCCT is on the 1 ppm level.

6.9 KEPCO and DCCT Tests

In addition to the DCCT tests described in section 6.8, we performed a test of a DCCT in a setup with a KEPCO power supply and a cRIO. This is similar to the setup planned for the tickle power supplies and potentially the analysis coils. The purpose of the test was to characterise the background signal measured by a DCCT in this setup. Knowledge of the background is important for monitoring the current and especially in relation to using the DCCT measurement as a feedback signal for the power supply current control.

The test setup consisted of a KEPCO BOP 20-10 power supply, which was shorted through a copper wire of a few meters. The DCCT head was placed around the wire approximately at the middle point of the wire. The DCCT itself was connected to a NI 9239 C-module in a cRIO-9039, which was grounded to mains ground. Both module and cRIO were of the types used in the MCS as described in section 6.5.1. Different cable types were tested to connect the cRIO to the DCCT. The KEPCO was in analogue control mode and driven by a "Wavefactory WF 1946 2CH" waveform generator.

The measurement scheme was as follows: use the waveform generator to make the KEPCO output a 1 A current through the DCCT, sample the DCCT output for 10 seconds, and then do a single-sided FFT (not a power spectrum) of the data set. The spectrum would reveal whether the setup was prone to pick up noise at any specific frequencies, which could potentially be filtered out before interpreting the current measurement and responding to it. The TOPACC DCCT is specified to have a bandwidth of 500 kHz, but the FPGA in the cRIO samples at 50 kHz at maximum, which translates to a Nyquist frequency of 25 kHz, so the recorded spectrum is in the 0–25 kHz range.

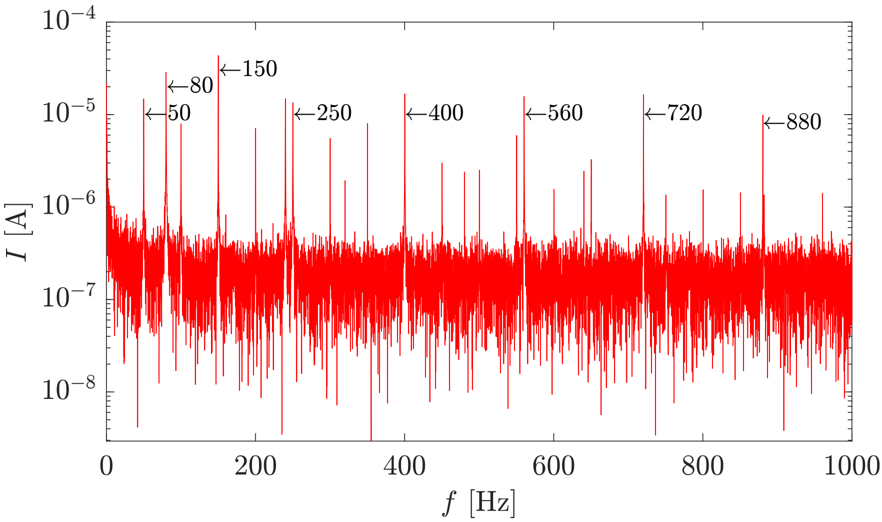

Figure 6.20: An FFT of the signal from the DCCT read by the cRIO. The spectrum is dominated by the 50 and 80 Hz peaks and their higher order peaks.

Figure 6.20 shows a typical FFT spectrum[15] in the 0–1 kHz region, where the most intense peaks are. It was recorded at a location less electrically noisy than the ALPHA experimental zone. The spectrum is dominated by the 50 and 80 Hz peaks and their higher order peaks. By carefully reconstructing the setup step by step, we were able to identify the KEPCO power supply as the source of the peaks associated with the 50 Hz signal. Hence, the 50 Hz signal would couple into the magnet and would be seen as noise on the magnetic field. On the other hand, the 80 Hz peak and its higher order peaks proved to be picked up somewhere in the DCCT circuit and are therefore not coupling into the magnet.

It is recommended by the manufacturer to read the output from the DCCT chassis with a 4-wire measurement. However, this proved more noisy than a two-wire measurement, which is what is recommended by the High Precision Measurements section at CERN. Other conclusions from the tests are:

- Powering the cRIO with a switch mode power supply does not introduce any noise on the input signals compared to a linear power supply.

[15]The 0 Hz/the DC component is removed from the figure.

- The noise on the input signal is practically independent on the position of the C-module. There is a negligible difference between slots 1–4 and 5–8.

- Shorting the unused BNC inputs on the C-module makes no difference for the noise on the input signal.

- The NI 9239 C-module applies a low pass filter at the Nyquist frequency. If the Nyquist frequency changes as a result of a change to a different sampling frequency, the low pass filter changes correspondingly.

- Something in the tested system induces noise at a frequency of about 6 kHz. It disappears after about 10 minutes of warm-up time.

- Referencing the COM of the cRIO to ground through a resistance (of 47 Ω) does not have an impact on the noise on the input signals compared to referencing it to ground through 0 Ω.

- Adding a 35 mH inductance between the terminals of the power supply (to simulate the presence of a magnet) did not reduce the induced noise on the signal read by the cRIO.

6.10 Magnetometry

To measure the magnetic field in ALPHA-g, there is a grid of NMR (Nuclear Magnetic Resonance) and Hall probes outside the magnet form and in the detector. A detailed description of these can be found in [105]. The probes allow measurements of the magnetic field simultaneously with a gravity measurement.

The magnets most important for gravity measurements are MAB, MGB, AnB, AnT, MGT, and MAT (see figure 6.5), so these all have an NMR probe glued onto their outside. There are also 10 NMR probes distributed throughout the body of the rTPC. Every NMR probe is paired with a Hall probe, which is less accurate but allows a higher readout rate. Of course, none of these probes measure the field seen by the particles in the trap, but the plan is to map the field measured by the NMR and the Hall probes to the field measured by ECR at specific locations in the trap. Note that ECR can not be done while ramping the magnets, as the field must be stable while the measurement is ongoing. There is also the risk that the

Material	Temperature [K]	$\gamma/2\pi$ [MHz/T]
Aluminium	4	11.112007(30)
Aluminium	300	11.112316(24)
Rubber	300	42.576490(88)

Table 6.5: Gyromagnetic ratios at relevant temperatures for aluminium and rubber. From [105].

microwaves used for the ECR measurements will induce positron spin flips in the antihydrogen and thereby bring it into an untrappable state.

NMR is based on having an ensemble of spins with a net magnetisation in a magnetic field. The ensemble is placed next to a small coil, which can be used to cause spin transitions by applying a resonant radio frequency pulse and thereby rotate the net magnetisation. The magnetisation will then precess back to equilibrium, which induces a signal in the coil. The relation between the spin precession frequency, f_0, and the external magnetic field, B, is given by

$$f_0 = \frac{\gamma}{2\pi} B \tag{6.26}$$

where γ is the gyromagnetic ratio, which is material and temperature dependent. Two different materials are used for the NMR probes in ALPHA-g. Rubber is used for the probes in the rTPC, which sit at 300 K, and aluminium is used for the probes that sit in the 4 K environment. Table 6.5 shows γ for aluminium and rubber for relevant temperatures. The relaxation time of the sample, which is what determines the readout rate, is about 0.5 s for aluminium at 4 K.

As the NMR probes are placed outside the internal magnets, they operate in a gradient magnetic field, which widens the transition linewidth and can make NMR spectroscopy impossible. To minimise the field gradient across the sample, the sample geometry is optimised for the field shapes. The expected accuracy of the NMR probes in ALPHA-g is 10 ppm, and the theoretical absolute precision for the rubber and aluminium probes are 270 nT and 31 μT respectively, and 34 μT precision has been demonstrated for the aluminium probes [105]. In ALPHA-g, the stability of the field of the external solenoid is the limiting factor to the precision.

Chapter 7

ALPHA-g Commissioning

The ALPHA-g apparatus, which is dedicated to measuring antigravity, is the newest addition to the ALPHA experiment. The construction of ALPHA-g began in the summer of 2018. The aim was to construct the experiment and have it operational before LS2 began in November. The construction was mostly successful; the results will be described in section 7.1. For most of LS2, the Babcock and BNL magnets were at the manufacturers as described in chapter 6. The magnets returned to CERN in spring 2021, with antiprotons being available from August.

Besides a brief description of the 2018 run, this chapter presents the status of the ALPHA-g magnet system as of spring 2022 as well as the particle work done in the 2021 run and partially in early 2022.

7.1 The 2018 Run

In 2018, ALPHA-g was assembled with the hardware and software required to trap antihydrogen in the bottom full strength region. About 60% of the total magnet control system was constructed and commissioned in order to achieve this. During the run, we managed to successfully operate the beamline in the ALPHA-g configuration, and to operate the external solenoid, the detector, the electrodes, and various diagnostics. This allowed us to trap antiprotons and positrons in ALPHA-g by the end of the run. Figure 7.1 shows the number of antiproton annihilation vertices reconstructed by the detector as a function of the z position, where the antiprotons were held in different potentials generated by the electrodes and annihilated against the background gas. The annihilation distributions are seen to be centred around the electrodes that generated the minimum of the potential.

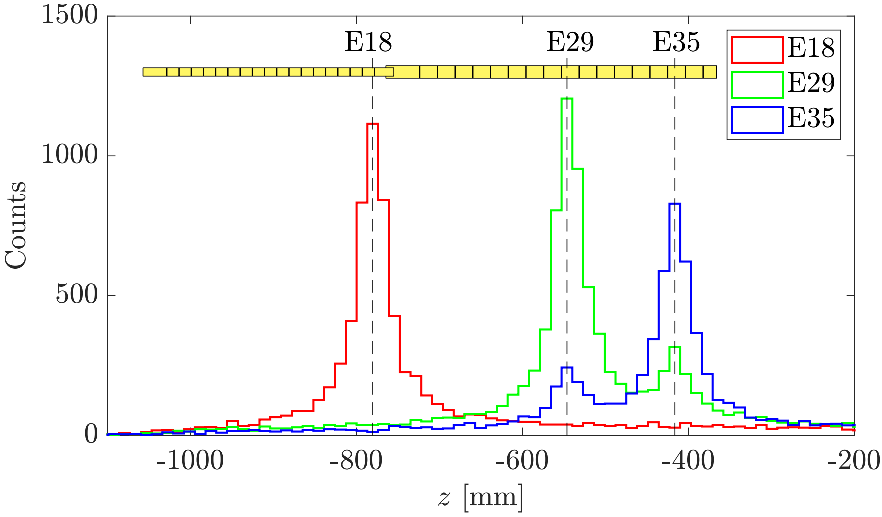

Figure 7.1: Number of antiproton annihilation vertices as a function of z-position to demonstrate antiproton trapping. The antiprotons were held in potentials at three different positions, before the antiprotons were allowed to escape and annihilate on the electrodes.

Unfortunately, the cryostat for the BNL magnets had a major leak from the liquid helium space, which meant that we were unable to keep the magnets at 4 K without an unsustainable helium consumption. Hence, the atom trap was not fully operational during the run, but the circumstances allowed us to commission MAB, MDB, and LOc.

7.2 Babcock Cryogenics

When the Babcock magnet returned to CERN in 2021 after having its shim coil installed, we found that the cryocoolers were unable to bring the magnet to 4 K. This was partially due to the static heat load introduced by the shim coil, which was later reduced. Furthermore, the current leads into the magnet had corroded while they were disconnected, which meant that the resistance across the connection had increased. The generated heat had a significant impact on the temperature increase during operation, which underlines the marginal cooling situation of the system. By polishing the connections, we reduced the resistance from the 1 mΩ to the 10 $\mu\Omega$ level, which significantly helped the thermal situation.

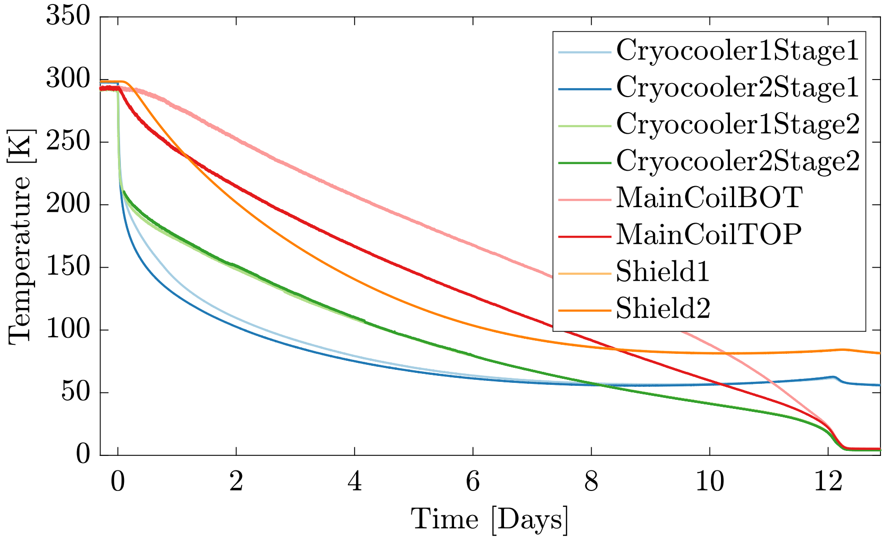

Figure 7.2: The temperatures of the cryocooler stages, the top and bottom of the main section of the coil, and the heat shields during a cooldown. It takes 12 days for the magnet to go from room temperature to its minimum temperature at 5–6 K.

The main variable parameter determining the cooling power is the helium supply pressure in the compressors. We found that we have to keep the supply pressure a few bars higher than what is specified by Sumitomo. As the system cools down, the helium pressure naturally decreases. To compensate, we found it necessary to increase the supply pressure post-cooldown[1], which should not be necessary according to Sumitomo. With these lessons in mind, we have achieved a reproducible cooldown time of 12 days. Figure 7.2 shows the temperatures of the cryocooler stages, the top and bottom of the main section of the coil, and the heat shields. When cooled down, stage 1 of the cryocoolers is around 55 K, stage 2 is at 4 K, while the magnet is around 5–6 K. The heat shields surrounding the magnet are about 80 K.

