Microwave Spectroscopy of Magnetically Trapped Atomic Antihydrogen

by

Mohammad Dehghani Ashkezari

M.Sc., Sharif University of Technology, 2008

B.Sc., University of Tehran, 2006

A THESIS SUBMITTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF

Doctor of Philosophy

in the
Department of Physics
Faculty of Science

© Mohammad Dehghani Ashkezari 2014
SIMON FRASER UNIVERSITY
Spring 2014

APPROVAL

Name: Mohammad Dehghani Ashkezari

Degree: Doctor of Philosophy

Title of Thesis: Microwave Spectroscopy of Magnetically Trapped Atomic Antihydrogen

Examining Committee: Dr. Eldon Emberly
Chair, Associate Professor

Dr. Michael Hayden
Senior Supervisor, Professor

Dr. Dugan O'Neil
Supervisor, Associate Professor

Dr. Paul Haljan
Supervisor, Associate Professor

Dr. Howard Trottier
Internal Examiner, Professor

Dr. Guy Savard
External Examiner, Professor
Department of Physics, University of Chicago
Senior Scientist, Argonne National Laboratory

Date Approved: April 03, 2014

Partial Copyright Licence

SFU

Abstract

We have every reason to believe that equal amounts of matter and antimatter were produced in the early universe. Moreover, theory predicts that the laws of physics make no distinction between the two. In this light, the fact that the observable universe is overwhelmingly dominated by matter is inexplicable.

ALPHA is an international project located at CERN involving approximately 40 physicists from 15 different institutions in 7 countries. The primary goal of the collaboration is to study the antihydrogen atom at the highest level of precision possible, and thereby enable comparisons between hydrogen and antihydrogen. Through these comparisons it hopes to improve our understanding of the distinction between matter and antimatter, and perhaps shed some light on the puzzle of why we live in a matter dominated universe. The hyperfine energy intervals of ground-state hydrogen and antihydrogen represent an opportunity for a precision comparison. A discrepancy between the energy levels of these two atomic systems would indicate a major revolution in physics, and in our understanding of the universe.

This thesis describes and interprets the first proof-of-principle spectroscopic measurements performed on magnetically trapped antihydrogen atoms. The experiments were performed by the ALPHA collaboration using microwave radiation tuned to induce transitions between hyperfine levels of ground state antihydrogen atoms. Our observations confirm that positron spin resonance transitions between hyperfine levels of ground state antihydrogen are consistent with expectations for hydrogen to within 4 parts in 10^3. The hyperfine splitting of ground state antihydrogen atoms is also constrained to 1420 ± 85 MHz.

To my loving Maman & Baba.

*"The first principle is that you must not fool yourself -
and you are the easiest person to fool."*

— Richard Feynman

Acknowledgments

Contents

List of Tables

List of Figures

xiv

Preface

In many respects the ALPHA experiment is an unique scientific project. It is a multi-disciplinary physics experiment involving high-energy, plasma, and atomic physics. The ALPHA collaboration is based at CERN, in Geneva, Switzerland. It has embarked upon a physics program that has the potential to address some of the most fundamental open questions in modern science. Having the opportunity of completing a PhD degree in connection with such an intriguing research study was a turning point in my academic life.

Due to the multi-dimensional nature of the experiment, contributions to the ALPHA project typically involve a wide range of skills and activities. I spent 7 months a year at CERN for four years to contribute to the operation of the apparatus and analysis of data. In the first two years that I worked on the project, I contributed to the production and first confinement of antihydrogen atoms. These experiments, which represent a critical milestone on the path to spectroscopy of the antihydrogen atom, are described in the second and third chapters of my thesis.

The real focus of my research, however, has been in the nascent area of anti-atomic spectroscopy. During the third and fourth years of my involvement with ALPHA, I concentrated on development of the apparatus and experimental techniques that enabled us to demonstrate the first-ever resonant microwave interaction with a pure antiatom. I also closely assisted in the analysis of data and preparation of the manuscript *Resonant Quantum Transitions In Trapped Antihydrogen Atoms*, that was published in Nature [1]. This work is discussed in the fourth and fifth chapters, and is further analyzed and interpreted in the sixth chapter of my thesis.

Finally, in the time since the publication of our work, I have contributed to the development of numerical models to better understand our data, and to begin the process of preparing for the next generation of antihydrogen spectroscopy experiments. This work is described in chapter 6 of the thesis. The analysis presented in chapter 6 is largely my own but I have benefited immensely from the advice and suggestions of Prof. Michael Hayden, Dr. Arthur Olin, Dr. Makoto Fujiwara, Dr. Simone Stracka, and Prof. Francis Robicheaux.

Chapter 1

Introduction

Do matter and antimatter truly obey the laws of physics identically? Why is our observable universe so strongly matter dominated? To date, these are open questions that have attracted the attention of physicists for many years. It is believed that the answer to these types of questions are intimately linked to fundamental principles underlying modern physics. An obvious approach to addressing them is to create antimatter in the laboratory and attempt to compare it with matter as precisely as possible.

Antihydrogen ($\bar{\text{H}}$), the bound state of an antiproton ($\bar{\text{p}}$) and a positron (e^+), is the simplest antiatomic system and is the antimatter counterpart of the very well-studied hydrogen atom. A central goal of antihydrogen research is now focused on precision comparison of antihydrogen and hydrogen through spectroscopic measurements. Currently, there are four active experiments involved in antihydrogen research: AEgIS, ALPHA, ASACUSA, and ATRAP. They are all located at the Antiproton Decelerator (AD) facility at the European Organization for Nuclear Research (CERN) in Geneva, Switzerland. World-wide, the AD is a unique facility; it is capable of providing low energy antiprotons in large quantities.

Recently, ALPHA announced the first successful antiatomic spectroscopy experiment, involving microwave driven transitions between the hyperfine levels of the ground state of the antihydrogen atom [1]. This thesis presents details of those experiments and their analysis.

1.1 Motivation for Antihydrogen Study

Tests of symmetries underlying fundamental interactions play a key role in modern physics. CPT (Charge, Parity, Time-reversal) is believed to be a perfect symmetry of Nature. This is supported by the broad success of quantum field theories. According to the CPT theorem any local Lorentz-invariant quantum field theory in a flat spacetime will be invariant under the combined operations of charge-conjugation (C), parity (P), and time-reversal (T). However, after experimental observations of C and CP violations under the weak interaction [2], it seems prudent to also investigate CPT symmetry through experiment.

In 2006 an analysis of the rotation of the cosmic microwave background polarization showed small non-zero signatures of Lorentz and CPT violations at the 1.5σ level [3]. Another search for sidereal time variations in B^0-$\bar{B}^0$ oscillation parameters conducted by the BaBar collaboration, indicated Lorentz and CPT-violating effects at the 2.2σ level [4]. Potentially, experimental tests of the CPT theorem can be further improved through precision comparisons of fundamental atomic parameters of hydrogen and antihydrogen.

The CPT theorem predicts that a particle and its antiparticle counterpart must have identical mass, opposite spin and opposite electric charge. Consequently, the energy levels and atomic spectra of hydrogen and antihydrogen atoms should be identical. This is a precise prediction that can be tested via spectroscopy.

The best candidate for precision optical spectroscopy of hydrogen and antihydrogen is likely the 1S-2S transition. The intrinsic linewidth of this transition is extremely narrow: $\Delta\nu/\nu = 4 \times 10^{-16}$. The most recent and accurate measurement of this transition frequency for atomic hydrogen, which involved comparison to a remote cesium fountain clock, is $\nu_{1S-2S} = 2\ 466\ 061\ 413\ 187\ 018\ (11)$ Hz, which has a relative uncertainty of 4.5×10^{-15} [5, 6, 7]. The ALPHA antihydrogen apparatus did not allow for laser access at the time of this study, and so I will not discuss optical spectroscopy further.

Another very precisely determined quantity is the zero-field ground state hyperfine splitting (HFS) of the hydrogen atom. This is the energy difference between the singlet (F=0) and triplet (F=1) states of the atom in the limit of no externally-applied magnetic field.

Here $\mathbf{F}$ is the quantum number for the total spin $\mathbf{F} = \mathbf{S}_e + \mathbf{S}_p$ of the 1S ground state which is split due to interaction of the electron spin $\mathbf{S}_e$ and the proton spin $\mathbf{S}_p$. The non-zero energy difference between these two states at zero external magnetic field corresponds to the well-known 21 cm line in radioastronomy. The frequency associated with this energy difference is known to an experimental precision of 10^{-12} [8, 9, 10], providing a solid benchmark against which the corresponding splitting in antihydrogen can be compared:

$$\nu_{HFS} = 1\ 420\ 405\ 751.\ 768 \pm 0.001\ \text{Hz}. \tag{1.1}$$

Measurement of the hyperfine splitting of atomic antihydrogen as an experimental probe for testing fundamental symmetries has been discussed for some time by the ALPHA [11] and ASACUSA [12, 17] collaborations. There is no expectation that these measurements will rival those performed on hydrogen any time soon, because of experimental challenges particular to antihydrogen. The most obvious challenges for experiments involving magnetic confinement are those associated with the inhomogeneous magnetic field, which broadens transitions and limits the duration of coherent interactions with applied microwave fields. Another challenge is that to date, experiments with antihydrogen have essentially involved one atom at a time. Despite these challenges, precisions at the 10 kHz level or better are anticipated in the next generation of experiments planned by the ALPHA collaboration.

Given the precision to which one might expect to measure ν_{HFS} for antihydrogen in the near future, one would also like to know at what level the measurement begins to yield useful information, such as imposing a constraint on CPT invariance. Alternatively, one might ask at what level of precision a measurement of ν_{HFS} might begin to rival constraints established by other experiments such as the mass anomaly for the neutral Kaon system (the $K^0 - \bar{K}^0$ mass difference is known to a relative precision $|\Delta m|/m \leq 6 \times 10^{-19}$ [13]). This is a complex and nuanced issue, which is ultimately beyond the scope of this thesis. However it has been argued that the appropriate figure of merit for such comparisons could be the absolute energy of any anomaly rather than the relative precision of the constraint [14, 15]. In this light, the absolute energy of the mass anomaly in the neutral Kaon system is 80 kHz in frequency units. Against this benchmark, one would thus already expect the next generation of ALPHA HFS experiments to push into unexplored territory. Even if these arguments prove to be incorrect, there is no reason to expect identical patterns of symmetry violation in meson-antimeson, baryon-antibaryon, and lepton-antilepton systems. Constraints

provided by hydrogen-antihydrogen comparisons are thus expected to be complementary to tests performed in other sectors.

Yet another motivation for pursuing measurements of ν_{HFS} for antihydrogen is presented in the next section, which enumerates the various factors that contribute to the hyperfine splitting of the hydrogen atom. Of these, the least well understood pertain to the internal structure of the nucleon. Corrections arising from internal structure effects contribute to ν_{HFS} at the 50 kHz level [16]. At this level, measurements of ν_{HFS} for antihydrogen begin to probe antibaryon structure effects.

1.2 The Ground State Hyperfine Separation of Atomic Hydrogen

The many significant figures required to specify the measured value of the hydrogen atom ground state hyperfine splitting reflect a history that is rich in physics. The first three digits (see Fig. 1.1) are predicted by the simple and elegant Fermi contact interaction formula:

$$
\begin{aligned}
\nu_{HFS} &= \frac{4}{3}\gamma_p\gamma_e \left|\psi(0)\right|^2 \hbar \\
&= \frac{16}{3}\left(\frac{M_p}{M_p + m_e}\right)^3 \frac{m_e}{M_p}\frac{\mu_p}{\mu_N}\alpha^2 c R_\infty
\end{aligned}
\tag{1.2}
$$

where $\psi(0)$ is the ground state wave function of the electron evaluated at the site of the nucleus, γ_p and γ_e denote the proton and electron gyromagnetic ratios, M_p and m_e the proton and electron mass, c the speed of light, μ_p the magnetic moment of the proton, μ_N the nuclear magneton, α the fine structure constant, and R_∞ the Rydberg constant. From Eq. 1.2, one sees that the ground state hyperfine splitting constant $a = h\nu_{HFS}$ of hydrogen is proportional in leading order to the proton magnetic moment μ_p [18]. This was in fact, one of the early motivations for pursuing precision measurements of ν_{HFS} for the antihydrogen atom. Until recently, the magnetic moment of the antiproton was only known to a precision of 3000 parts per million (ppm) [19, 20], and so even a relatively crude measurement of ν_{HFS} for antihydrogen would have improved this value. However, the ATRAP collaboration

announced a new measurement of the antiproton magnetic moment in 2013 with an uncertainty of 4.4 ppm [21]. This result is now sufficiently precise for it to be used in experiments that probe departures of ν_{HFS} from theory, as described below.

$$\boldsymbol{\nu}_{HFS} = \boxed{142}\ \boxed{0}\ \boxed{40}\ \boxed{575}\ \boxed{1.768}\ \text{Hz}$$

Figure 1.1: The ground state hyperfine splitting of the hydrogen atom. The various digits have been color-coded to indicate historical progress in measurement precision.

To four digits of precision Eq. 1.2 predicts $\nu_{HFS} = 1.419$ GHz. The first experimental observation of the hyperfine separation of ground state hydrogen was performed by Nafe, Nelson and Rabi (1947) in an atomic-beam experiment [22]. A similar but more precise experiment conducted by Nagel, Julian and Zacharias [23] yielded $\nu_{HFS} = 1.420\ 410(6)$ GHz, in clear disagreement with the prediction of Eq. 1.2.

Equation 1.2 is valid only if Dirac theory is strictly correct. In particular if the electron g-factor is not exactly -2, then the right-hand side of Eq. 1.2 must be multiplied by $-g/2$. Today we know that radiative corrections cause g to depart from -2, significantly [18]. Conclusive experimental evidence of the anomalous electron g-factor was announced by Kusch and Foley in 1947 [24]. Accounting for this factor brings Eq. 1.2 into accord with experiments out to the 4$^{\text{th}}$ significant figure (Fig. 1.1). Additional corrections to Eq. 1.2 can be made, including further contributions from QED, the internal structure of the proton, the hadronic vacuum polarization, and the weak interaction. The best calculations today reproduce ν_{HFS} out to the 6$^{\text{th}}$ digit [25]. Interestingly, it is the corrections related to the internal structure (charge and magnetic moment distributions) of the nucleon (proton/antiproton) that presently limit calculations of ν_{HFS}. Even though the charge of the proton is well known, its internal charge distribution is not. Measurements obtained by different methods vary by as much as 4% [26, 27]. As noted at the end of the previous section, experiments that probe ν_{HFS} for antihydrogen at this level of precision will inform our understanding of the internal structure of the antiproton. In this sense, experiments that probe ν_{HFS} for

antihydrogen provide access to information that is not obtained from high precision measurements of individual particle parameters such as their masses and magnetic moments.

By 1960 magnetic resonance experiments employing beams of hydrogen atoms had resolved ν_{HFS} out to approximately the 9^{th} digit in Fig. 1.1. At this point in time the hydrogen maser was invented. This dramatically increased the coherence time for atom-field interactions, and enabled much more precise experimental determination of ν_{HFS}. Today it is known to 13 digits of precision, limited by the reproducibility of coatings applied to the storage bulbs in which atoms are confined.

The standard model of particle physics predicts that ν_{HFS} must be identical for the hydrogen and antihydrogen atoms. In recent years Kostelecký and various collaborators have published a series of papers that outline an extension to the standard model (SME) of particle physics [28, 29, 30, 31, 32]. This extension amounts to a framework that allows for both CPT and Lorentz invariance violating terms in the Lagrangian of effective quantum field theories. This leads to corrections to the 1S-2S transition and various transitions between hyperfine sublevels. Some of these corrections are asymmetric for hydrogen and antihydrogen atoms. While the SME does not make predictions about the size or even the existence of these corrections, it does provide guidance as to where one ought to look in a search for discrepancies. Moreover, once experimental constraints are established for a particular measurement, it provides a framework for understanding the implications of these constraints for other systems.[1]

1.3 Thesis Overview

The ALPHA collaboration has a goal of performing precision microwave spectroscopy experiments on trapped antihydrogen atoms. Even as this thesis is being written an entirely new trapping apparatus (which I helped to construct) is being commissioned. While it is unlikely that antihydrogen spectroscopy experiments will challenge hydrogen spectroscopy

---

[1]The SME allows one to set stringent limits on CPT violation using data from matter-only experiments (through the intrinsic link to Lorentz invariance [33]). The value of matter-antimatter comparisons, such as those envisioned for hydrogen and antihydrogen, is that they are direct and model independent.

for precision any time soon, it is likely that they will explore very valuable regions of parameter space. In particular, direct comparisons of matter-antimatter systems will provide model independent tests of CPT invariance.

In this thesis I report a series of experiments in which we drive microwave transitions between hyperfine levels of ground state antihydrogen atoms. These experiments were the first of their kind ever performed and set the stage for precision microwave spectroscopy experiments on antihydrogen in the near future.

The organization of this thesis is as follows. Chapter 2 reviews general background information about the ALPHA apparatus. This is needed to understand the experimental context. Readers who are already familiar with the ALPHA apparatus can skip this chapter. Chapter 3 then discusses the experimental methods used for production and confinement of antihydrogen in the ALPHA apparatus. Again this is background information insofar as the main focus of the thesis is concerned, but I did contribute to many of the experiments that are described. Chapter 4 introduces methods for injecting microwaves into the apparatus. Also described is a novel diagnostic tool for in-situ measurement of static magnetic fields along the axis of the ALPHA trap. Chapter 5 is dedicated to the microwave spectroscopy experiments in which positron spin-flip transitions between hyperfine levels of ground state antihydrogen atoms are induced. Finally, Chapter 6 presents further analysis of the experimental observations reported in Chapter 5. It sheds light on systematic uncertainties associated with our proof of principle spectroscopy experiments. This analysis is based on a numerical model that accounts for atom dynamics, trap fields, and microwave power levels to determine the probability of a spin-flip transition occurring. Chapter 7 starts with a brief discussion of future microwave spectroscopy experiments that could be performed in the next-generation ALPHA apparatus, and then concludes with a summary of the thesis.

Chapter 2

ALPHA Apparatus

This chapter surveys key aspects of the apparatus used to perform the experiments reported in this thesis. First, an overview of the apparatus is presented to familiarize the reader with the fundamental components of the machine. Next, the traps designed for confinement of charged particles and neutral antihydrogen atoms are discussed. Finally, various diagnostics and detection systems that are employed in our experiments are described. A much more extensive review of the apparatus and its capabilities can be found in Ref. [43].

2.1 Overview

In a nutshell, the ALPHA apparatus is a machine specifically designed for synthesizing and trapping antihydrogen atoms. It involves two basic components: a Penning trap that is used to confine and manipulate the (charged) antiprotons ($\bar{p}$) and positrons (e^+) needed to make antihydrogen ($\bar{H}$), and a magnetic minimum trap to confine (some of) the synthesized atoms. Confinement of neutral antihydrogen atoms is accomplished by exploiting the interaction of their magnetic dipole moments with an appropriately designed inhomogeneous magnetic field. A schematic drawing of the main section of the apparatus is shown in Fig. 2.1.

Antiptotons are extracted from the Antiproton Decelerator (AD) ring and positrons are provided by an adjacent machine, a positron accumulator (dedicated to the ALPHA apparatus). These constituents are injected axially into the experiment from either end. We refer to these as the upstream and downstream ends, respectively, in reference to the end from which antiprotons enter. The heart of the apparatus provides the Ultra-High Vacuum
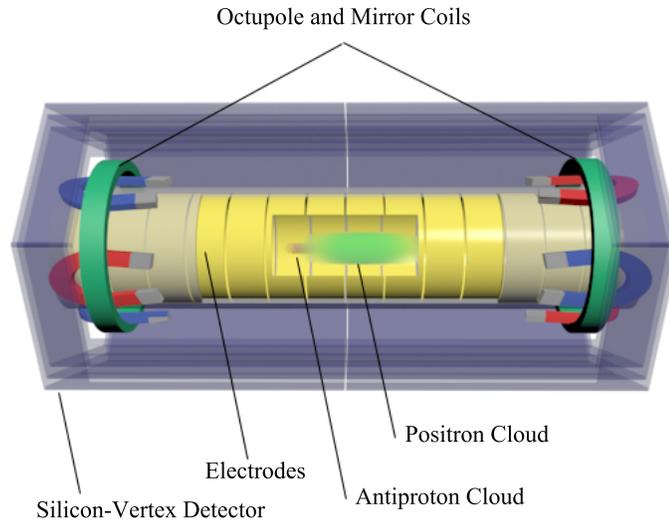
Figure 2.1: A schematic cutaway view of the trapping section of the ALPHA apparatus, showing the Penning trap electrodes, magnetic minimum trap, and the annihilation vertex detector. The magnetic minimum trap consists of a pair of mirror coils, an octupole and a solenoid (not shown).

(UHV) and cryogenic environments needed to confine charged plasmas and antihydrogen atoms for long periods of time. Figure 2.2 shows an illustration of the entire apparatus. All of the main components of the apparatus will be described in this chapter.

2.2 Penning Trap: Charged Particle Confinement

Penning traps confine charged particles using electric and magnetic fields. In this section a brief review of the theory of an ideal Penning trap is presented. This is followed by a description of the ALPHA Penning trap.

2.2.1 Theory of an Ideal Penning Trap

Charged particles can be confined provided that an appropriate three-dimensional potential is established. The Penning trap accomplishes this using static electric and magnetic fields. The forces to which charged particles are exposed inside the trap can depend on the coordinates in many ways; however, analysis of the particle motion is simplified when a harmonic potential is used. A general quadratic form for the potential energy (in Cartesian
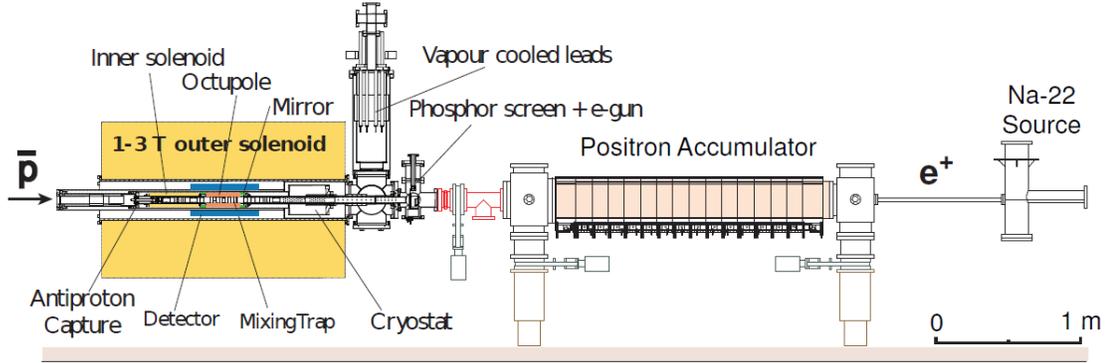
Figure 2.2: The full ALPHA apparatus, including the positron accumulator. Antiprotons enter from the left (upstream end) and positrons enter from the right (downstream end). Antihydrogen is formed in the mixing trap (image from ALPHA collaboration).

coordinates) is thus:

$$U = \zeta \left(\alpha x^2 + \beta y^2 + \gamma z^2 \right),$$ (2.1)

where the parameters ζ, α, β, and γ are constants. For confinement of a charged particle using electrostatic fields, U can be expressed in terms of the electrostatic potential Φ:

$$U = Q\Phi.$$ (2.2)

Moreover, in order to satisfy Laplace's equation

$$\nabla^2 \Phi = 0$$ (2.3)

the parameters in Eq. 2.1 must satisfy $\alpha + \beta + \gamma = 0$. For cylindrically symmetric potentials, this leads to $\alpha = \beta = 1$ and $\gamma = -2$, reducing U to the form:

$$U = \zeta \left(x^2 + y^2 - 2z^2 \right).$$ (2.4)

An ideal Penning trap is formed by superimposing a homogeneous axial magnetic field $\mathbf{B} = (0, 0, B_0)$ and the electric field $\mathbf{E} = -\nabla\Phi$ associated with a cylindrically symmetric quadratic electrostatic potential:

$$\Phi = \frac{U_0}{2d^2} \left(2z^2 - x^2 - y^2 \right),$$ (2.5)

where U_0 is the depth of the potential well and d is the characteristic trap length. The motion of a particle of charge Q and velocity $\mathbf{v} = (v_x, v_y, v_z)$, in the combined electric $\mathbf{E}$ and magnetic $\mathbf{B}$ fields is governed by the Lorentz force:

$$\mathbf{F} = Q\left(\mathbf{v} \times \mathbf{B} - \nabla\Phi\right). \tag{2.6}$$

It is characterized by three distinct oscillatory motions. An oscillation parallel to the magnetic field at frequency f_z, an oscillation perpendicular to the magnetic field at f'_c (a 'cyclotron' motion) and another oscillation around the trap axis at frequency f_m known as 'magnetron' motion. The modified cyclotron frequency f'_c is closely related to the free space cyclotron frequency f_c. In terms of magnitude, the axial and magnetron frequencies tend to be much smaller than the cyclotron frequency. For instance, in a 2 cm long, 20 V deep well in a 1 T magnetic field the free space cyclotron frequency of an electron ($f_c = \frac{QB}{2\pi m}$) is 28 GHz while the axial ($f_z = \sqrt{\frac{QU_0}{2m\pi^2 d^2}}$) and magnetron ($f_m = \frac{f_z^2}{2f'_c} \approx \frac{f_z^2}{2f_c}$) frequencies are 15 MHz and 4 KHz, respectively. These values are characteristics of those encountered in ALPHA experiments, although as described below we do not use a precise quadratic potential and we work with many particles in the trap. A full treatment of the Penning trap can be found in [34, 35].

2.2.2 The ALPHA Penning Trap

ALPHA employs a Penning trap that comprises an axial array of thirty-four cylindrical aluminum electrodes (Fig. 2.3). Technically, this arrangement is known as Penning-Malmberg trap. The potentials it produces are not perfectly quadratic, but the particle motions are largely similar to those described above. To avoid oxidation the electrodes are plated with a thin layer of gold. The electrodes range in (inner) diameter from 36.6 mm to 44.6 mm and have varying lengths. They are mounted inside a vacuum chamber surrounded by a liquid helium cryostat. The electrodes are electrically isolated from one another by ruby spheres and from the surrounding vessel by ceramic spacers. The electric potentials applied to these electrodes are individually controlled. They are set by a programmable 16-bit National Instruments PXI-6733 digital-to-analogue converter. The output of the PXI-6733 (± 10 V) is then amplified using external amplifiers. Depending on the electrode, the amplification factor is either 7.2 or 14. The entire stack of electrodes is immersed in a 1 T magnetic field directed along the axis of the electrodes (see Fig. 2.2), providing radial confinement of

charged particles. Then, by manipulating the electric potentials applied to the electrodes, clouds of antiprotons and positrons can be confined independently and transfered from one axial location to another.
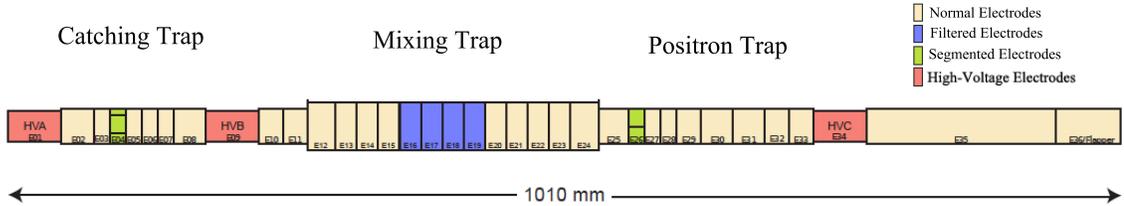
Figure 2.3: Electrodes of the ALPHA Penning trap. Antiprotons and positrons are confined and prepared in the catching and positron traps, respectively. Antiproton and positron clouds are then transfered into the mixing trap at the center of the array. Some of the electrodes are electronically filtered more heavily than others (blue). A couple electrodes are segmented to apply rotating potentials (green), and a few electrodes are designed so that they can sustain high-voltages (red) (figure from ALPHA collaboration).

The thirty-four electrodes are grouped in three sections. The 'catching trap' where antiprotons delivered by the AD are caught and cooled; the 'positron trap' where positrons from the accumulator are caught and cooled; and the 'mixing trap' where antihydrogen atoms are formed by gentle mixing of antiproton and positron clouds. The mixing trap is located between the catching and positron traps.

The catching trap consists of eleven electrodes (electrodes 1-11). Two of the electrodes (HVA, HVB) are designed so that high voltages (in the kV range) can be applied. These electrodes are used to catch the injected antiproton bunch. In the catching trap there is also one electrode that is azimuthally segmented into six independently controlled sectors. The potentials applied to these sectors are manipulated so as to apply a 'rotating wall' potential, which results in radial compression of the confined particle cloud [36].

At the heart of the apparatus, thirteen electrodes form the mixing trap. The diameter of the electrodes in this region is larger (44.6 mm) than in the catching and positron traps. This is done to allow antihydrogen atoms to approach the windings of the octupole magnet

as close as possible without hitting a wall, and thereby obtain the maximum possible magnetic trap depth. Using a novel ultra-thin design, the maximum thickness of the electrodes in the mixing trap is 1.5 mm [43].

2.3 Magnetic Trap: Neutral Atom Confinement

A neutral atomic system like the antihydrogen atom cannot be confined using static electric fields. However, any atom with a non-zero magnetic moment can be trapped using appropriately designed static magnetic fields. The energy of an atom with a magnetic dipole moment $\boldsymbol{\mu}$, in a magnetic field $\boldsymbol{B}$ is:

$$U = -\boldsymbol{\mu} \cdot \boldsymbol{B}. \tag{2.7}$$

Accordingly, the force experienced by the atom is:

$$\mathbf{F} = \nabla(\boldsymbol{\mu} \cdot \mathbf{B}) = \mu_z \nabla \mathbf{B} \tag{2.8}$$

where μ_z is the projection of the magnetic moment $\boldsymbol{\mu}$ onto the field direction. The Maxwell equations do not allow for a static magnetic field maximum in free space; however a minimum is allowed and the force expressed by Eq. 2.8 can be used to trap atoms in the vicinity of this minimum. This is possible when the direction of the magnetic moment is opposite to that of the magnetic field; ie. $\mu_z < 0$. As will be discussed later, this is the case for two of the hyperfine levels of the ground state antihydrogen atom. Such states are known as 'low field seeking states.'

In the ALPHA apparatus, the three-dimensional minimum in the magnetic field is created by superposition of the fields produced by a pair of mirror (or 'pinch') coils with those produced by an octupolar magnet. Figure 2.4 shows the magnetic field profile when the magnetic trap is energized and superimposed on a uniform 1 Tesla axial magnetic field. The mirror coils are situated either side of the center of the trap at an effective distance of 137 mm, generating a minimum along the axis of the Penning trap (Fig. 2.5).