During the 2021 run, the Babcock was kept cold for 34 consecutive days. Unfortunately, we experienced a gradual loss of cooling power, as illustrated by figure 7.3, which limited the cold time. Besides illustrating how the temperatures of the system behave during operation, figure 7.3 shows that the temperature of cryocooler 1 stage 1 increases

[1] Or begin the cooldown at an equivalently higher starting pressure.

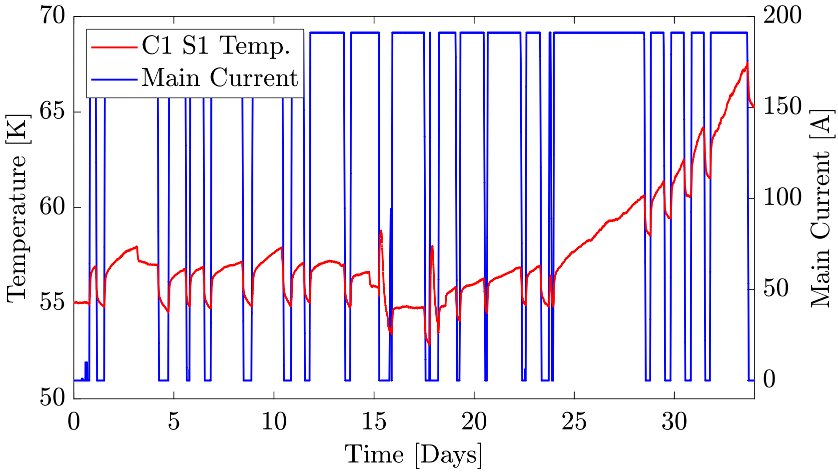

Figure 7.3: The temperature of cryocooler 1 stage 1 overlaid with the current in the main coil. The cooling power of the cold head is seen to decrease over time causing the temperature to increase. The periodic temperature increases of a few K are caused by resistive heating in the magnet leads.

significantly over time. The temperature increase propagated to the magnet, which approached the critical temperature (around 6.5–7.0 K) after 34 days, and we were forced to stop operating.

Figure 7.3 also shows how the energy dissipated by the current in the resistive sections of the leads warms up stage 1. The temperature curve is overlaid with the current in the main coil to illustrate this point. The temperature increase is also seen by the rest of the system, although the increase on the magnet itself is only 100–200 mK. The temperature spikes after 15 and 18 days are caused by quenches. After a quench from full current (191 A), it takes about 12 hours for the magnet to cool down again.

To counter the apparent loss of cooling power in cryocooler 1, we increased the helium supply pressure in the compressor. Figure 7.4 shows the stage 1 temperature overlaid with the helium pressure. The effect of the increase in pressure is clear after 13 and 17 days, where the increased cooling power is seen to overcome/equal the heating rate. However, as the temperature accelerates after about 20 days, the increase in cooling power is not enough to counter the heating. As the

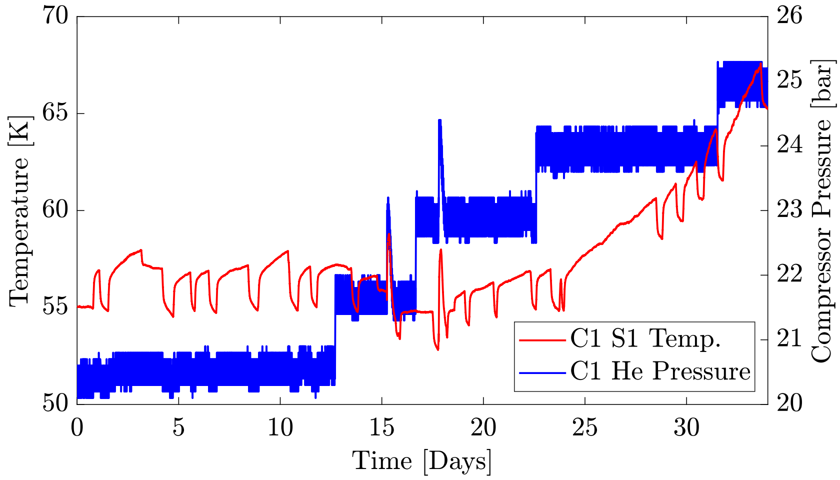

Figure 7.4: The temperature of cryocooler 1 stage 1 overlaid with the helium supply pressure in the compressor. Increasing the pressure is seen to overcome/equal the heating after 13 and 17 days, but when the heating accelerates after 20 days, it overcomes the effect of increasing the pressure.

maximum supply pressure is 26 bar, the temperature could no longer be controlled.

Based on our observations and conversations with Sumitomo, we believe that the loss in cooling power is due to contaminants in the helium that accumulate in the cold head. The compressors are filled with ultra pure 6.0 helium (99.9999% pure), which does not contain any significant contaminants, but the contaminants could be introduced to the system via a tiny leak – perhaps from another part of the compressor system. To test the theory, we disconnected the cold heads from the rest of the system, warmed up stage 2 to either 100 or 300 K, changed the helium, and cooled down. This method proved to restore the cooling power, but warming to 300 K seems to give a better result, as the cooling power quickly deteriorated after warming to 100 K. The 100 K heat cycle takes about one week, and the 300 K cycle takes about 3 weeks.

Although the magnet is equipped with 40 W internal heaters to accelerate a warm-up, it takes about a month to get from 4 K to room temperature. The process can be accelerated by adding gas into the vacuum chamber to increase the thermal conductivity. The

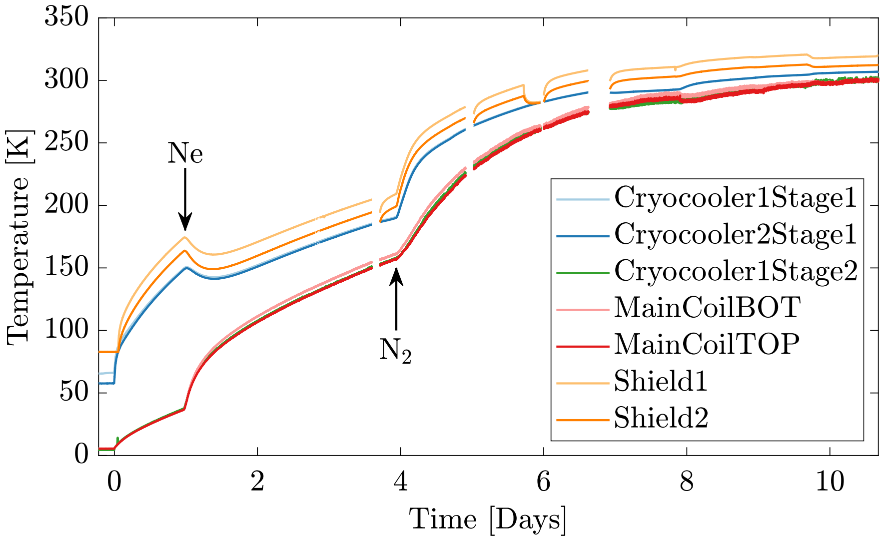

Figure 7.5: The temperatures of the cryocooler stages, the top and bottom of the main section of the coil, and the heat shields during a warm-up. Neon and nitrogen gas were added to the vacuum chamber to accelerate the warm-up after 1 and 4 days respectively. The temperature for cryocooler 2 stage 2 has been omitted.

temperature of the interior must be above the boiling point of the introduced gas, as it would otherwise just stick to the surfaces. Hence, neon can be used above 27 K, and nitrogen can be used above 77 K.

Figure 7.5 shows the temperatures during an accelerated warm-up, where neon was added after 1 day, and nitrogen was added after 4 days. The warm-up time is seen to be 10 days. The temperature of cryocooler 2 stage 2 is omitted in the figure due to an unstable temperature reading caused by a bad connection. If too much gas is added, the outside of the vacuum chamber will freeze over, which could damage the detector in the bore of the magnet, so there is a limit to, how much the warm-up can be accelerated with this method.

7.3 Babcock Field Maps

To characterise the field produced by the Babcock magnet, we recorded multiple field maps along the empty bore. The field is measured by mounting an on-axis NMR probe, four off-axis hall probes, and one

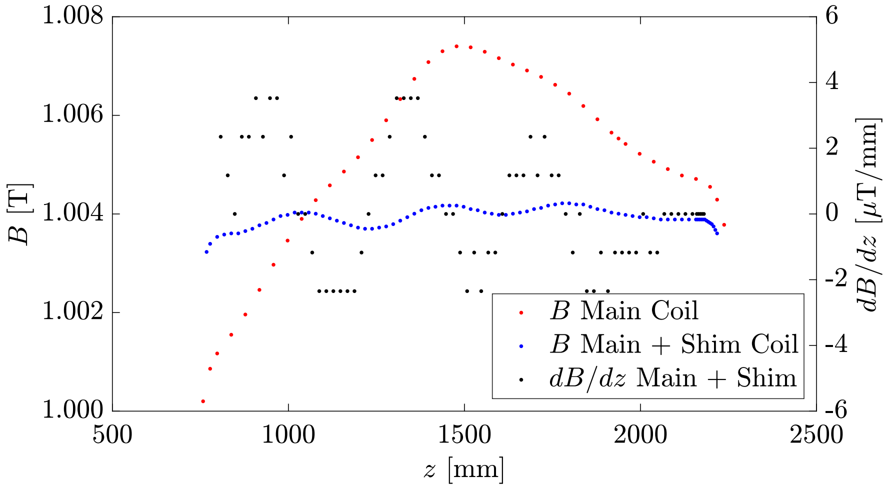

Figure 7.6: The field maps recorded with the on-axis NMR probe, with 191 A in the main coil and with/without 5 A in the shim coil. The gradient of the total field with both magnets energised is shown as well. The resolution of the field and the field gradient is 23.5 μT and 1.2 μT/mm respectively.

off-axis high resolution hall probe on a custom-made jig, which can be positioned in the bore. In 2018 and 2019, we used this method to record the field maps that were used to design the shim coil, before we sent the magnet to the manufacturer to have the shim coil installed.

When the magnet returned to ALPHA with the shim coil installed, we remeasured the fields to verify the effect of the shim coil and to determine, whether to adjust the position of the coil. The field maps obtained with the NMR probe are seen in figure 7.6. The field gradient, which is calculated from the two nearest points, of the field with both magnets energised is shown as well. The resonance frequencies are recorded with a 1 kHz (23.5 μT) resolution at positions spaced 20 mm apart, which gives a field gradient resolution of 1.2 μT/mm.

Figure 7.6 shows that with 5 A in the shim coil, the difference between the maximum and minimum field in the region of interest is reduced by about a factor of 10. As described in section 6.3, the total field gradients cannot be more than $1.77 \cdot 10^{-6}$ T/mm for the up/down measurement and $1.77 \cdot 10^{-8}$ T/mm for the 1% measurement without introducing a lower energy cut-off for the antihydrogen considered for

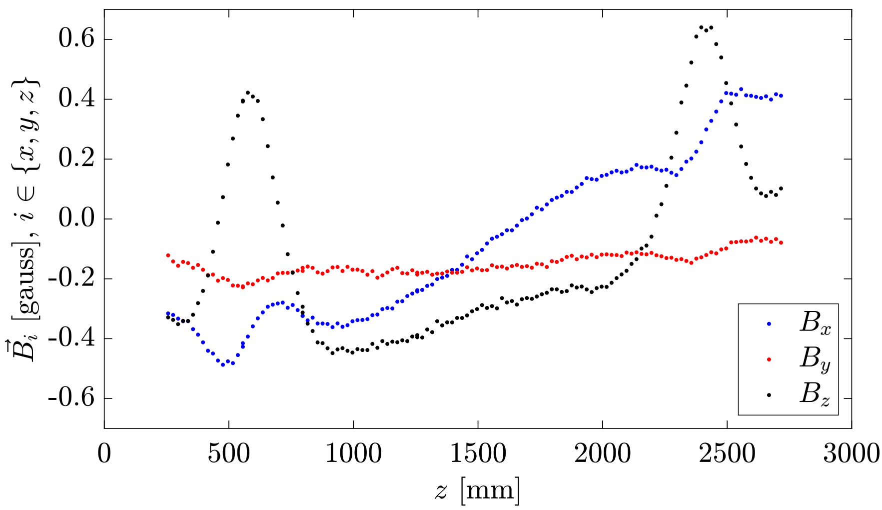

Figure 7.7: The field as measured with the high-precision hall probe through the bore of the magnet, with both the main and the shim coil at 0 A.

a gravity measurement. One could use the internal magnets to smooth the field further or add additional shim coils. Near the ends of the main coil, the field quickly decreases, and the large gradient prevents the NMR probe from measuring the field.

Figure 7.7 shows a map of the background field (both main and shim coil at 0 A) as measured with the high-precision hall probe. The background map was recorded shortly after the map shown in figure 7.6. The background is seen to be well below the 1 gauss level. There are some interesting features around the ends of the magnet, which might be explained by persistent currents, but further examination would be required to determine their origin.

7.4 BNL Magnet Commissioning

Although MAB, MDB, and LOc were commissioned in 2019, we considered all the internal magnets to be "new", when they returned from BNL, after they had installed the remaining magnets in 2021, so we restarted the commissioning process. When a superconducting magnet is ramped for the first time, it might quench, as the magnetic forces on the wires make them cause small cracks in the magnet form, which can release enough energy to make the magnet normalconducting. It might take multiple attempts to reach full current for the first time. This process is known as training.

The commissioning procedure is as follows: the magnet is gradually ramped to higher current levels at a conservative ramp rate. At each level, the quench response is tested by manually triggering it, if it has not already been triggered by an actual quench. The process is repeated, until the maximum current is reached. Simultaneously, the gains and balances described in section 6.5 are tuned. During commissioning, the external solenoid is off, and only a single magnet is energised at a time.

The ramp rate used in the initial part of the commissioning is typically 1 A/s, but after demonstrating that the magnet can reach full current, the rate can be increased. A higher ramp rate will generate a higher inductive voltage across the magnet, and the voltage thresholds for detecting quenches across the inductive part must therefore be increased[2]. At high ramp rates, V_{avg} can exceed the size of its data type (I16) on the FPGA, which is critical to the quench detection. To avoid the issue, we installed voltage dividers in the signal conditioners as described in section 6.5. In this way, we have achieved ramp rates of 95 A/s for MAB and MGB.

Both capture solenoids have proven to quench consistently around 40–50 A, although they are designed to handle a maximum current of 140 A. As the behaviour is consistent, the magnets do not appear to be training. The problem seems independent of the ramp rate, which has been as low as 0.1 A/s. When a quench is detected, the software triggers a "fast readout", which stores the complete set of voltage tap readings around the time of the quench recorded with a 140 μs time resolution. For such a quench of SoB, the voltage drops on the LTS are shown in figure 7.8.

Initially in figure 7.8, the voltage drops across the two halves are seen to be constant, as the magnet is ramping at a constant rate. The sign is different for the two halves due to the orientation of the taps in the SC. Around $t = -5$ ms, there is an event, which causes the voltage drops to go more negative. This is consistent with part of

[2]Although the voltage drops across the magnet should cancel each other, the imperfect balance results in a non-zero signal during normal operation. The threshold for quench detection must be higher than this signal.

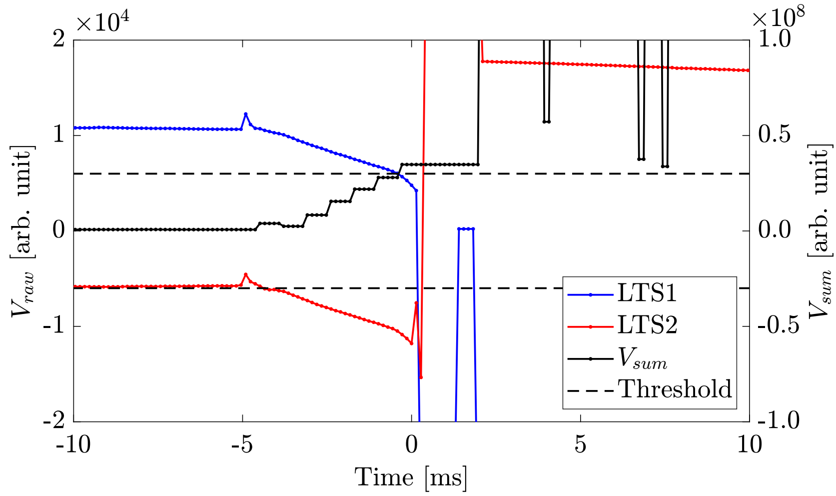

Figure 7.8: The fast readout triggered by the QD at around 50 A while SoB was ramping. The voltage drops across the LTS are overlaid with the V_{sum}. The time is relative to the detection of the quench. The event at $t = -5$ ms is consistent with part of LTS2 going resistive.

LTS2 going resistive. V_{sum} starts to increase, and when the threshold is crossed[3], the QPS fires. This of course has dramatic consequences for the voltage drops, but these are no longer relevant, as the QPS has fired. It is currently not understood what causes the quench. As described in section 6.3, the capture solenoids should be far from their critical current.

Another issue was observed with the octupoles, as a quench would normally be detected within an hour after reaching full current. We observed that the power supply voltage would sometimes oscillate or spike, although it was set to keep the current constant. Most of the time, these features would appear on both halves of the magnet, and the increase in V_{sum} would stay below the threshold. However, a quench would eventually be detected. Figure 7.9 shows the fast readout on the LTS for OcB during such an event, while the magnet was energised to 800 A.

What figure 7.9 shows is a sudden increase in voltage around $t = -20$ ms, which causes the amplifier in the SC to saturate, before the

[3]In this case, 5 samples go into calculating V_{avg}, so V_{sum} is only considered to have crossed the threshold after 5 samples.

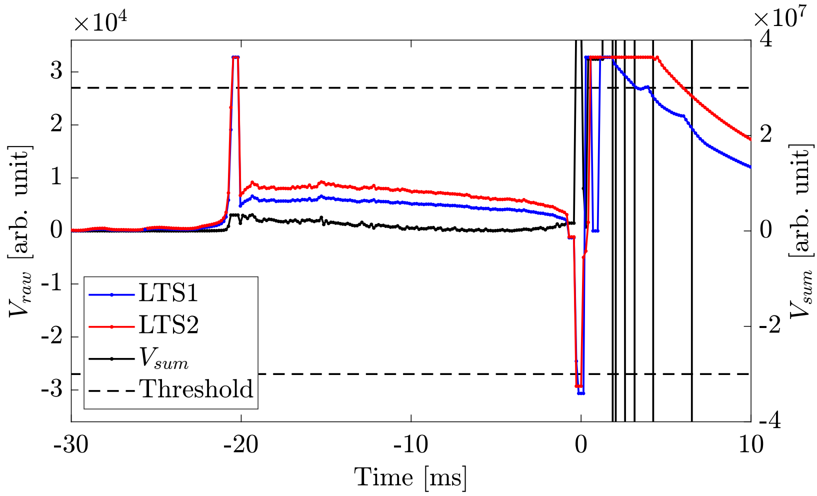

Figure 7.9: The fast readout triggered by the QD while OcB was kept at 800 A. The voltage drops across the LTS are overlaid with the V_{sum}. The time is relative to the detection of the quench. At $t = 0$, the required gain switching time becomes too short for the system to keep up, and the QPS triggers as described in the text.

SC switches to a lower gain. The voltage then decreases slowly over the next 20 ms. Just before $t = 0$ ms, the voltage becomes low enough that the software decides to switch to a higher gain, but as soon as it switches, the signal grows to a large negative value, which is amplified by the higher gain. This time, V_{sum} exceeds the threshold, and before the gain can be adjusted, the QD part of the software decides that a quench has happened.

The problem with the octupoles was mitigated by reducing the number of samples that enters the average and by increasing the number of samples that must exceed the threshold for a quench to be detected, but the cause of the issue was identified to be poor connections in the octupole circuits. However, the problem displays the inability of the code to deal with current changes that are faster than the gain switching time. Such rapid changes in current are not observed during normal operation in ALPHA-2, which has a similar magnet system, so there is no reason to expect them to happen during normal operation in ALPHA-g.

Table 7.1 shows the status of the magnet commissioning as of spring 2022. The consequences of the issues experienced with the capture

Magnet	Commissioned	Max Current w/wo Bab [A]	Ramp time up/down [s]
Bottom			
SoB	Partially	38/50	25/25
BgB	No	–	–
MAB	Yes	60/95	1/1
MBB	Yes	–/95	5/5
MCB	Yes	–/95	5/5
MDB	Yes	–/95	5/5
MFB	Yes	–/95	5/5
MEB	Yes	–/95	5/5
MGB	Yes	60/95	1/1
TrB	No	–	–
Analysis			
AnB	No	–	–
CCB	No	–	–
COB	No	–	–
AnT	No	–	–
CCT	No	–	–
COT	No	–	–
Top			
TrT	No	–	–
MGT	No	–	–
MFT	No	–	–
MET	No	–	–
MDT	No	–	–
MCT	No	–	–
MBT	No	–	–
MAT	No	–	–
BgT	No	–	–
SoT	Partially	–/37	37/37
Other			
LOc	Yes	800/975	30/2.5
OcB	Yes	700/1080	30/2.5
OcT	No	–	–

Table 7.1: Table of the status of the magnet commissioning as per spring 2022. The maximum current that has been confirmed to work for operation, and the shortest ramp time to full current confirmed to work without the Babcock energised are shown.

solenoids will be discussed in section 7.6. For the internal solenoids, we have not yet experienced any limit to how long they can stay at full current.

To understand the context of the 2021 run described in section 7.6, the status of the internal magnets for the majority of the run were as follows: SoB could be ramped to 33 A (1 T on axis field) at 0.1 A/s and MAB and MGB to 60 A (0.6 T) at 1 A/s. The octupoles could be ramped to 800 A in 60 s, but it was questionable for how long they could stay on.

At the time of the submission of this thesis, SoB, MAB, MGB, and the octupoles are operational, they are fully integrated in the experimental protocols, and they are being used in attempts to trap antihydrogen.

7.4.1 Filtering the Power Supply Output

The way the magnet power supplies are currently operated, when a power supply is requested to begin a ramp, it will try to use a constant ramp rate, which in principle causes a discontinuity in the beginning and at the end of the ramp. As a result, the current output will oscillate around the requested current as shown by the blue curve in figure 7.10. The figure shows the voltage drop across the inductive part of the magnet for a ramp from 0 A to 5 A followed by a ramp to 0 A. The time resolution is only 250 ms, but it is enough to see the oscillating behaviour of the voltage.

It is possible that the oscillations could be reduced by tuning the PID parameters of the power supply, but we initially mitigated the issue by filtering the power supply output by installing a high power resistor in parallel with the magnet, which works as a high-pass filter. With the resistor in parallel, the voltage drop for an identical ramp is shown by the red curve in figure 7.10. The time data has been manipulated to make the two ramp downs begin at the same time, which has arbitrarily been chosen to be 8 s.

7.5 Magnet Inductances

When running multiple magnets at once, we observed that their mutual inductances can cause two kinds of issues: firstly, the oscillating behaviour of the current at the beginning and at the end of a ramp

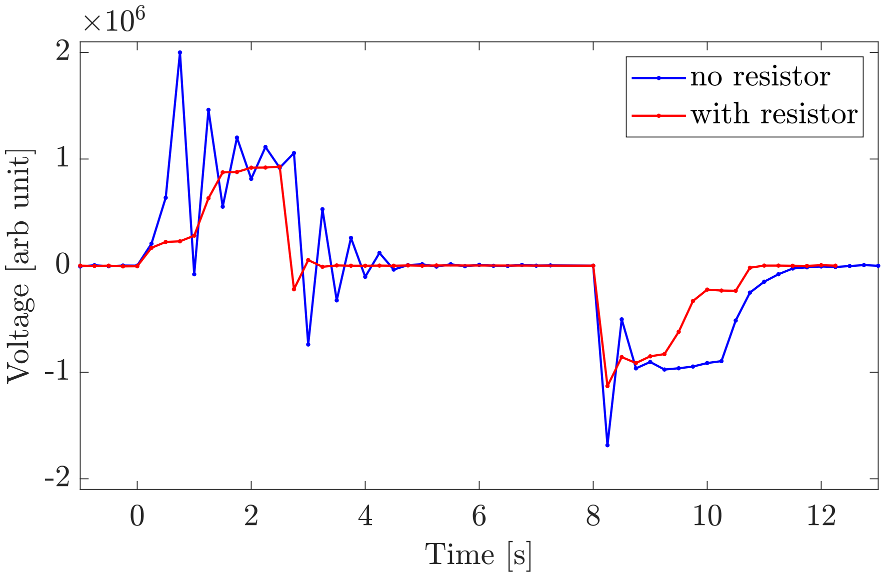

Figure 7.10: The voltage drop across LTS during a ramp from 0 to 5 A starting at $t = 0$ s, and a ramp from 5 to 0 A starting at $t = 8$ s. The time data has been offset to make the second ramp downs start at the same time to help the comparison.

can induce signals that are detected as quenches in the neighbouring magnets. This has proven to be an issue for the octupoles in particular, as they have high mutual inductance. Secondly, when a magnet quenches or its QPS fires, it can cause quenches to be detected in other magnets. This is not an issue in itself, but since the recovery time for the Babcock magnet is 12 hours, a lot of time is lost as a consequence of a detected quench.

Common to both types of issues is that the QD system falsely identifies the induced signal as quenches. Since the total voltage across an inductive section is

$$V = IR + L\frac{dI}{dt} - \sum_n M_n \frac{dI_n}{dt} \tag{7.1}$$

where n indexes all the other magnets, it should be possible to filter out the induced signal in the software, if we can measure $M_n dI_n/dt$ for the other magnets. Figure 7.11 shows the voltages across the boost section of the Babcock main coil induced by a quench in SoB. Of the different sections in the main coil, the largest difference in coupling between the two halves of the section and SoB is the boost section, so

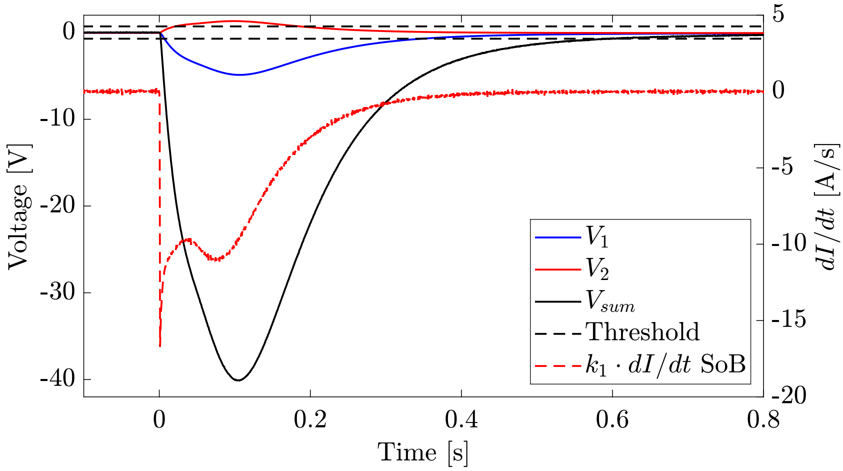

Figure 7.11: The voltage (solid red, blue) across the boost section of the Babcock induced by a quench in SoB. As V_{sum} (solid black) violates the threshold (dashed black), a quench is detected. The time derivative of the current in SoB (dashed red) is plotted as well.

this is where the quench is detected first. V_{sum} is seen to drastically violate the threshold.