In general, the radial minimum in the magnetic field can be constructed from a multipolar field $B_\perp$ oriented perpendicular to the uniform axial field:
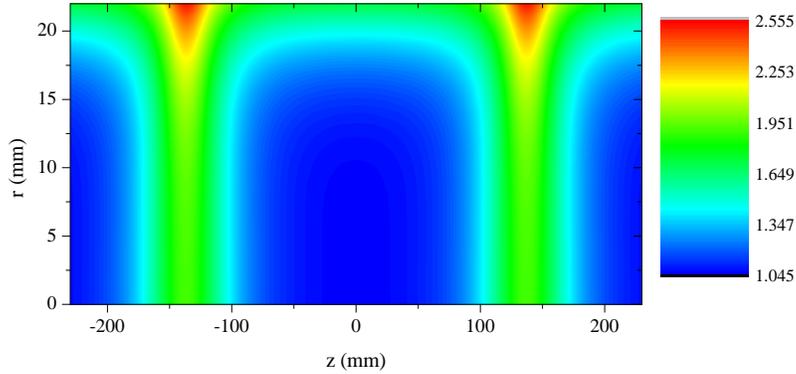
Figure 2.4: Magnetic field strength in the ALPHA apparatus (in Tesla) when the magnetic minimum trap is energized. The absolute minimum field occurs near the centre of the trap. Note that the radial (r) and axial (z) axes are plotted at very different scales.

$$B_\perp = B_w \left(\frac{r}{R_w}\right)^l \tag{2.9}$$

where B_w is the magnitude of the field at radius R_w. The multipolar order l characterizes both the azimuthal and radial variation of the magnetic field. In the case of an octupole $l = 3$. The magnetic fields in the ALPHA apparatus are very similar to a typical Ioffe-Pritchard configuration [37] except that in the classic Ioffe-Pritchard configuration the transverse magnetic field is generated by a quadrupolar magnet ($l = 1$) [38, 39]. The transverse magnetic fields associated with the multipolar magnet can produce undesirable perturbative effects on the dynamics of charged particles, such as radial transport [40, 41]. The key motivation for using an octupole magnet instead of a quadrupole is that the transverse magnetic field near the axis is much smaller and thus perturbative effects are less pronounced. An added benefit that will become apparent later in this thesis is that the magnetic field profile associated with an octupole provides a magnetic environment that is more conducive to spectroscopic measurements on trapped atoms.

The magnetic well depth, ΔB, is determined by the field at bottom of the well (B_{min}), and the field at the electrode wall. Therefore, ΔB is given by

$$\Delta B = \sqrt{B_w^2 + B_{min}^2} - B_{min} \tag{2.10}$$

At typical current settings of 900 A for the octupole and 650 A for the mirror coils, the trap depth, ΔB, is approximately 0.8 T or 0.5 K (in temperature units) for ground state antihydrogen atoms. The magnetic well is constructed so that the direction of the octupole field (responsible for radial confinement) is perpendicular to that of the axial field while the direction of the mirror coil fields (responsible for axial confinement) are parallel to the axial field. So the octupole and axial fields add in quadrature while the mirror and axial fields add linearly. Consequently, when the trapping fields are ramped down the radial depth of the well collapses faster than the axial depth of the well. This means that atoms almost always escape towards the walls rather than through the ends.
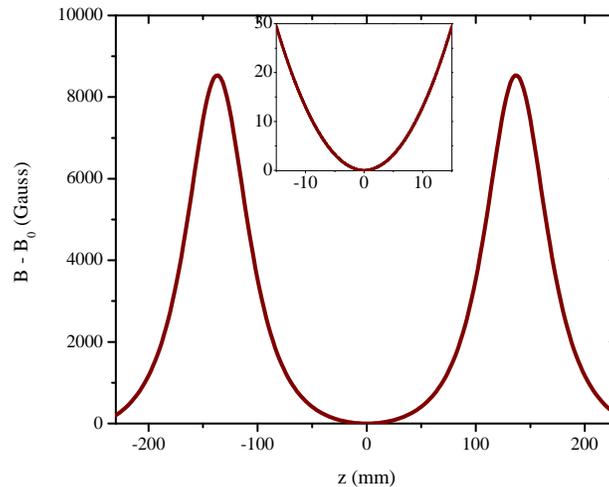
Figure 2.5: Variation of magnetic field strength along the axis of the ALPHA trap. The field B_0 at the centre of the trap has been subtracted, and so the field shown here comes predominantly from the mirror coils. There is some contribution from the ends of the octupole coil; see Fig. 2.1. The inset is similar to the main figure applies to a shorter axial range around the centre of the trap.

.

The magnets are wound using a superconducting alloy of niobium-titanium and are immersed in a bath of liquid helium at 4 K [38]. If the temperature of the superconductors exceeds their critical temperature, they become resistive resulting in Joule heating. This effect, known as a 'quench,' can potentially cause serious damage to the magnets because of

the enormous stored energy in the field. To avoid this, a quench protection system (QPS), based on a Brookhaven design by Ganetis [42], continuously monitors the voltage across the magnets. If the measured voltage goes beyond a predefined threshold, the QPS turns off the magnets and extracts the stored energy in a controlled fashion. Using an isolated gate bipolar resistor (IGBT) switch, the high current flowing in the magnet is rapidly diverted to a bank of resistors.

Controlled rapid shut down of the magnetic trapping fields is a key feature of the AL-PHA apparatus [38]. The magnets can be ramped down with a time constant of 9 ms [43]. This enables detection of the released antihydrogen atoms within a very short time interval, and hence, reduces the background from cosmic rays.

2.3.1 Magnetic Field Homogeneity

The primary design target for the magnetic field configuration in the ALPHA apparatus is the efficient confinement of antihydrogen. From a spectroscopist's point of view, however, exposing atoms to a highly non-uniform magnetic field imposes limitations on the precision to which atomic energy intervals can be determined. Figures 2.5 and 2.6 illustrate the level of magnetic field inhomogeneity near the trap center. Figure 2.5 shows the magnetic field strength along the axis of the ALPHA trap with respect to the field strength at the centre (B_0). Note the rapid increase in field strength as one moves away from the centre in the axial direction. The magnetic field 5 mm away from the centre is of order 3 Gauss higher than at the centre (see inset to Fig. 2.5). The homogeneity of the static field near the trap centre is dominated by the two mirror coils, which give rise to a saddle point in $|\mathbf{B}|$ at the origin ($r = z = 0$):

$$|\mathbf{B}| = B_0 + \alpha \left(z^2 - \frac{r^2}{2} \right) \tag{2.11}$$

where $\alpha \sim 16$ T/m^2 when the currents in all magnets are set at their nominal operating values. Superposition of the octupole, mirror, and solenoidal fields leads to the formation of a minimum in the trapping field situated on a circle of radius of order 4 mm centred on the trap axis. Figure 2.6 shows the magnetic field as a function of radius near the trap

centre. As is seen, the minimum magnetic field in the radial direction occurs 4 mm away
from the trap axis. The magnetic field at this point is approximately 0.7 Gauss less than
the magnetic field at the centre.
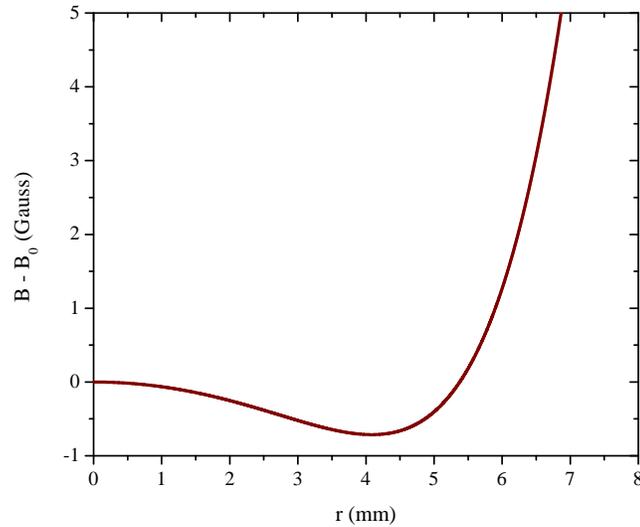
Figure 2.6: Deviation of the static magnetic field B from its value B_0 at the trap centre,
along the radial direction (z=0 mm). The minimum field B_{min} occurs on a circle of radius
4 mm around the trap axis. Above this radius the octupole field starts to dominate.

2.4   Cryostat and Trap Vacuum

Electromagnetic manipulation of plasmas is only possible under high or ultra-high vacuum
(UHV) conditions, where the collision rate with background gases is low. This situation is
exacerbated in the ALPHA apparatus, where antiparticle plasmas are studied. Even lower
pressures are required since collisions with background gas molecules result in annihilation.

The ALPHA cryostat consists of two intersecting cylinders, forming an L-shaped volume
(Fig. 2.7). The horizontal section encloses the Penning trap and superconducting magnets,
and is positioned so that its axis coincides with the axis of the external solenoid (not shown
in Fig. 2.7).
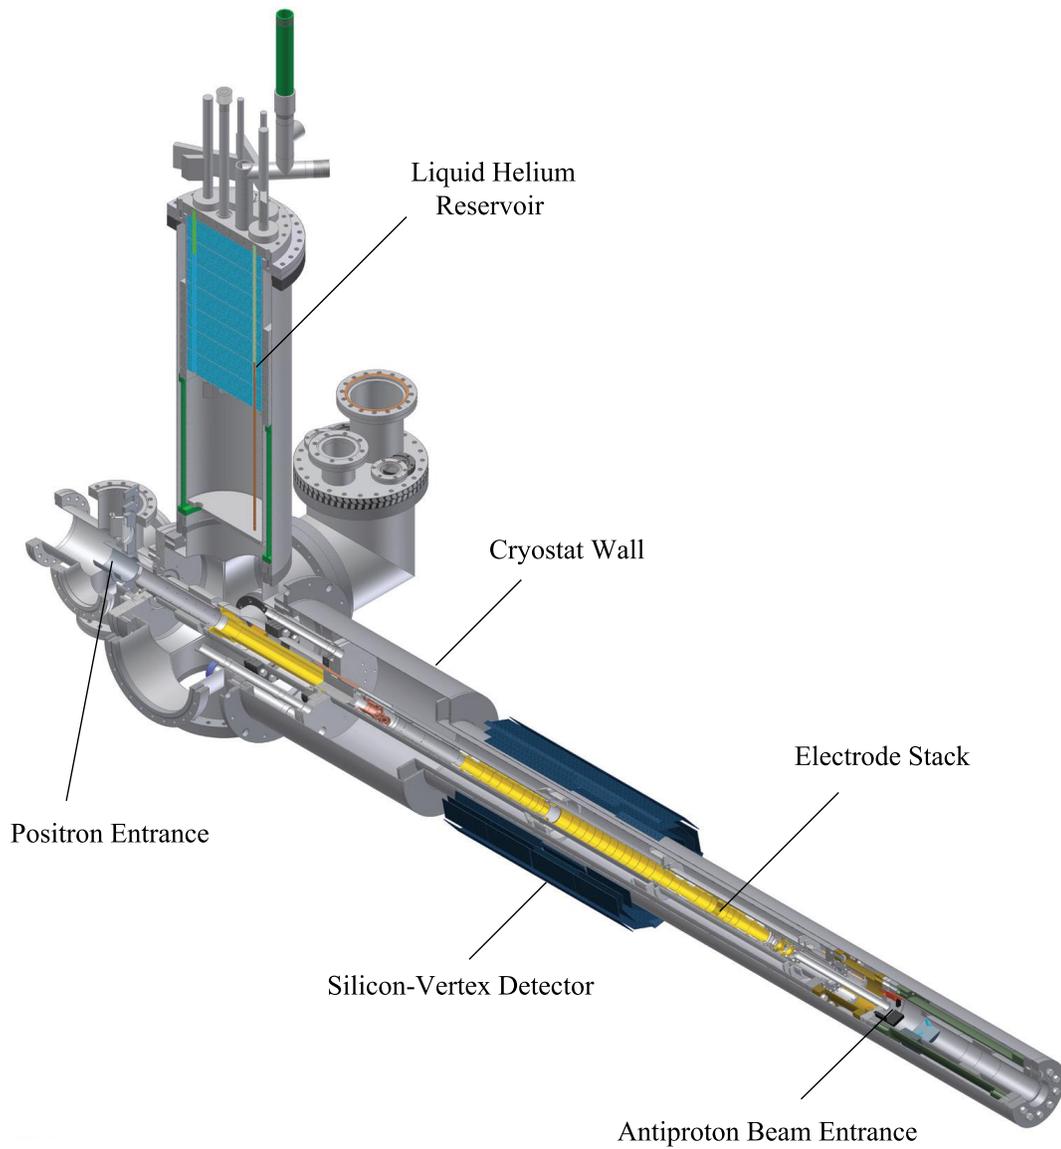
Figure 2.7: Mechanical drawing of the ALPHA cryostat. Antiprotons are injected from the upstream (right) side of the apparatus and positrons are transfered from the downstream (left) side (Adapted from ALPHA collaboration).

The horizontal section of the cryostat consists of three concentric tubes which create three distinct regions: the trap vacuum chamber, a liquid helium reservoir and an outer vacuum chamber (OVC). The innermost chamber is a 47 mm diameter tube that forms the trap vacuum chamber. It accommodates the Penning trap. The middle tube is filled with liquid helium and houses the superconducting magnets. The outermost tube (OVC) provides thermal isolation for the helium vessel, separating it from the outside world at room temperature. The silicon vertex detector is placed on the external surface (140 mm diameter) of this tube.

The Penning trap electrodes are in weak thermal contact with the walls of the liquid helium vessel, which acts as a heat sink. Normally the electrodes reach an equilibrium temperature of ~ 8 K within a few hours of filling the liquid helium vessel. The cold surface of the helium reservoir and the electrodes act as a powerful cryopump. Background gas molecules that collide with the cold walls readily adsorb to the surface. It is thus only under cryogenic conditions that the trap vacuum reaches UHV levels. Physical constraints make it impossible to introduce a pressure gauge in the Penning trap for direct pressure measurements. However, based on measurements of antiproton lifetimes (storage time) in the Penning trap [44], the trap pressure is inferred to be in the $10^{-13} - 10^{-14}$ mbar range.

2.5 Particle Sources

Clearly, positrons and antiprotons are needed for antihydrogen production. However, in the ALPHA antihydrogen synthesis procedure, electrons are also used. They act as a coolant for the antiprotons. In this section I briefly explain how the key particles used in the antihydrogen production procedure are provided in the ALPHA experiment.

2.5.1 Antiprotons

The nucleus of antihydrogen, the antiproton, is the antimatter counterpart of the proton. It is both difficult and expensive to acquire. Antiprotons can be produced in a high-energy collision between protons through the process:

$$p + p \rightarrow 3p + \bar{p} \tag{2.12}$$

This process does not occur unless the proton beam exceeds a threshold energy of 6 GeV in the lab frame [45]. The Antiproton Decelerator (AD) facility at CERN is the only machine currently capable of producing antiprotons through process 2.12. The procedure it uses is as follows: First, a bunch of $\sim 10^{13}$ protons is supplied by the Proton Synchrotron (PS). These particles then collide with an iridium target at an energy of 26 GeV. The antiprotons that are created are then magnetically extracted and injected into the AD storage ring with a momentum of 3.5 GeV/c [46, 47].
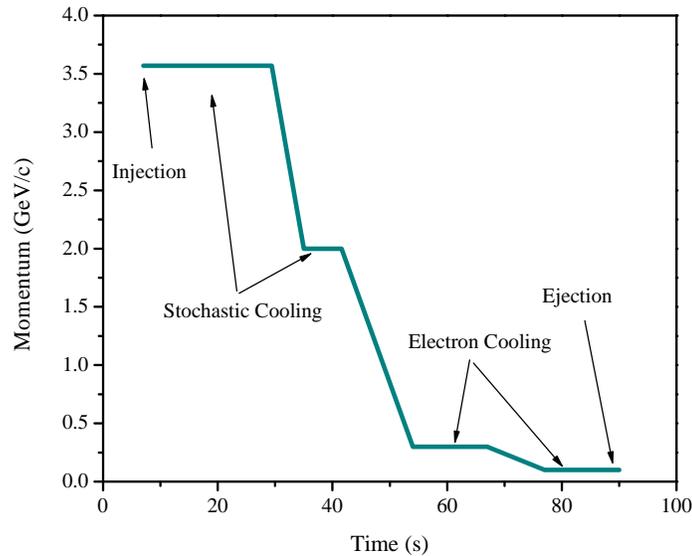
Figure 2.8: Antiproton deceleration and cooling stages in a typical AD cycle. Deceleration occurs where the particle momentum decreases. Stochastic and electron cooling occur during the flat sections. This process is periodic. Here the cycle period is less than ~ 100 s.

Once inside the AD ring, antiprotons go through a series of deceleration and cooling processes, summarized in Fig. 2.8. Antiprotons are decelerated as they pass through a series of RF-cavities with oscillating electric fields that oppose the motion of the circulating beam. This procedure increases the transverse broadening (emittance) of the beam. Stochastic and electron cooling techniques are thus interspersed with the deceleration stages to reduce

the antiproton beam size and energy spread.

Stochastic cooling involves a series of diagnostic and beam tuning devices to correct for the deviation of particle momenta from the mean value. Corrections are applied in the form of electric field kicks [48]. During electron cooling stages a beam of electrons co-circulate with the antiproton beam. The transverse energy of the antiprotons is transfered to the electrons through Coulomb interactions. The electrons are then extracted from the antiproton beam.

At the end of this sequence of processes the energy of the antiprotons is reduced to 5.3 MeV (longitudinal momentum 0.1 GeV/c). A ~ 200 ns bunch of approximately 4×10^7 particles is then delivered to one of the experiments located in the AD hall. Antiprotons at 5.3 MeV are still too energetic to be captured by the high-voltage (kV) potentials that are applied in the ALPHA catching trap. Therefore, at the entrance of the apparatus a 218 μm thick degrading aluminum foil is located. The antiprotons lose further energy as they pass through the foil via interaction with the degrader. Even after the energy degrading process, fewer than 0.1% of the antiprotons delivered to the ALPHA apparatus are captured. The rest strike the surrounding materials of the apparatus and annihilate.

2.5.2 Electrons

The primary reason to introduce electrons into the ALPHA apparatus is to use them as a coolant for antiprotons. Once antiproton and electron plasmas are mixed together, the antiprotons transfer their energy to the electrons via elastic collisions, and the electrons can efficiently radiate their energy through cyclotron radiation. In strong magnetic fields, the rate at which electrons radiate their energy is orders of magnitude higher than it is for antiprotons.

Electrons are supplied from the downstream side of the apparatus. A barium-oxide filament integrated into an electron gun is used to generate and accelerate electrons. The electron gun assembly is mounted to a vertically moveable structure (see Sec. 2.6) which is located outside of the main solenoid. This structure also holds the MCP/phosphor screen assembly and microwave horn antenna, both of which will be discussed later on.

The filament emits electrons when heated by an electric current. An electrode placed in front of the filament is biased to -15 V to generate a collimated electron beam. The electrons then follow the fringe field ($\sim$0.024 T) of the solenoid and enter the main trap. The number of electrons introduced into the apparatus can be controlled by the current in the filament, the bias potential, the position of the electron gun assembly in the fringe field and by the depth of the potential well that is prepared by biasing the Penning trap electrodes.

2.5.3   Positrons

Positrons are also transfered into the apparatus from the downstream side of the trap. They are provided by a dedicated Surko-type positron accumulator that was originally developed by the positron group at the University of California at San Diego [50]. This device was used by the ATHENA collaboration [51, 52] during the early 2000s and then passed on to ALPHA. A schematic drawing of the accumulator is shown in Fig. 2.9.
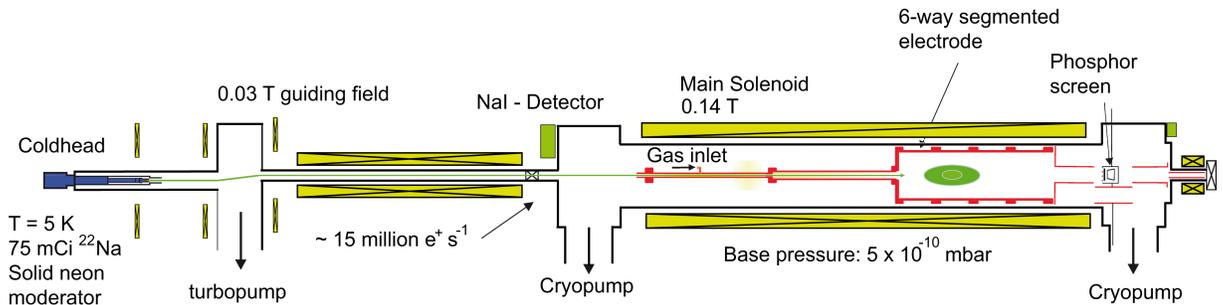
Figure 2.9: The ATHENA/ALPHA positron accumulator. Positrons emerging from the source (left end of the accumulator) are guided to a Penning trap where they cool through collisions with Nitrogen buffer gas. The positrons are then transferred to the main trap of the apparatus located on the right side of the accumulator (not shown in the figure). Figure adapted from [43].

Positrons are obtained from a $^{22}_{11}$Na radioactive source through the beta decay process:

$$^{22}_{11}\text{Na} \;\rightarrow\; ^{22}_{10}\text{Ne} +^{0}_{1}\text{e}^+ + \nu_\text{e} + \gamma \tag{2.13}$$

wherein a proton in the sodium nucleus decays to a neutron via the weak interaction, releasing a positron (e^+) and a neutrino (ν_e). A gamma ray (γ) of energy 1.274 MeV follows shortly afterward (3.7 ps) as the nucleus of the Neon daughter relaxes to its ground state. The half-life of $^{22}_{11}$Na is about 2.6 years [53].

ALPHA uses a 2.8 GBq ($\sim$ 75 mCi, peak activity, at the time of installation in 2007) $^{22}_{11}$Na source with a thin layer of solid neon directly deposited on its surface. The positrons in reaction 2.13 have a wide energy spread, the most energetic being in the MeV range. The solid neon layer acts as a moderator; only 0.5% of the positrons radiated from the source are able to pass through the neon layer, the positrons that do escape have a fairly low energy ($\sim$ 80 eV) [54].

Positrons released from the source are guided to a Penning trap by a magnetic field of 0.14 T. In this field the positrons have a long cyclotron cooling time constant, and some other mechanism is required to further cool them. Cooling is accomplished using a nitrogen buffer gas; the positrons lose kinetic energy through inelastic collisions with the nitrogen molecules. One of the trap electrodes is azimuthally segmented into six pieces. This electrode is used to apply a rotating electric field (typically at a frequency of 600 kHz), which in turn exerts a net torque on the positron plasma. Via conservation of angular momentum the positron plasma is then compressed or expanded in the radial direction [36, 55, 56]. Using this technique the density and radius of the positron plasma is controlled.

The positron plasma is then transfered to the main trap. To transfer the positrons, the nitrogen gas is first evacuated using two cryopumps. Then a mechanical valve that separates the positron accumulator vacuum from the main UHV trap opens and positrons are transfered ballistically into the main trap. The efficiency of this transfer procedure is about 50% and typically a population of about 2×10^7 positrons is re-captured in the main ALPHA trap.

2.6   The Vertically Moveable Translator

A vertically moveable translator is located between the positron accumulator and the mixing (main) trap. The translator structure functions as a support structure for a number of

Microwave Horn

Micro-Channel Plate
(MCP)

Pass-Through

Electron Source

Metallic Reflector
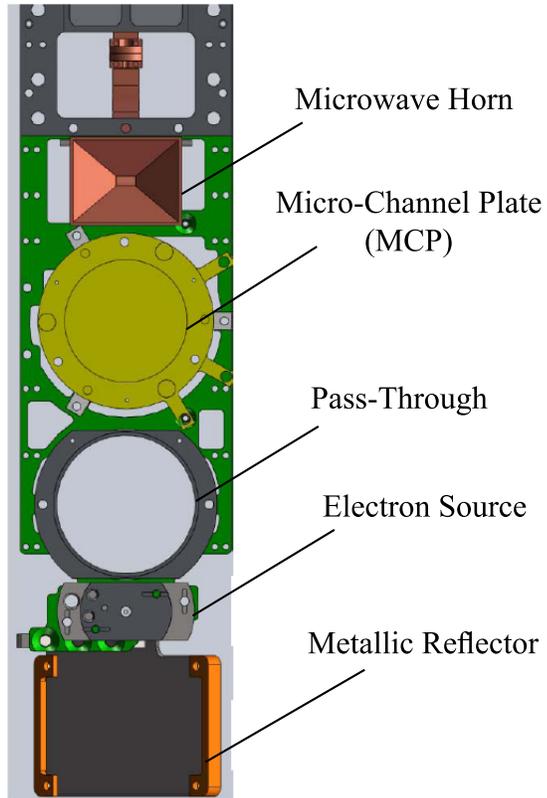
Figure 2.10: A drawing showing placement of devices mounted on the moveable translator (back view, looking toward the upstream end of the apparatus). From top to bottom, the components are the microwave horn, MCP, a pass-through cylinder which is simply a tube that allows positrons to move from the accumulator to the trap, an electron source and a $45°$ metallic reflector plate that is used to direct K_a band microwaves from an external source down the bore of the apparatus (Picture adapted from [43]).

devices and diagnostics, including an electron gun assembly, a phosphor screen and micro-channel plate (MCP), a microwave horn antenna, a 45° metallic reflector plate (see Sec. 4.1), and a pass-through cylinder (Fig. 2.10). During each antihydrogen synthesis, trapping, and physics sequence, the translator will typically move a number of times positioning various devices into the beamline, where the fringe field of the main solenoid is approximately 0.024 T.

The vertical movement of the translator is controlled by a stepper motor, gearbox, and an optical encoder to ensure positioning reproducibility. The typical time for the translator to move from one device to another is about 10 s. Usually, this movement occurs in parallel with other operations during the course of an experiment.

2.7 Detection and Diagnostic Systems

The ALPHA apparatus is equipped with a number of particle detectors and plasma diagnostics. The particle detectors function as counting and/or event location devices to identify the annihilation of antiparticles. The plasma diagnostic devices are primarily used to characterize the properties of ensembles of (anti)particles. Both types of systems are briefly discussed below.

2.7.1 Faraday Cup

A Faraday cup is a device used to measure the quantity of charge that is deposited when a packet of charged particles is intercepted. The device has an intrinsic capacitance that is usually known to good precision. Measurement of changes in its potential can be directly interpreted in terms of charge.

The Faraday cup used in ALPHA is a thin Al foil that is electrically isolated from its surroundings. It is located on the upstream end of the catching trap and simultaneously serves as the first degrader for the incoming antiproton beam.

2.7.2   Micro-Channel Plate, Phosphor Screen, and CCD Camera

The combination of a Micro-channel Plate (MCP) and a Phosphor screen is a well-established diagnostic tool for destructive characterization of non-neutral plasmas. An MCP is a plate made from highly resistive material with many tiny tube like holes passing through it. Each hole acts like an electron multiplier [58].

ALPHA uses a type E050JP47 MCP device which is manufactured by El-Mul Technologies [59]. It is mounted on the vertically moveable translator (see Sec. 2.6) in a region of the fringe field of the 1 T solenoid where B $\sim$ 240 Gauss. The ALPHA MCP has a circular active face with a 41.5 mm diameter and is covered with an regular array of holes 12 μm in diameter spaced by 15 μm in a hexagonal array. When an energetic particle strikes the walls of a hole, secondary electrons are emitted. A large potential difference is applied across the plate which causes the emitted electrons to accelerate toward one side. As electrons accelerate, they strike the walls many times and generate a shower of electrons. The magnitude of the potential difference that is applied between the faces of the MCP varies from 900 V for low numbers of trapped particles ($< 10^5$ leptons or $< 10^3$ antiprotons), down to 400 V for large numbers ($< 10^8$ leptons or $< 10^5$ antiprotons).

The electron shower is ejected from the back face of the MCP and is further accelerated onto a phosphor screen. The impact of the electrons on the phosphor screen excites phosphor atoms which results in emission of visible light. The light from the phosphor screen travels through a vacuum window and is finally captured by a charge coupled device (CCD) camera mounted outside the vacuum system (Fig. 2.11). Typically the camera shutter is left open for a period of 1 ms to collect and integrate the light. When the maximum potential difference (1 kV) is applied across the micro-channel plate, the device has a gain of 8×10^5. The gain behaviour of the MCP has been investigated for each of the particle species used in ALPHA (antiprotons, positrons and electrons) over a range of operating parameters [60].

This combination of MCP, Phosphor screen, and CCD camera is used in ALPHA for temperature measurement and imaging of the electron, positron, and the antiproton plasmas. Also we can (destructively) measure the plasma's integrated radial density profile [60].
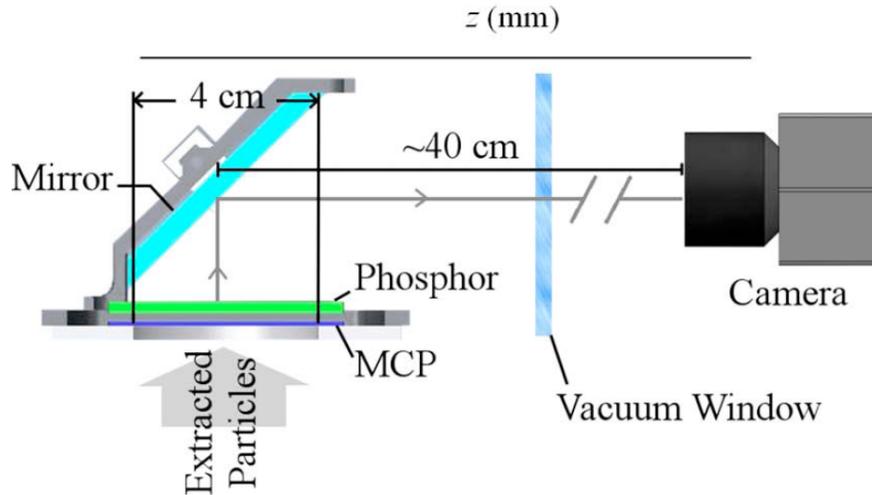
Figure 2.11: A simplified schematic of the MCP/phosphor screen/CCD arrangement. The incident particles produce a shower of electrons which then excite the phosphor screen. The light from decay of the phosphor atoms is captured by a CCD camera (Adapted from [60]).

2.7.3 Plastic Scintillator Paddles

Scintillators are commonly used for single particle detection, particularly in high energy and astro-particle physics. When an energetic particle (or radiation) passes through the detector, the scintillating material is ionized and produces light. The emitted light is collected by a Photo-Multiplier Tube (PMT) and is converted to a potential difference. If the voltage exceeds a predefined threshold then a 'count' is identified. Any environmental radiation or cosmic rays (with enough energy) can also trigger the scintillators. To reduce background noise, these detectors are usually used in pairs. A count is considered valid only if it is observed in both detectors and within a predefined time window.

In ALPHA, scintillator detectors are mainly used to identify the annihilation of antiprotons (or antihydrogen). Three pairs of rectangular scintillator paddles (40 cm wide by 60 cm high) are vertically positioned next to one another alongside the ALPHA apparatus. The first pair is close to the upstream end of the apparatus and sits around the beam degrader. The second pair covers the central region of the ALPHA trap and the last pair is positioned close to the downstream of the apparatus. It is more sensitive to annihilation events on the MCP.

The scintillators do not cover a full 4π steradians around the ALPHA trap, and so their efficiency for annihilation detection is limited. However, by simultaneous comparison of counts recorded by the scintillators with events identified by the silicon vertex detector (see Sec. 2.7.4) we estimate the detection efficiency of the scintillators to be $\sim 25\%$.

2.7.4 Silicon Vertex Detector

To detect and locate the position of an annihilation event, ALPHA utilises a silicon vertex detector which consists of 60 double-sided silicon wafer modules. The modules surround the ALPHA trap in a three-concentric-layer arrangement. They are located outside the vacuum vessel in a (controlled) room temperature environment [61, 62]. The ALPHA silicon vertex detector is similar to one used previously in the ATHENA experiment [63, 64, 65] except for the fact that the latter had only two layers of detectors, and that the ALPHA system does not include CsI crystals for γ-ray detection.