Figure 7.11 also shows the dI/dt signal for SoB (multiplied with a constant, k_1), whose shape resembles the shape of V_{sum}. In the figure, dI/dt is simply calculated as $(I_{i+1} - I_i)/\Delta t$ based on the DCCT signal, where i indexes the data points. In an attempt to make the signals match, we apply an infinite impulse response (IIR) filter to the dI/dt signal:

$$y_{i+1} = \alpha x_i + (1 - \alpha) y_i \tag{7.2}$$

with $y_0 = 0$, and $\alpha \in [0, 1]$ is a tunable parameter. The filtered dI/dt (multiplied with a constant, k_2) is plotted in figure 7.12 with V_{sum} and their difference. V_{sum} has been divided by a factor of 10 to improve the readability of the plot. The difference is seen to almost not violate the threshold (which could be increased) for periods of more than 50 ms. The method needs fine tuning but seems like a possible way of filtering out signals induced by other magnets.

The Danfysik QD system for the Babcock (see section 6.6), does currently not have the capability to do any of the advanced quench detection described above. Hence, we are working towards replacing

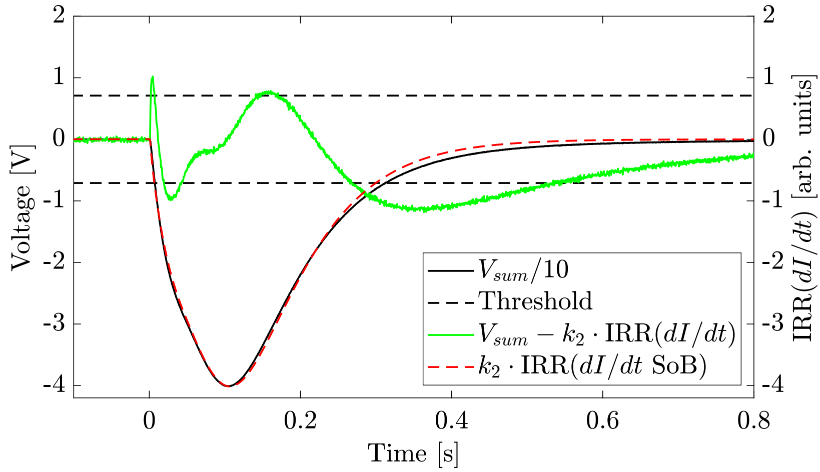

Figure 7.12: The filtered dI/dt signal (dashed red) is seen to match V_{sum} (solid black), and their difference (green) is seen to almost not violate the threshold (dashed black) for a period longer than 50 ms.

the Danfysik system with a QD system similar to the system for the BNL magnets (see section 6.5). Both the new QD system for the Babcock and the QD systems for the BNL magnets would need a live reading of the DCCT signals to calculate the time derivatives of the currents. The DCCT signals are currently fed into the PS cRIO, so connecting the DCCTs to multiple cRIOs requires a more complicated connection scheme, but it cannot alter the value of the signal. Alternatively, lower quality DCCTs dedicated quench detection could be added.

In a general campaign to determine the self- and mutual inductances of the magnets, we measured the voltage drops across the magnets during a ramp with a constant ramp rate. At the time, it was only possible to get a quality measurement for a few of the magnets. The results can be seen in table 7.2. To determine the self-inductance, we ramped the magnets slowly to ensure a stable ramp rate, but to determine the mutual inductances, we ramped the magnets quickly to get as big an induced signal as possible. During commissioning, we verified that the mutual inductance with the beamline magnets is negligible.

Note that it is not possible to do an isolated measurement of the self-inductance, as there will be second order mutual inductance effects from the surrounding magnets. However, the self-inductances

M_{ij} [mH]	SoB	MAB	MGB	LOc	OcB
SoB	649	2.0	0.15	–	–
MAB	2.0	54.5	0.26	–	–
MGB	0.18	0.23	55.3	–	–
LOc	–	–	–	1.9	0.84
OcB	–	–	–	0.91	2.9

Table 7.2: Measured self- and mutual inductances in mH for selected magnets. No value is given, if the mutual inductance measured was less then 0.1 mH.

measured for MAB and MGB are very similar (as they should be), which indicates that the second order mutual inductance effect is not significant. The inductances in table 7.2 agrees with the calculated inductances in section 6.3.6 within about 10 %.

7.6 Particle work

This section describes some of the work done with particles in ALPHA-g during the 2021 run and in early 2022, before the antiproton beam became available. Besides preparing the plasmas and attempting antihydrogen production in the bottom full strength region, this section touches upon the efforts on recording antiproton annihilations in the detector and on ECR measurements.

7.6.1 Plasma Preparations

The process of particle preparation and antihydrogen production in ALPHA-g is intended to be similar to the process in ALPHA-2 described in chapter 4. However, unlike in ALPHA-2, where the antiprotons and positrons enter from different sides, in ALPHA-g both particle species are loaded from below. The positrons enter the SoB region first, where they are cooled and compressed, before they are moved to a deep well under MGB (as far from SoB as possible). The antiproton plasma then enters the SoB region, where it is cooled sympathetically with electrons and compressed. To produce antihydrogen, the antiprotons are moved to E28 and the positrons to E29 (see figure 6.10), where they are merged.

With SoB operating at 33 A and the external solenoid at full current, the total field in the initial trapping region under SoB is about 2 T. In comparison, the charged particles in ALPHA-2 are trapped in a

3 T field, so the cyclotron cooling power in ALPHA-g is less than half of that in ALPHA-2. About 3 million positrons were loaded into ALPHA-g, where they cooled under SoB for 60 s, before being transferred to the deepest possible well under MGB generated by 150 V on the electrodes (E34–E36) and 1.6 T generated by the magnets. In this well, they would wait until the antiprotons were loaded.

By dumping the positrons to an MCP from E35, we determined their temperature to be 50 K in the coldest cases. When the antiprotons were ready, the positrons were moved to a shallow pre-mix well under E29. However, in this well we measured their temperature to be 400 K in the best cases. In ALPHA-2, the positron temperature prior to mixing is less than 20 K, and the mixing only yields 10–30 trappable antihydrogen atoms, so the creation of trappable antihydrogen with 400 K positrons did not seem promising. One possible explanation for the high positron temperatures could be heating caused by patch potentials – small patches of charge – on the electrodes [106, 107]. As of early 2022, a campaign to minimise patch potentials in ALPHA-g is ongoing, and the results are promising.

The 2021 run was the first time we received antiprotons from ELENA (see chapter 3). As part of a major upgrade to the catching trap, we installed a new degrading foil primarily made of aluminium, which should be compatible with the 5 keV beam from ELENA. Simulations have determined the optimal foil thickness to be about 1 μm, so we tested a few different foils with a thickness in that range. The best performing foil allowed us to catch about 5 %, which translated to 180000 antiprotons. After transferring them to ALPHA-g and cooling them there, about 10000 antiprotons could be used for mixing.

7.6.2 Antihydrogen Production Attempts

Despite the high positron temperatures, which might make production of trappable antihydrogen infeasible, we attempted to produce antihydrogen near the end of the 2021 run. The antiprotons and positrons were brought to the premix well under E28 and E29, where they were mixed by slowly making the two wells shallower as illustrated by figure 7.13. This type of “slow merge”, which is similar to the mixing procedure in ALPHA-2, unfortunately did not yield any signs of antihydrogen.

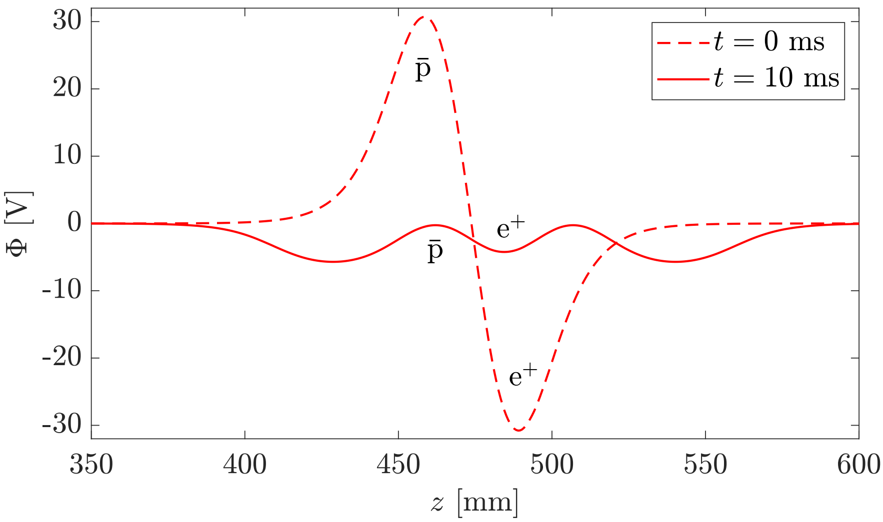

Figure 7.13: A ”slow merge” of antiprotons and positrons in an attempt to produce antihydrogen by making the two plasmas overlap.

Another mixing method to create antihydrogen, which increases the chance of producing untrappable antihydrogen but decreases the chance for trappable antihydrogen, is the so-called “ATHENA style” mixing [108]. In this type of mixing, the positrons are held under E29, and the antiprotons are held in a well under E25 as shown in figure 7.14. The potential confining the antiprotons to the right is then lowered, so the antiprotons are injected into the deeper well around the positrons. The injection energy is matched to the positrons, so all antiprotons will have enough energy to traverse the positron plasma and interact/mix with it. The potential shown in figure 7.14 does not take the positron self-potential into account.

In ATHENA mixing, the static positrons will interact with the traversing antiprotons via the Coulomb force to slow them down and potentially form antihydrogen. In previous experiments, this method yielded a clear increase in annihilation signals due to formed antihydrogen [108]. Slowed antiprotons will not have enough energy to traverse the positron plasma and will thus be trapped in one of the wells on either side of the positron plasma.

By first lowering the potential barrier for either of the wells, the hot antiprotons and any cold antiprotons in the given well will escape, which leaves only the cold antiprotons in the other well trapped. The

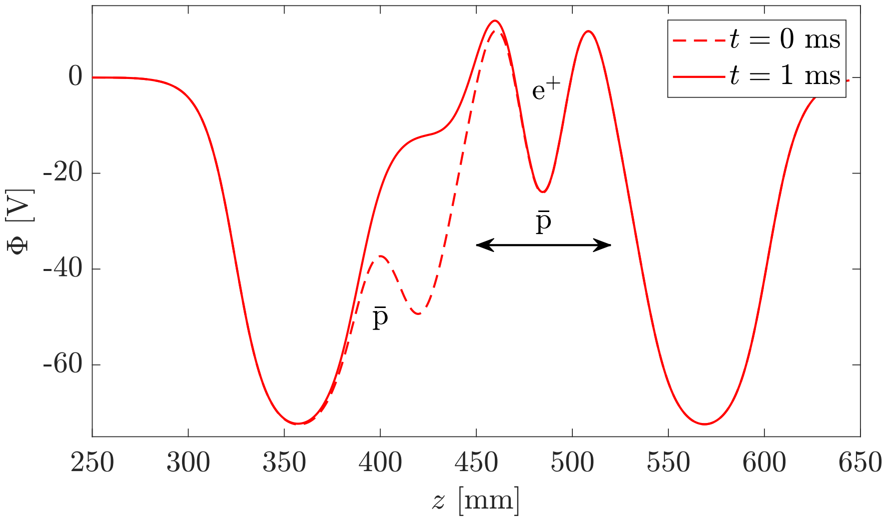

Figure 7.14: "ATHENA style" mixing of antiprotons and positrons in an attempt to produce antihydrogen. The antiprotons are injected at a potential energy matching the positrons, so the antiprotons can traverse the positrons plasma and interact with it to form antihydrogen.

number of antiprotons left in the second well indicates the interaction strength with the positron plasma. Although we did detect antiprotons in both wells, proving the interaction with the positron plasma, we did not see any clear annihilation signal indicating antihydrogen production.

7.6.3 Antiproton Triggers

As described in section 6.2, the machine learning-based algorithm to identify annihilations in the detector needs data with clearly defined events for training. Releasing high numbers of antiprotons over short periods of time produces very clean data sets. This is done automatically as a result of antihydrogen production attempts, but we also did dedicated antiproton annihilation runs, where different detector settings could be studied. Ideally in such runs, the annihilation rate is as high as possible without saturating the detector, which currently has a maximum readout rate of a few hundred Hz.

To let antiprotons escape the trap at a desirable rate, we hold them in an axially long well (i.e. six electrodes or more). It is not understood why expanding to a long well will let the antiprotons escape, but they

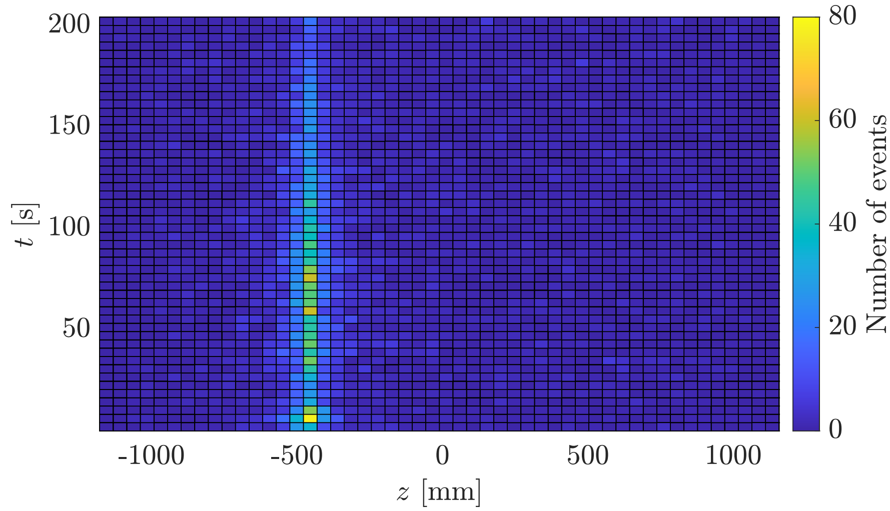

Figure 7.15: The number of antiproton annihilations as a function of axial position and time. The antiprotons are held in a long shallow well from $z = -400$ to -600. The data indicates a preferred axial escape position around $z = -500$.

have been observed experimentally to slowly heat and escape. Figure 7.15 shows the number of annihilations as a function of time and axial position for an antiproton plasma held in a well generated by a 15 V potential on electrodes E27–E36.

E27–E36 extend over the region from -400 to -600, but the antiprotons are seen to escape at a well defined axial position. This “hotspot” behaviour is not understood, but it is consistently observed. 200 s after expanding to the long well, the electric trapping potential is turned off to let any remaining antiprotons escape, but they proved to have all been lost from the long well. Figure 7.16 shows the (z, ϕ) distribution of the annihilations of the escaped antiprotons. The data shows an azimuthal preference for the annihilations, which is different from what is normally observed with this type of losses. The asymmetry could be explained by the existence of the patch potentials causing the hot positrons.

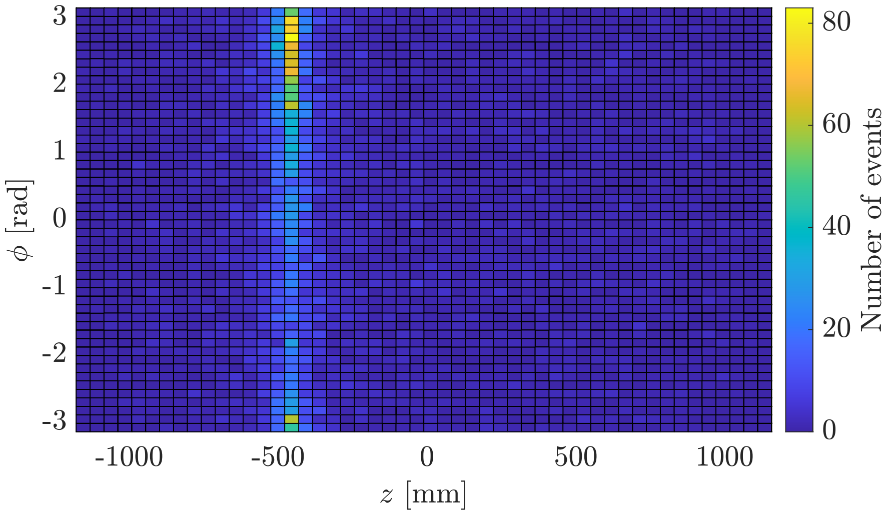

Figure 7.16: The number of antiproton annihilations as a function of axial and azimuthal position. The antiprotons are seen to have a preferred azimuthal escape direction, which could be explained by the existence of patch potentials.

7.6.4 ECR Measurements

ECR measurements can be used to precisely determine the field produced by the inner magnets at a given current (see section 2.4). It is of particular interest to know the magnetic field at the positions of the mirrors that will be used for gravity measurements. By keeping a large electron plasma under SoB, small electron bunches can be extracted to make multiple measurements of the heating of the electron plasma at different microwave frequencies in quick succession and thereby determine the resonance frequency at a given position for a given field. The temperature of the plasma is determined by dumping it to a MCP. We aim to achieve a precision of the ECR measurements in ALPHA-g that is better than the required control of the fields of $5 \cdot 10^{-6}$ T for the 1% measurement.

By using a three-electrode well, the axial position of the plasma can be controlled with sub-millimetre precision. Figure 7.17 shows the strength of the on-axis magnetic field measured as a function of the axial position with MGB at 59.52 A as measured by the DCCT. The background field ($I_{MGB} = 0$) has not been measured. The centre of E35, which sits under MGB, is defined to be at $z = 0$. A second order polynomial is fitted to the data to determine the central axial position

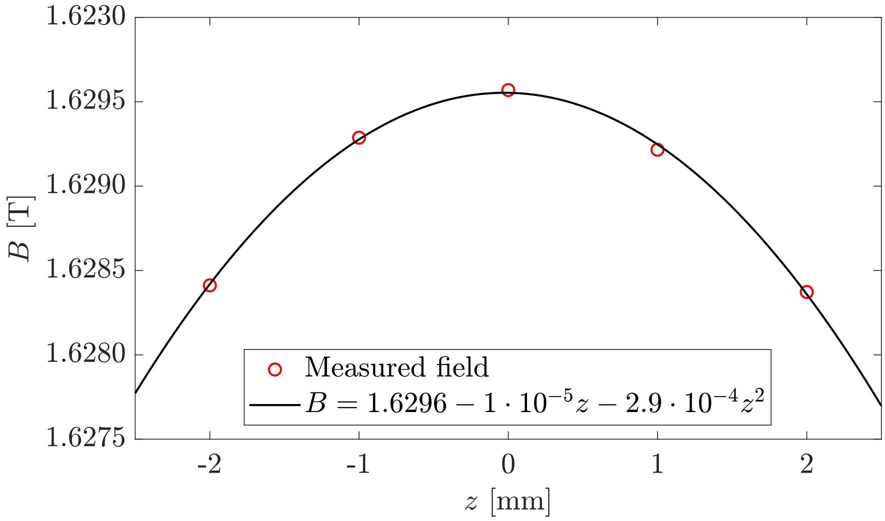

Figure 7.17: A fit to the axial magnetic field strength measured with ECR under MGB relative to the axial centre of E35. At a current of 59.52 A, the total maximum field is 1.6296 T. The precision is better than indicated by the number of significant digits given, but it has yet to be carefully characterised.

of MGB and the maximal axial field strength. The fit gives a maximum axial field of 1.6296 T (including the field from the Babcock) with five significant digits. However, the precision of the ECR measurements is better than the number of significant digits would suggest, but it has yet to be carefully characterised.

Chapter 8

Magnetic Field Calculations

Multiple computer codes to simulate antihydrogen behaviour etc. exist within the ALPHA collaboration. One of these is currently being developed with the main goal of simulating adiabatic expansion and cooling of antihydrogen [54]. This code will be referred to as the AE code. The code is written in C++, and is capable of simulating the behaviour of 0.2 million antihydrogen atoms in a magnetic trap. The simulation time is 24 s with a 3.5 μs timestep, and a computation time of 40 s per antiatom. With a 100 jobs running in parallel, the total simulation time is about a day.

Simulating the behaviour of trapped antihydrogen naturally involves calculating the magnetic field. As several magnets with different geometries contribute to the field, the total field must be calculated, before the potential can be derived. To keep the computation time down, the magnetic field must be calculated quickly. In general, there are two approaches to this: either the field is calculated at the relevant position in situ, or the field the precalculated on a grid, and the field at the relevant position is determined by extrapolating between the grid points. The method used for the AE code is the former.

The method used to calculate the field in the AE code in situ left room for improvement, so we conducted a study of the optimal way to calculate the field with both accuracy and computation speed in mind [2]. The resulting paper was published as a general study. A selection of the results will be summarised in this chapter with the ALPHA experiment in mind.

8.1 Calculating Magnetic Fields of Solenoids

When calculating the field of a solenoid (or any other type of magnet), one can choose to represent the magnet by an idealised *magnet model*. For idealised magnet models, there are typically analytic expressions to calculate the field available, as well as approximative *field models*. The alternative to an idealised model is a detailed representation of the wire configuration. The field of such magnet models are typically calculated with a Biot-Savart field model.

In the following section, a selection of the field models presented in [2] are discussed. In general, analytic models and Biot-Savart models are accurate but slow, so while they are good for calculating fields accurately, they are poor for simulations. The TAVP model, which was originally used in the AE code, is fast but inaccurate. Together with the aforementioned models, a promising alternative for the AE code and for simulations in general – the McDonald model – is presented here.

Analytic Models

When representing a solenoid with an idealised current distribution, three models are particularly common: a single current loop, an infinitely thin cylindrical shell with finite length, and a cylindrical shell with finite length and thickness. The current distribution is often assumed uniform, which is a decent assumption for magnets with a high uniform wire density, but not for magnets like the capture solenoids described in chapter 6. For all three idealised magnet models, there are exact analytic expressions for the magnetic field [109, 110]. The simplest is for the single current loop with radius R and current I, which in cylindrical coordinates (ρ, ϕ, z) is

$$B_\rho = \frac{Cz}{2\alpha^2\beta\rho}\left[(R^2+\rho^2+z^2)E(k^2) - \alpha^2 K(k^2)\right] \tag{8.1}$$

$$B_\phi = 0 \tag{8.2}$$

$$B_z = \frac{C}{2\alpha^2\beta}\left[(R^2-\rho^2+z^2)E(k^2) + \alpha^2 K(k^2)\right] \tag{8.3}$$

where

$$\alpha^2 = R^2 + \rho^2 + z^2 - 2R\rho \tag{8.4}$$

$$\beta^2 = R^2 + \rho^2 + z^2 + 2R\rho \tag{8.5}$$

$$k^2 = 1 - \frac{\alpha^2}{\beta^2} \tag{8.6}$$

$$C = \frac{\mu_0 I}{\pi} \tag{8.7}$$

and $K(x)$ and $E(x)$ are the complete elliptic integrals of the first and second kind, which are available in libraries of most programming languages. The exact analytic expressions for the magnetic field components of a cylindrical shell are more complicated, but they are also based on elliptic integral functions and are thus calculable [110], whereas the expressions for a finite length and thickness solenoid require numerical integration. Hence, the expression is only as exact as the numerical integration method[1].

While the analytic models give accurate values for the magnetic field, to the extend that the idealised magnet model is a good representation of the magnet, they are typically slow compared to approximative field models. Although the computation time of a field model can be hard to quantify, as it depends on the computer processing system and the input parameters, the computation times of the analytic models were found to be 1–100 times longer than those of approximative models. These comparisons are elaborated further in [2].

Biot-Savart Models

Given a small line segment, $d\vec{l}$, with current, I, the magnetic field given at a point, P, is given by the Biot–Savart law [111]

$$d\vec{B} = \frac{\mu_0 I}{4\pi} \frac{d\vec{l} \times \vec{r}}{|\vec{r}|^3} \tag{8.8}$$

where $\vec{r}$ is the coordinate vector from the line segment to P. By dividing a magnet model into small line segments and adding their contributions to the field together, the field of any magnet type can be calculated using the Biot-Savart model. Hence, the model lends itself well to calculating the field of the complicated geometries like

[1] It is noted that all calculations done on a computer are of course limited by *machine precision*, but since such levels of precision are typically irrelevant, this limit is ignored.

the capture solenoids and the octupole end-turns. The computation time is unfortunately very long, so Biot-Savart models are unsuited for in situ calculations. As mentioned in section 6.3, the winding of the BNL magnets was carefully documented with photographs of each wire layer. Based on these photos, we have constructed a Biot-Savart model, which is the most accurate model we have available. It is from this model that the values in table 6.1 and appendix A.1 are extracted.

TAVP Model

A model, which we will refer to as the Truncated Approximate Vector Potential (TAVP) model, was previously used in ALPHA to calculate solenoid fields in ALPHA-2 (i.e. in the AE code and [112, Appendix A.1]). In spherical coordinates, (r, ϕ, θ), the TAVP model gives the vector potential for a solenoid, $\vec{A} = (0, A_\phi, 0)$, with radius R

$$A_\phi = \frac{C}{2R\lambda}\left((R^2 + r^2 - 2R\lambda\rho)^{-1/2} - (R^2 + r^2 + 2R\lambda\rho)^{-1/2})\right) \quad (8.9)$$

where $C = \frac{\mu_0 I R^2}{4}$, and λ is a tunable parameter. The TAVP model is inspired by the following expression from Jackson [111, eq. 5.40]

$$A_\phi(r,\theta) = C\frac{r\sin(\theta)}{(R^2 + r^2)^{3/2}}\left(1 + \frac{15R^2r^2\sin(\theta)^2}{8(R^2 + r^2)^2} + \cdots\right) \quad (8.10)$$

which assumes that the point of evaluation is near the axis of symmetry or far away from the magnet, so $R \gg r$, $R \ll r$, or $\theta \ll 1$. For $\lambda = 0.866$ the first and the second terms of equation 8.10 are reproduced, while $\lambda = 0.902$ is the best approximation for the ALPHA-2 mirror coils [112].

Figure 8.1 shows the relative deviation of the TAVP model compared to a Biot-Savart model of an ALPHA-2 mirror coil along four different paths in space[2]. The Biot-Savart model is considered to be exact. The details of the comparison are described in [2]. The coordinates are normalised to the inner radius of the ALPHA-2 solenoid, R_1. The TAVP is seen to deviate from the true field by a few to 10s of percent, especially away from the axis of symmetry.

In ALPHA-2, the electrode wall is at about $R_1/2$, so within the relevant region, the deviation of the TAVP model is at the percent level.

[2] On the axis of the solenoid, half a radius off the axis, in the centre-plane of the solenoid, and half a radius off the plane [2].