Each silicon wafer module features 256 readouts and an active area of 6 cm $\times$ 23 cm. This area is covered by two sets of perpendicular micro-strips. Each set consists of 128 micro-strips. The pitch widths in the $R - \phi$ and z directions are 227 μm and 875 μm, respectively (where R, ϕ, and z are cylindrical coordinates). For more details refer to [66, 67].

When antiproton annihilation byproducts (principally charged pions) pass through the silicon module, they ionise atoms and deposit charge on perpendicular micro-strips of the module. The charge is monitored by a VF48 ADC module, and if it exceeds a defined threshold level, the strip is considered to be affected by passage of a charged particle. The intersection of two orthogonal charged strips defines the passage location of the particle and is called a 'hit'. By fitting a helix to three hits in different layers of the detector a 'track' is constructed, representing a candidate pion trajectory.

The intersection of at least two tracks defines a 'vertex', localizing the position of a candidate antiproton (or antihydrogen) annihilation event. An example of a reconstructed antiproton annihilation is shown in Fig. 2.12(a). The detector is also triggered by charged particles from cosmic rays, which pass though the silicon modules in a straight line. Usually these events consists of two tracks which tend to line up (Fig. 2.12(b)). The total axial length

of the detector is 46 cm, which provides a solid angle coverage of 90% for annihilation events occurring at the center of the trap. The rate at which annihilation events can be read out is about 470 Hz and the overall annihilation trigger efficiency is estimated to be $(90\pm10)\%$ [66].
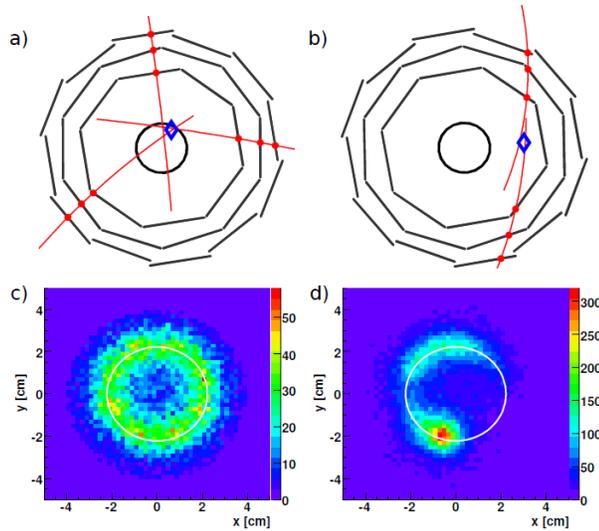
Figure 2.12: (a) A typical reconstructed antiproton/antihydrogen annihilation event on the electrode walls. The blue diamond is the identified vertex, which is an estimate of the annihilation location. (b) An example of reconstructed tracks created by cosmic rays which traverse the detector in an approximately straight line. (c) The spatial distribution of approximately 2×10^4 antihydrogen atom annihilations projected along the z-axis. The distribution is approximately azimuthally uniform and concentrated around the surface of the electrodes, indicated by the white circle. Small non-uniformities are interpreted to be due to the escape of field-ionized antihydrogen. (d) The escape of bare antiprotons tends to produce highly nonuniform distributions (Adapted from [68]).

The characteristic distribution of detected annihilation events from antiprotons and antihydrogen atoms in the combined electric and magnetic fields in the ALPHA trap are different. Antihydrogen atoms give rise to an azimuthally uniform annihilation distribution because they are neutral (Fig. 2.12(c)). Antiprotons, which are charged, tend to follow magnetic field lines and generate localized annihilation hot spots (Fig. 2.12(d)). The position sensitive silicon vertex detector plays critical role in the ALPHA experiment, since it provides a means to discriminate between the antiproton and antihydrogen annihilation events [68].

Chapter 3

Antihydrogen Synthesis and Trapping

Laboratory production of antihydrogen atoms first occurred in 1996 at CERN [69] and then at Fermilab in 1998 [70]. These atoms, however, were highly relativistic and annihilated through collisions with surrounding matter only a tiny fraction of second after they formed. The ATHENA collaboration was the first to produce non-relativistic antihydrogen atoms at the CERN Antiproton Decelerator (AD) facility in 2002 [71]. Shortly thereafter, ATRAP, another collaboration based at the CERN Antiproton Decelerator also succeeded in producing cold antihydrogen atoms [72]. Although the atoms produced by ATHENA were considered 'cold' compared to those created in 1996, their kinetic energy was approximately 1 eV (or 10^4 K in temperature units) which is far too energetic to be trapped in a laboratory-scale magnetic potential well.

The ALPHA collaboration evolved from ATHENA, and has a goal of performing precision experiments that probe the antihydrogen atom. The ALPHA approach is to create and magnetically trap antihydrogen atoms, and then use electromagnetic radiation (laser or microwaves) as tools to measure atomic energy intervals. This chapter reviews the period of atomic antimatter history during which the methods for production and confinement of antihydrogen in ALPHA were developed [73, 74]. Although I contributed to this series of experiments [66, 68, 79, 96, 101], they are not the primary focus of this thesis. Nevertheless, it is both useful and instructive to review them because they set the stage for the remainder

of the thesis: the first measurement of an anti-atomic energy interval.

3.1   Antiproton 'Catch and Cool'

The procedure for antihydrogen production begins with extraction of a bunch of antiprotons from the AD ring. The energy of the extracted antiprotons is of order 5.3 MeV, which is several orders of magnitude too high for antihydrogen production. A number of energy reduction and cooling procedures are therefore performed. The first stage of energy reduction takes place at the entrance to the ALPHA apparatus. The antiprotons pass through a 218 μm thick sheet of aluminum foil (degrader) and thereby lose energy. The thickness of this degrader was determined experimentally (guided by stopping power calculations based on the SRIM code [75]) by optimizing the number of antiprotons that are slowed down and caught in the Penning trap.

Antiprotons pass through the degrading foil and enter the apparatus (catching trap) where a 3 T magnetic field produced by superconducting solenoids is present. Before antiprotons arrive in the catching trap region of the apparatus, a 4 keV blocking potential is raised. Most of the antiprotons have enough kinetic energy to overcome this potential barrier and escape, but a few of them have low enough energies that they bounce back. In order to catch these reflected antiprotons a second potential barrier is quickly raised just upstream of the first. This sequence is outlined in Fig. 3.1. This method for catching antiprotons was first demonstrated at the CERN LEAR facility [76], and has been used ever since by all of the experiments located in the AD hall.

The time delay between the arrival of the antiprotons and erection of the second blocking potential has to be long enough to make sure that the entire bunch of antiprotons has entered the apparatus, and it has to be short enough that reflected particles are not missed. This time interval, which is known as the 'closing time' is experimentally determined to maximize the number of particles that are caught. Figure 3.2 shows data from a measurement of the number of captured antiprotons as function of closing time.

The 4 keV energy of antiprotons in the catching trap is still too high for forming trappable antihydrogen, and so further cooling is necessary. In order to do this, before catching
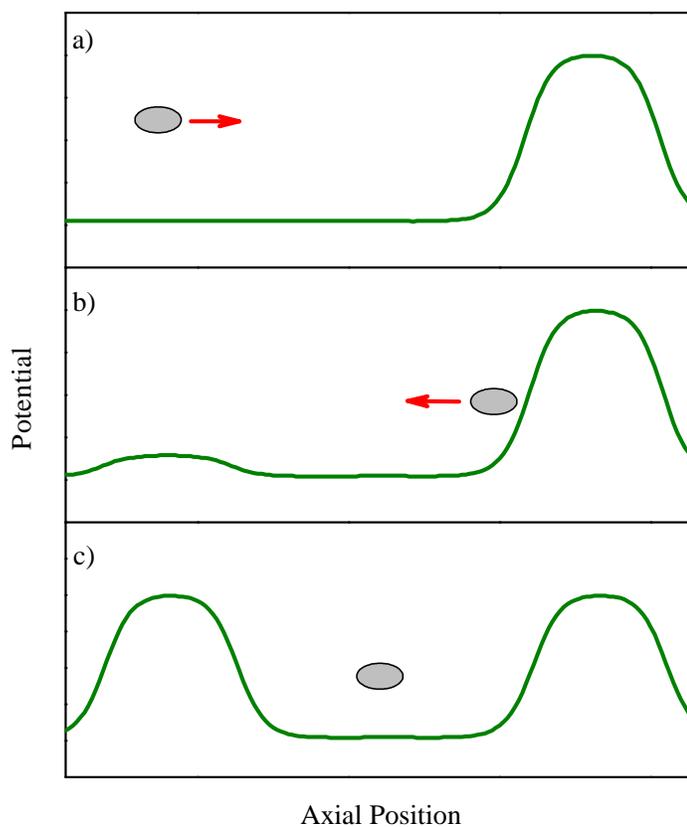
Figure 3.1: Graphical illustration of antiproton catching. a) An electrostatic blocking potential is erected to reflect a fraction of the incident particles. b) A second potential barrier is then raised to catch reflected particles. c) Antiprotons are trapped between the blocking potentials.
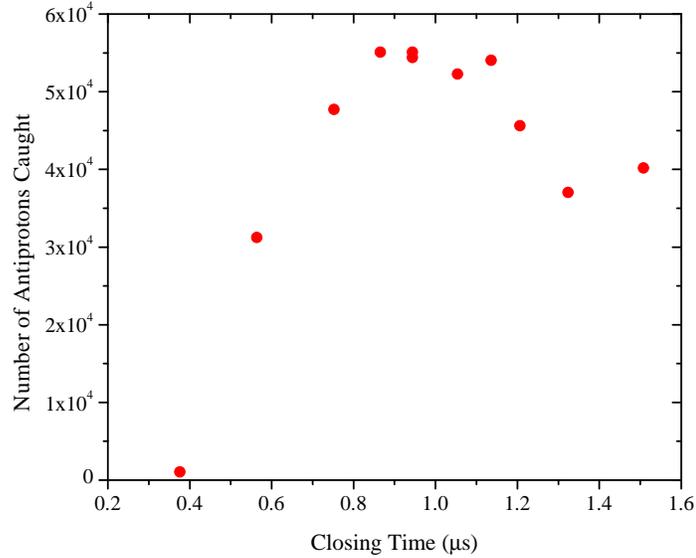
Figure 3.2: Example of the number of antiprotons caught as a function of the delay between their arrival time and the time at which the second blocking potential is raised.

antiprotons, an electron plasma of approximately 20×10^6 particles and a plasma radius of 0.55 mm is loaded into a shallow potential well positioned between the high-voltage blocking potentials (not shown in Fig. 3.1). Once antiprotons have been caught, they transfer their energy to the electrons through elastic Coulomb collisions. Meanwhile the electrons self-cool toward equilibrium with the surrounding cryostat via synchrotron (or cyclotron) radiation. The cooling time for electrons is much shorter than that for antiprotons and is given by [77]:

$$\tau_e = 3\pi\epsilon_0 \frac{c^3 m_e^3}{e^4 B^2} \tag{3.1}$$

where e is the elementary charge, m_e is the electron mass and B is the magnetic field. At $B = 3$ T, the electron cooling time is approximately 0.3 s while the antiproton cooling time is nearly 60 years! Therefore, as long as the magnetic field is high enough the electrons act as a coolant for antiprotons [78], enabling us to lower the temperature of the antiproton plasma down to approximately 200 K. If the antiprotons and electrons reach thermal equilibrium, one would expect the antiproton temperature to be close to the electron temperature, which in turn is expected to be comparable to that of the cryogenic surroundings.

However, several mechanisms cause the electron temperature to be higher than expected. These include electronic noise, plasma instabilities, and thermal radiation from surfaces at higher temperatures than the electrodes.

Antiprotons that reach equilibrium with the electron plasma end up being trapped in the same shallow electrostatic potential. After some time period the high-voltage potential barriers of the catching trap are lowered and hot antiprotons that were not cooled by the electrons are released. The remaining antiprotons and electrons are allowed to interact and further cool for a period of time that is referred to as the 'cooling time'.
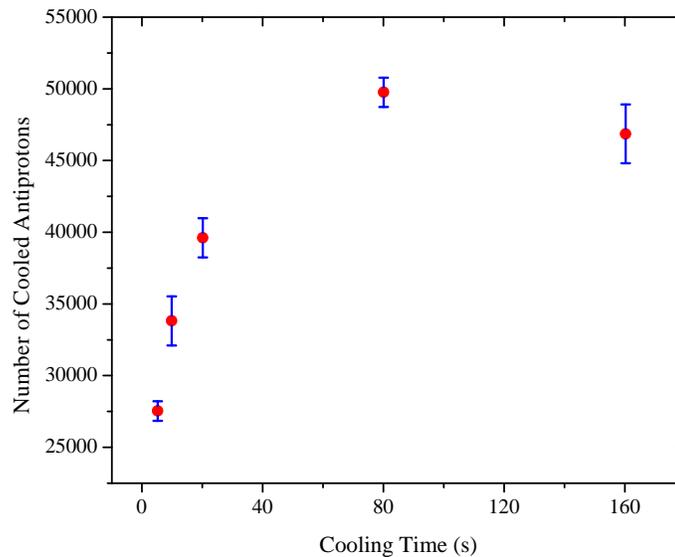
Figure 3.3: The final number of captured antiprotons after being allowed to interact with an electron plasma for different cooling time periods. Error bars indicate one standard deviation from the mean for multiple measurements.

Confinement of antiprotons for long periods of time is limited by losses caused by annihilations on residual background gas. The optimum efficiency for cooling antiprotons is determined experimentally. The efficiency can be characterized by the ratio of the number antiprotons that are initially caught to the number of antiprotons left over after the cooling procedure. An example data set showing the number of antiprotons remaining after electron

cooling is presented in Fig. 3.3. From this figure it is apparent that the optimum cooling time lies close to 80 s, which is the time that is routinely used in our generic experimental sequence.

3.1.1 Electron Ejection

Before bringing the antiprotons into contact with positrons to form antihydrogen, the electrons have to first be removed from the antiproton plasma. The inadvertent presence of electrons when the antiproton and positron clouds are brought together can result in electron-positron annihilations that in turn heat the antiprotons. Additionally, electrons and positrons may combine to form positronium atoms (a bound state of an electron and a positron) which in turn can destroy antihydrogen atoms through charge exchange.
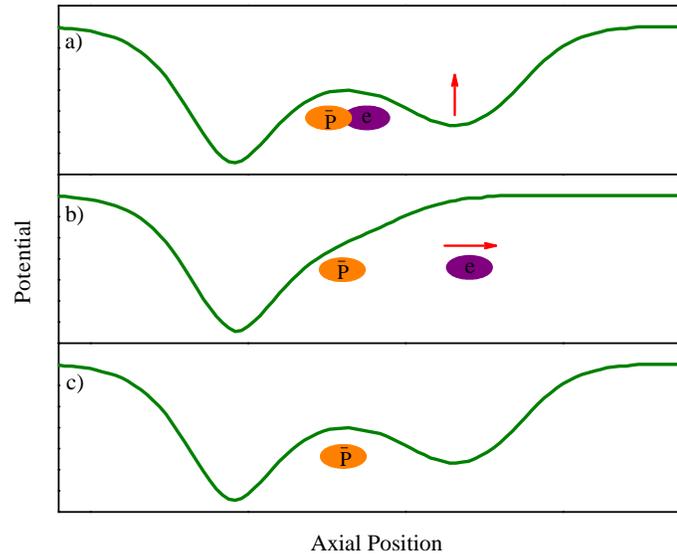
Figure 3.4: Electron ejection process. a) The negative electrons and antiprotons are initially trapped in the vicinity of the local maximum in applied potential. The barrier to the right is then lowered. The resulting slope in potential reflects the presence of an electric field. b) Electrons are ejected over a time period of 100 ns while antiprotons effectively remain stationary. c) Once the electrons have been ejected the potential barrier is reconstructed.

Electrons are separated from the antiprotons by taking advantage of the large mass difference between the two species. To remove electrons that are in equilibrium with antiprotons, one can quickly lower one of the side walls of the potential well. Under the influence of the electric field (indicated by the slope of the potential in Fig. 3.4) the electrons escape; being more massive, the antiprotons do not escape as quickly. The lowered wall is then raised again before the antiprotons move appreciably.

Typically the time scale for lowering the potential barrier, waiting for electrons to leave, and raising the barrier again is of order 100 ns. Usually, this electron ejection process has to be repeated a number of times to remove all of the electrons. A cartoon illustrating the procedure is shown in Fig. 3.4.

3.1.2 Evaporative Cooling of Charged Species

The process of removing electrons from the antiproton cloud using fast electric field pulses perturbs the antiprotons and increases their average temperature. Consequently, when the antiprotons are transfered to the mixing trap, where they are confined in a 1 T magnetic field, their temperature is observed to be about 400 K. Antihydrogen atoms synthesized with such energetic antiprotons will not be confined in the ALPHA magnetic trap. Recall that the ALPHA atom trap depth is approximately 0.5 K (in temperature units) or 4.3×10^{-5} eV for ground state antihydrogen atoms. So, after the electrons have been ejected, further cooling of the antiprotons is necessary. At the expense of losing some of the antiprotons, we use evaporative cooling to further cool them prior to antihydrogen synthesis [79].

Evaporative cooling is a process in which energetic particles are selectively expelled from an ensemble, lowering the average kinetic energy (and hence temperature) of those that remain. This is a common procedure that is used, for example, for neutral atoms in preparation of Bose-Einstein condensates [80, 81]. However, it is not a procedure that is routinely used for charged particles.

Naively one might expect evaporative cooling of charged particles to be fundamentally different from cooling of neutral atoms, leading to unexpected results. Neutral atoms are typically confined in magnetic potential wells while charged particles are confined in electrostatic potential wells. Charged particles are subject to long-range Coulomb interactions

with trapping fields and with each other while atoms trapped in magnetostatic wells are subject to short-range dipolar interactions associated with their magnetic moments. The coupling between the magnetic moment of an atom and typical laboratory magnetic fields is several orders of magnitude weaker than the coupling between an elementary charge and typical electrostatic fields. Moreover, charged particles that are confined in Penning traps are tightly bound to magnetic field lines by the Lorentz force, which is not the case for neutral atoms.

Proper understanding of the evaporative cooling process for charged particles requires a theoretical framework that is distinct from that used for neutral atoms. The energy distribution of an ensemble of particles that are in global thermal equilibrium involves a Maxwell-Boltzmann distribution (Fig. 3.5). In a finite potential well, atoms (or particles) with more kinetic energy than that associated with the maximum well depth can escape over the top of the potential barrier and carry away extra energy.
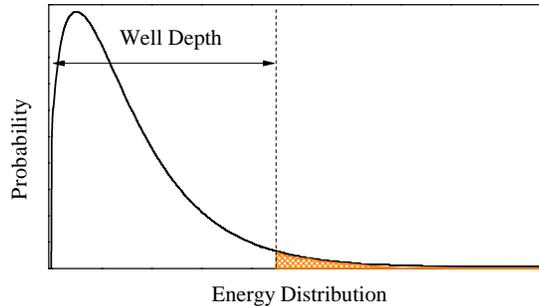
Figure 3.5: Maxwell-Boltzmann energy distribution for an ensemble of particles in equilibrium. In a finite potential well, particles in the high energy tail of the distribution can escape.

In the standard theory of evaporation, the temperature T and the number of atoms N are linked together via two first-order differential equations [82]:

$$\frac{dN}{dt} = -\frac{N}{\tau_{ev}} \tag{3.2}$$

$$\frac{dT}{dt} = -\alpha \frac{T}{\tau_{ev}} \tag{3.3}$$

where t represents time, τ_{ev} is an evaporation time constant, and α is the ratio of the average excess energy carried away by an escaping atom to the average energy of a trapped atom. These equations have to be modified for charged particles in a Penning trap, to include effects such as high collision rates (due to long-range Coulomb interactions) and the self-field of the charged particles which can dramatically change the potential well. Further discussion of these effects may be found in [79, 83, 84]. The modified versions of Eqs. 3.2 and 3.3 are:

$$\frac{dN}{dt} = -\frac{N}{\tau_{ev}} - \gamma N \tag{3.4}$$

$$\frac{dT}{dt} = -\alpha \frac{T}{\tau_{ev}} + P \tag{3.5}$$

where γ characterizes the antiproton loss rate due to annihilation on residual background gas (experimentally measured to be $\gamma = 1 \times 10^{-4}$ Hz in the ALPHA apparatus), and P is a heating term associated with expansion driven Joule heating.

In practice, we hold the antiprotons in a 1.5 V potential well. One side of the potential well is then lowered by linearly ramping down the voltage on one of the electrodes. The potential well depth is thus reduced from 1.5 V to some pre-determined value. As the well becomes shallower, lower temperatures are achieved but more antiprotons are lost.

Figure 3.6 shows data that demonstrate agreement between the model described above and measurements. A cloud of antiprotons with a radius of 0.6 mm containing 4.5×10^4 particles at a density of 7.5×10^6 cm^{-3} was used in this demonstration. The potential well depth is reduced to a number of different final values, and then, after a 10 s wait (to allow the remaining particles to re-equilibrate), the temperature and the number of particles remaining are measured.

The lowest temperature in this series of experiments was observed with a final well depth of 10 ± 4 mV; only 6 ± 1 % of the particles remained in the trap, with a measured temperature of 9 ± 4 K [79]. Note that in our typical antihydrogen production sequence the
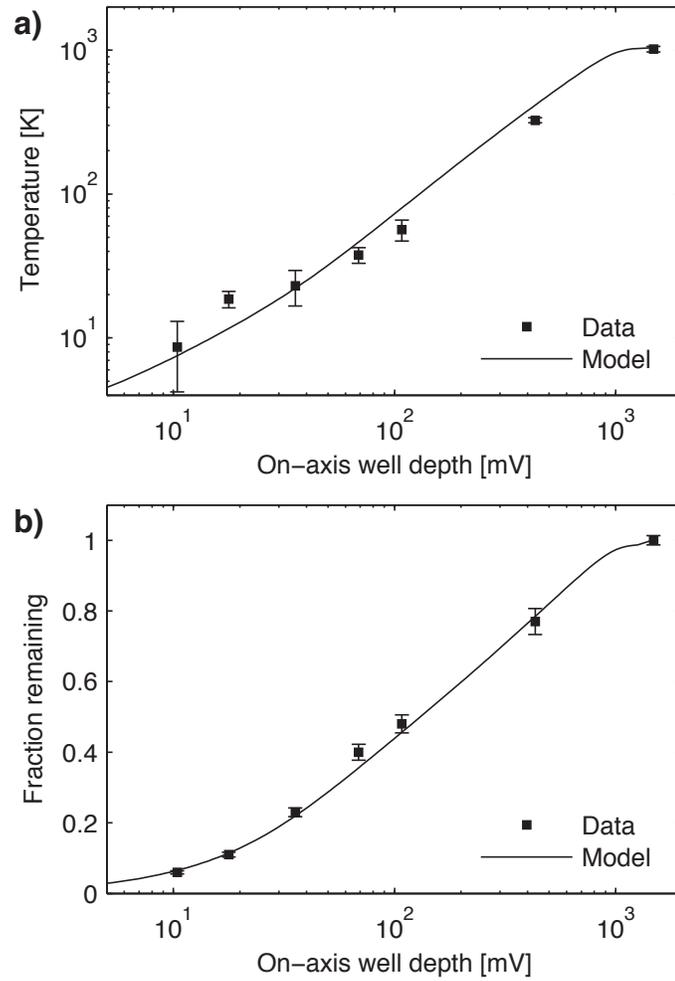
Figure 3.6: The temperature (a) and fraction (b) of antiprotons remaining in the trap as a function of the final well depth during evaporative cooling. The solid lines represent the prediction of the evaporative cooling model, modified to account for particle charge (From [79]; Copyright (2009) by the American Physical Society).

antiproton well depth is reduced to ~ 70 mV and the final antiproton cloud temperature is at best approximately 40 K. This final temperature may vary depending on the density of the antiproton cloud.

3.2 Antihydrogen Production

In parallel with the antiproton preparation procedures described above, a cloud of 2×10^7 positrons is transferred to the mixing trap where it cools down through the emission of cyclotron radiation in a 1 T magnetic field for tens of seconds (see Sec. 2.5.3). The evaporative cooling procedures outlined for antiprotons are similarly applied to the much denser positron plasma in order to achieve temperatures of order 40 K [84].

After evaporative cooling of both antiproton and positron clouds, the next step is to bring them together to synthesize antihydrogen, a process that is called 'mixing.' Before mixing, both of the clouds have to be positioned close to each other in the centre of the Penning trap (the mixing trap). Simultaneous confinement of negatively charged antiprotons and positively charged positrons is realized using a potential configuration known as a 'nested potential well' [85].

Figure 3.7 shows an example of a nested potential that is used in connection with ALPHA. It consists of a long potential well in which we can confine antiprotons and a shorter inner well in which oppositely charged positrons are held. The antiprotons are confined in one of the two 'side wells' of the nested potential. The next step would be to give the antiprotons enough energy to overcome the potential barrier and pass through the positron plasma. There are a few different methods that are used to inject antiprotons into the positron plasma, which will be discussed later. For now we briefly go through different possible antihydrogen formation scenarios, as antiprotons are brought into contact with the positrons.

3.2.1 Antihydrogen Formation Scenarios

The combination of an antiproton and a positron can lead to a bound atomic state only if a 'third body' that can carry away the excess energy is involved. Otherwise, the law of
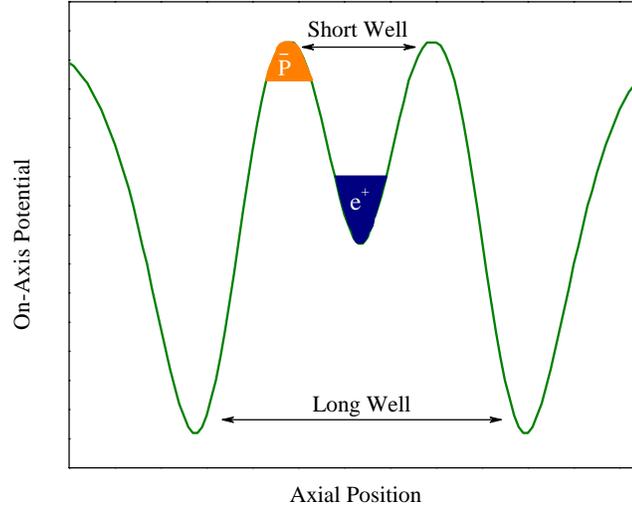
Figure 3.7: An example of a nested potential well, in which oppositely charged species are confined close to one another. Only the vacuum potential is shown here.

linear momentum conservation would be violated. The third body can be a positron, an antiproton, or a photon.

In the simplest scenario a positron is captured by an antiproton and the binding energy is released via the emission of a photon. This process is known as 'spontaneous radiative recombination'; It has been studied in detail [87]:

$$\bar{p} + e^+ \rightarrow \bar{H} + \gamma. \tag{3.6}$$

This is an allowed electric dipole transition that favours the formation of deeply bound (typically $n < 10$, where n is the principal quantum number) antihydrogen. It has a weak dependence on positron temperature, T_{e^+}, close to $1/\sqrt{T_{e^+}}$ [86]. The antihydrogen formation rate expected for this process is much lower than that observed for ALPHA, and thus it is believed that this process is not the dominant mechanism. In theory, the antihydrogen formation rate to a particular quantum state can be enhanced by the presence of an

appropriate photon (laser stimulation) [88]. This process was attempted by the ATHENA collaboration but no increase in formation rate was observed [89].

Another possible antihydrogen formation scenario involves a charge exchange process in which antiprotons collide with positronium in an excited state (Ps*):

$$\bar{p} + Ps^* \to \bar{H} + e^-. \tag{3.7}$$

This type of formation scenario was studied during the 1990s [90, 91] and in 2004 the ATRAP collaboration reported the creation of antihydrogen using this method [92]. In a similar manner, a double Rydberg charge exchange method for antihydrogen production has also been proposed [93].

The most probable antihydrogen formation process is thought to be a 'three-body recombination' mechanism in which the excess energy is removed by a spectator positron:

$$\bar{p} + e^+ + e^+ \to \bar{H} + e^+. \tag{3.8}$$

This reaction resembles an elastic collision between an antiproton and a positron in the presence of another positron. The thermal energy of the positron, $k_B T_{e^+}$ is typically dominated by the binding energy of the synthesized antihydrogen atom. Detailed cross-section calculations for this reaction show that, under the assumption of global equilibrium between the positrons and the antiprotons, the antihydrogen formation rate is strongly correlated with the positron temperature, scaling as $(1/T_{e^+})^{9/2}$ [94, 88].

3.2.2 Mixing

To induce the antiprotons and positrons to interact, the antiprotons have to be given enough energy to escape the side well in which they are trapped and transit the positron plasma. The easiest way to do this, in the nested well configuration, is to trap the antiprotons in a well that is at a higher potential than that of the positrons.

Figure 3.8 illustrates this type of mixing scheme, which was implemented by the ATHENA
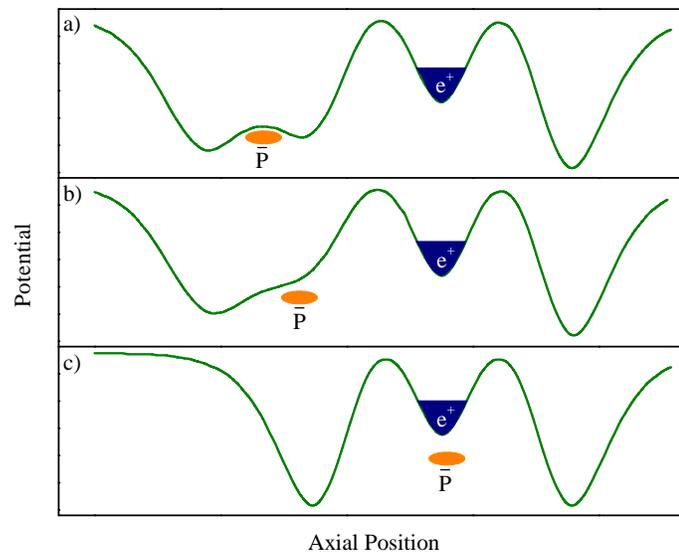
Figure 3.8: Illustration of a generic mixing sequence. (a) Antiprotons are trapped in a potential well to the side of the positron well (b) The antiprotons are released and (c) enter the positron plasma. In the original application of this method [71] the antiprotons were injected into the positron plasma with a kinetic energy of about 30 eV.

collaboration, prior to the ALPHA collaboration [71]. The antiprotons are held in a potential well situated above the positron space charge energy and then are released to interact with positrons. The antiprotons enter the positron cloud with a kinetic energy of about 30 eV. Within a few tens of milliseconds they reach thermal equilibrium with the positrons through Coloumb interactions [95, 86].
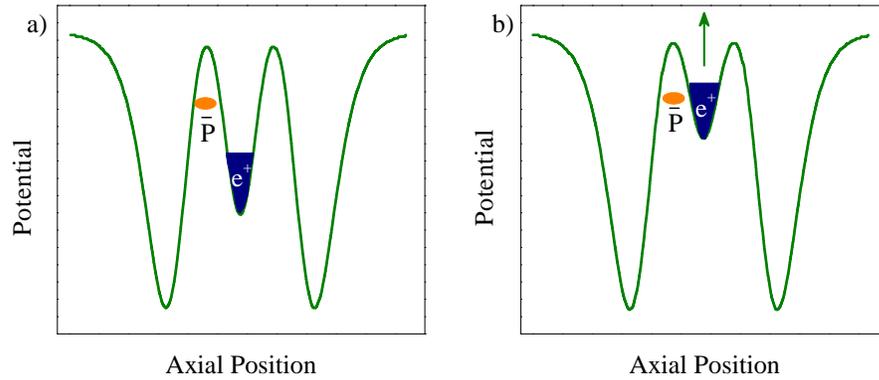
Figure 3.9: Incremental mixing scheme. (a) Antiprotons and positrons are held in neighboring wells. (b) The relative energy difference between two wells is modified until antiprotons start to escape and pass through the positron well.