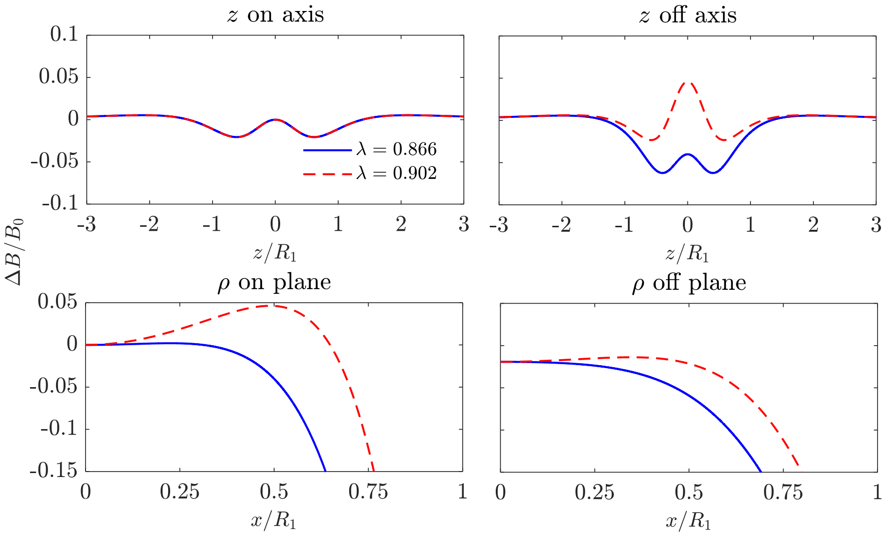

Figure 8.1: The relative deviation of the TAVP model compared to a Biot-Savart model of an ALPHA-2 mirror coil. Plots for two different values of λ are shown. The TAVP model is seen to deviate a few percent in the relevant region. From [2].

As described in section 6.2.1, the magnetic fields in ALPHA-g must be known and controlled to better than 5 μT for the 1% measurement. This number does not directly translate into a requirement on the magnetic fields in a simulation, as the errors in the simulation are systematic and to some extend symmetric around the antihydrogen trapping region. However, for simulations of a gravity measurement, it seems preferable to have a more accurate field model available.

McDonald Model

Of the approximative field models studied in [2], a model based on off-axis expansion of an analytic expression for the on-axis field seems particularly promising. The model is referred to as the McDonald model [113]. It is assumed that the magnetic field is azimuthally symmetric, in which case the field components in cylindrical coordinates,

(ρ, ϕ, z) are given as

$$B_z(\rho, z) = \sum_{n=0} (-1)^n \frac{a_0^{(2n)}(z)}{(n!)^2} \left(\frac{\rho}{2}\right)^{2n} \tag{8.11}$$

$$B_\rho(\rho, z) = \sum_{n=0} (-1)^{n+1} \frac{a_0^{(2n+1)}(z)}{(n+1)(n!)^2} \left(\frac{\rho}{2}\right)^{2n+1} \tag{8.12}$$

with $B_\phi = 0$, and

$$a_0^{(n)} = \frac{d^n a_0}{dz^n} \quad \text{and} \quad a_0(z) = B_z(0, 0, z) \tag{8.13}$$

For a single current loop with radius R placed at $z = 0$, the on-axis field is given as

$$B_z(0, 0, z) = \frac{\mu_0 I}{2} \frac{R^2}{(R^2 + z^2)^{3/2}} \tag{8.14}$$

for a finite length shell solenoid with ends at z_1 and z_2 the on-axis field is

$$B_z(0, 0, z) = \frac{\mu_0 IN}{2L} \left(\frac{z - z_1}{\sqrt{R^2 + (z - z_1)^2}} - \frac{z - z_2}{\sqrt{R^2 + (z - z_2)^2}} \right) \tag{8.15}$$

and for a finite thickness and length solenoid with inner and outer radius R_1 and R_2, the field is

$$B_z(0, 0, z) = \frac{\mu_0 IN}{2L(R_2 - R_1)} \left[(z - z_1) \ln \left(\frac{\sqrt{R_2^2 + (z - z_1)^2} + R_2}{\sqrt{R_1^2 + (z - z_1)^2} + R_1} \right) - (z - z_2) \ln \left(\frac{\sqrt{R_2^2 + (z - z_2)^2} + R_2}{\sqrt{R_1^2 + (z - z_2)^2} + R_1} \right) \right] \tag{8.16}$$

Hence, the McDonald field model is applicable to all three idealised current distributions. The deviation of the field of a finite thickness solenoid as calculated by the McDonald model relative to a detailed Biot-Savart model can be seen in figure 8.2 and 8.3 for different orders in n. Figure 8.3 does not show the on-axis deviation, as the field calculated by the McDonald model here is exact[3]. As the deviation

[3]Hence, any difference between the field models is caused by the difference between the idealised magnet model and the Biot-Savart model, and is hence not a result of poor field model performance.

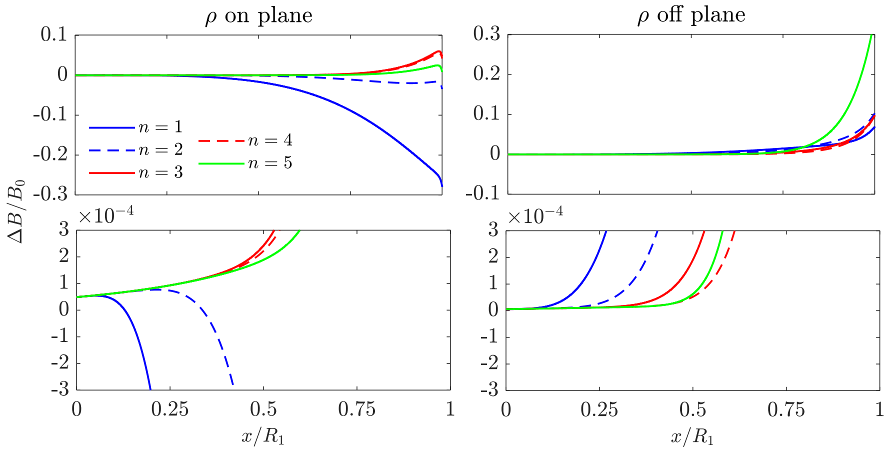

Figure 8.2: The deviation as defined in [2] of different orders of the McDonald model relative to a detailed Biot-Savart model along a path in the centre-plane of the solenoid (ρ on) and in the plane half a radius off the centre-plane (ρ off). From [2].

decreases significantly with the order, figures 8.2 and 8.3 include a magnified version of each plot. The deviations are seen to stay below $2 \cdot 10^{-4}$ for the third and higher order expansions.

The computation time of the TAVP model falls between the computation times of the first and second order expansion of the McDonald model, but the McDonald model is more accurate. Both the TAVP and the McDonald model are orders of magnitude faster than the detailed Biot-Savart model. As we show in [2], it is possible to derive an algorithm to calculate $a_0^{(n)}$, so the McDonald model can easily be expanded to arbitrary orders. It has now been implemented in the AE code as an alternative to the TAVP model.

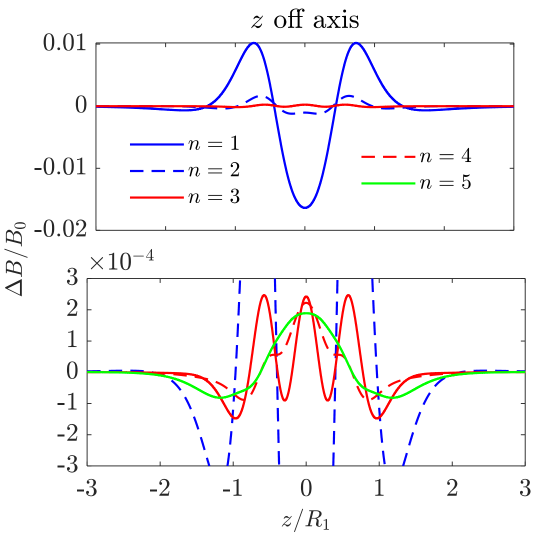

Figure 8.3: The deviation as defined in [2] of different orders of the McDonald model relative to a detailed Biot-Savart model along a path half a radius off the axis (z off). The on-axis plot is not shown, as the McDonald model is exact on-axis. From [2].

Chapter 9

Conclusion and Outlook

The construction of the ALPHA-g experiment [1] with emphasis on the magnet system, and how the experiment will measure antigravity was described in chapter 6. The experiment contains two types of atom traps, which will be used for an up/down and a 1% measurement respectively. The initial focus of the collaboration has been on synthesising and trapping antihydrogen in the bottom full strength region, where an up/down measurement could be made.

The results obtained during the commissioning of ALPHA-g were described in chapter 7 with a special focus on the magnet system, and they can be summarised as follows: already in its first year on construction, before LS2, we managed to demonstrate trapping of positrons and antiprotons in ALPHA-g. With the lack of antiprotons during LS2, it was not until august 2021 that ALPHA-g would be operational again. The experiment was then in an highly improved operational state compared to 2018, which quickly allowed us to trap antiprotons and positrons and attempt antihydrogen creation. Unfortunately, the attempts did not yield any positive results. This can be blamed on the high positron temperatures, which were likely caused by the existence of patch potentials on the electrodes. As per summer 2022, antihydrogen is being created in ALPHA-g, and SoB, MAB, MGB and the octupoles are operational and used in attempts to trap it.

Although the capture solenoids in ALPHA-g have proved unable to reach full current, which worsens our abilities to prepare charged particles for antihydrogen production compared to ALPHA-2, the rest of the internal magnet system is in a good state to trap antihydrogen and later perform an up/down measurement of antigravity. Regarding the 1% measurement, the experiment will have to be improved based on our future experience with the up/down measurement, but we have not yet encountered any problems that will definitively prevent us from achieving 1% precision. The required heat cycles of the external solenoid complicate the development of particle work in ALPHA-g, but it is not a fundamental issue.

The recent results obtained with the ALPHA-2 experiment, described in chapter 5, can be summarised as follows: the latest measurements of the 1S–2P transition [3] improve on the first measurement of the 1S–2P$_c$ transition [9] and demonstrate the first measurement of the 1S–2P$_f$ transition. The measurements agree with the hydrogen spectrum, and the average of the deviations tests CPT symmetry to the $16 \cdot 10^{-9}$ level. The 1S–2P$_f$ transition is suitable for laser cooling of antihydrogen, which we demonstrated recently [4], but its full potential has yet to be explored. It is expected that laser cooling will contribute significantly to improve the precision of future measurements of antihydrogen as described below. Finally, we have demonstrated how laser cooled beryllium ions can be used to sympathetically cool positrons and reduce their temperature from 17 to 7 K in ALPHA-2 [84]. When implemented with antihydrogen production, the colder positrons are expected to yield an increase in the number of trapped antihydrogen atoms by a factor of 5.

The ALPHA-2 atom trap is currently in the process of being upgraded to ALPHA-3. The upgrade includes the addition of a metrology laboratory, where a caesium fountain clock and a hydrogen MASER will improve the laser system significantly. Together with the demonstrated laser cooling and the production of more and colder antihydrogen, there is potential to improve the precision of the 1S–2S measurement with a factor of 500 and thereby achieve hydrogen-like precision. In addition, the developed techniques would contribute to the 1% gravity measurement. Both hydrogen-like 1S–2S precision and the 1% gravity measurement are future goals of the ALPHA collaboration. Finally, the ALPHA apparatus has the potential to be loaded with hydrogen atoms [114], which would enable direct comparisons of hydrogen and antihydrogen.

Appendix A

Appendix

A.1 Detailed Magnet Tables

Table A.1 and A.2 list detailed information about the BNL magnets. The tables contain the following information

- I_{max}: the maximum current that can be sent through the magnet.
- B_{max}: the field determining the trap depth generated at I_{max} as calculated by the detailed Biot-Savart model referred to in section 8.1. For the solenoids, it is the field at the centre of the magnet, and for the octupoles, it is the field at the electrode surface. The values are calculated with a 1 T axial background field and then subtracted 1 T.
- L: the self-inductance of the magnets calculated with equation 6.9 as described in section 6.3.6. The inductance listed is for the individual magnets, and not the total inductance of the top-bottom pair.
- Series: whether the magnet is nominally wired in series with its top/bottom partner. It does not reflect the current setup.
- Length: the axial length of the magnet. These values are derived from the number of turns per layer and the turn spacing.
- R_1: the radius of the centre of the wires in the innermost layer.
- Layers: the number of wire layers.
- Turns per layer: the number of wire turns per layer

- $\Delta\rho$: The radial spacing between the centres of wires in neighbouring layers.
- Δz: The axial spacing between the centre of neighbouring turns in a layer.

For the geometrical properties, the values listed are the "as designed" values. More detailed information should be extracted from the photographs taken during the winding process. For magnets with complicated geometries, it is not possible to quote a single value for some of the parameters. See section 6.3 for further information.

Magnet	I_{max} [A]	B_{max} [T]	L [mH]	Series [y/n]	Length [mm]	R_1 [mm]	Layers	Turns per layer	$\Delta\rho$ [mm]	Δz [mm]
SoB	140	4.6	768	n	249.30	29.5100	9/1/8	–	1.227	–
BgB	140	1.48	75.1	y	43.92	43.0658	8	115	0.5332	0.3853
MAB	140	1.46	51.8	n	34.29	43.0658	8	90	0.5332	0.3853
MBB	140	1.47	51.8	y	34.29	43.0658	8	90	0.5332	0.3853
MCB	140	0.792	17.0	y	16.95	43.0658	8	45	0.5332	0.3853
MEB	140	0.794	17.0	y	16.95	43.0658	8	45	0.5332	0.3853
MFB	140	0.790	17.0	y	16.95	43.0658	8	45	0.5332	0.3853
MEB	140	1.48	51.8	y	34.29	43.0658	8	90	0.5332	0.3853
MGB	140	1.48	51.8	n	34.29	43.0658	8	90	0.5332	0.3853
TrB	140	1.79	75.1	y	43.92	43.0658	8	115	0.5332	0.3853
TrT	140	1.83	75.1	y	43.92	43.0658	8	115	0.5332	0.3853
MGT	140	1.49	51.8	n	34.29	43.0658	8	90	0.5332	0.3853
MFT	140	1.50	51.8	y	34.29	43.0658	8	90	0.5332	0.3853
MET	140	0.778	17.0	y	16.95	43.0658	8	45	0.5332	0.3853
MDT	140	0.799	17.0	y	16.95	43.0658	8	45	0.5332	0.3853
MCT	140	0.795	17.0	y	16.95	43.0658	8	45	0.5332	0.3853
MBT	140	1.50	51.8	y	34.29	43.0658	8	90	0.5332	0.3853
MAT	140	1.49	51.8	n	34.29	43.0658	8	90	0.5332	0.3853
BgT	140	1.53	75.1	y	43.92	43.0658	8	115	0.5332	0.3853
SoT	140	4.7	768	n	249.30	29.5100	9/1/8	–	1.227	–

Table A.1: Table of the properties of the inner magnets in the bottom and top region in ALPHA-g.

Magnet	I_{max} **[A]**	B_{max} **[mT]**	L **[mH]**	**Series [y/n]**	**Length [mm]**	R_1 **[mm]**	**Layers**	**Turns per layer**	$\Delta\rho$ **[mm]**	Δz **[mm]**
AnB	15	332	249	n	71.25	40.7413	12	166	0.508	0.4318
CCB	1	2.27	2.02	n	139.00	33.2015	2	140	0.508	1.0000
COB	1	0.11	0.28	n	71.25	34.2175	2	25	0.508	–
AnT	15	336	249	n	71.25	40.7413	12	166	0.508	0.4318
CCT	1	2.27	2.02	n	139.00	33.2015	2	140	0.508	1.0000
COT	1	0.11	0.28	n	71.25	34.2175	2	25	0.508	–
LOc	1100	110	1.99	–	1472.50	–	2	13	1.227	–
OcB	1100	610	3.37	y	386.40	–	6	11/15/15	1.227	–
OcT	1100	610	3.37	y	386.40	–	6	11/15/15	1.227	–

Table A.2: Table of the properties of the inner magnets in the analysis region and the octupoles in ALPHA-g.

A.2 SC Circuit Diagram

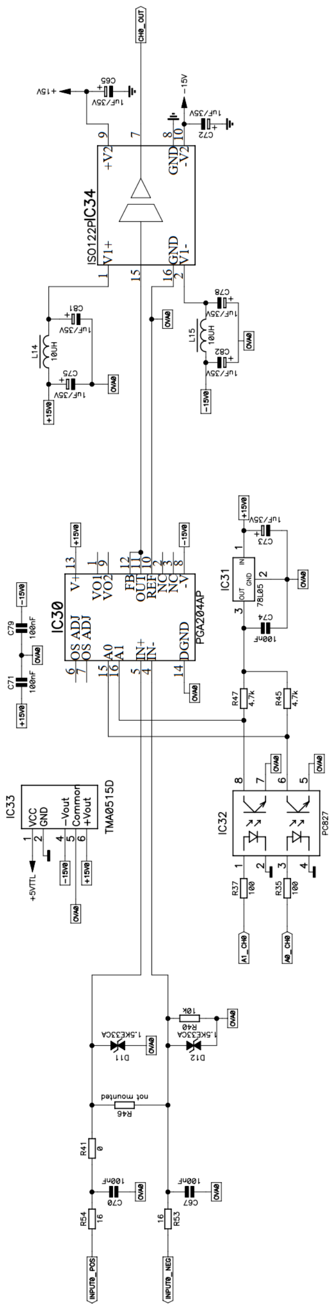

Figure A.1: Circuit diagram for an SC channel.

A.3 Babcock Diagram

Figure A.2 shows a detailed voltage tap diagram of the Babcock magnet, and figure A.3 shows a QD channel in the Danfysik QD system. The mapping from the Danfysik QD system to the Babcock magnet sections and the taps on the connector is as follows:

- Channel 1 Upper, pin 1 and 3, main coil
- Channel 1 Lower, pin 3 and 5, main coil
- Channel 2 Upper, pin 7 and 8, outer coil
- Channel 2 Lower, pin 8 and 9, outer coil
- Channel 4 Upper, pin 9 and 10, boost/shim coil
- Channel 4 Lower, pin 10 and 11, boost/shim coil

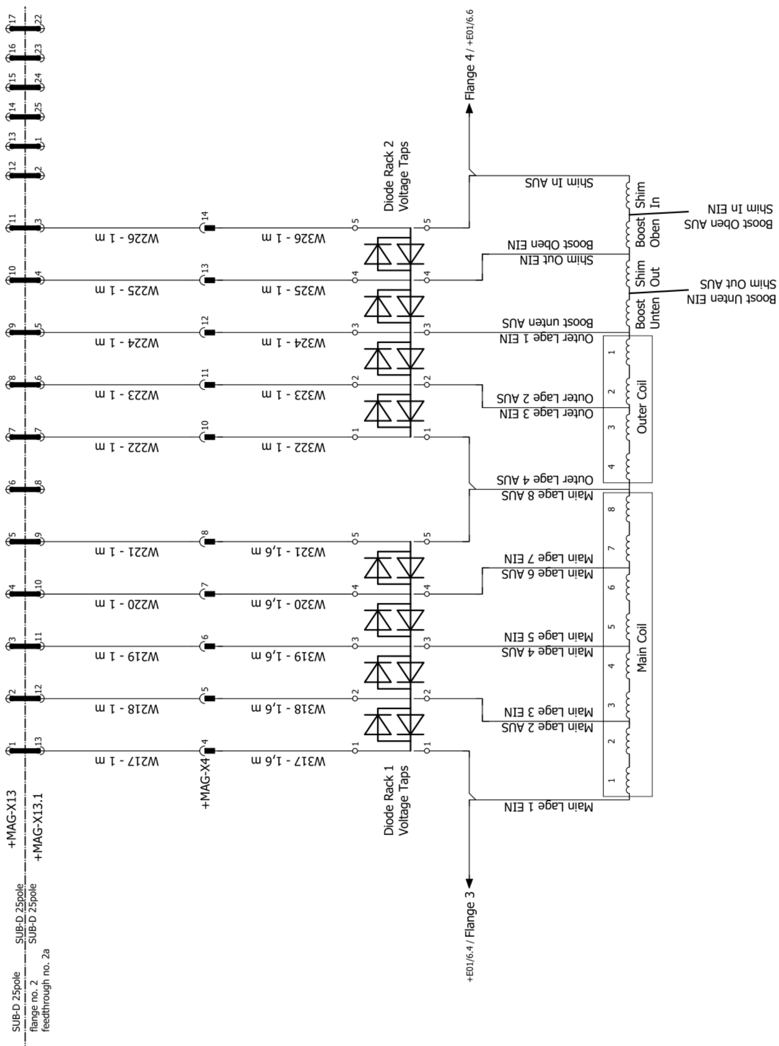

Figure A.2: Diagram of the location of the voltage taps used for QD on the Babcock magnet. The overvoltage protection diodes and the voltage tap connector pin mapping are shown as well.

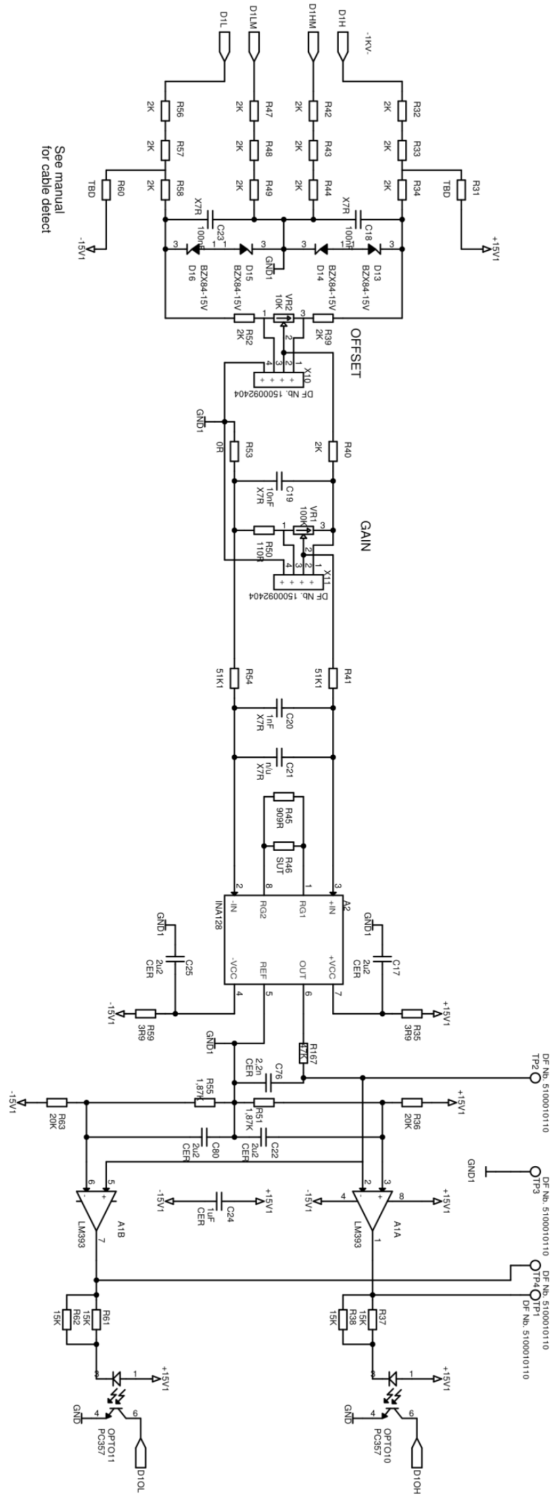

Figure A.3: Circuit diagram of a Danfysik QD channel. Credit: Bilfinger Noell GmbH.

Bibliography

[1] J. S. Hangst et al. Addendum to the ALPHA proposal. 2016.

[2] P. Granum, M. L. Madsen, J. T. K. McKenna, D. L. Hodgkinson, and J. Fajans. Efficient calculations of magnetic fields of solenoids for simulations. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 1034:166706, 2022.

[3] M. Ahmadi et al. Investigation of the fine structure of antihydrogen. *Nature*, 578(7795), 2020.

[4] C. J. Baker et al. Laser cooling of antihydrogen atoms. *Nature*, 592(7852):35–42, 2021.

[5] C. D. Anderson. The positive electron. *Physical Review*, 43(6):491–494, 1933.

[6] G. B. Andresen et al. Trapped antihydrogen. *Nature*, 468:673–676, 2010.

[7] M. Ahmadi et al. An improved limit on the charge of antihydrogen from stochastic acceleration. *Nature*, 529(7586):373, 2016.

[8] M. Ahmadi et al. Observation of the hyperfine spectrum of antihydrogen. *Nature*, 548(7665):66, 2017.

[9] M. Ahmadi et al. Observation of the 1S–2P Lyman-α transition in antihydrogen. *Nature*, 561(7722):211, 2018.

[10] M. Ahmadi et al. Characterization of the 1S–2S transition in antihydrogen. *Nature*, 557(7703), 2018.

[11] P. A. M. Dirac. The quantum theory of the electron. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, 117(778):610–624, 1928.

[12] M. E. Peskin and D. V. Schroeder. *An Introduction To Quantum Field Theory*. Taylor & Francis Inc, 1995.

[13] A. Rubbia. Phenomenology of particle physics I, lecture notes, 2017. Lecture notes, ETH Zürich.

[14] O. Chamberlain, E. Segrè, C. Wiegand, and T. Ypsilantis. Observation of antiprotons. *Physical Review*, 100(3):947–950, 1955.

[15] G. Baur et al. Production of antihydrogen. *Physics Letters B*, 368(3):251–258, 1996.

[16] A. D. Sakharov. Violation of CP variance, C asymmetry, and baryon asymmetry of the universe. *Soviet Physics Uspekhi*, 34(5):392–393, 1967.

[17] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson. Experimental test of parity conservation in beta decay. *Physical Review*, 105(4):1413–1415, 1957.

[18] J. H. Christenson, J. W. Cronin, V. L. Fitch, and R. Turlay. Evidence for the 2π decay of the k_2^0 meson. *Physical Review Letters*, 13(4):138–140, 1964.

[19] A. G. Cohen and D. B. Kaplan. Spontaneous baryogenesis. *Nuclear Physics B*, 308(4):913–928, 1988.

[20] A. Kostelecky and A. J. Vargas. Lorentz and CPT tests with hydrogen, antihydrogen, and related systems. 2015.

[21] M. Charlton, S. Eriksson, and G. M. Shore. Testing fundamental physics in antihydrogen experiments, 2020.

[22] C. G. Parthey et al. Improved measurement of the hydrogen 1S-2S transition frequency. *Physical Review Letters*, 107(20), 2011.

[23] A. Matveev et al. Precision measurement of the hydrogen 1S-2S frequency via a 920-km fiber link. *Physical Review Letters*, 110(23), 2013.

[24] N. Kolachevsky et al. 2S–4S spectroscopy in hydrogen atom: The new value for the Rydberg constant and the proton charge radius. In *AIP Conference Proceedings*. Author(s), 2018.

[25] M. M. Nieto and T. Goldman. The arguments against “antigravity” and the gravitational acceleration of antimatter. *Physics Reports*, 205(5):221–281, 1991.

[26] P. Morrison. Approximate nature of physical symmetries. *American Journal of Physics*, 26(6):358–368, 1958.

[27] P. Scampoli and J. Storey. The AEgIS experiment at CERN for the measurement of antihydrogen gravity acceleration. *Modern Physics Letters A*, 29(17):1430017, 2014.

[28] C. Amsler et al. Pulsed production of antihydrogen. *Communications Physics*, 4(1), 2021.