In order to reduce the kinetic energy of the antiprotons as they are released into the positrons, a modified mixing scheme known as 'incremental mixing' was proposed (Fig. 3.9) [84]. The antiprotons are held in the side well next to the positron plasma while the depth of the central well in which the positrons are situated, is reduced. Eventually antiprotons escape from their well and pass through the positron plasma. In this case the kinetic energy of the antiprotons is controlled by the depth of the central potential well. In typical experimental sequences this scale is about 1 eV. Although this is substantially less than 30 eV, it is still large compared to the magnetic potential well intended for antihydrogen trapping (a few tens of μeV).

The incremental mixing scheme is also fairly susceptible to the initial conditions of the antiproton and positron plasmas. For instance, the energy of the antiprotons must match the positron space charge potential, but this in turn varies depending on the number of

positrons in the well which changes from run to run. If the antiproton injection energy falls below the positron space charge potential the antihydrogen formation rate is significantly reduced.

3.2.3   Autoresonance Mixing

The mixing technique developed and ultimately used by ALPHA leaves the trapping potentials unaltered. Instead a sinusoidal axial electric field drive with a negative frequency chirp is applied to the confined antiprotons. Before the mixing procedure is initiated, the full magnetic trap (mirror coils and octupole) is energised. The antiproton plasma, which acts like a single particle during the excitation, is phased-locked to the electric field drive. The antiprotons then gain just enough energy to be autoresonantly injected into the positron plasma [96].

The autoresonance phenomenon used in this scheme appears in a variety of dynamical systems, ranging from plasma modes [97] to orbital dynamics [98]. In general a swept frequency drive is applied to a nonlinear oscillator, with the result that the oscillator can become phase-locked to the drive. The oscillator's amplitude can then be precisely controlled simply by adjusting the frequency of the drive.

In the ALPHA apparatus, antiprotons are confined in an anharmonic potential well next to the positron well (see Fig. 3.9.a) and oscillate back and forth (axially) at a natural 'bounce' frequency, ω_b. Autoresonant injection of antiprotons involves applying a swept-frequency drive of the form $v_0 \sin(\omega t)$ to one of the electrodes used to create the antiproton potential well. The drive frequency starts above the bounce frequency and is swept down to a frequency that corresponds to a pre-determined final energy. This energy needs to be larger than the space charge potential of the positrons.

In a typical antihydrogen synthesis procedure we employ a 0.75 ms long drive, swept from 325 kHz to 235 kHz. Antihydrogen annihilation events are observed shortly after the sweep is initiated (Fig. 3.10) after which the production rate decreases by an order of magnitude over a few tens of ms. A possible reason for this decrease is the positron plasma temperature increase when the antiprotons are injected. Measurements indicate that the
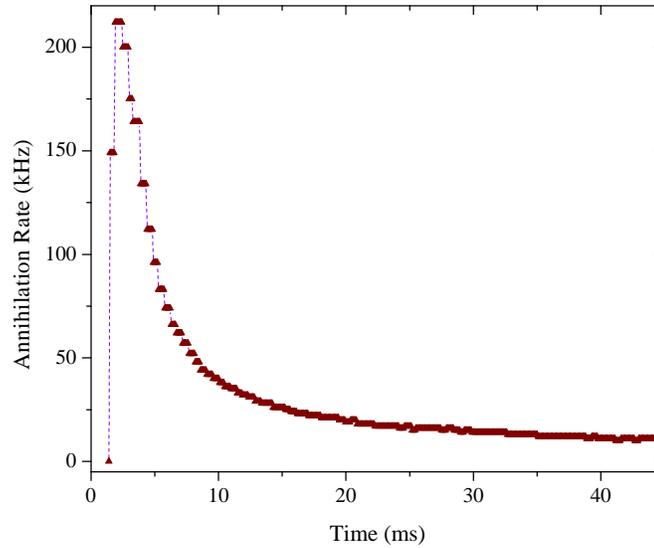
Figure 3.10: Antihydrogen annihilation rate during autoresonant injection of antiprotons into a positron cloud, which starts at $t = 0$.

positron temperature changes from 40 K to 80 K during the antiproton injection procedure.

Autoresonant injection of antiprotons is now a well-established procedure in ALPHA and is routinely used for antihydrogen production.

3.3 Antihydrogen Detection

Most of the antihydrogen atoms synthesized in ALPHA are too energetic to be trapped. They strike the trap walls (Penning trap electrodes) and annihilate. ALPHA uses a position sensitive silicon-vertex detector as an imaging device to locate these annihilation events (see Sec. 2.7.4). During mixing, cosmic rays can also trigger the detector. It is therefore important to be able to distinguish between triggers caused by annihilation events and those caused by cosmic rays.

Typically cosmic rays produce two linear tracks as a single particle transits from one side of the apparatus to the other. On the earth's surface the main source of cosmic rays

are muons with a mean energy of 4 GeV [99]. The bending radius of these particles in a 1 T magnetic field is more than 10 m, and so their tracks are expected to be essentially straight over distances comparable to the detector diameter. On the other hand, antihydrogen (or antiproton) annihilation produces two or more tracks; one for each pion. The average energy of these pions is close to 120 MeV [57], which implies a bending radius of ~ 0.7 m/sin(θ) in a 1 T magnetic field. Here θ is the angle between the particle momentum and the magnetic field. Figure 3.11 compares examples of reconstructed annihilation event by-products and cosmic ray tracks projected onto the $r - \phi$ plane. The former are curved while the latter are not.
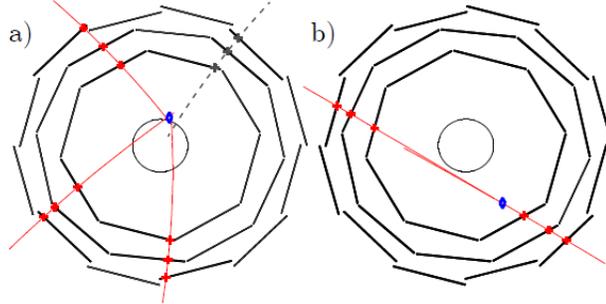
Figure 3.11: Example reconstruction of a) an annihilation event, and b) a background cosmic ray event. The solid red curves show the reconstructed tracks and the red dots show the identified hit points on the silicon modules. The dashed track illustrates a track that was considered, but ultimately excluded from the vertex determination. The blue diamonds illustrate the position of the reconstructed vertices (Adapted from [100]).

The algorithm used to reconstruct vertex positions and to identify cosmic ray events involves the following [100, 67]: First the tracks associated with a given event are reconstructed and their total number is determined (N_{Tracks}). If $N_{\text{Tracks}} < 2$ the event is discarded because at least two tracks are required to reconstruct a vertex. Next the mean distance of closest approach (D) is calculated for this configuration of tracks. A loop is then run $N_{\text{Tracks}} - 1$ times. In each iteration one of the tracks is excluded and D is recalculated. The minimum value generated by this process is then compared to the original. If the new value is more than 60% smaller than the original value (with all tracks), the new configuration is adopted and the appropriate track is excluded. Otherwise the original configuration is considered to be valid (Fig. 3.11). If the number of tracks in the new configuration still

involves more than two tracks the process is repeated until the result converges, or the set is reduced to two tracks.

Once the tracks and vertex locations have been identified, annihilation and cosmic ray events are systematically distinguished in a manner that is not subject to experimenter bias. A number of parameters are set as cuts (thresholds) to classify annihilation and cosmic events. These parameters are determined using training data sets that are independent of data used in eventual measurements.

The training data sets used in our initial report of antihydrogen trapping [73] consisted of a) a total of 165,520 events collected during 335 s of antiproton-positron mixing with all trapping fields engaged (the annihilation event training set) and b) a total of 109,824 events collected during about 3 hours of operation with the trapping fields engaged but no positrons or antiprotons present (the cosmic ray training set).

Three parameters were chosen as selection criteria: the number of tracks, the radial location of the vertex relative to the trap axis, and the residual from linear fits to the hit patterns from which tracks are reconstructed. The number of tracks is important because typically events with two tracks are produced by cosmic rays while events with more than two tracks are associated with antihydrogen (or antiproton) annihilation. The event location is also important; The distribution of antihydrogen annihilation events is expected to be evenly distributed on the electrode walls. Finally, cosmic rays are typically high energy particles that cross the detector in almost a straight line, thus they exhibit a low residual value.

After a full analysis of the training data sets, the following cuts were selected: for two track events a vertex radius < 4 cm and a linear squared residual of > 2 cm^2 was defined as an annihilation event. For more than two track events, a vertex radius of < 4 cm and a linear squared residual of > 0.05 cm^2 is required for acceptance as an annihilation event. Figure 3.12 shows the distribution of reconstructed vertices for cosmic rays and annihilation events, along with the cut parameters that were established. Extensive discussion of this analysis can be found in Refs. [67, 68, 84]
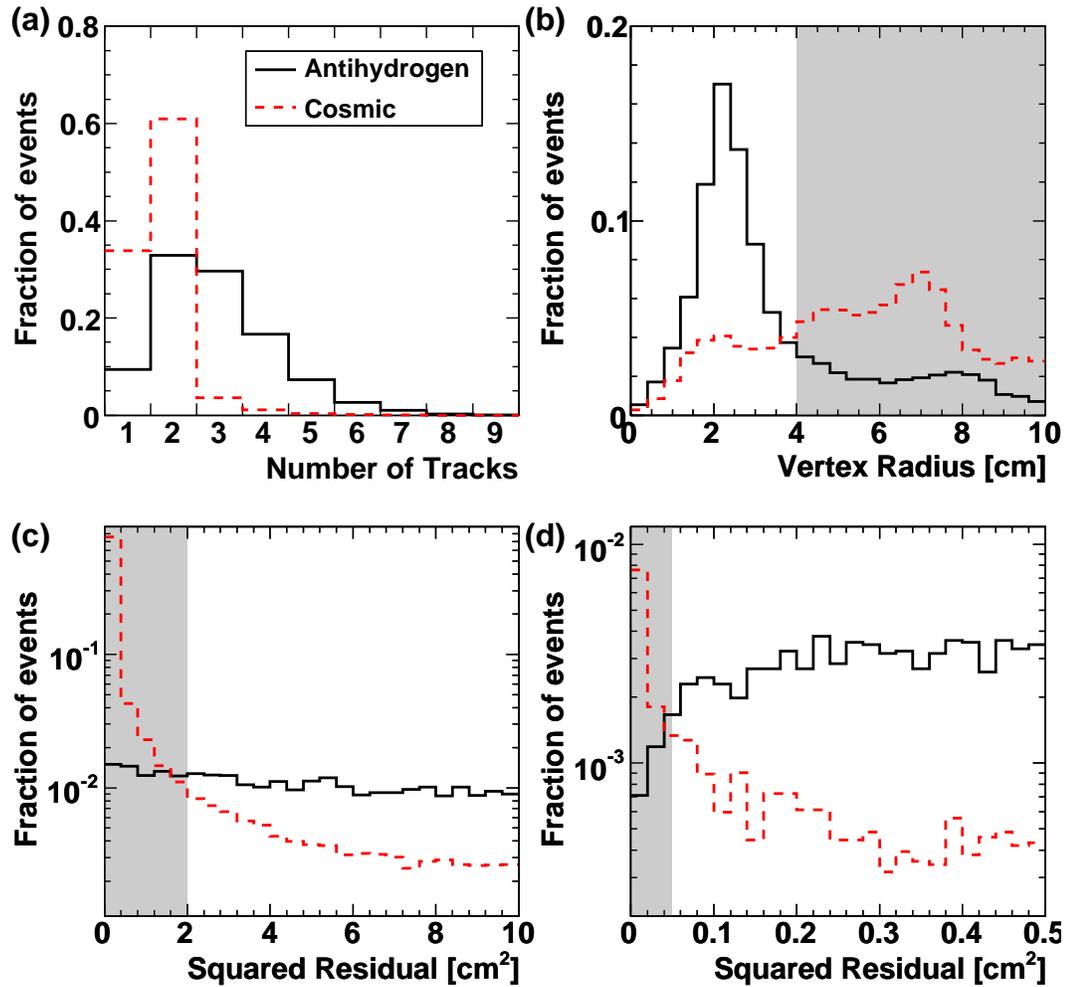
Figure 3.12: Distribution of reconstructed vertices for the training data sets; the gray regions indicate cut parameters. (a) Distribution of the number of reconstructed tracks per event for the cosmic ray data set (red) and annihilation event data set. (b) Distribution of radial coordinate for both training data sets. (c) Two-track event distribution and (d) more than two-track event distribution (Reproduced from [83]).

These cuts were applied to the training data sets to examine their efficiency. A $(99.54 \pm 0.02)\%$ cosmic rejection rate was observed when the cuts were applied to the cosmic ray data set. This corresponds to an acceptance rate (approving a cosmic ray event as annihilation event) of 0.047 ± 0.002 events per second. Applying the cuts to the annihilation event data set, $(64.4 \pm 0.1)\%$ of the events pass the cuts. If one combines these efficiencies with the 90% trigger efficiency of the detector, an overall annihilation event detection efficiency of $(58 \pm 7)\%$ is obtained [100].

The algorithm outlined above is the default algorithm that is used to identify annihilation events associated with trapped antihydrogen atoms. For identification of annihilation events that are caused by microwave radiation, a complementary analysis is performed; It will be discussed in Sec. 5.4.3 and App. A.

3.4   Antihydrogen Confinement

The ALPHA apparatus was designed to demonstrate trapping of neutral antihydrogen atoms. The first step in the experimental sequence is the production of antihydrogen, as discussed earlier in this chapter. These atoms are produced at the midpoint of the fully energised magnetostatic trap. Characteristics of the typical antiproton and positron plasmas used for experimental demonstration of antihydrogen trapping are summarized in Tab. 3.1.

	e^+ plasma	$\bar{\text{p}}$ plasma
Number of particles	2×10^6	$(30 \pm 5) \times 10^3$
Density (cm^{-3})	5.5×10^7	6.5×10^6
Radius (mm)	0.9	0.8
Temperature (K)	40 ± 15	200 ± 40

Table 3.1: Positron and antiproton plasma characteristics prior to mixing.

The two plasmas were confined in a nested well configuration (similar to Fig. 3.7) and then antihydrogen atoms were produced by autoresonant injection of antiprotons into the positron well. The particles were allowed to interact for 1 s during which time we observed 5000 ± 400 annihilation events in the silicon detector. Most of the atoms that were created

were energetic enough to escape the magnetic trap, strike the electrode walls, and annihilate. However, atoms with kinetic energies less than the magnetic well depth (0.5 K, in temperature units) are confined.

After mixing and before checking to see if any antihydrogen atoms have been trapped, all unbound charged particles have to be ejected. To do this, we axially 'dump' the antiprotons and positrons onto matter targets where they annihilate. The dumps are accomplished using a series of electric field pulses, and are designed to facilitate counting of the charged particles. After the dumps, all the Penning trap electrodes are grounded; any antiprotons that remain in the trap at this point must be trapped by the mirror and octupole fields alone (see Sec. 3.4.1) [101].

Next, in order to remove any remaining charged particles we apply four electric field pulses - known as 'clearing' pulses - up to 500 V/m. Once all charged particles have been removed, the neutral trap fields are ramped down (very quickly) and we look for annihilation events associated with any atoms that might have been trapped. The sequence of manipulations after mixing and before the trap can be shut down takes 172 ms to complete. This then sets the scale for the minimum hold time for trapped atoms.

The ALPHA apparatus has two key features that play an important role in identifying the signatures of trapped antihydrogen atoms; the first feature is the rapid turn off of the trap magnets [38, 39] and the second feature is the position sensitive silicon vertex detector used to locate antihydrogen atom annihilation events [66].

The rapid turn off feature of the magnets minimizes the chance of misinterpreting cosmic rays as annihilation events. Recall that the background rate is 0.047 ± 0.002 s^{-1}. The high current flowing in the magnets (up to 650 A in the mirror coils and up to 900 A in the octupole) is delivered to a bank of resistors resulting in a decay time constant of about 9 ms (see Fig. 3.13). The time window over which we search for the release of antihydrogen atoms is chosen to be 30 ms, or slightly more than three e-folding times for the currents. Evidence for trapped antihydrogen atoms was first observed in 2009 [68]. It was fully established in the following year [73].
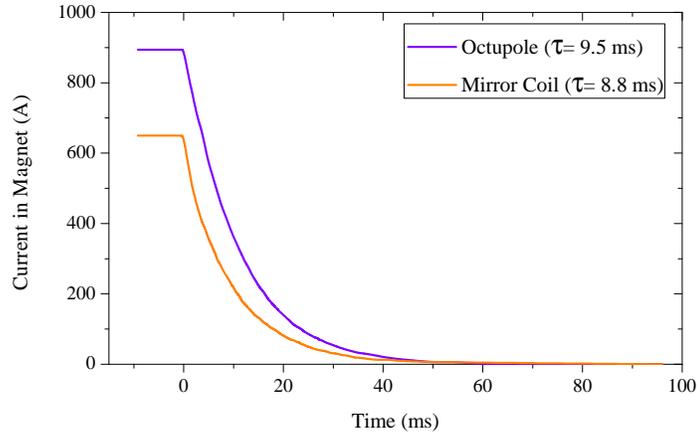
Figure 3.13: Current in the octupole and mirror coil magnets as they are rapidly switched off. The decay time for the octupole is $\tau = 9.5$ ms and the decay time for the mirror coil is $\tau = 8.8$ ms.

3.4.1 Mirror-Trapped Antiprotons

Under certain conditions, the ALPHA magnetic trap can also confine bare charged particles, particularly unbound antiprotons, even when all of the Penning trap electrodes are grounded. These antiprotons will then be released and annihilate when the magnets are switched off, mimicking the annihilation of antihydrogen atoms.

The unwanted confinement of antiprotons is associated with conservation of the magnetic moment (μ) of the particle as it gyrates in the nonuniform magnetic field:

$$\mu = \frac{E_\perp}{B} = \frac{mv_\perp^2}{2B} \tag{3.9}$$

where $E_\perp$ ($v_\perp$) is the kinetic energy (velocity) of the particle in the plane perpendicular to the magnetic field $\mathbf{B}$. As an antiproton moves into a region of higher magnetic field, $E_\perp$ has to increase in order to conserve angular momentum. Thus energy is converted between the degrees of freedom perpendicular and parallel to the local magnetic field.

As the particle travels toward regions of higher magnetic fields, more and more energy is transferred to the motion perpendicular to the field, until the parallel component of velocity reaches zero and the particle reverses its direction. This process is known as

'mirror-trapping.' The kinetic energy associated with motion parallel to the field can be written [102]:

$$E_{||} = E_0 \left(1 - \frac{E_{0\perp}}{E_0} \frac{B}{B_0} \right) \tag{3.10}$$

where $E_0 = E_{||} + E_\perp$ is the total kinetic energy of the particle and $E_{0\perp}$ is the perpendicular energy at a point where $B = B_0$. The trapping condition is then satisfied if $(E_{0\perp}B) / (E_0 B_0)$ is larger than one.

As previously mentioned, we apply four axial electric field pulses up to 500 V/m in order to clear any mirror trapped antiprotons. The motion of these antiprotons while the electric field pulses are applied can be understood in terms of a pseudo-potential Φ that combines the energy of interaction between the magnetic dipole moment and the magnetic field with the electrostatic potential energy $-e\phi(z, t)$, where $\phi(z, t)$ is the electric potential [68]:

$$\Phi(z, t) = E_{0\perp} \left(\frac{B - B_0}{B_0} \right) + (-e)\phi(z, t). \tag{3.11}$$

Figure 3.14.a shows the electrostatic clearing potential and the axial magnetic field in the ALPHA apparatus when the magnetic trap is energised. Fig. 3.14.b shows the calculated pseudo-potential (Eq. 3.11) for three different antiproton transverse energies $E_{0\perp} = 0$, 10, and 50 eV. In order to mirror trap an antiproton a well needs to exist in the pseudo-potential. The well depth is a function only of perpendicular energy $E_{0\perp}$.

For our parameters, wells in the pseudo-potential only develop for antiprotons with $E_{0\perp} > 24.4$ eV. Any antiproton with $E_{0\perp} < 24.4$ eV will necessarily be expelled when the strong clearing fields are applied even if it has $E_{0||} = 0$ eV [101]. The off-axis pseudo-potential has a complicated form because of contributions from the octupolar field. The trajectory of trapped particles can thus only be properly calculated numerically.

In our numerical simulations of trapped antiprotons we use a distribution that is spatially uniform throughout the trap region. We choose a velocity distribution that is isotropic and flat up to a total energy of 75 eV. This is far more extreme than the actual energy of antiprotons in the trap. However, our knowledge of the relevant distribution is limited
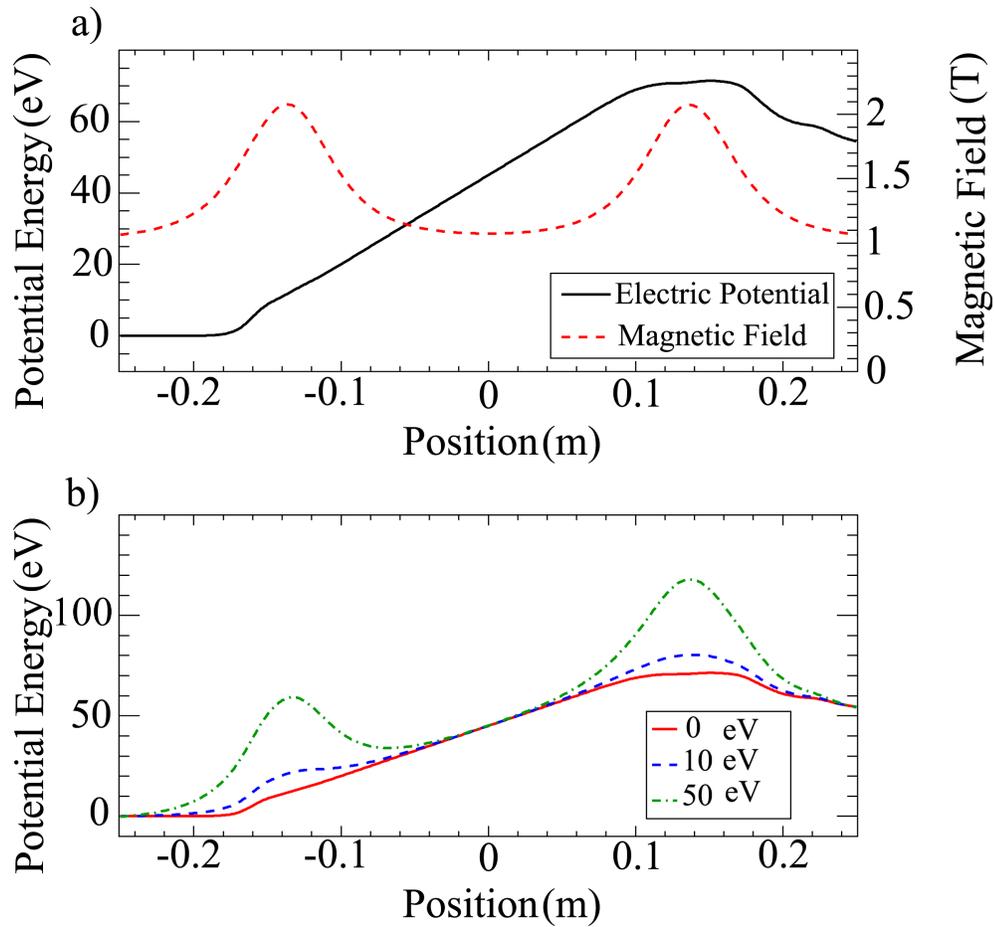
Figure 3.14: (a) The solid line shows the clearing electric potential applied to remove antiprotons from the magnetic trap during shut down. The dashed line shows the on-axis magnetic field. (b) Pseudo-potential for antiprotons, combining the interaction of the magnetic moment of the gyrating particles with the magnetic field and the electrostatic potential energy (cf. Eq. 3.11). Examples for three different values of the initial perpendicular energy $E_{0\perp}$ are shown. (Adapted from [68])

and so we have taken the conservative approach of modeling a very broad range of initial conditions (and many more trapped antiprotons than we have observed in all trapping experiments that we have performed).

Simulations show that fewer than 2% of antiprotons survive the clearing cycles, and all that survive have $E_\perp > 50$ eV [101]. Based on these calculations we conclude that the minimum transverse energy of antiprotons that is required for mirror trapping is several orders of magnitude higher than the energy of antiprotons that are typically used for antihydrogen production.

3.4.2   Measurement Variations

Once the antiproton clearing process is complete the neutral trap is ramped down and we look for antihydrogen annihilation events within a 30 ms window. The magnetic trap depth 30 ms after initiation of the ramp down procedure is less than 0.1% of its initial value, hence it is believed that all trapped atoms should have been released by this time.

In order to deflect any charged particles that might not already have been cleared from the trap (perhaps some exotic mirror trapped antiprotons) we apply a precautionary electric field before switching off the magnets, and we allow it to be present until the end of the experiment. This electric field has the same strength as the clearing fields: 500 Vm^{-1}. The use of this bias field allows us to use the annihilation imaging detector to distinguish between the release of trapped antihydrogen (which is neutral and is therefore unaffected by electric fields) and that of mirror-trapped antiprotons.

During our initial demonstrations of antihydrogen trapping we conducted the experimental cycle 335 times in six distinct variations. In one variation of the experiment, referred to as the 'left bias' configuration, just before the magnet shut down we erect a static electric field to deflect any remaining antiprotons to the left (negative z or upstream direction) of the apparatus as they are released. In the second variation, referred to as the 'right bias' configuration, a static electric field intended to deflect any remaining antiprotons to the other end of the device is employed. In the third variation, referred to as the 'no-bias' configuration, all Penning trap electrodes are held at ground potential during the magnet shut down.

To ensure that detected events are in fact associated with the annihilation of anti-hydrogen atoms we repeated the above measurements with deliberately heated positrons (up to 1100 K). Under these conditions the antihydrogen production rate is significantly suppressed (cf. Sec. 3.2.1). Moreover, any antihydrogen that is formed is unlikely to be trapped because it is too energetic (the antiprotons approach thermal equilibrium with the hot positrons through Coulomb interactions). The number of annihilation events detected during 1 s of mixing with heated positrons was 97 ± 16 which is remarkably smaller than the number of annihilation events during mixing with cold positrons (5000 ± 400).

A summary of all the measurement sequences and corresponding observations for our initial demonstration of antihydrogen trapping [73] is presented in Tab. 3.2. In total, the experiment was repeated 335 times; 38 annihilation events were observed corresponding to a rate of 0.11 events per attempt. The experiments conducted with heated positrons (which act as null experiments) were repeated 246 times; 1 annihilation event was observed corresponding to a rate of 0.47×10^{-2} events per attempt. The total observation time for experiments with cold positrons was $335 \times 30 \times 10^{-3}$ s $=10.05$ s during which time we expected to count 0.46 ± 0.01 cosmic ray events misidentified as antihydrogen annihilation.

Measurement Variation	Experimental Cycles	Events
No bias	137	15
Left bias	101	11
Right bias	97	12
No bias, heated positrons	132	1
Left bias, heated positrons	60	0
Right bias, heated positrons	54	0

Table 3.2: A summary of antihydrogen trapping attempts. The number of detected annihilation events is shown in last column [73].

Using numerical calculations, we have studied the trajectories of antihydrogen atoms and hypothetical mirror trapped antiprotons in the time dependent fields of the trap as it is shut down. In these simulations the particles propagate via classical force equations: the Lorentz force for antiprotons and the dipole-gradient force for the antihydrogen atoms. We also consider the effect of the bias electric field that is applied during the magnet shut down.

These simulations are intended to test the plausibility of the time- and location-distributions of detected annihilation events and to provide insight into the antihydrogen release process.
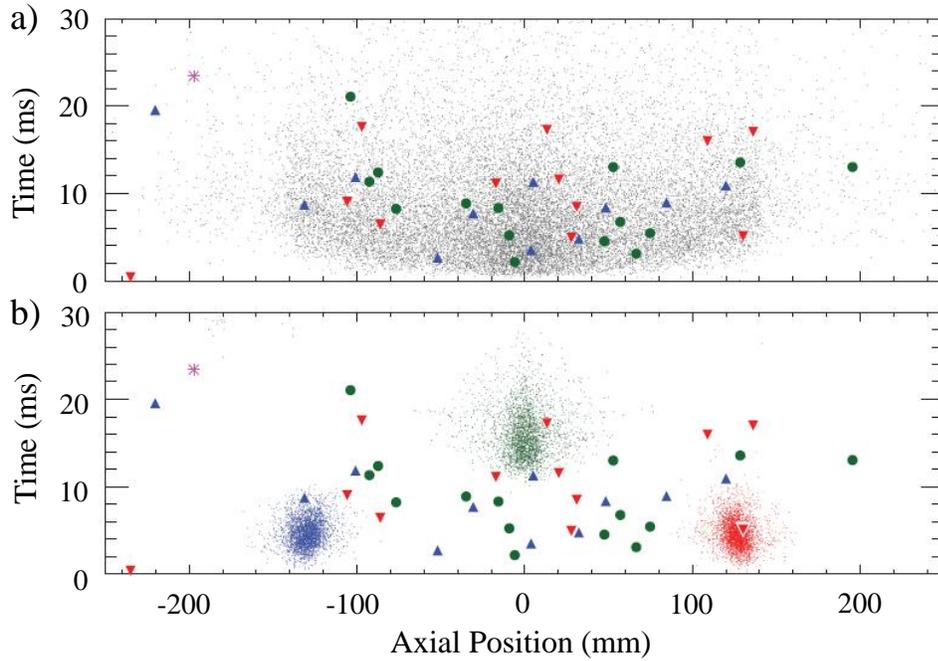
Figure 3.15: Temporal (t) and spatial (z) distribution of observed annihilation events after magnet shut down. (a) Measured $t - z$ distribution for annihilations obtained with no bias (green circles), left bias (blue triangles), right bias (red triangles) and heated positrons (violet star). The grey dots are from a numerical simulation of antihydrogen atoms released from the trap during the magnet shut down. The simulated atoms were initially in the ground state, with a maximum kinetic energy of 0.1 meV. The typical kinetic energy is larger than the depth of the neutral trap, ensuring that all trappable atoms are considered. (b) The coloured dots represent the $t - z$ distribution results for a numerical simulation of mirror trapped antiprotons being released from the trap. The colour codes are the same in both figures. Notice that the axial resolution of the silicon detector is ~ 5.6 mm. (Adapted from [73])

Figure 3.15 shows the time (t) and axial position (z) distribution for the 38 annihilation events observed using cold positrons and the one annihilation event observed using heated positrons (violet star). These events are superimposed on the $t - z$ distributions expected for released antihydrogen and antiprotons on the basis of numerical calculations. Figure 3.15.a

illustrates the consistency between the experimental annihilation event distribution with that predicted by simulations of antihydrogen atoms, while Fig. 3.15.b shows the lack of consistency between the observed event distribution and the prediction for simulated mirror trapped antiprotons.

We continued to repeat the first three variations of the antihydrogen trapping experiments listed in Tab. 3.2 after submitting our initial report for publication [73]. By the end of the 2010 experimental run, we had conducted 1243 trapping experiments and observed 320 annihilation events for an effective rate of 0.26 events per attempt [101]. An updated $t - z$ distribution of events is shown in Fig. 3.16.
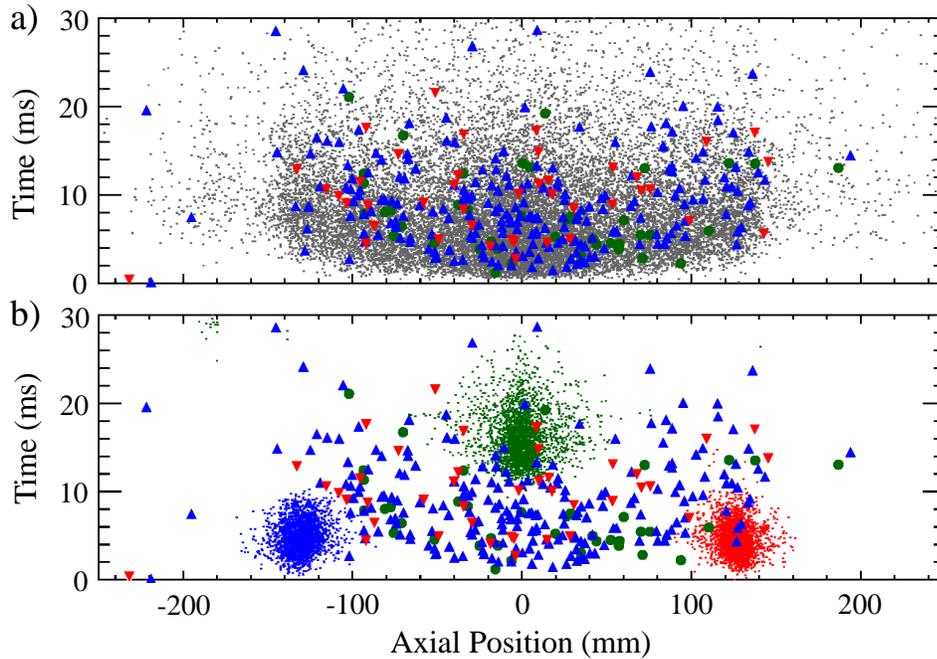
Figure 3.16: An updated $t - z$ distribution of annihilation events, reflecting the situation at the end of the year 2010. The colour codes and descriptions are the same as in Fig. 3.15.

3.4.3   Antihydrogen Storage Lifetime

After successful demonstration of antihydrogen confinement, the first major question to address was to find the upper limit on the storage lifetime of the trapped antiatoms. Sufficiently long confinement times are critical to the success of precision laser or microwave spectroscopy experiments on trapped antihydrogen atoms. Atomic hydrogen is stable [99], and in theory, antihydrogen is as well. However even matter atoms cannot be confined indefinitely in a magnetic trap.

Magnetic confinement of matter atoms in cryogenic environments has already been demonstrated for very long times (10-30 minutes) [103, 104]. Likewise, one should be able to confine antihydrogen but there are a number of mechanisms that may lead in annihilation of the trapped antihydrogen atoms. In principle antihydrogen atoms in ALPHA can undergo spin-exchange collisions and escape from the trap. However, at present the atom density (at most a few trapped atoms in any attempt) is far too low for this loss mechanism to be significant. Instead elastic collisions or annihilation on background gas is the most likely loss mechanism, simply because of the much higher density of the background gas.

In our initial trapping experiments (Sec. 3.4.2) the trapping time was set to 172 ms, the minimum time required to complete all of the necessary manipulations. In those experiments the goal was to demonstrate magnetic confinement of antihydrogen and so a short trapping time is beneficial because it minimizes the likelihood of misidentifying a cosmic ray event.

To study the storage lifetime of trapped antihydrogen atoms, a series of 201 measurements similar to a generic trapping experiment were carried out. In these experiments the time between the clearing pulses and the magnet shutdown was varied. This time window sets the confinement time for the trapped atoms.

The cosmic ray rejection algorithm used for these experiments was identical to that discussed previously in Sec. 3.3. We observed that even at a 1000 s confinement time there is strong evidence for survival of trapped atoms. The probability that the events observed
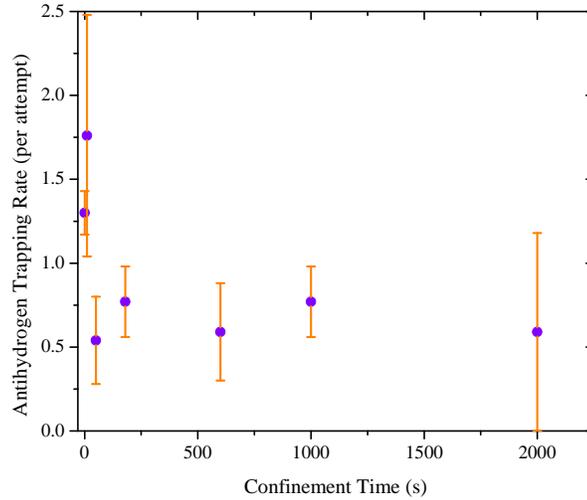
Figure 3.17: The number of annihilation events detected (per attempt) after magnet shut down, as a function of confinement time. The confinement time is varied from 0.4 to 2000 seconds. Error bars indicate one standard deviation from the mean for multiple measurements.

after a 1000 s confinement time are caused by cosmic rays is 10^{-5}, corresponding to a statistical significance of 8σ. After 2000 s the statistical significance of the observed annihilation events is reduced to 2.6σ [74]. A summary of the various experiments that were performed is presented in Tab. 3.3.

Confinement time (s)	0.4	10.4	50.4	180	600	1000	2000
Experimental Cycles	119	6	13	32	12	16	3
Detected Events	76	6	4	14	4	7	1
Background Estimate	0.17	0.01	0.02	0.05	0.02	0.02	0.004
Statistical significance (σ)	$\gg 20$	8.0	5.7	11	5.8	8.0	2.6
Trapped $\bar{\text{H}}$ Rate	1.3 ± 0.13	1.76 ± 0.72	0.54 ± 0.26	0.77 ± 0.21	0.59 ± 0.29	0.77 ± 0.29	0.59 ± 0.59

Table 3.3: A summary of antihydrogen trapping experiments conducted with different confinement times [74].

A 1000 s confinement time is already close to four orders of magnitude longer than our initially reported value (172 ms). As far as laser and microwave spectroscopy of antihydrogen is concerned, these confinement times are essentially infinitely long and open the door to a

variety of possible spectroscopic measurements. In chapter 5 a series of measurements in which microwave radiation is used to drive transitions between hyperfine levels of ground-state trapped antihydrogen atoms will be discussed. These experiments represent the first-ever spectroscopic interrogation of an antimatter atom.

Chapter 4

Microwave Injection and Diagnostics

The primary goal for ALPHA during the 2011 experimental run was to embark upon the study of ground state antihydrogen via resonant interaction with microwave radiation. To accomplish this, microwaves need to be injected into the mixing trap where antihydrogen atoms are synthesized and stored. The minimum magnetic field in this volume is of order 1 Tesla, and so radiation at frequencies in the range 28-30 GHz (lying in K_a band) is required to induce PSR spin flip transitions. An Agilent 8257D PSG synthesizer is used as the source of this radiation; it is delivered to the atom trap via methods (circuits) described in this chapter.

The microwave circuit employed in our experiments was designed to encompass two independent injection paths. In one mode of operation a low power microwave beam radiated from an antenna outside of the vacuum system is focused through a window and then reflected down the bore of the apparatus. This low power microwave radiation is mainly used as a diagnostic tool for measuring the magnetic field via electron cyclotron resonance (see Sec. 4.4). This method of injection will be referred to as external injection. In the other mode of operation higher power microwave radiation enters the vacuum system via waveguide and is then launched from an antenna into the electrode stack. This method of injection will be referred to as internal injection. This higher power radiation is mainly used for PSR spin flip experiments in which the internal states of trapped antihydrogen atoms

are manipulated.

Propagation of microwaves down the electrode stack is a complex problem, and one that is largely intractable from a modeling perspective. The geometry bears some resemblance to an overmoded circular waveguide into which many azimuthal breaks in current paths and several abrupt changes of diameter have been introduced (see Fig. 2.3). In detail, however, the situation is further complicated by complex electromagnetic coupling to the thin annular space between the electrodes and the trap vacuum wall (which houses the many conductors and dielectric materials associated with the electrode biasing system).

In this chapter both of the injection methods (external and internal) are first described and compared. This is followed by a brief presentation of the theory of collective modes of oscillation of a charged plasma and the electron cyclotron resonance methods we use to determine the magnetic field in the vicinity of its minimum, at the centre of the ALPHA trap. In a complimentary mode of operation we are able to make in-situ measurements of time average microwave electric field amplitudes inside the electrode stack. Knowing the electric field, we are able to make reasonable estimates of in-situ microwave magnetic fields.

4.1 External Injection

External injection was the first method by which we introduced microwave radiation into the apparatus. It was relatively simple to implement and did not require any modifications that might compromise the vacuum integrity of the apparatus. It involves two metallic 'mirrors' or reflectors. One of these reflectors is a simple stainless steel plate mounted at the bottom end of the vertical translator (see Sec. 2.6). It can thus be positioned vertically so that it is aligned with the bore of the electrode stack. It is oriented such that vectors normal to its surface lie in the horizontal plane and make a 45° angle with respect to the trap axis and the normal to a 4″ diameter glass viewport. The surface area of this reflector is approximately 40 cm^2 (Fig. 4.1). It is cut in the shape of the letter L simply because of physical constraints imposed by the vertical translator.

The second reflector consists of a 5×10^2 cm^2 section of the inner surface of a prolate aluminum ellipsoid of revolution with the following parameters: semi-major axis=19.52 cm,
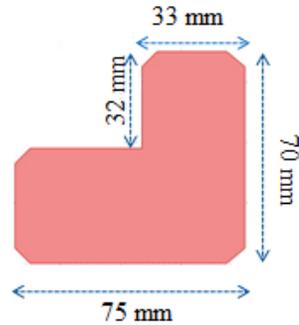
Figure 4.1: The L-shaped stainless steel reflector. This reflector is mounted in the vacuum system on the vertical translator at a 45° angle with respect to the trap axis. Microwave radiation passing through a glass viewport and striking this mirror is reflected down the bore of the apparatus into the mixing region.

semi-minor axis=17.00 cm, foci=9.59 cm, and eccentricity=0.49. The surface area of this mirror corresponds to approximately $1/8^{\text{th}}$ of a full prolate ellipsoid [43]. It is located outside of the apparatus, in front of a $4''$ diameter glass view port mounted on a CF-100 vacuum flange. It is positioned such that one of its focal points is located outside of the apparatus while the other sits on or near the surface of the L-shaped internal reflector. The throat of a 20 dB horn antenna is positioned at the external focal point of the ellipsoidal mirror and oriented so that microwave radiation is focused onto the internal reflector and down the bore of the electrode stack. The horn antenna is fixed with respect to the ellipsoidal mirror and the ellipsoidal mirror is fixed on an X-Y translator plate. Using the X-Y translator we are able to precisely locate the other focal point of the ellipsoidal mirror on the surface of the internal reflector. A schematic drawing showing a mock-up of the external injection path is shown in Fig. 4.2. In this arrangement, which was used to study efficiency prior to implementation, a microwave detector is used to characterize the injected microwave intensity.

In the final configuration of the external injection circuit, the insertion loss of the microwave components (coaxial cables and switches) situated between the synthesizer and the external horn antenna was measured to be 10.8 dB at a frequency of 28.233 GHz. In addition, studies conducted using electron cyclotron resonance (described later in this chapter)
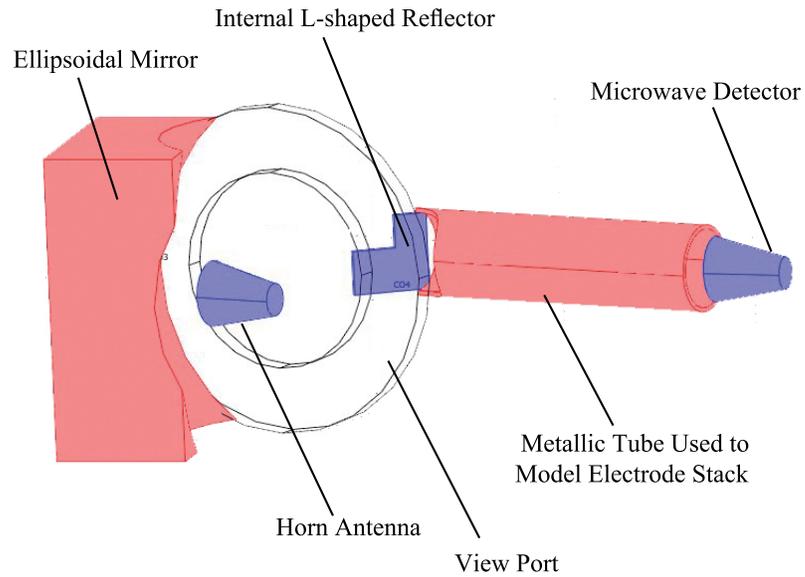
Figure 4.2: A schematic drawing showing a mock-up of the external injection scheme. The curved surface of the ellipsoidal mirror faces to the right. The horn antenna is positioned at one of the focal points while the internal reflector is positioned at the other. Microwaves propagate from the horn antenna to the ellipsoidal mirror, pass through the window, strike the internal reflector and finally enter the electrode stack (the X-Y translator plate is not shown). In this setup, a second horn antenna feeding a microwave detector is placed at the far end of the tube used to mimic the electrode stack.

indicate a further 20 dB attenuation in microwave power from the external horn antenna (attached to the ellipsoidal mirror) to the centre of the mixing trap. The external injection scheme was first implemented in June 2010 (Fig. 4.3).
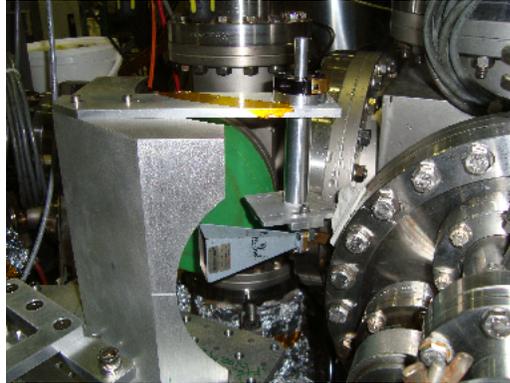
Figure 4.3: Photograph showing the external injection setup. The horn antenna and the ellipsoidal mirror are located in front of the viewport flange (covered with a green cap). The internal reflector is inside the apparatus and cannot be seen.

4.1.1   Positron Cyclotron Resonance Heating (PCRH)

The ability to introduce microwave radiation had an immediate potential application to the antihydrogen trapping sequence. As was described in Sec. 3.4.2, in order to establish the validity of the original antihydrogen trapping signal, a series of control experiments with intentionally heated positrons (up to 1100 K) was performed. The goal of this exercise was to demonstrate suppression of antihydrogen confinement when the kinetic energy of the antiatoms is increased.

In our normal sequence, heating of the positron plasma is accomplished by applying a white noise drive potential to one of the adjacent electrodes. It is likely that the antiproton plasma, which is confined next to the positron plasma, is also influenced by the noise drive. In other words, our normal methods are not able to produce a situation in which only the positron plasma temperature is modified prior to antihydrogen synthesis and trapping.

As an alternative approach, microwave radiation can be employed to heat the positron

plasma. Once the positron and antiproton plasmas are confined at the center of the apparatus, and immediately before mixing, microwaves tuned to the positron cyclotron frequency are injected. This heats the positron plasma but has no effect on the antiprotons, since the antiproton cyclotron resonance frequency is orders of magnitude lower than that of the positrons. The positron plasma temperature can be controlled by varying the duration and amplitude of the applied microwave radiation.

Figure 4.4 shows the effect of heating a standard positron plasma used for antihydrogen trapping (i.e. $\sim 2 \times 10^6$ particles and density of 5.5×10^7 cm^{-3}) at different microwave power levels. At each power level, a train of 4 μs duration microwave pulses is injected for 1 second. The microwave frequency is swept over a 5 MHz range centred on the positron cyclotron frequency, in 50 KHz steps [11]. The manner in which the electron/positron cyclotron frequency in a static magnetic field is determined will be discussed later, in Sec. 4.4.1.
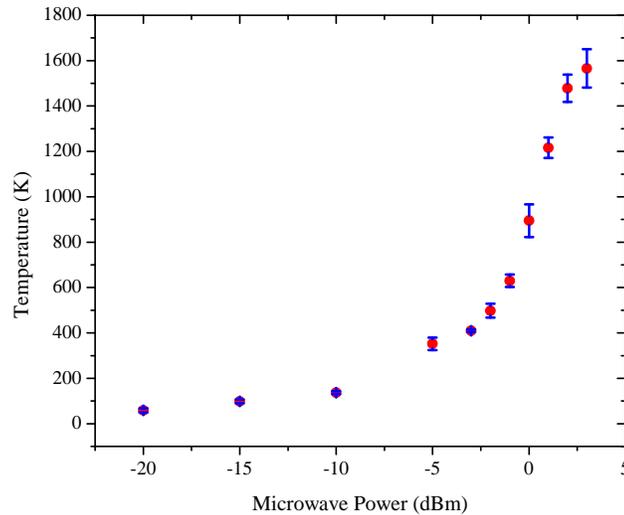
Figure 4.4: Temperature increase for a standard positron plasma as used in ALPHA trapping experiments, when the particles are subjected to microwaves that are swept through the cyclotron frequency. Here the external injection scheme was employed. The microwave power corresponds to the output of the synthesizer (Reproduced from [11]).

Figure 4.5 shows the number of particles that are still confined in the electrostatic potential well, after the 1 s microwave heating procedure. In order for this to be useful for antihydrogen trapping control experiments, the positron plasma needs to be heated but no particles should be lost. The data shown in Fig. 4.5 indicate that the threshold above which particle loss becomes significant is about 0 dBm (at the output of the synthesizer).
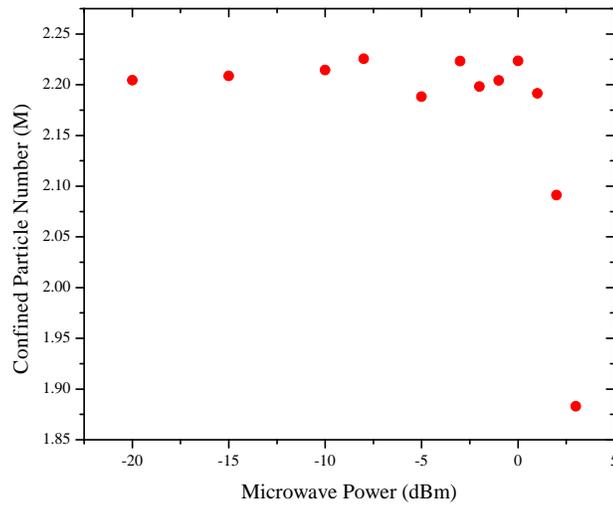
Figure 4.5: Number of positrons remaining in the electrostatic potential well after 1 s of microwave heating (cf. Fig. 4.4).

When the microwave radiation is turned off, the positrons cool down through cyclotron radiation. There are well-established theories that predict the radiative damping time of leptons in an ideal Penning trap. In a magnetic field of 1 T the cooling time is calculated to be 2.6 s [77, 105]. Figure 4.6 shows the measured radiative cooling curve for a positron plasma subjected to a 1 s microwave heating sweep at 0 dBm, as described above. Fitting an exponential decay function to the data yields a cooling time of 2.8 s, which is in good agreement with theory.
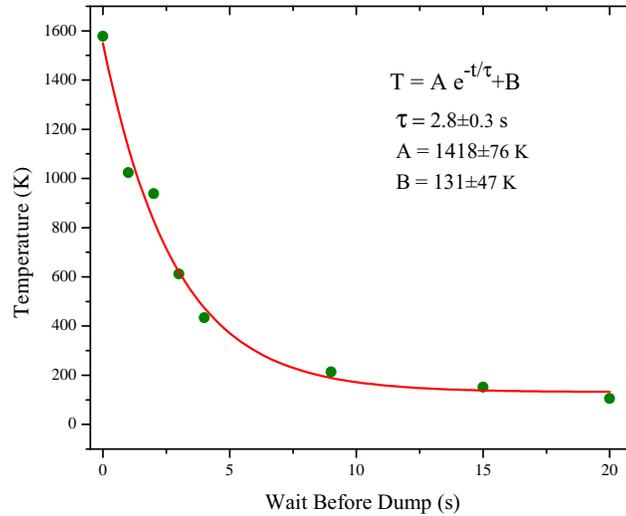
Figure 4.6: Radiative cooling curve for a positron plasma subjected to a 1 s microwave heating sweep at 0 dBm.

4.2 Internal Injection

As previously mentioned, data from our first round of experiments using microwaves introduced via the external injection method (July 2010) suggested that the total attenuation from the external horn antenna to the center of the trap was $\sim$ 20 dB. In 2007, when the apparatus was open, Prof. Walter Hardy and Sarah Seif el Nasr (University of British Columbia) measured the transmission of microwaves down the full length of the electrode stack. Their measurements suggested that the insertion loss to the centre of the stack at 28.5 GHz was about 10 dB. This implies that an additional $\sim$ 10 dB insertion loss occurs between the external microwave horn and the electrode stack. The most effective method for reducing this loss was to put the horn inside the apparatus. By mounting the horn antenna on the vertical translator it can be moved until its axis coincides with the axis of the Penning trap. This modification was implemented in time for the 2011 experimental run. A decision was made to leave the internal reflector in place, in order to provide an independent path for microwave injection.

4.2.1   New Vertical Translator Design

Accommodating the microwave horn on the vertical translator required many design modifications, including the addition of a custom waveguide window to bring microwaves from outside into the UHV environment of the apparatus. In its new configuration, the electron gun, MCP, internal microwave reflector plate, and the pass-through tube are still present. In addition, a vertical length of rigid WR28 waveguide connects a waveguide window (on the top of the vertical translator) to a 20 dB gain rectangular microwave horn that is mounted along with the rest of the instruments at the bottom end of the vertical translator (as shown in Fig. 2.10). The microwave horn is aligned so that its axis is horizontal and points toward the center of the apparatus (mixing trap). Flexibility to move this rigid assembly is provided by a length of high frequency coaxial cable outside of the vacuum space.

The waveguide window is a home-made hermetically-sealed block of quartz (length = 3.27 mm, width = 5.46 mm, and height = 2.67 mm) that separates the vacuum from the outside environment. The length of the window was chosen to be close to a half wavelength over the frequency band of interest. It is glued (using Torr-Seal) into a rectangular stainless steel plate of thickness 3.27 mm with a slightly larger (tapered) rectangular hole. There is also a 0.127 mm thick Kapton film between the window assembly and the mating microwave flange. The transmission of the window over the frequency range 28-30 GHz was measured to lie between 60% and 80%.

4.2.2   Circuit

The external injection path described in Sec. 4.1 is suitable for injecting low power microwaves into the apparatus for electron cyclotron resonance experiments. However higher microwave power levels are desirable for driving transitions between hyperfine levels of ground-state antihydrogen atoms because the probability of inducing a spin flip is directly related to the microwave power amplitude (further discussed in Chapter 6). An independent microwave circuit was developed to deliver microwaves for the internal injection scheme. In the process, we included an option to amplify the microwaves using a broad-band Miteq AMF-4B amplifier prior to injection into the apparatus. The nominal small-signal gain of the amplifier is close to 20 dB over the frequency range 26-31 GHz. It delivers a maximum

of approximately 4 W of power. The amplified microwaves then pass through the window and feed the horn mounted on the vertical translator (Sec. 4.2.1 and Fig. 2.10). A schematic of this circuit is shown in Fig. 4.7.

Just before the waveguide enters the vacuum system, a microwave detector (an HP R422A crystal detector) is used to monitor the power level. It is attached to the coupled port of a 10 dB directional coupler. The output of the microwave detector is then amplified using a custom-made amplifier with a measured gain of 99.5. The amplified output of the microwave detector is then digitized and recorded by a computer. Alternately the amplified detector output is fed to an oscilloscope or precision voltmeter for calibration purposes.

Figure 4.8 shows the measured calibration between microwave power and the amplified detector readout at $f = 28.5$ GHz. Once calibrated, the detector is used to monitor the power delivered by the microwave circuit to the cryostat over the frequency range of interest. An example of microwave power delivered to the apparatus for a few frequencies is shown in Fig. 4.9. Here power refers to the forward power entering the apparatus as inferred from the microwave detector. Significant variations in forward power are observed.

Figure 4.10 shows the measured reflection coefficient at the point where the microwaves enter the apparatus, over the frequency range 27.5 GHz to 30.0 GHz. Based on these data we identified a pair of frequencies (that are separated by the hydrogen hyperfine splitting frequency) for PSR spin flip experiments such that the microwave injection efficiency was high (reflection coefficient low). We found a pair of frequencies such that the sum of their reflection coefficients was less than that for any other pairs in this range. The two frequencies that we identified are 28.275 GHz and 29.695 GHz, and the sum of the reflection coefficients at these frequencies was 0.1. These two frequencies correspond to the target frequencies for the PSR transitions that will be discussed in Chapter 5. The sum of the reflection coefficients for other pairs of frequencies in this range can be as high as 0.7. The efficiency of microwave transmission into the apparatus at these frequencies was subsequently verified using electron cyclotron resonance experiments.

The attenuation in microwave power through the high power injection path is ~ 13.6 dB, measured from the source to the point where the microwaves enter the apparatus. The
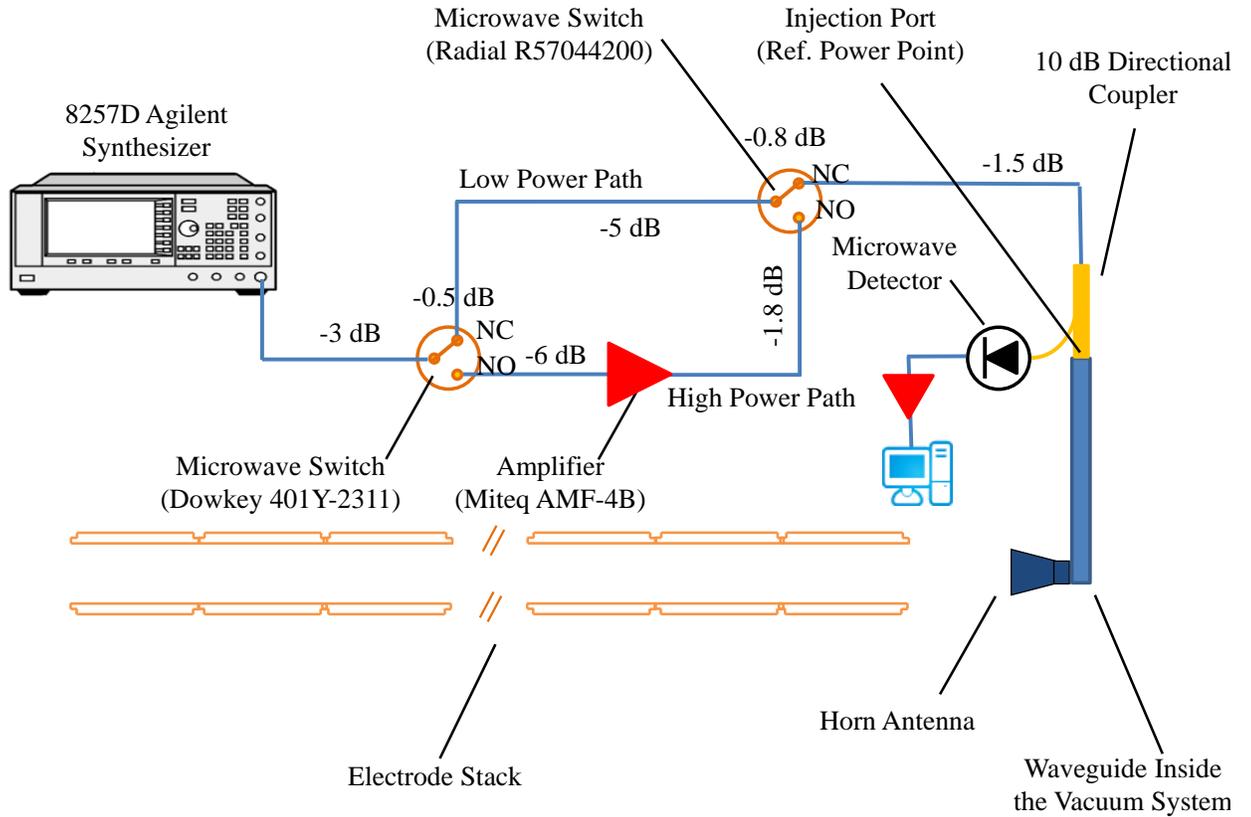
Figure 4.7: An outline of the microwave circuit used to introduce microwaves into the ALPHA apparatus via the internal horn antenna. Coaxial cables are indicated as solid lines. The numbers specify the insertion loss of various components over the frequency range of interest.
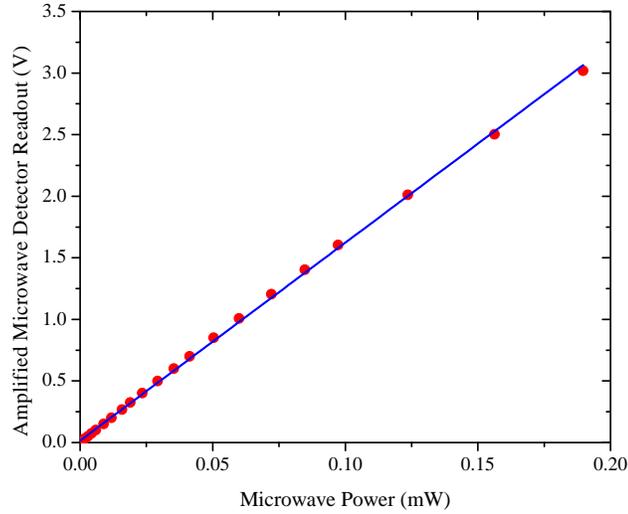
Figure 4.8: Calibration of the amplified microwave detector readout with low power microwaves (without amplification) at f= 28.5 GHz.
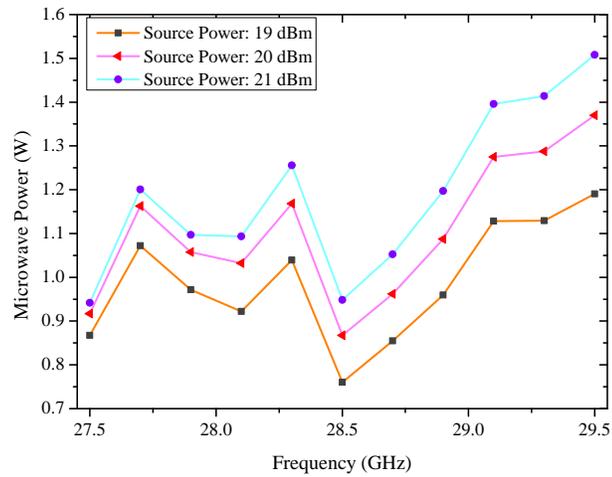
Figure 4.9: Microwave power delivered to the apparatus, as inferred from the microwave detector. From the microwave's point of view, the apparatus has a complicated form of input impedance which varies as a function of frequency. Higher degrees of impedance mismatch between the source and apparatus lower the power transmission efficiency. Data for three different source (synthesizer) power levels are shown. The straight lines connecting data points are only meant as guides for the eye.
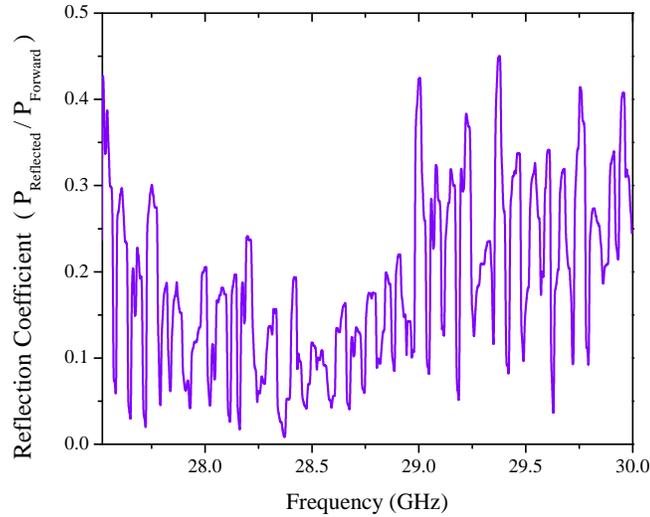
Figure 4.10: Reflection coefficient at the point where microwaves are injected into the apparatus. The frequency is swept from 27.5 GHz to 30.0 GHz in 108 equal steps.

primary factor contributing to this loss is the attenuation in the 2.92 mm coaxial cables (-12.3 dB) used to interconnect components of the circuit as shown in Fig. 4.7. Additional contributions come from various connectors, adapters, and microwave switches. Note that the practical low frequency limit for using the internal injection scheme is 22 GHz. This is set by the cut off frequency for WR28 waveguide.