[29] A. Sótér, H. Aghai-Khozani, D. Barna, A. Dax, L. Venturelli, and M. Hori. High-resolution laser resonances of antiprotonic helium in superfluid 4He. *Nature*, 603(7901):411–415, 2022.

[30] M. Hori, H. Aghai-Khozani, A. Sótér, A. Dax, and D. Barna. Laser spectroscopy of pionic helium atoms. *Nature*, 581(7806):37–41, 2020.

[31] C. Malbrunot et al. The ASACUSA antihydrogen and hydrogen program: results and prospects. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 376(2116):20170273, 2018.

[32] C. Smorra et al. BASE – The Baryon Antibaryon Symmetry Experiment. *The European Physical Journal Special Topics*, 224(16):3055–3108, 2015.

[33] M. J. Borchert et al. A 16-parts-per-trillion measurement of the antiproton-to-proton charge–mass ratio. *Nature*, 601(7891):53–57, 2022.

[34] G. Chardin et al. Proposal to measure the gravitational behaviour of antihydrogen at rest. Technical report, CERN, 2011.

[35] P. Blumer et al. Positron accumulation in the GBAR experiment, 2022.

[36] T. Aumann et al. Experiment proposal for puma: antiprotons and radioactive nuclei. Technical report, CERN, 2019.

[37] J. Walz and T. W. Hänsch. A proposal to measure antimatter gravity using ultracold antihydrogen atoms. *General Relativity and Gravitation*, 36(3):561–570, 2004.

[38] P. Perez and A. Rosowsky. A new path toward gravity experiments with antihydrogen. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 545(1-2):20–30, 2005.

[39] M. Knoop et al. *Trapped Charged Particles: A Graduate Textbook with Problems and Solutions*. World Scientific Publishing Co. Inc., 2016.

[40] L. S. Brown and G. Gabrielse. Geonium theory: Physics of a single electron or ion in a penning trap. *Reviews of Modern Physics*, 58(1):233–311, 1986.

[41] E. D. Hunter, A. Christensen, J. Fajans, T. Friesen, E. Kur, and J. S. Wurtele. Electron cyclotron resonance (ECR) magnetometry with a plasma reservoir. *Physics of Plasmas*, 27(3):032106, 2020.

[42] J. Larmor. On the theory of the magnetic influence on spectra and on the radiation from moving ions. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 44(271):503–512, 1897.

[43] T. M. O'Neil. Cooling of a pure electron plasma by cyclotron radiation. *Physics of Fluids*, 23(4):725, 1980.

[44] J. H. Malmberg, C. F. Driscoll, B. Beck, D. L. Eggleston, J. Fajans, K. Fine, X. P. Huang, and A. W. Hyatt. Experiments with pure electron plasmas. In *AIP Conference Proceedings*. AIP, 1988.

[45] X.-P. Huang, F. Anderegg, E. M. Hollmann, C. F. Driscoll, and T. M. O'Neil. Steady-state confinement of non-neutral plasmas by rotating electric fields. *Physical Review Letters*, 78(5):875–878, 1997.

[46] G. B. Andresen et al. Compression of antiproton clouds for antihydrogen trapping. *Physical Review Letters*, 100(20):203401, 2008.

[47] J. R. Danielson and C. M. Surko. Torque-balanced high-density steady states of single-component plasmas. *Physical Review Letters*, 94(3):035001, 2005.

[48] G. B. Andresen et al. Evaporative cooling of antiprotons to cryogenic temperatures. *Physical Review Letters*, 105(1):013003, 2010.

[49] C. Carruth. *Methods for plasma stabilization and control to improve antihydrogenproduction.* PhD thesis, University of California, Berkeley, 2018.

[50] M. Ahmadi et al. Enhanced control and reproducibility of non-neutral plasmas. *Physical Review Letters*, 120(2):025001, 2018.

[51] D. J. Griffiths. *Introduction to Quantum Mechanics (2nd Edition).* Pearson Prentice Hall, 2004.

[52] D. E. Pritchard. Cooling neutral atoms in a magnetic trap for precision spectroscopy. *Physical Review Letters*, 51(15):1336–1339, 1983.

[53] W. Bertsche et al. A magnetic trap for antihydrogen confinement. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 566(2):746–756, 2006.

[54] D. L. Hodgkinson. *On the Dynamics of Adiabatically Cooled Antihydrogen in an Octupole-Based Ioffe-Pritchard Magnetic Trap.* PhD thesis, University of Manchester, pending publication.

[55] M. Zhong, J. Fajans, and A. F. Zukor. Axial to transverse energy mixing dynamics in octupole-based magnetostatic antihydrogen traps. *New Journal of Physics*, 20(5):053003, 2018.

[56] C. J. Foot. *Atomic Physics.* Oxford University Press, 2005.

[57] M. Horbatsch and E. A. Hessels. Tabulation of the bound-state energies of atomic hydrogen. *Physical Review A*, 93(2):022513, 2016.

[58] G. Breit. An interpretation of Dirac's theory of the electron. *Proceedings of the National Academy of Sciences*, 14(7):553–559, 1928.

[59] W. E. Lamb and R. C. Retherford. Fine structure of the hydrogen atom by a microwave method. *Physical Review*, 72(3):241–243, 1947.

[60] B. R. Holstein. *Topics in Advanced Quantum Mechanics.* Dover Publications, 2013.

[61] V. A. Yerokhin, K. Pachucki, and V. Patkos. Theory of the Lamb shift in hydrogen and light hydrogen-like ions. 2018.

[62] N. Bezginov, T. Valdez, M. Horbatsch, A. Marsman, A. C. Vutha, and E. A. Hessels. A measurement of the atomic hydrogen Lamb shift and the proton charge radius. *Science*, 365(6457):1007–1012, 2019.

[63] J.-P. Karr and D. Marchand. Progress on the proton-radius puzzle. *Nature*, 575(7781):61–62, 2019.

[64] C. Ø. Rasmussen, N. Madsen, and F. Robicheaux. Aspects of insert spectroscopy of trapped antihydrogen atoms. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 50(18):184002, 2017.

[65] S. A. Baird et al. Design study of the antiproton decelerator: AD, 1996.

[66] M. Hori and J. Walz. Physics at CERN's antiproton decelerator. *Progress in Particle and Nuclear Physics*, 72:206–253, 2013.

[67] B. Autin et al. The antiproton decelerator (AD), a simplified antiproton source (feasibility study), 1995.

[68] D. Gamba, L. Bojtár, C. Carli, B. Dupuy, A. Frassier, L. Jørgensen, L. Ponce, and G. Tranquille. AD/ELENA electron cooling experience during and after CERNs Long Shutdown (LS2). *Proceedings of the 13th International Workshop on Beam Cooling and Related Topics*, COOL2021:Russia, 2021.

[69] V. Chonan et al. Extra Low ENergy Antiproton (ELENA) ring and its transfer lines: Design report, 2014.

[70] C. Amole et al. The ALPHA antihydrogen trapping apparatus. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 735:319–340, 2014.

[71] M. Ahmadi et al. Observation of the 1S–2S transition in trapped antihydrogen. *Nature*, 541(7638):506, 2017.

[72] M. Amoretti et al. The ATHENA antihydrogen apparatus. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 518(3):679–711, 2004.

[73] C. Amole et al. Resonant quantum transitions in trapped antihydrogen atoms. *Nature*, 483(7390):439, 2012.

[74] C. Amole et al. An experimental limit on the charge of antihydrogen. *Nature Communications*, 5(1):1, 2014.

[75] C. Amole et al. Description and first application of a new technique to measure the gravitational mass of antihydrogen. *Nature Communications*, 4(1), 2013.

[76] M. Ahmadi et al. Antihydrogen accumulation for fundamental symmetry tests. *Nature Communications*, 8(1):1, 2017.

[77] X. Mougeot. Reliability of usual assumptions in the calculation of β and ν spectra. *Physical Review C*, 91(5):055504, 2015.

[78] C. M. Surko and R. G. Greaves. Emerging science and technology of antimatter plasmas and trap-based beams. *Physics of Plasmas*, 11(5):2333–2348, 2004.

[79] M. Sameed. *Laser-Ablated Beryllium Ions for Cold Antihydrogen in ALPHA*. PhD thesis, Swansea University, 2017.

[80] A. Capra et al. Lifetime of magnetically trapped antihydrogen in ALPHA. *Hyperfine Interactions*, 240(1), 2019.

[81] G. Bendiscioli and D. Kharzeev. Antinucleon-nucleon and antinucleon-nucleus interaction. A review of experimental data. *Riv. Nuovo Cim.*, 2007.

[82] C. Amole et al. Silicon vertex detector upgrade in the ALPHA experiment. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 732:134–136, 2013.

[83] A. Hoecker et al. TMVA - Toolkit for Multivariate Data Analysis. 2007.

[84] C. J. Baker et al. Sympathetic cooling of positrons to cryogenic temperatures for antihydrogen production. *Nature Communications*, 12(1), 2021.

[85] J. Jones. *Sympathetically Laser-Cooled Positron Plasmas for Antihydrogen Formation*. PhD thesis, Swansea University, 2022.

[86] G. B. Andresen et al. Antiproton, positron, and electron imaging with a microchannel plate/phosphor detector. *Review of Scientific Instruments*, 80(12):123701, 2009.

[87] E. D. Hunter, J. Fajans, N. A. Lewis, A. P. Povilus, C. Sierra, C. So, and D. Zimmer. Plasma temperature measurement with a silicon photomultiplier (SiPM). *Review of Scientific Instruments*, 91(10):103502, 2020.

[88] S. Eriksson. Precision measurements on trapped antihydrogen in the ALPHA experiment. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 376(2116):20170268, 2018.

[89] R. Pohl et al. The size of the proton. *Nature*, 466(7303):213–216, 2010.

[90] A. Beyer et al. The Rydberg constant and proton size from atomic hydrogen. *Science*, 358(6359):79–85, 2017.

[91] A. P. Mills and M. Leventhal. Can we measure the gravitational free fall of cold rydberg state positronium? *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 192(1-2):102–106, 2002.

[92] A. Antognini et al. Studying antimatter gravity with muonium. *Atoms*, 6(2), 2018.

[93] M. Johnson. *Design and Commissioning of Beamlines for the ALPHA Antihydrogen Experiment.* PhD thesis, The University of Manchester, 2019.

[94] C. J. Baker et al. Design and performance of a novel low energy multi-species beamline for the alpha antihydrogen experiment. Pending publication.

[95] A. Capra et al. Design of a radial TPC for antihydrogen gravity measurement with ALPHA-g, 2016.

[96] G. Smith. Characterization and analysis development for the ALPHA-g barrel scintillator. Master's thesis, 2021.

[97] C. So, J. Fajans, and W. Bertsche. The ALPHA-g antihydrogen gravity magnet system. *IEEE Transactions on Applied Superconductivity*, 30(4):1–5, 2020.

[98] H. Nagaoka. Note on the self-inductance of solenoids. *Proceedings of the Tokyo Mathematico-Physical Society. 2nd Series*, 4(16):314–325, 1908.

[99] L. Lorenz. Ueber die fortpflanzung der electricität. *Annalen der Physik und Chemie*, 243(6):161–193, 1879.

[100] A. T. A. M. de Waele, P. P. Steijaert, and J. Gijzen. Thermodynamical aspects of pulse tubes. *Cryogenics*, 37(6):313–324, 1997.

[101] A. T. A. M. de Waele, P. P. Steijaert, and J. J. Koning. Thermodynamical aspects of pulse tubes II. *Cryogenics*, 38(3):329–335, 1998.

[102] A. T. A. M. de Waele. Basic operation of cryocoolers and related thermal machines. *Journal of Low Temperature Physics*, 164(5-6):179–236, 2011.

[103] G. Hudson and K. Bouwknegt. 4–13 kA DC current transducers enabling accurate in-situ calibration for a new particle accelerator project, LHC. In *2005 European Conference on Power Electronics and Applications*. IEEE, 2005.

[104] M. Bastos et al. High accuracy current measurement in the main power converters of the large hadron collider: tutorial 53. *IEEE Instrumentation & Measurement Magazine*, 17(1):66–73, 2014.

[105] N. Evetts. *Solid-state nuclear magnetic resonance magnetometry at low temperature with application to antimatter gravity experiments by ALPHA*. PhD thesis, University of British Columbia, 2021.

[106] J. H. Malmberg and C. F. Driscoll. Long-time containment of a pure electron plasma. *Physical Review Letters*, 44(10):654–657, 1980.

[107] C. J. Baker et al. The effects of patch potentials in penning-malmberg traps. Pending publication.

[108] G. B. Andresen et al. Production of antihydrogen at reduced magnetic field for anti-atom trapping. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 41(1):011001, 2007.

[109] J. Simpson, J. Lane, C. Immer, and R. Yuongquist. Simple analytic expressions for the magnetic field of a circular current loop. https://ntrs.nasa.gov/citations/20010038494 (accessed on May 1st, 2021), 2001.

[110] J. T. Conway. Exact solutions for the magnetic fields of axisymmetric solenoids and current distributions. *IEEE Transactions on Magnetics*, 37(4):2977–2988, 2001.

[111] J. D. Jackson. *Classical Electrodynamics*. John Wiley & Sons Inc, 1998.

[112] C. Amole et al. Discriminating between antihydrogen and mirror-trapped antiprotons in a minimum-B trap. *New Journal of Physics*, 14(1):015010, 2012.

[113] K. T. McDonald. Expansion of an axially symmetric, static magnetic field in terms of its axial field. https://physics.princeton.edu/ mcdonald/examples/axial.pdf (accessed on May 1st, 2021), 2011.

[114] S. A. Jones. An ion trap source of cold atomic hydrogen via photodissociation of the BaH^+ molecular ion. *New Journal of Physics*, 24(2):023016, 2022.