4.3   Collective Modes of a Plasma

If a perturbation is applied to a trapped cylindrically symmetric non-neutral plasma in thermal equilibrium, its subsequent motion may be described as a superposition of natural (collective) modes of oscillation [106]. These modes are classified by two indices (l, m), where l is the axial mode number and m is the azimuthal mode number. Measurements of mode characteristics can be used as a non-destructive diagnostic for the properties of the plasma [107]. In our experiments, we exclusively employ axially symmetric modes for which $m = 0$.

The lowest-order mode has $l = 1$. In this mode the plasma acts like a single particle that executes an axial oscillating motion in the electric potential well. This mode is known as the 'dipole' mode (Fig. 4.11; left). The next mode has $l = 2$ and is known as the quadrupole or 'breathing' mode. In this mode the plasma undergoes an elongation-compression behavior along the axis of the magnetic field (Fig. 4.11; right) [107]. In other words, this mode involves an oscillation of the plasma aspect ratio (defined as the ratio of plasma length to diameter) in time.

Figure 4.11: Cartoon showing the two lowest-order modes of an oscillating plasma. The dipole mode (left) involves a center-of-mass oscillation of the plasma about its equilibrium position. The quadrupole mode (right) involves an oscillation of the plasma aspect ratio.

An analytic theory for the frequencies of the modes of a plasma in global thermal equilibrium has been developed [106]. The angular frequency of the axisymmetric mode, for cold ($T = 0$ K) and strongly magnetized plasmas ($\omega_c \gg \omega_p$) is given by:

$$\omega_\ell = \omega_p \left[1 - \frac{k_2}{k_1} \frac{P_\ell(k_1) Q'_\ell(k_2)}{P'_\ell(k_1) Q_\ell(k_2)} \right]^{-1/2} \tag{4.1}$$

where P_ℓ and Q_ℓ are Legendre functions of the first and second kinds, respectively and P'_ℓ and Q'_ℓ are their derivatives. Additionally, ω_c is the cyclotron frequency, α is the aspect ratio of the plasma, and ω_p is the plasma frequency:

$$\omega_p = \sqrt{\frac{nq^2}{\epsilon_0 m}}. \tag{4.2}$$

The arguments of the Legendre functions of the first and second kinds and their derivatives are given by:

$$k_1 = \frac{\alpha}{\sqrt{\alpha^2 - 1 + \frac{\omega_p^2}{\omega_\ell^2}}} \tag{4.3}$$

and

$$k_2 = \frac{\alpha}{\sqrt{\alpha^2 - 1}}. \tag{4.4}$$

For a plasma at temperatures $T > 0$ confined to a perfect harmonic potential (and in the limit that the Debye length is small compared to the plasma size) one finds that the angular quadrupole frequency ω_2 is shifted upward such that [108]:

$$\omega_2^2 = (\omega_2^c)^2 + 20 \left[3 - \frac{\omega_p^2 \alpha^2}{2(\omega_2^c)^2} \frac{\partial^2}{\partial \alpha^2} f(\alpha) \right] \frac{k_B T}{mL^2} \tag{4.5}$$

where

$$f(\alpha) = \frac{Q_1 \left(\frac{\alpha}{\sqrt{\alpha^2 - 1}} \right)}{\alpha^2 - 1} \tag{4.6}$$

and ω_2^c is the angular quadrupole frequency as T $\to$ 0, k_B is the Boltzmann constant, and L is the length of the plasma. The typical electron plasmas that are used in ALPHA for cyclotron resonance measurements consist of 3×10^6 to 4×10^7 of electrons. These plasmas have a radius of 1 mm and a length of 20 to 40 mm. The quadrupole mode frequency for these plasmas is in the range 24 MHz to 28 MHz [109].

Equation 4.5 reveals that the frequency of the quadrupole mode shifts linearly with temperature in the limit of small temperature changes, which is generally the case in our experiments. That is

$$\Delta f_2 \approx \beta \Delta T \tag{4.7}$$

where f_2 is the quadrupole frequency and β is the calibration coefficient between the quadrupole frequency change and the temperature change, which is obtained experimentally. For instance, $\beta^{-1} = 3.7 \pm 0.3$ K/kHz for typical electron plasmas that are used in ALPHA experiments involving 12×10^6 electrons, $\alpha = 16$, $L = 26$ mm, and a base plasma

temperature of ~ 150 K [109]. This correlation is used extensively in the microwave spectroscopy experiments described later in this thesis. Plasma mode frequencies are measured and employed as a diagnostic to deduce the amplitude of microwave fields in-situ. Moreover, when they are correlated with the frequency of microwave fields applied at the cyclotron resonance, we obtain an absolute measurement of local (static) magnetic fields.

4.3.1   Instrumentation

Plasma mode frequency measurements were largely performed on electron plasmas confined to a harmonic three-electrode potential well (Fig. 4.12). The plasma is excited using an arbitrary waveform generator (National Instruments PXI-5421) to deliver a Gaussian-modulated sinusoidal pulse to one of the electrodes adjacent to the plasma. This drive signal is gated by a TTL pulse with a typical duration of 1 μs. The plasma responds by oscillating, or 'ringing', at one (or more) of its natural frequencies (Fig. 4.13). This measurement is normally performed by monitoring the potential on the central electrode used to form the potential well. Note that the typical time scale for the ringing motion of the plasma to die away is less than 100 μs. As the plasma oscillates, it induces a potential on the electrodes that can be measured.
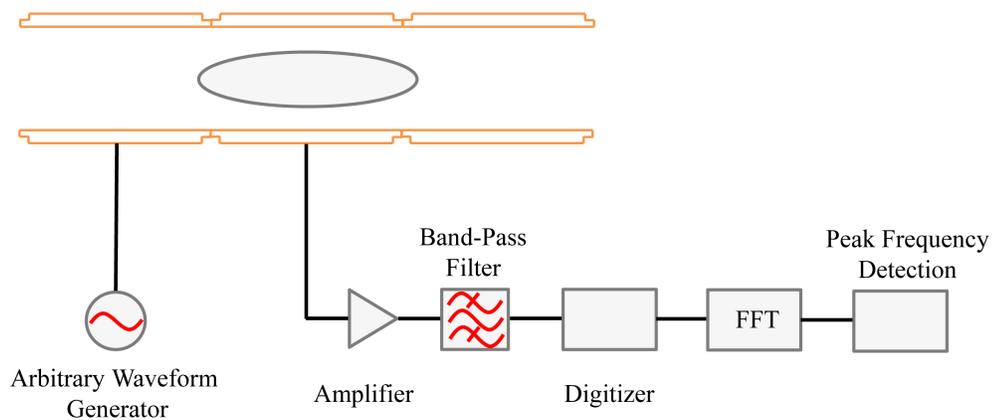
Figure 4.12: Schematic of setup employed to drive plasma oscillations and detect the response frequency.

The oscillating voltage induced by plasma motion is picked up from the central electrode

by a room temperature broadband amplifier (20 dB gain, HP 8447A), it is then filtered and digitized by a high-speed digitizer (100 MS/s National Instruments PXI-5122). A Fast Fourier Transform is then applied to the response (cf. Fig. 4.13) to obtain a frequency spectrum. Typically responses from a series of pulses are averaged and used to produce the spectrum. This spectrum is then analysed using a peak detection routine to determine the frequency of the mode.
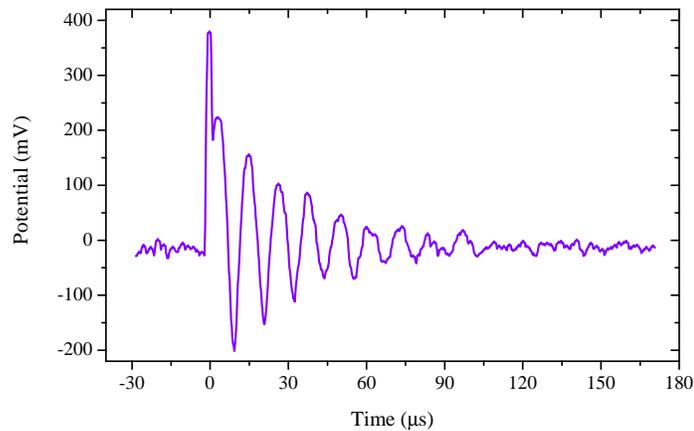
Figure 4.13: Oscillation of a plasma after being driven at its quadrupole mode frequency (here approximately 26 MHz). The response is detectable by monitoring the potential on the central electrode.

4.4 *In-Situ* Measurement of Magnetic Fields

The primary application of the plasma mode diagnostic in ALPHA is to infer the magnetic field along the axis of the Penning trap. This is accomplished by resonant heating of an electron plasma when microwave radiation is injected at the cyclotron frequency. This heating is driven by the microwave electric field. The methods used to perform this measurement are described below.

4.4.1 Electron Cyclotron Resonance: Static Magnetic Field Measurement

Measurement of the electron cyclotron resonance frequency is typically conducted as follows: an electron plasma comprising approximately 7×10^7 particles (density of 6.5×10^{14} m^{-3}) is loaded at the center of the apparatus and then a series of 4 μs duration microwave pulses is injected via either the internal or external path (see Secs. 4.1 and 4.2) at a rate of one pulse every 30 s. The microwave frequency is incremented every time a pulse is injected, in order to scan through the cyclotron resonance.
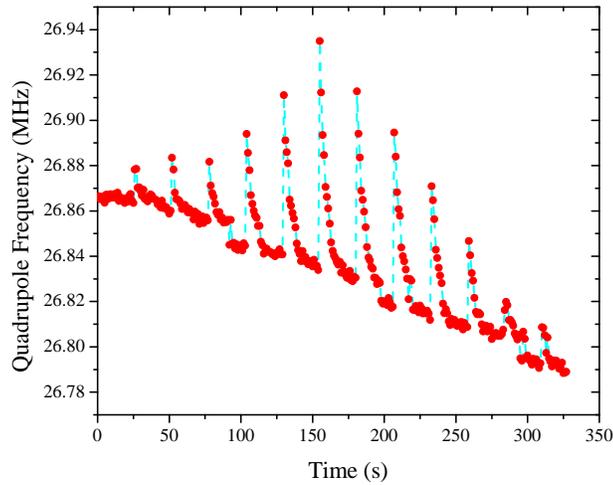
Figure 4.14: Quadrupole mode frequency of an electron plasma in response to a series of microwave pulses, as the microwave frequency is stepped through the cyclotron resonance. The steady background drift in the mode frequency is caused by gradual expansion of the plasma leading to smaller aspect ratios and consequently lower quadrupole frequencies (Reproduced from [11]).

Meanwhile, collective modes of the electron plasma are excited (via potentials applied to the trap electrode adjacent to the plasma), and the frequency of the quadrupole mode response is determined and recorded. As previously described (Sec. 4.3) the frequency of this particular mode is a function of aspect ratio and plasma temperature [108]. Every time a microwave pulse is injected the plasma temperature increases and the quadrupole mode frequency shifts upwards. The time between microwave pulses is sufficient for the plasma to cool down and return to equilibrium (via cyclotron radiation). When the frequency of

the applied microwave electric field matches the cyclotron frequency, the plasma response is maximal (Fig. 4.14) [11, 109]. By plotting the quadrupole frequency shift as a function of microwave frequency, the cyclotron resonance frequency is obtained and the corresponding magnetic field is inferred. Figure 4.15 shows a generic cyclotron resonance lineshape acquired from a plasma in a uniform magnetic field.
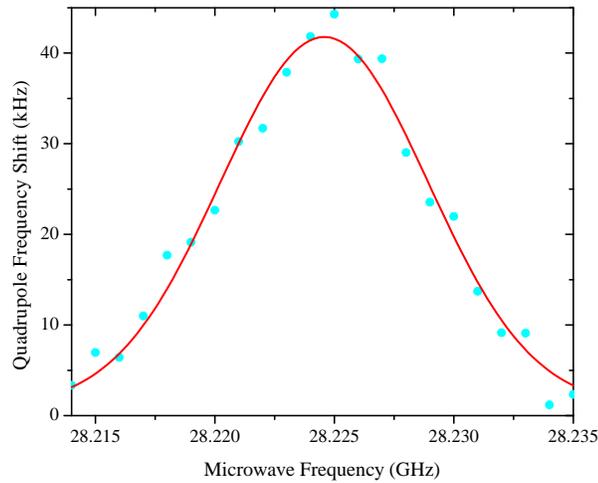
Figure 4.15: A generic cyclotron resonance lineshape in a uniform magnetic field. In this example the resonance frequency is $f_c = 28.224$ GHz and the width (FWHM) of the distribution is 7.8 MHz.

4.4.2 Calibration of Solenoid and Mirror Coil Fields

Clearly, using the plasma mode diagnostic one can measure the contribution of individual magnets to the total axial magnetic field at the centre of the ALPHA trap where antihydrogen is synthesized and confined. Thus, one can calibrate the magnetic field along the trap axis with, for instance, the current flowing in each magnet. These calibrations are then useful for setting magnetic fields in control sequences. Figure 4.16 shows an example of a measurement in which the electron cyclotron frequency is measured as function of current in the solenoid. These data directly reflect the on-axis magnetic field at the center of the trap. As expected, a linear correlation is observed.
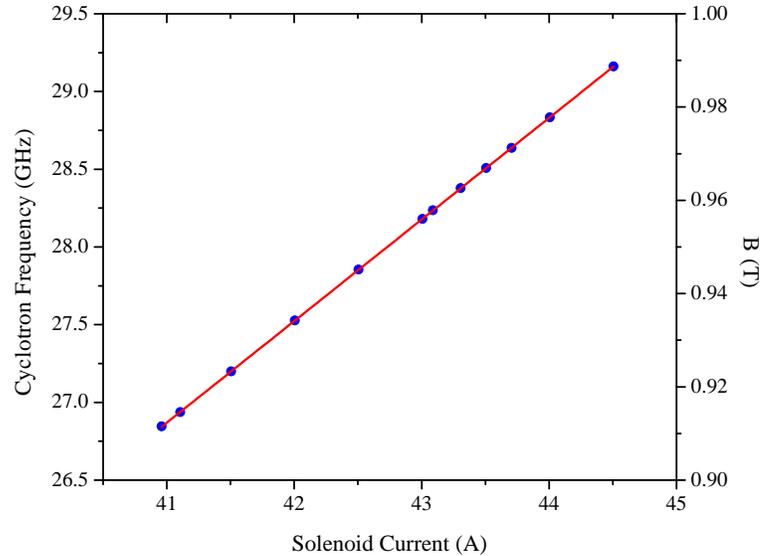
Figure 4.16: Calibration between solenoid current and on-axis magnetic field as determined using the plasma mode diagnostic to measure the cyclotron frequency.

Figure 4.17 shows data from a similar measurement, as the current in the mirror coils is varied. In this case the solenoid is turned on and its current is set to its normal operating value, giving a background field of 1 Tesla. A linear trend is observed between the electron cyclotron frequency and the current in the mirror coils. Using these calibrations we are able to determine the electron cyclotron resonance frequency at the center of the ALPHA apparatus with an uncertainty of 10 MHz, corresponding to a relative magnetic field measurement of $\Delta B/B = 3.4 \times 10^{-4}$ [109]. The thermal motion of electrons parallel to the static magnetic field leads to broadening of the cyclotron lineshape (Fig. 4.15). This broadening is the primary parameter limiting the uncertainty in magnetic field determination.

4.5 Microwave Field Calibration, Power Meter

Knowledge of the microwave fields in the region where spin flip transitions are induced plays a crucial role in understanding the spectroscopy experiments reported in the next chapter. Methods similar to those described in Sec. 4.4.1 can be used to obtain estimates
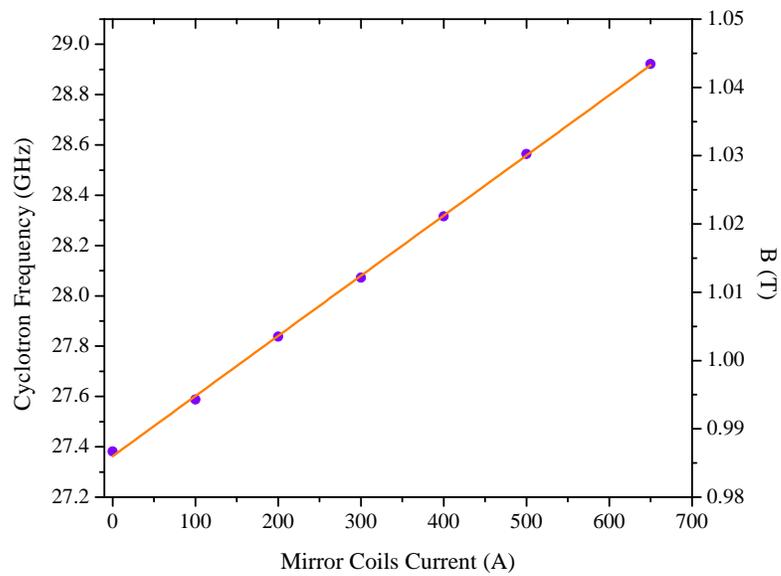
Figure 4.17: Calibration between Mirror coil currents and on axis magnetic field as determined using the plasma mode diagnostic to determine the cyclotron frequency. The solenoid field was held fixed at a value just below 1 T for these measurements. The mirror coils contribute about 6% to the total field at the centre of the trap, under normal operating conditions.

of microwave electric field. We start by loading an electron plasma (approximately 12×10^6 electrons with a typical density of 8×10^{10} m^{-3}) in the center of the apparatus. Then a short pulse (80 ns width) of microwave radiation at the cyclotron resonance frequency is injected, causing the plasma temperature to increase by an amount ΔT (determined by monitoring the quadrupole mode frequency of the trapped plasma). A single particle theory has been developed by Profs. Hayden, Robicheaux, and Hardy to relate this temperate change to the microwave fields (see App. B for further details). In the limit where the duration τ of a rectangular microwave pulse is short compared to damping/collisional times in the plasma [109, 110],

$$\Delta T = \frac{q^2 \tau^2 E_+^2}{12 m k_B},\tag{4.8}$$

where E_+ is the component of the microwave electric field co-rotating with the cyclotron motion, and q is the elementary charge. Thus, by measuring ΔT one can infer the amplitude of the microwave electric field. Estimates of the microwave magnetic field amplitude B$_{\mu\text{W}}$ and time-average power P propagating down the bore of the apparatus can then be made in various limits. For example, an analogy to plane wave propagation in free space would imply

$$P = \frac{AE^2}{2Z_0}.\tag{4.9}$$

where $A \approx 16$ cm^2 is the cross sectional area of the electrode stack and $Z_0 = 377$ Ω is the impedance of free space. Clearly the interpretation of the electric field E in this expression (in relation to E_+) depends on factors such as the standing wave ratio. Nevertheless, such estimates provide useful insight. In the same spirit (i.e. using the free space plane wave approximation) the microwave magnetic field is related to the microwave electric field such that $H_{\mu W} = E_{\mu W}/Z_0$ or $B_{\mu W} = \mu_0 E_{\mu W}/Z_0 = E_{\mu W}/c$. Table 4.1 lists estimates of microwave field amplitudes at the center of the ALPHA trap under typical operating conditions during microwave spectroscopy experiments. These results are inferred from a microwave heating experiment at the cyclotron resonance frequency involving 80 ns rectangular pulses.

As indicated above, we relate the microwave magnetic field to the microwave electric field using the impedance of free space. It is worth examining the validity of this approximation. Assuming that the ALPHA electrode stack can be modeled as a circular waveguide with a

Frequency (GHz)	$E_{\mu W}$ (V/m)	P_{Source} (mW)	$B_{\mu W}$ Gauss
28.276	110	700	3.7×10^{-3}

Table 4.1: Estimates of microwave field amplitudes at the centre of the trap under typical operating conditions for spectroscopy experiments. $E_{\mu W}$ refers to the amplitude of the microwave electric field.

radius of 22 mm, the cut off frequency of the dominant (and lowest order propagating) TE_{11} mode is approximately 4 GHz. For typical experimental conditions in ALPHA, the (free space) wavelength of the microwaves used for these measurements is approximately 1 cm. The wave impedance of TE modes in an ideal circular waveguide is given by [111]:

$$Z_{nm} = \frac{k_0}{\beta_{nm}} Z_0, \tag{4.10}$$

while the wave impedance of TM modes is:

$$(Z_{nm})_{TM} = \frac{\beta_{nm}}{k_0} Z_0, \tag{4.11}$$

where k_0 is the free space wave number, β_{nm} is the propagation factor:

$$\beta_{nm} = \left[k_0^2 - k_{c,nm}^2\right]^{1/2}, \tag{4.12}$$

$k_{c,nm}$ is the cut off wave number, which for a TE_{nm} mode is:

$$k_{c,nm} = \frac{p_{nm}'}{a}, \tag{4.13}$$

where p_{nm}' are roots of the first derivative of the Bessel function, $J_n'(x)$. The guide wavelength is thus:

$$\lambda_{g,nm} = \frac{2\pi}{\beta_{nm}} = \frac{\lambda_0}{\sqrt{1 - \left(\frac{f_{c,nm}}{f}\right)^2}}, \tag{4.14}$$

where λ_0 is the free space wavelength and $f_{c,nm}$ is the cut off frequency for the particular mode. For the TE_{11} mode: $n = 1, m = 1$, and $p_{11}' = 1.841$. The wave impedance for this mode at frequencies of order 30 GHz as employed in the spin flip experiments discussed in the next chapter is equal to Z_0 to within 1%. The wave impedances for higher order modes exhibit larger deviations from Z_0, but even up to $f_c \sim$19 GHz (roughly 2/3 of the operating

	TE$_{11}$	TE$_{21}$	TE$_{12}$	TE$_{22}$	TE$_{13}$	TE$_{23}$	TE$_{14}$	TE$_{24}$
Cut off Frequency (GHz)	4.0	6.6	11.6	14.6	18.5	21.8	25.4	28.3
Guide Wavelength (m)	0.010	0.010	0.011	0.011	0.013	0.015	0.019	0.037
Z_{nm}/Z_0	1.0	1.0	1.1	1.1	1.3	1.5	1.9	3.7

Table 4.2: Cut off frequency, guide wavelength, and wave impedance for a selected number of TE modes in a 44 mm diameter circular waveguide at a frequency $f = 29.696$ GHz. The free space wavelength is approximately 1 cm.

frequency) corrections to the free space plane wave approximation are only at the 30% level.

This is illustrated in Tab. 4.2 which shows the cut off frequency, guide wavelength, and the ratio of the wave impedance to the free space impedance for a few TE modes in an ideal circular waveguide with the same dimensions as the ALPHA mixing trap, at a frequency $f = 29.696$ GHz. Additionally, higher order modes (which are generated by mode conversion at geometric discontinuities) will tend to be attenuated more quickly than low order modes.

These arguments suggest that the uncertainty in determining B$_{\mu\text{W}}$ from E_+ is ultimately linked to the unknown mixture of standing and traveling waves present in the electrode stack, rather than the choice of Z_0 as the wave impedance.

Chapter 5

Antihydrogen Positron Spin-Flip Transitions

This chapter is dedicated to microwave spectroscopy experiments. I start with a short discussion of potential approaches for microwave spectroscopy of atomic antihydrogen. Next, I describe our successful attempts to induce PSR (Positron Spin Resonance) transitions between hyperfine levels of ground state atomic antihydrogen. Data from these experiments set a crude limit on the hyperfine splitting of the antihydrogen atom. Finally, I discuss systematic errors associated with our measurements and their potential interpretations. Experiments that build on those described here are anticipated, and are expected to yield constraints on the zero-field hyperfine splitting of antihydrogen at much higher level of precision, perhaps one day approaching that of the hydrogen atom.

5.1   Approaches to Microwave Spectroscopy

To date, there have been two major proposals for pursuing microwave spectroscopy of antihydrogen. In one scheme a beam of antihydrogen atoms is produced [112]. These atoms would then pass through a microwave cavity in which interactions with an RF field take place on the fly. In the other approach antihydrogen atoms are first confined in a magnetic potential well. Interactions then take place as atoms pass through regions of space in which transitions come into resonance with applied RF fields [1]. The ASACUSA collaboration is pursuing the former approach while the ALPHA collaboration has taken the latter. Here, I

briefly discuss the general strategy for both approaches.

5.1.1 Spectroscopy of Trapped Antihydrogen Atoms

The advent of the hydrogen maser brought about a remarkable leap in the precision to which intervals between ground-state hyperfine levels of the hydrogen atom could be measured. In large part these advances are due to the fact that atoms are stored inside the microwave cavity for long periods of time, leading to long coherent atom-field interaction times. Simultaneously these atoms are stored in a very homogeneous magnetic field, close to a turning point where the atomic transition is first-order independent of field. The situation for antihydrogen is quite different. A physical (matter) enclosure cannot be used to confine antimatter atoms because they would simply annihilate at first contact with the walls. A magnetic potential well can be used instead, but unless the well is very shallow (necessitating the use of very cold atoms), atom-field interactions are limited by transit time effects.

The Breit-Rabi diagram for the ground state of the (anti)hydrogen atom is shown in Fig. 5.1. The four eigenstates are labeled $|a\rangle$, $|b\rangle$, $|c\rangle$ and $|d\rangle$ in order of increasing energy in low magnetic fields. The 'low-field seeking' states $|c\rangle$ and $|d\rangle$ have magnetic moments $\boldsymbol{\mu}$ antiparallel to $\boldsymbol{B}$ and can be magnetically trapped. The 'high-field seeking' states $|a\rangle$ and $|b\rangle$ have $\boldsymbol{\mu}$ parallel to $\boldsymbol{B}$ and cannot be trapped. The relative energies of these states are given by:

$$E_a = -\frac{a}{4} - \frac{a}{2}\sqrt{1 + \left[\frac{\hbar(\gamma_e + \gamma_p)B}{a}\right]^2} \tag{5.1}$$

$$E_b = \frac{a}{4} - \frac{\hbar}{2}(\gamma_e - \gamma_p)B \tag{5.2}$$

$$E_c = -\frac{a}{4} + \frac{a}{2}\sqrt{1 + \left[\frac{\hbar(\gamma_e + \gamma_p)B}{a}\right]^2} \tag{5.3}$$

$$E_d = \frac{a}{4} + \frac{\hbar}{2}(\gamma_e - \gamma_p)B \tag{5.4}$$

where a is the zero-field hyperfine splitting constant and γ_e and γ_p are the gyromagnetic ratios of the electron and proton. Similar expressions hold for antihydrogen, involving the gyromagnetic ratios of the positron (γ_{e^+}) and antiproton ($\gamma_{\bar{p}}$).
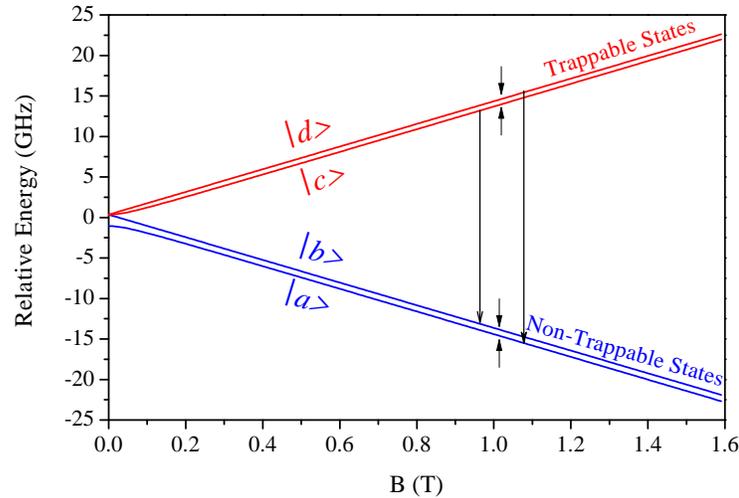
Figure 5.1: Relative hyperfine energy levels (in frequency units) and allowed transitions for ground state (anti)hydrogen when a time-varying magnetic field is applied perpendicular to the static field.

The underlying idea for microwave spectroscopy of magnetically trapped antihydrogen atoms is to induce transitions from trapped (low-field seeking) to non-trapped (high-field seeking) states by applying time-varying magnetic fields. The consequence of such transitions would be the ejection of atoms from the trap, followed by their annihilation on nearby electrodes; products of these annihilation events would then be recorded by the ALPHA annihilation detector. Knowing the frequency at which transitions are induced would then enable one to extract fundamental parameters such as the zero-field hyperfine splitting constant of the antihydrogen atom and the gyromagnetic ratio of the antiproton. The biggest challenge presented by this scheme is the fact that the inhomogeneous magnetic fields required for trapping limit the time for which a (moving) atom can interact with an RF field, making high-precision spectroscopy measurements difficult.

Figure 5.1 also shows allowed transitions between hyperfine levels when a time-varying magnetic field is applied perpendicular to the static magnetic field.[1] In the high-field limit

---

[1]The $|c\rangle \rightarrow |a\rangle$ transition is also allowed, but requires application of a time-varying magnetic field parallel

(B≫ 0.05 T), two of these transitions ($|c\rangle \to |b\rangle$ and $|d\rangle \to |a\rangle$) amount to a positron spin flip. We refer to these as PSR (Positron Spin Resonance) transitions (cf. ESR transitions for the hydrogen atom). In the same limit, the other two transitions ($|d\rangle \to |c\rangle$ and $|b\rangle \to |a\rangle$) amount to an antiproton spin flip, and are thus referred to here as NMR transitions.

5.1.2   Spectroscopy of Antihydrogen in a Beam

The ASACUSA collaboration plans to use a method similar to classical atomic beam (Stern-Gerlach type) experiments to measure the antihydrogen ground-state splitting [113].
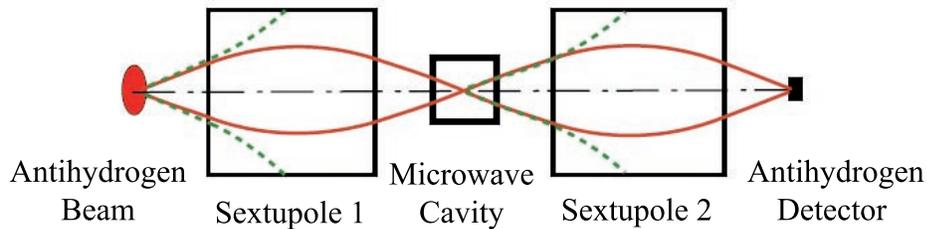
Figure 5.2: Antihydrogen beam spectroscopy experiment envisioned by the ASACUSA collaboration. The setup consists of two sextupole magnets and a microwave cavity. The solid lines represent trajectories of antihydrogen atoms in low-field seeking states while the dashed lines represent trajectories of atoms in high-field seeking states (Adapted from [12]).

Figure 5.2 shows the antihydrogen beam spectroscopy experiment envisioned by the ASACUSA collaboration. The antihydrogen beam is generated as atoms escape from a Paul or a cusp trap [112, 113, 114]. When these atoms pass through the first sextupole magnet, those which are in high-field seeking states will be deflected towards regions of higher magnetic field and will thus be defocused and hit the walls. In contrast, atoms in low-field seeking states will be attracted towards the sextupole axis and will thus be focused into a radiofrequency cavity. Without any interaction at the cavity, the low-field seeking atoms will traverse the second sextupole magnet and reach the detector. However, by tuning the RF cavity to induce transitions between low-field seeking and high-field seeking states, atoms entering the second sextupole magnet will be defocused and will fail to reach the detector. In January 2014 the ASACUSA collaboration reported detection of 80 antihydrogen atoms at a distance 2.7 m downstream of their production region (in a setup similar but not identical to that shown

---

to the static magnetic field.

in Fig. 5.2), where the perturbing influence of the magnetic fields used to produce the antihydrogen atoms is small [112]. This is an important step towards an eventual microwave spectroscopy experiment on antihydrogen atoms in a beam.

A potential advantage of the experiment described above is that it should work with a relatively high-temperature $(50-100 \text{ K})$ source of atoms, in contrast to the very cold $(< 1 \text{ K})$ atoms that are required by ALPHA. However, the challenge faced by ASACUSA to date seems to have been that essentially all of the atoms produced in their source are in highly excited Rydberg states. This thesis is focused on the magnetic confinement approach to antihydrogen spectroscopy and so the beam approach will not be described further. Interested readers may refer to the publications of the ASACUSA collaboration for further information.

5.2 Features and Challenges: ALPHA Microwave Spectroscopy Apparatus

The ALPHA apparatus has a number of distinct features that are conducive to a successful spectroscopy experiment. Silicon vertex detectors capable of detecting the byproducts of matter-antimatter annihilation events are among the most powerful and efficient diagnostic tools in antimatter experiments. In principle they should enable one to conduct spectroscopic measurements even if only very few atoms are confined and interrogated. Also, the very long antihydrogen confinement times that have been demonstrated in the ALPHA apparatus [74] are significant for spectroscopy. First, atoms in excited states have plenty of time to cascade down to their ground state prior to interrogation. More importantly, long confinement times provide ample opportunity for atoms to interact with microwave fields, even if the volume over which those fields are resonant is small because of the intrinsic inhomogeneity of the trapping fields.

At the same time, the ALPHA trapping apparatus presents significant challenges for a successful microwave spectroscopy experiment. The depth of the magnetic well for $|d\rangle$-state antihydrogen atoms is about 0.5 K, in temperature units. This implies that the mean velocity of atoms passing through the centre of the trap is approximately 70 m/s, and so the time scale for the atoms to traverse the trap minimum (a distance of order 1 cm in the radial
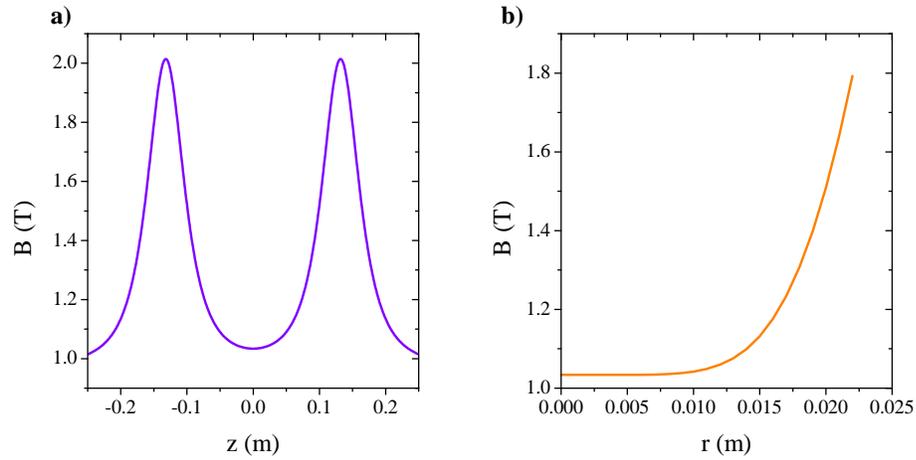
Figure 5.3: Magnetic field strength (a) along the axis and (b) along the radial direction (at z=0 mm) of the ALPHA apparatus. The field strength at the position of the mirror coils is about 1 T higher than the field at center of the trap.
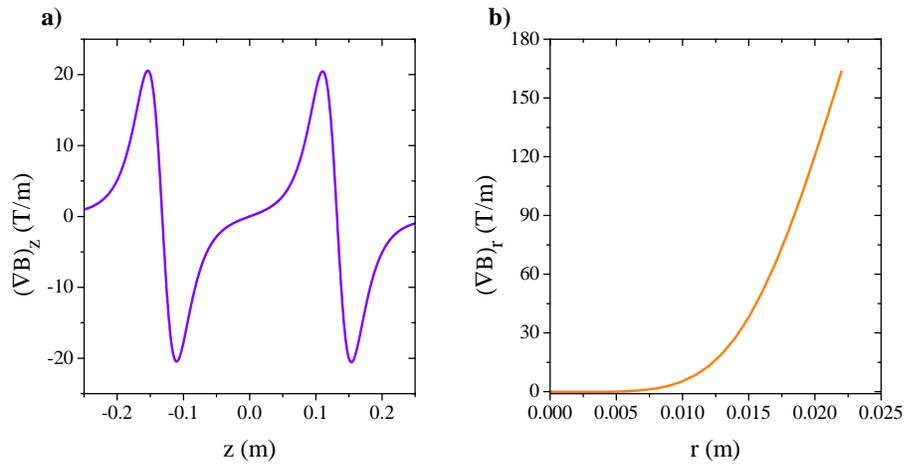
Figure 5.4: Magnetic field gradient (a) along the axis of the ALPHA trap and (b) and along the radial direction at z=0 mm.

direction[2]) and experience resonant interaction with time-varying fields ($B_{\mu W}$) is of order 10^{-4} s. In the reference frame of an individual atom passing through resonance the spectral width of the time-varying field ($B_{\mu W}$) is thus at least of order 10 kHz. This transit-time broadening effect is almost certainly going to be the limiting factor for any first-generation *precision* antihydrogen microwave spectroscopy experiment. Laser cooling atoms to ultra cold temperatures and simultaneously lowering the depth of the magnetic potential well may help to overcome this issue.

Given the existing ALPHA antihydrogen trapping apparatus, the limiting factor in any microwave spectroscopy experiment is the homogeneity of the magnetic field used to trap atoms. Figures 5.3.a and 5.3.b show the magnetic field strength along the axis and along the radial direction (at z=0 mm) of the ALPHA trap; Figs. 5.4.a and 5.4.b show the magnetic field gradient along the axial and radial directions, respectively. Notice that the field gradient in both directions is strong. Deeper well depths are desirable for antihydrogen confinement, but unless the field profile near the minimum is somehow flattened, increasing the trapping fields also increases the field inhomogeneity. The uniformity of the magnetic field in the vicinity of the minimum magnetic field (near $r = 0$ and $z = 0$) is thus the dominant limiting factor for microwave spectroscopy experiments in the ALPHA apparatus.

5.3   Measurement Scenarios

Ground state antihydrogen atoms trapped in the ALPHA apparatus must be in one of the trappable states, $|c\rangle$ or $|d\rangle$. Given this fact, several different approaches to measuring energy intervals are possible. One of the simplest approaches is discussed here. Another approach which will be pursued in the future is discussed in Chapter 7.

---

[2]The field homogeneity is substantially worse in the axial direction, and so these order-of-magnitude estimates are overly optimistic. The field homogeneity of the ALPHA-II apparatus is expected to be significantly improved: see Chapter 7.

5.3.1 PSR Transitions

The simplest approach to measuring hyperfine intervals involves inducing PSR transitions. Low-field seeking atoms ($|c\rangle$ and $|d\rangle$) are first confined in the atom trap. Then a transverse microwave magnetic field $B_{\mu W}$ resonant with the $|c\rangle \rightarrow |b\rangle$ transition is applied to convert $|c\rangle$ state atoms into $|b\rangle$ state atoms. These $|b\rangle$ state atoms are high-field seeking and are thus ejected from the trap and annihilate when they strike the electrode walls. Once the $|c\rangle$ state atoms have been removed from the trap the process can be repeated with the microwave radiation tuned to drive the $|d\rangle \rightarrow |a\rangle$ transition. If both PSR transition frequencies (f_{bc} and f_{ad}) are measured under precisely the same conditions, their difference yields the zero-field hyperfine splitting $\Delta\nu_{HFS} = a/h$ independent of magnetic field; that is

$$(E_d - E_a) - (E_c - E_b) = a \tag{5.5}$$

$$\text{or} \quad f_{ad} - f_{bc} = \Delta\nu_{HFS} \tag{5.6}$$

where h is the Planck constant.
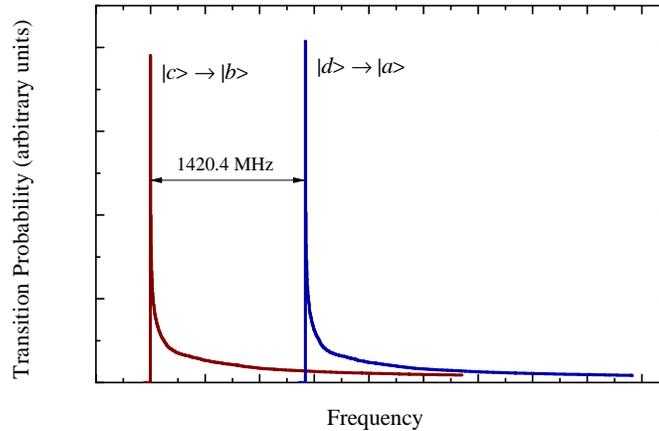
Figure 5.5: Distribution of PSR transition frequencies in the trap, as determined from a Monte Carlo simulation of trapped atom dynamics. In this example, atoms were loaded from the low energy tail of a 50 K thermal distribution. The central magnetic field is 1 Tesla (Reproduced from [11]).

Figure 5.5 shows the distribution of PSR transition frequencies in the ALPHA trap, as determined from a Monte Carlo simulation in which the trap is loaded with equal numbers of $|c\rangle$ and $|d\rangle$ state atoms. The abrupt low-frequency edges in these distributions are associated with the minimum in the static field near the center of the trap. The long high-frequency tails reflect the fact that everywhere else in the trap, the static field is inhomogeneous. Importantly, these low frequency edges represent a sharp spectroscopic feature that can be probed by varying the frequency of the microwave fields used to induce transitions. With many cold trapped atoms available for study (and a carefully designed and controlled magnetic field profile near the trap minimum), one might expect measurements of PSR transition frequencies to eventually yield the hyperfine splitting constant a to about 1 part in 10^7.

What makes this measurement intriguing for ALPHA is the fact that the frequency of the PSR transitions in a 1 T magnetic field (the minimum field in the trap under normal operating conditions) lie in the range 28-30 GHz. The wavelength at these frequencies is of order 1 cm which is smaller than the $\sim$ 4 cm diameter of the trap electrodes. Microwave radiation at these frequencies propagates down the stack of electrodes. An experimental program focused on PSR transitions was thus identified as the top priority for the 2011 run.

5.4 Experimental Sequence

An oscillating magnetic field $B_{\mu W}$ applied perpendicular to the static trapping field can drive positron spin-flip transitions between the trappable and untrappable states, *i.e.* $|c\rangle \to |b\rangle$ and $|d\rangle \to |a\rangle$. This leads to the escape and subsequent annihilation of the trapped atom as it strikes the trap walls. An experimental demonstration of this process, and of the *resonant* nature of this interaction is described below.

Prior to each attempt at inducing a spin-flip transition, the minimum magnetic field in the trapping apparatus was tuned to one of two different predetermined values using the plasma mode diagnostic discussed in Sec. 4.4. These two settings differed by 36 Gauss, or the equivalent of 100 MHz in terms of the electron cyclotron frequency. Experiments were further subdivided into 'on-resonance' and 'off-resonance' attempts. For the former, applied

microwave fields were tuned so as to be nominally resonant with antihydrogen atoms situated at the minimum magnetic field. For the latter, applied microwave fields were detuned from the nominal resonance condition by 100 MHz, toward lower frequencies. Technically there are still locations in the trap where $|c\rangle$ state atoms can pass through resonant conditions during 'off-resonance' attempts.

An experimental cycle, regardless of being on-resonance or off-resonance, begins with the antihydrogen trapping sequence described in Sec. 3.4. Next the trapped atoms are simply held for 60 s. The purpose of this hold period is to let the magnetic fields stabilize. Next, microwave radiation is injected into the apparatus for 180 s. This time period is referred to as the 'microwave window'. Since trapped antihydrogen atoms are either in the $|c\rangle$ state or the $|d\rangle$ state, both of the PSR transitions need to be driven to ensure a spin-flip.

The 180 s microwave window is divided into six 30 s intervals. Each of these intervals is subdivided into two 15 s periods. During the first period, the microwave frequency is swept linearly over a 15 MHz band near the target $|c\rangle \rightarrow |b\rangle$ transition frequency.[3] The sweep extends from 5 MHz below the target frequency to 10 MHz above. Likewise, the second period involves a 15 MHz frequency sweep near the target $|d\rangle \rightarrow |a\rangle$ transition frequency. It starts precisely 1420.405 MHz higher than the lower sweep. Figure 5.6 illustrates the full microwave frequency sweep procedure. Note that in practice stepped frequency sweeps were employed with a step size of 100 KHz.

After the 180 s microwave irradiation period the microwaves are turned off. Then after a 1 s hold period the trapping field currents are rapidly ramped down (with a time constant of 9 ms [73]), just as they are in our standard antihydrogen trapping experiments. If trapped antihydrogen atoms survive the microwave radiation, they would be released and the subsequent annihilation events are detected by the silicon vertex detector. Ideally, one would not expect to see annihilation events after an 'on-resonance' experiment, because trapped atoms should have already undergone PSR transitions and should have been kicked

---

[3]Nominally this target frequency corresponds to the frequency at which $|c\rangle$ state atom sitting at the minimum trap field would be in resonance with the $|b\rangle$ state.
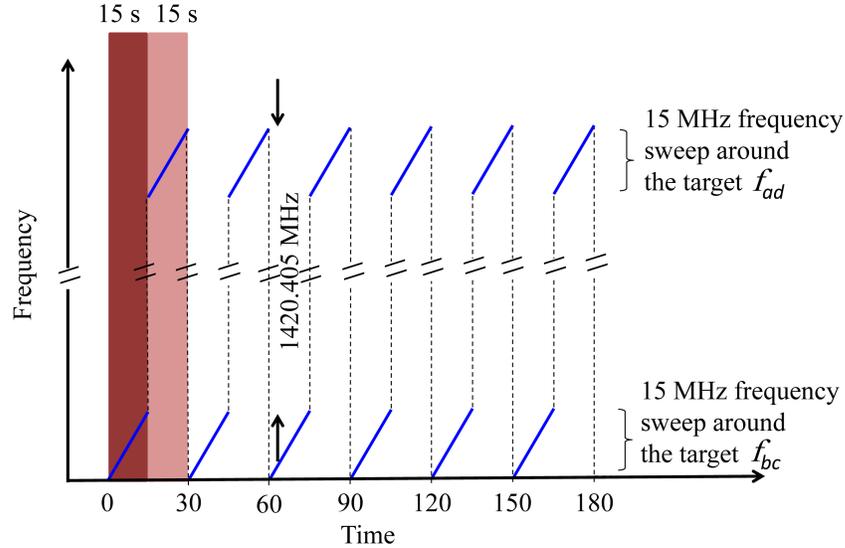
Figure 5.6: The 180 s microwave frequency sweep pattern employed in our spin-flip demonstration experiments. The two coloured bars highlight the first two 15 s periods during which microwaves are swept near the two target transition frequencies.

out of the trap during the microwave irradiation window. Similarly, one expects to see annihilation events after an 'off-resonance' experiment because trapped atoms have not undergone PSR transitions. Many factors cause departures from this idealized behavior. Most obviously we do not trap atoms with 100% efficiency; many trapping attempts are made with no atoms in the trap. Even if an atom is in the trap, its trajectory and the microwave fields that are established in the trap do not guarantee that a spin flip will occur. Finally, as noted above, the high frequency tail of the $|c\rangle \rightarrow |b\rangle$ transition extends well above the onset of the $|d\rangle \rightarrow |a\rangle$ transition. There is thus a finite probability that PSR transitions will be induced during 'off-resonance' attempts.

5.4.1 Experimental Variations

The experiments described above were conducted over a one-month period, during October and November 2011. Microwave radiation was injected at three different power levels:

- Full Power: 20.5 $\mathrm{dB_m}$ at the microwave source and ~ 700 mW referenced to the

window where the waveguide enters the UHV volume of the apparatus.

- 1/4 Power: 14.48 dB$_\mathrm{m}$ at the microwave source and $\sim$ 175 mW referenced to the window where the waveguide enters the UHV volume of the apparatus.

- 1/16 Power: 8.5 dB$_\mathrm{m}$ at the microwave source and $\sim$ 44 mW referenced to the window where the waveguide enters the UHV volume of the apparatus.

At each power level data were acquired under at least three distinct and interspersed experimental conditions: an on-resonance condition, an off-resonance condition, and a no-microwaves (zero-power) condition. The on-resonance and off-resonance conditions were identical except for the minimum magnetic field in the trap or the target frequencies at which microwaves were injected. Background (zero-power) acquisition sequences were run for each magnetic field setting; i.e., if the difference between on-resonance and off-resonance condition was a change in magnetic field, then zero-power data were acquired at both fields.

Two microwave frequency sweeps were employed. We label these Sweep$_A$ and Sweep$_B$. They differ by 100 MHz: i.e., Sweep$_B$ =Sweep$_A$+ 100 MHz. The precise specifications for these sweeps are given in Tab. 5.1.

Sweep	Lower Range (GHz)	Upper Phase Range (GHz)	f_{bc} (GHz)	f_{ad} (GHz)
Sweep$_A$	[28.270600, 28.285600]	[29.691000, 29.706000]	28.275500	29.695900
Sweep$_B$	[28.370600, 28.385600]	[29.791000, 29.806000]	28.375500	29.795900

Table 5.1: Ranges for the two microwave frequency sweeps. The last two columns give the target frequencies about which the scans are positioned.

Similarly two magnetic field configurations were employed. We label these B$_A$ and B$_B$. They differ such that the electron cyclotron frequencies at the trap centre are separated by 100 MHz.

- B$_A$: The minimum on-axis magnetic field with the solenoid and mirror coils on is B$_\mathrm{min}^\mathrm{axis}$ = 1.0322 $\pm$ 0.0003 T. The electron cyclotron frequency at this field corresponds to f_A^{cyc} = 28.895 $\pm$ 0.010 GHz.

- B_B: The minimum on-axis magnetic field with the solenoid and mirror coils on is $B_{min}^{axis} = 1.0358 \pm 0.0003$ T. The electron cyclotron frequency at this field corresponds to $f_B^{cyc} = 28.995 \pm 0.010$ GHz. In order to reach this magnetic field from B_A, the mirror field currents are increased.

Figures 5.7, 5.8, and 5.9 summarize the different combinations of sweep and field configurations that were employed. The manner in which the lineshapes are calculated is described in Chapter 6.
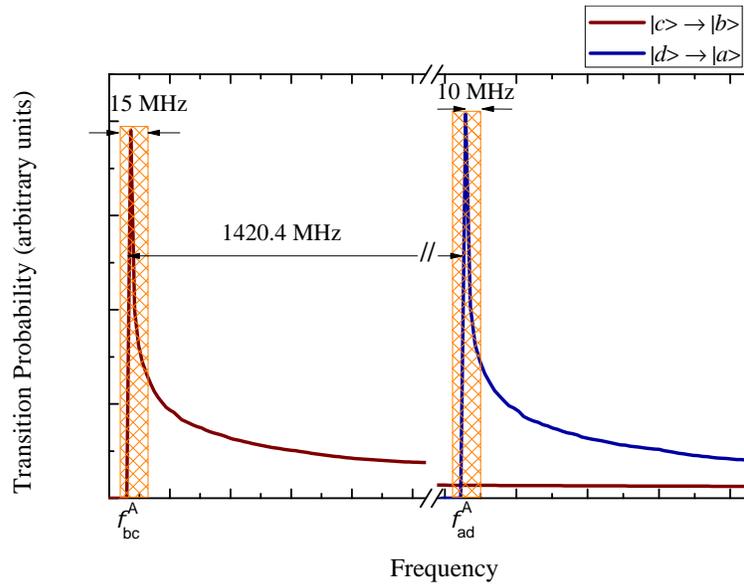
Figure 5.7: Nominal 'on-resonance' conditions when the magnetic field is set to configuration B_A and frequency $Sweep_A$ is employed. The orange bars represent the 15 MHz bands over which the microwave frequency is scanned.

Table 5.2 summarizes all of the conditions under which data were acquired. The amplitudes of microwave electric fields are inferred from ECR heating measurements performed
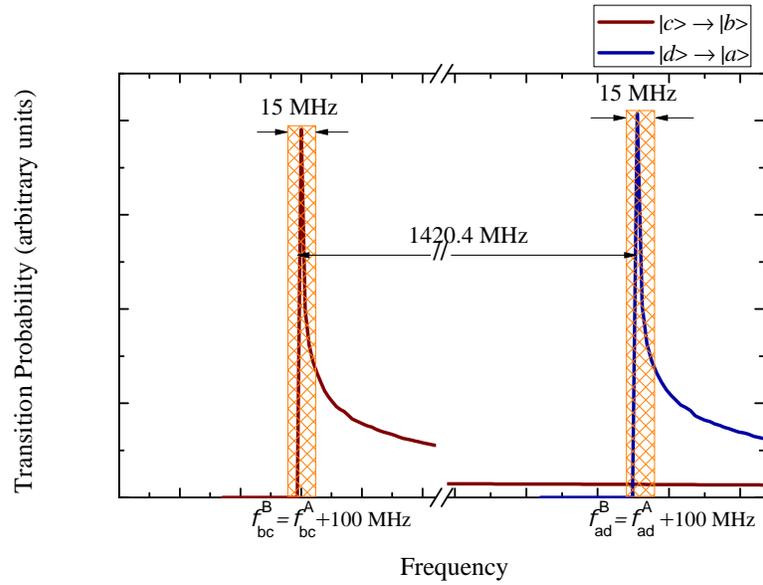
Figure 5.8: Nominal 'on-resonance' conditions when the magnetic field is set to configuration B_B and frequency Sweep$_B$ is employed. The orange bars represent the 15 MHz bands over which the microwave frequency is swept.
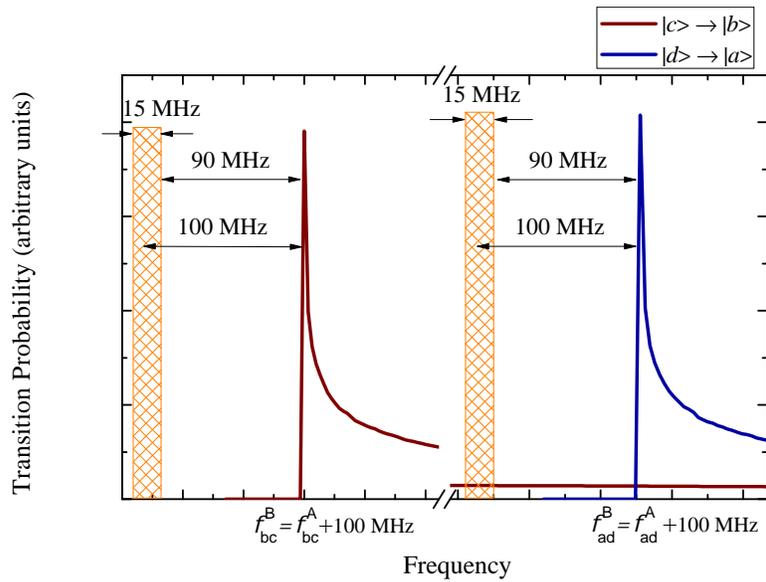
Figure 5.9: Nominal 'off-resonance' conditions when the magnetic field is set to configuration B_B and frequency Sweep$_A$ is employed.

under the same experimental conditions and at the target frequency f_{bc} (ECR heating measurements are described in Sec. 4.5). The amplitudes of microwave magnetic fields are calculated from the inferred microwave electric fields based on the results expected for a plane wave propagating in free space.

Two independent sets of data are obtained when the experiments described above are performed. Antihydrogen annihilation events are detected (a) during the microwave irradiation period and (b) when the trap fields are ramped down at the end of each measurement cycle. We refer to these as 'Appearance' and 'Disappearance' mode data sets, respectively.

Relative μW Power	Magnetic Field	Frequency Sweep	*In-Situ* μW E-Field (V/m)	*In-Situ* μW B-Field (mGauss)	Type
Full	B_A	Sweep$_A$	110	3.7	On-Resonance
Full	B_B	Sweep$_A$	110	3.7	Off-Resonance
Full	B_B	Sweep$_B$	130	4.3	On-Resonance
1/4	B_A	Sweep$_A$	55	1.8	On-Resonance
1/4	B_B	Sweep$_A$	55	1.8	Off-Resonance
1/16	B_A	Sweep$_A$	27.5	0.9	On-Resonance
1/16	B_B	Sweep$_A$	27.5	0.9	Off-Resonance
Zero	B_A	OFF	0	0	No Microwaves
Zero	B_B	OFF	0	0	No Microwaves

Table 5.2: Conditions under which data were acquired. The fourth and fifth columns give estimates of the microwave (μW) field amplitudes at the target frequency f_{bc}. These are inferred from ECR heating measurements.

5.4.2 Results: Disappearance Mode

Table 5.3 summarizes the number of times (attempts) each variation of the experiment was performed, along with the number of antihydrogen atoms that were detected in a 30 ms window when the trap fields were ramped down. The algorithm employed to identify annihilation events was precisely the same as that employed in our initial demonstration of antihydrogen trapping. The rate at which cosmic ray events are misinterpreted as annihilation events by this selection scheme is $(4.7 \pm 0.2) \times 10^{-2}$ s^{-1} and so the probability of misidentifying an annihilation event is $(14.1 \pm 0.6) \times 10^{-4}$ per attempt.

Relative μW Power	Magnetic Field	Frequency Sweep	Number of Attempts	$\bar{\text{H}}$ Events after Trap Release	Type
Full Power	B_A	Sweep$_A$	36	0	On-Resonance
Full Power	B_B	Sweep$_A$	61	13	Off-Resonance
Full Power	B_B	Sweep$_B$	19	0	On-Resonance
1/4 Power	B_A	Sweep$_A$	15	1	On-Resonance
1/4 Power	B_B	Sweep$_A$	15	1	Off-Resonance
1/16 Power	B_A	Sweep$_A$	33	1	On-Resonance
1/16 Power	B_B	Sweep$_A$	34	9	Off-Resonance
Zero Power	B_A	OFF	52	17	No Microwaves
Zero Power	B_B	OFF	48	23	No Microwaves

Table 5.3: Disappearance mode data set, classified by relative microwave (μW) power.

Table 5.4 summarizes the same data collected into three major classes: on-resonance, off-resonance, and no-microwaves attempts. A clear decrease in survival rate is observed when on-resonance data are compared to off-resonance (or no-microwaves) data. This is precisely the effect one would expect to observe if spin-flip transitions are being induced. By comparing the rate at which antihydrogen atoms are detected during on-resonance attempts with the corresponding rate for off-resonance attempts, one can evaluate the probability (p-value) that the observed number of outcomes could have occurred by chance [116]; Here $p = 1.0 \times 10^{-5}$.

	Number of Attempts	$\bar{\text{H}}$ Events After Trap Release	Rate (events per attempt)
On-Resonance	103	2	0.02 ± 0.01
Off-Resonance	110	23	0.21 ± 0.04
No-Microwaves	100	40	0.40 ± 0.06

Table 5.4: Aggregate disappearance mode data set [1].

It is clear from Tab. 5.4 that the number of atoms surviving after the no-microwaves attempts is greater than the case in which microwaves are injected but are off-resonance. The p-value for this being a chance occurrence is 6×10^{-3}. This observation can be explained by far off-resonant interactions with $|c\rangle$ state atoms. Notice that the upper half of the sweep

shown in Fig. 5.9 overlaps with the tail of the $|c\rangle \rightarrow |b\rangle$ transition. This scenario requires sufficient microwave power to induce spin flip transitions as atoms pass through resonant conditions at positions in the trap that are distant from the trap centre. This scenario is discussed further in Chapter 6. The only other plausible scenario would be the annihilation of antihydrogen atoms on residual gas atoms. Microwave radiation does increase the trap electrode temperature, which could in turn lead to the desorption of matter atoms [1]. This scenario is discussed further in Sec. 5.5.1.

5.4.3 Results: Appearance Mode

The observation interval for data acquired in appearance mode is the full 180 s microwave window. The rate at which cosmic rays are misinterpreted as annihilation events using our conventional algorithm is 0.047 ± 0.002 s^{-1}, and so over 180 s more than 8 events will be registered on average. This is significantly greater than the typical number of antihydrogen atoms trapped per experimental cycle (typically less than 0.5 trapped atoms per attempt).

To overcome this challenge we introduced an alternative set of annihilation event acceptance criteria, based on a bagged decision tree classifier using the random forest approach [117, 118, 119]. These criteria were established and fixed using training without microwaves (They are further discussed in App. A). In addition, we tightened up the cut on the axial (z) event positions used in our alternative event acceptance criteria, from $|z| < 10$ cm to $|z| < 6$ cm. The motivation for this is that the cosmic ray background has a uniform axial distribution over the entire length of the ALPHA trap while spin-flipped antihydrogen atoms are expected to come from a small region near the center of the trap, and therefore, they tend to annihilate near the center of the trap. The net result of these modifications is approximately a factor of 30 improvement in rejecting the cosmic ray background (to $(1.7 \pm 0.3) \times 10^{-3}$ s^{-1}), while retaining 75% of the signal.

Figure 5.10 shows the time distribution of valid events recorded during the microwave radiation window, as identified by the alternative acceptance criteria. The bins here are 15 s to correspond with the two halves of each sweep (cf. Figs. 5.6, 5.7, 5.8, and 5.9). Thus the bin for 0 s < t < 15 s corresponds to irradiation at the $|c\rangle \rightarrow |b\rangle$ transition and the bin for 15 s < t < 30 s corresponds to irradiation at the $|d\rangle \rightarrow |a\rangle$ transition. Recall that the period -60 s < t < 0 s corresponds to a period during which magnetic fields are allowed to

stabilize. Microwave injection does not start until t = 0 s.
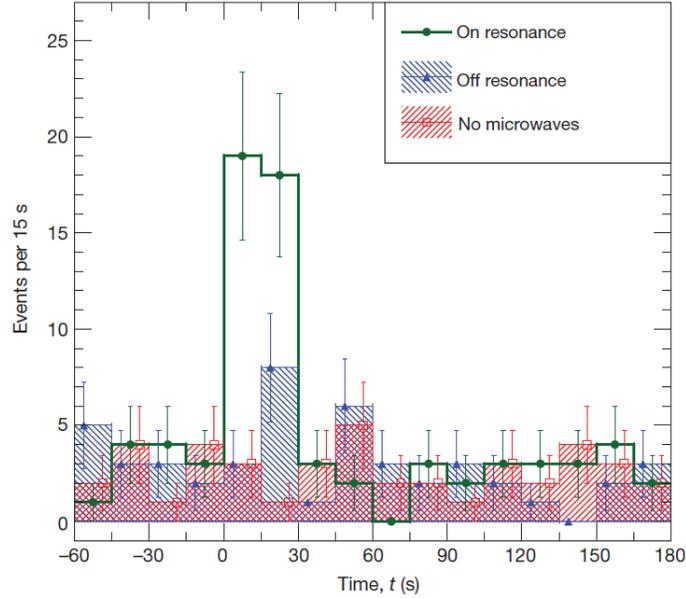
Figure 5.10: Time distribution of valid events before and during the microwave irradiation window. Aggregated data are shown for attempts in which microwaves were injected on resonance, off resonance, or not at all. The error bars are associated with counting statistics. A clear excess of counts is observed in on-resonance experiments as soon as microwaves are injected at t = 0 s (Reproduced from [1]).

In the first 30 s (the first two bins) of the microwave window (when both PSR transitions are probed sequentially), a clear excess of counts is observed for on-resonance attempts relative to off-resonance and no-microwaves attempts, with a p-value of 2.8×10^{-5}, as soon as microwave radiation is injected. Again, this is what one would expect to observe if antihydrogen atoms undergo spin flip transitions leading to ejection from the trap. Note that cosmic ray backgrounds have not been subtracted from the data shown in Fig. 5.10. The anticipated contribution of cosmic rays to the data over a 15 s time window are listed in Tab. 5.5. The apparent cosmic ray background rates are slightly different simply because the number of attempts for each sequence type are not identical.

The second off-resonance bin (blue) after the start of microwave irradiation in Fig. 5.10, exhibits a slight excess of counts relative to the case where no microwaves are injected $(p = 1.3 \times 10^{-2})$. As discussed in Sec. 5.4.2, this excess may be associated with $|c\rangle$ state atoms undergoing spin flip transitions when microwaves are injected at frequencies that are 100 MHz below the onset of the $|d\rangle \rightarrow |a\rangle$ transition (Fig. 5.9) [1]. As one would expect, this excess appears during the second half of the first 30 s sweep when the microwaves are tuned just below the $|d\rangle \rightarrow |a\rangle$ transition.

Type	Cumulative Cosmic Ray Background (per 15 s)
On-Resonance	2.7
Off-Resonance	2.8
No Microwaves	2.5

Table 5.5: Anticipated cumulative contribution of cosmic rays to the data shown in Fig. 5.10, expressed as the number of events in a 15 s period. These events are associated with 103 on-resonance, 110 off-resonance, and 100 no microwave attempts.

5.5   Systematic Effects

Taken at face value, the observations summarized in the previous section are consistent with a successful experimental demonstration of interactions between trapped antihydrogen atoms and a resonantly tuned microwave field. In this section I address systematic effects pertaining to these observations. I start with the possibility that the experimental observations arise from some effect other than a resonant interaction with a microwave field. I then discuss issues with spectroscopic relevance. That is, how well we can determine the magnetic field and other experimental parameters.

5.5.1   Microwave Effects on the Trap Vacuum

Microwave radiation does heat up the trap electrodes, causing desorption of cryo-pumped material from cold surfaces. It is thus plausible that confined antihydrogen atoms encounter released gases and annihilate. However, our numerical models of this process indicate that the axial distribution of annihilation events one expects to observe is different from that caused by spin-flip interactions at the centre of the trap. Figure. 5.11 compares the axial

distribution of annihilation events expected for these two mechanisms. Annihilation events caused by spin-flip interactions are highly localized around the trap centre (gray histogram) while those caused by a background of matter atoms are much more broadly distributed, extending out the trap ends (dotted black curve). Moreover, estimates made by Prof. Svante Jonsell indicate that the rate at which antihydrogen atoms are expected to annihilate on released gases (in the ALPHA trap) is very low compared to the rates shown in Fig. 5.10. This assumes that the released gases are mostly cold H_2 molecules at temperatures of order 10 K (the electrode temperature only changes by about a degree during microwave injection). It also assumes the cross section for $\bar{H} - H_2$ collisions is not very different from that for $\bar{H} - He$ [120]. Scaling appropriately to account for the $H_2 - He$ mass difference one obtains an annihilation rate constant of 10^{-16} m^3/s (the pressure of the released gas is assumed to be less than 10^{-11} mbar). This implies that the expected number of annihilation events associated with background gas during a 15 s period (frequency sweep interval) is more than two orders of magnitude smaller than the expected cosmic background rate. And finally and most importantly, if the desorption of cryo-pumped materials were the source of annihilation events, we should have observed a similar time distribution of annihilation events in both on-resonance and off-resonance experiments, which was not the case.

In principle this unwanted effect (desorption of matter atoms from trap walls caused by microwave heating) occurs during all of our experiments except those in which no microwaves were injected. However, the magnitude of the effect could be different for Sweep$_A$ and Sweep$_B$, because the frequencies for these sweeps differ by 100 MHz. We continuously monitor the temperature of one of the trap electrodes as well as the pressure inside the trap. Figure 5.12 shows the evolution of the electrode temperature over the 180 s time period during which microwaves are injected for Sweep$_A$ and Sweep$_B$. Similarly, Fig. 5.13 shows the vacuum pressure inside the apparatus. The data are recorded over the time period that microwaves are injected. On average the electrode temperature during Sweep$_A$ is only about 0.7% higher than that during Sweep$_B$ while the trap pressure is 2.2% higher. These measurements lead to the conclusion that the thermal response of the trap electrodes, and the unintended release of absorbed atoms during the two sweeps is very similar.
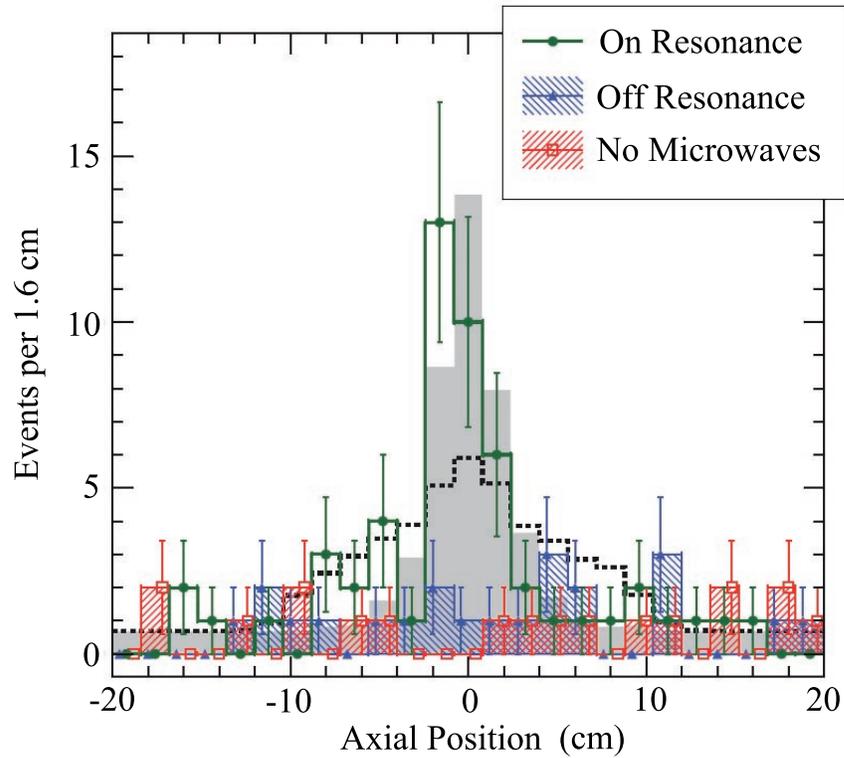
Figure 5.11: Axial distribution of annihilation events during the time period 0 s < t < 30 s. The gray histogram shows the distribution of events caused by microwave induced spin-flips, predicted by simulations. The dashed black histogram shows the distribution of events expected for a background distribution of (residual) matter atoms. The green, blue, and red histograms show the experimental data. Both simulations (gray and dashed black histograms) have been normalized to the number of events observed on-resonance. Reproduced from [1].
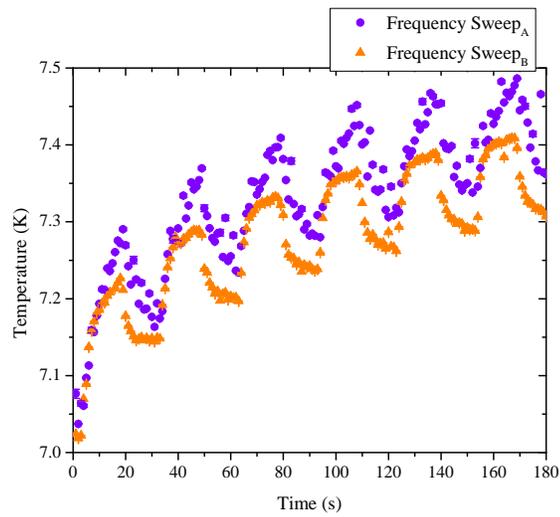
Figure 5.12: Temperature evolution of trap electrodes during Sweep$_A$ and Sweep$_B$. In both cases, during the first 15 s period (microwaves tuned in the vicinity of the $|c\rangle \rightarrow |b\rangle$ transition) the electrodes warm up. During the second 15 s (microwaves tuned in the vicinity of the $|d\rangle \rightarrow |a\rangle$ transition) the electrodes cool down somewhat (less energy is deposited). This pattern is repeated six times over the 180 s microwave injection window.
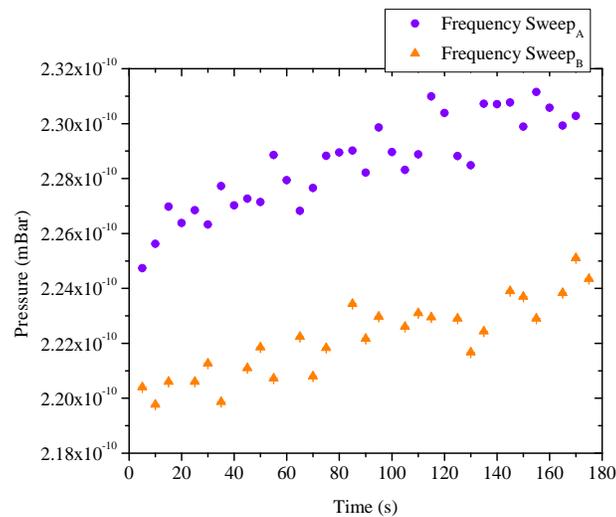
Figure 5.13: Vacuum pressure evolution inside the apparatus during Sweep$_A$ and Sweep$_B$, over the same time period as that shown in Fig. 5.12. This measurement is made in a room temperature portion of the vacuum chamber that is distant from the cryogenic trap, and so the pressure shown here is only an upper limit for that to which antihydrogen is exposed.

5.5.2 Stability and reproducibility of the Trapping Magnetic Field

Looking forward, the most critical experimental parameter for microwave spectroscopy experiments using the ALPHA apparatus is the trapping magnetic field, which sets the energy levels of trapped atoms and hence transition frequencies. With this in mind our experimental protocol involves measuring the absolute magnetic field at the centre of the trap prior to all sequences in which microwaves are injected, as well as monitoring the stability of currents in all of the superconducting magnets.

The main superconducting solenoid that provides the uniform and constant 1 T magnetic field exhibits a slow drift of approximately -0.8 Gauss per day (corresponding to a ~ 2 MHz decrease in electron cyclotron frequency). To deal with this we elected to set the current in the solenoid before each experimental sequence.

Once the currents in the mirror coils are set to their nominal value of 650 A, time is required for them to settle down (Fig. 5.14). After 60 s these currents are stable to $\sim 0.5\%$, or 1-2 MHz in electron cyclotron resonance frequency.

For technical reasons, we are not able to measure electron cyclotron frequencies with the current in the octupole magnet at its nominal value (900 A). Using numerical models we estimate the contribution of the octupole to the axial magnetic field at the centre of the trap. A perfect octupole magnet would contribute nothing to the field, but we estimate that the finite extent and distributed winding pattern of our octupole contributes approximately 14 G to the total axial field at the centre of the trap. In anticipation of this contribution we assumed that the electron cyclotron frequency would be 40 MHz higher with the octupole current set to its nominal value, than it would be with only the solenoid and mirror fields turned on. This assumption (linear superposition of fields) opens the door to a potential systematic error in the registration of antihydrogen PSR frequencies with respect to measured electron cyclotron frequencies, at the level of 40 MHz.

We infer the magnitude of the magnetic field at the trap centre by measuring the electron cyclotron resonance frequency during each experimental sequence (see Sec. 4.4.1 for a description of the plasma modes diagnostic). This procedure indicates that the field we
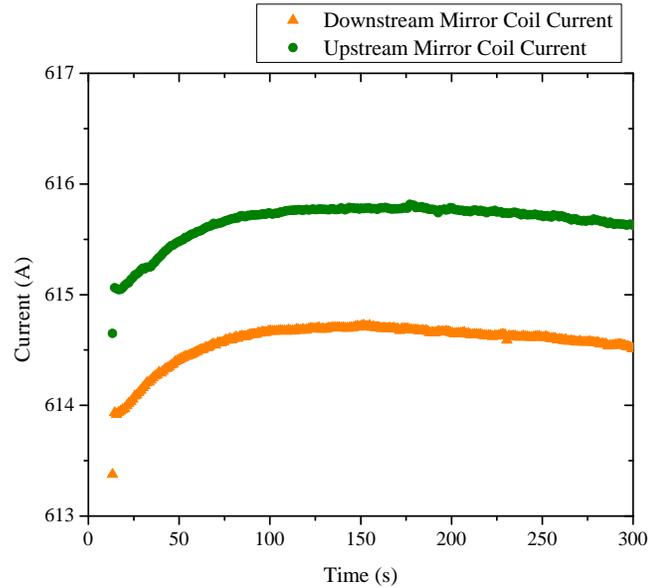
Figure 5.14: Mirror coil currents as a function of time, after nominally being set to fixed values.

establish is reproducible to within an equivalent of ± 2 MHz and that its absolute value is known to an equivalent of 10 MHz. These uncertainties are limited by the size of the electron plasma that is used. In theory, using a smaller number of electrons at lower temperatures should improve the precision of this measurement. Further discussion of this may be found in Refs. [109, 123].

Table 5.6 lists the various factors that currently dominate the uncertainty in the measurement and inference of the magnetic field on the axis of the ALPHA apparatus, and which thus contribute to uncertainty in the registration between PSR transition frequencies and ECR frequencies. Assuming that these factors are independent, we can infer the magnetic field to 15 G or an equivalent electron cyclotron frequency of 41 MHz.

5.6 Summary

Viewed as a whole, our appearance and disappearance mode datasets paint a compelling picture in which resonant interactions between microwave fields and the internal quantum

Source	Contribution to the Inferred Magnetic Field Error Budget
Mirror Coils	2 MHz
Octupole	40 MHz
Plasma Mode Diagnostic	10 MHz

Table 5.6: Factors that contribute to uncertainties in the inferred magnetic field, expressed in equivalent ECR frequency.

states of trapped ground state antihydrogen atoms have been induced and observed. As soon as microwaves are injected, PSR transitions are induced and annihilation events are observed as confined atoms are released. During the experiments in which the frequency of microwave fields are detuned by 100 MHz from the anticipated resonance, the annihilation events largely vanish. The small signal that does remain can be understood in terms of much weaker interaction with $|c\rangle$ state atoms as they pass through resonance 'far' from the trap centre.

These observations constrain the frequencies at which the onset of the $|c\rangle \rightarrow |b\rangle$ and $|d\rangle \rightarrow |a\rangle$ transitions occur to within 100 MHz. Similarly the frequency difference between the two lines and hence $\Delta\nu_{HFS}$ (Eq. 5.6) is constrained to 100 MHz. In the next chapter these ideas are revisited and a more precise statement of the bounds on $\Delta\nu_{HFS}$ is made. A model that accounts for the probability of spin-flip transitions occurring is introduced and used to reproduce the time distribution of annihilation events (similar to Fig. 5.10). This then allows for a better understanding of the systematic uncertainties associated with our proof-of-principle spectroscopy experiments.

Chapter 6

Numerical Modeling of Spin-Flip Experiments

6.1 Introduction

The PSR line shape shown in Fig. 5.5 was generated using a numerical model developed by Prof. Francis Robicheaux in connection with our publication of Ref. [1]. I have independently developed a numerical model that generates results that are essentially identical to those of Prof. Robicheaux's model, and which I have since used to extend the understanding of our experimental data. These models and the manner in which various simulations are performed are discussed in this chapter.

6.2 Numerical Simulation of Microwave Interactions with Ground-State Antihydrogen

Our simulations of microwave-induced spin flip transitions of magnetically trapped antihydrogen atoms use a mixture of classical and quantum mechanics. We use classical mechanics to calculate the motion of atoms in the trap and we use quantum mechanics to compute the transition probability each time an atom encounters a resonance condition.

A detailed description of the classical part of the calculation can be found in Ref. [124].

The force acting on an antiatom can be calculated from the potential energy of a magnetic dipole moment in a magnetic field: $U = -\boldsymbol{\mu} \cdot \boldsymbol{B}$. The magnetic moment of the antiatom is approximately one Bohr magneton. For low-field seeking states, the magnetic moment is directed opposite to the magnetic field so that the potential energy is $U = \mu B$ where B is the magnitude of the magnetic field. The classical force on the atom is then $\boldsymbol{F} = -\nabla U = -\mu \nabla B$.

The magnetic field consists of a superposition of four separate components: a uniform field, two mirror fields and an octupole field (as previously discussed in Chapter 2). The contribution of each component was computed using the techniques described in Ref. [125]. Essentially, each field was accurately represented by appropriate analytic functions. The superposition of these fields yields the profile shown in Fig. 2.4.

Once the field is known, we compute the force using a symmetric finite difference formula. An adaptive step-size Runge-Kutta algorithm [126] was then used to track the motion of the antiatom in the trap.

The initial conditions for our simulations mimic experimental conditions. We know that the positron plasma has a temperature of approximately 50 K and that the antihydrogen atoms are formed through three-body recombination. At this temperature it is expected that the antiprotons come into thermal equilibrium with the positrons before recombination occurs [127]. Thus in our simulations we launch antihydrogen atoms randomly from a volume defined by the size and location of the positron plasma, with velocities consistent with a 50 K thermal distribution. As in the experiment, most of the simulated antihydrogen atoms are too energetic and almost immediately hit the walls of the trap. Only a small fraction of the simulated atoms have low enough kinetic energy (≤ 0.54 K) to be confined by the magnetic potential well. We allow the trapped atom to propagate in the trap potential for approximately 1 s before turning on the microwaves. This 1 s period is to ensure that the atom trajectory is not biased by its initial conditions (position and momentum). The adaptive Runge-Kutta algorithm is not a symplectic integrator and is thus intrinsically susceptible to energy drift [128]. We reject any atom trajectory in which the energy of the atom changes by more than 1% relative to its initial value.

To determine the probability of a microwave-induced spin flip transition we use the Landau-Zener approximation [129, 130, 131, 132]. As atoms propagate in the trap, we keep track of the energy detuning, $\Delta E = E_u - (E_l + hf)$, where E_u is the energy of the upper (trapped) state, E_l is the energy of the lower (untrapped) state, and f is the frequency of the applied microwave field. This formalism implicitly assumes the microwave field is spatially uniform and orthogonal to the trapping field. Resonance conditions are identified by quadratic interpolation of ΔE between three successive time steps to identify points at which $\Delta E = 0$. The probability of a spin flip transition occurring as an atom passes through resonance is:

$$P = 1 - \exp\left(-2\pi \frac{|V|^2}{\hbar \frac{d\Delta E}{dt}}\right) \tag{6.1}$$

where V is the coupling matrix element between the upper and lower state and all parameters are evaluated at the time when $\Delta E = 0$ [133]. In the limit where $|V|^2/\hbar$ is small compared to the slew rate $d\Delta E/dt$ (diabatic limit), Eq. 6.1 can be approximated as $P = 2\pi V^2 / \left(\hbar \frac{d\Delta E}{dt}\right)$. To good approximation in large magnetic fields, $V = \frac{1}{2}\mu B_{\mu W}$ for PSR transitions, where $B_{\mu W}$ is the amplitude of a linearly polarized microwave magnetic field and μ is the magnetic moment of antihydrogen atom.

Any time that a resonance condition is met, the location and momentum of the atom is recorded and the spin-flip probability is estimated using the Lanadu-Zener approximation. Figure 6.1 shows an example of the distribution of resonance locations that are identified when the applied microwave fields are detuned by +1320 MHz with respect to the onset of the $|c\rangle \rightarrow |b\rangle$ transition (or by -100 MHz with respect to the onset of the $|d\rangle \rightarrow |a\rangle$ transition) in the minimum trap field. The resonance condition for a given frequency f is satisfied on a surface over which the magnetic field is constant. When the frequency is swept over some range the surface becomes a shell with finite thickness. When the frequency is set to zero detuning (resonance at the minimum magnetic field in the trap), the resonance surface collapses to a region close to the trap centre. Interaction between the radial components of the mirror and octupole fields results in this minimum being a circle of radius 4 mm about the trap midpoint, under typical trapping conditions.
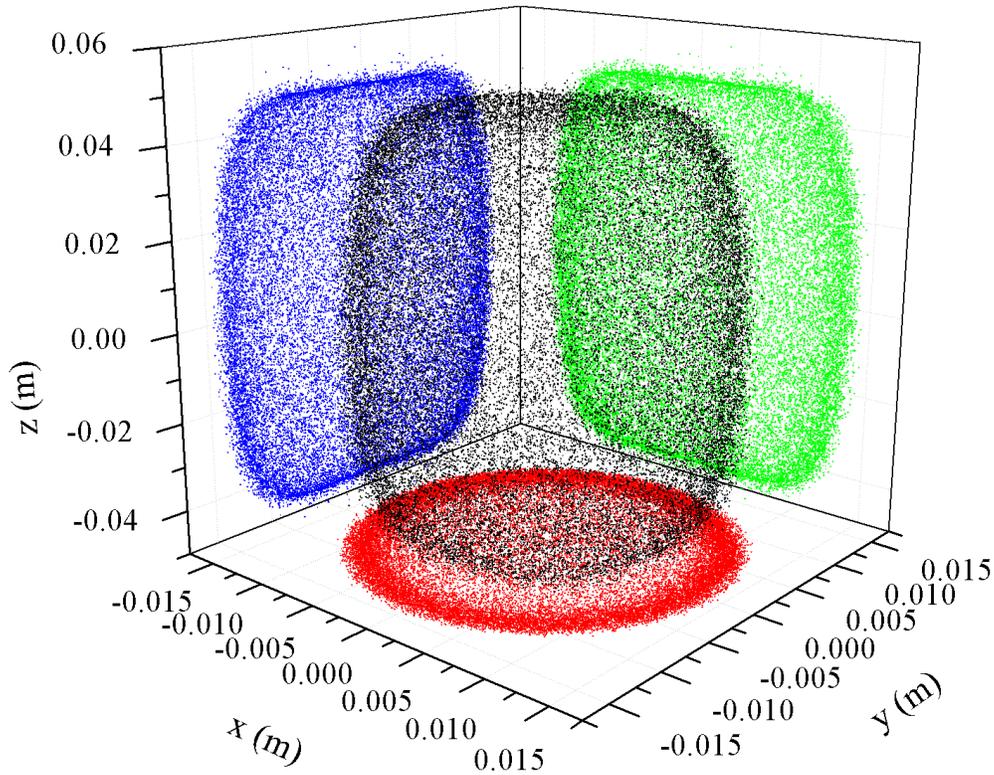
Figure 6.1: Distribution of resonance locations. Black dots represent the locations at which the resonance condition for the $|c\rangle \rightarrow |b\rangle$ PSR transition is met when the microwave frequency is set 1320 MHz above the minimum frequency at which this transition occurs in the minimum magnetic field (detuning= 1320 MHz). This distribution is a result of 5000 trajectories. The apparent fuzziness of the surface is an artifact of the numerical algorithm (step size) used to generate the distribution. The red, blue and green dots show the projection of this distribution on the xy, xz, and yz planes, respectively. Note that the scale for the transverse (x and y) coordinates is magnified relative to the scale for the axial (z) coordinates.

Figure 6.2 shows the distribution of axial (z) and radial (r) spin-flip locations and corresponding probability for three different detunings. In the top row the detuning is set to 1 MHz, in the next row it is set to 10 MHz and in the bottom row it is set to 1320 MHz (which corresponds to the off-resonance conditions employed in our experiments).

In all three cases the microwave magnetic field is set to 5.3 mGauss, which for a plane wave propagating in free space corresponds to a microwave electric field of 160 V/m. The distribution of spin flip locations depends strongly on detuning. For example, when the detuning is 1320 MHz (bottom row of Fig. 6.2), this distribution is spread out over a 10 cm long, 2.5 cm diameter shell (Fig. 6.1). When the detuning is small (eg, top row of Fig. 6.2) the distribution is much closer to a 4 mm thick, 1 cm diameter washer with a 5 mm hole. Ultimately these distributions reflect the static magnetic fields used for trapping.

The rightmost column of distributions in Fig. 6.2 (panels c, f, and i) show the spin-flip probabilities for the three detunings. The peak of the spin-flip probability distribution shifts to the left as the detuning increases. As the detuning increases the resonance condition is satisfied at higher and higher magnetic field, where the spatial gradient of the magnetic field is also increasingly larger. Since $\frac{d\Delta E}{dt}$ is proportional to $|\nabla B|$ ($\frac{d\Delta E}{dt} \propto \mathbf{v} \cdot \nabla B$, where v is the atom speed), increasing $|\nabla B|$ decreases the spin flip probability. Notice that the spin-flip probability distributions have been shown on a semi-log plot to illustrate the contrast between conditions that lead to spin flips with near-unit probability (conditions that rarely occur) and conditions that lead to spin flips with low probability (conditions that occur often). Note that rare interactions that lead to spin flips with near unit probability occur when atom trajectories are nearly tangential to the resonant surface.

6.2.1 Spatial Distribution of Annihilation Event Locations

We compute a uniformly distributed random number between 0 and 1 each time a simulated antihydrogen atom passes through a resonance condition; if the random number is less than the spin-flip probability P, evaluated using the Landau-Zener approximation, then we flip the magnetic moment of the atom. If the atom is initially in a low-field seeking state, a flip converts it into a high field seeker that is repelled from the magnetic minimum, and eventually hits the walls of the trap. We allow for the possibility of a spin flip every time
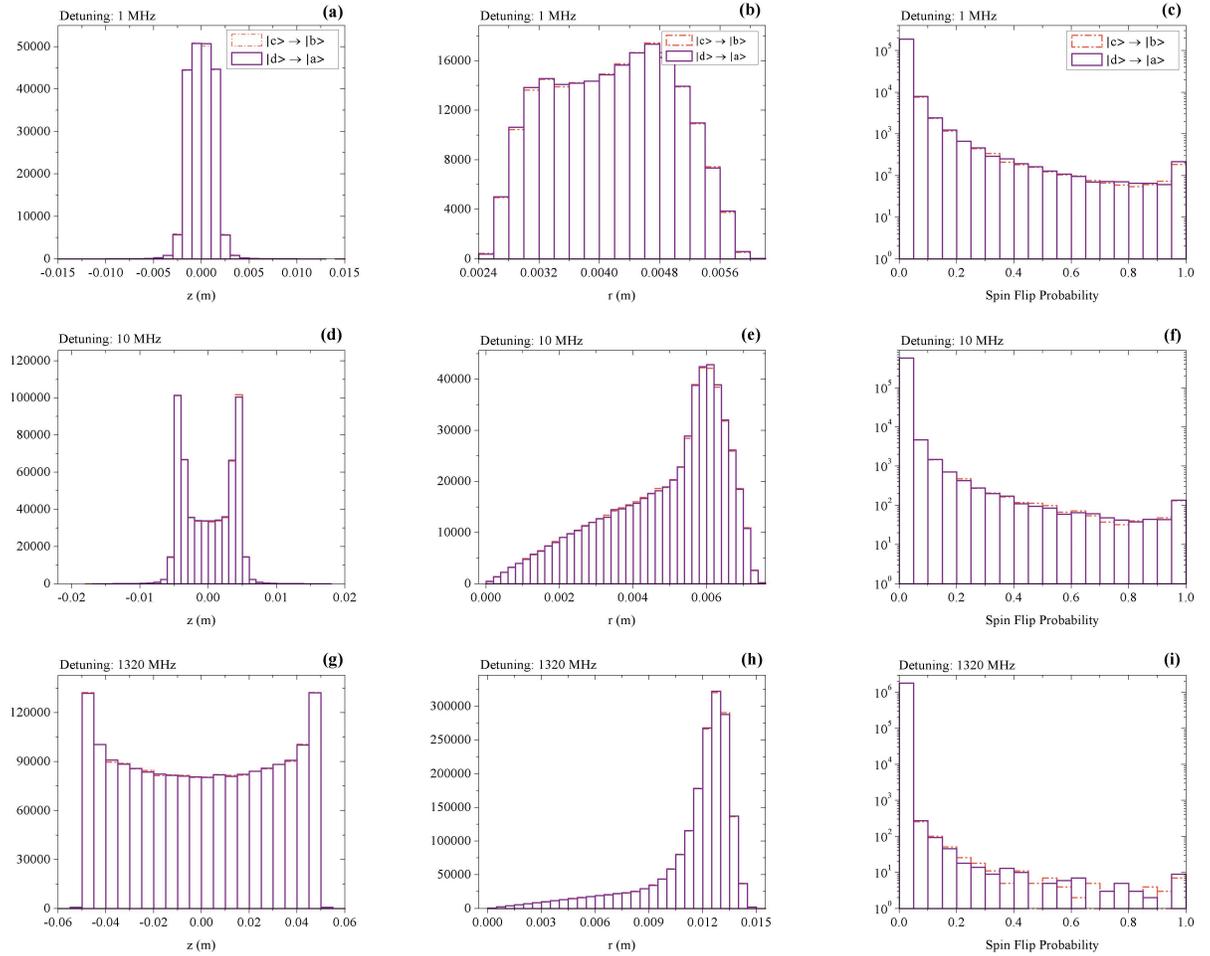
Figure 6.2: Spin flip locations and probabilities. Panels a, d, and g show the distribution of axial (z) spin flip locations for three detunings of 1, 10, and 1320 MHz, respectively. Panels b, e, and h show the radial (r) distribution of these locations for the same detunings. Panels c, f, and i show the distribution of spin flip probabilities for these detunings (1, 10, 1320 MHz). In each plot two distributions are shown, with orange and violet representing the results calculated for $|c\rangle$ and $|d\rangle$ state atoms, respectively. The distributions are the result of 5000 trajectories and 1 s of interaction with microwave fields. In all cases the microwave magnetic field is set to 5.3 mGauss, which for a plane wave propagating in free space corresponds to a microwave electric field of 160 V/m. Care should be exercised comparing the distributions shown in the leftmost and central columns, as the ranges are quite different.

an atom encounters a resonant condition. Therefore, it is possible for a high field seeker to be converted back to a low field seeker, although the probability for this to occur is low. Once the antihydrogen atom hits the wall of the trap, the time and position of the simulated annihilation event are recorded.
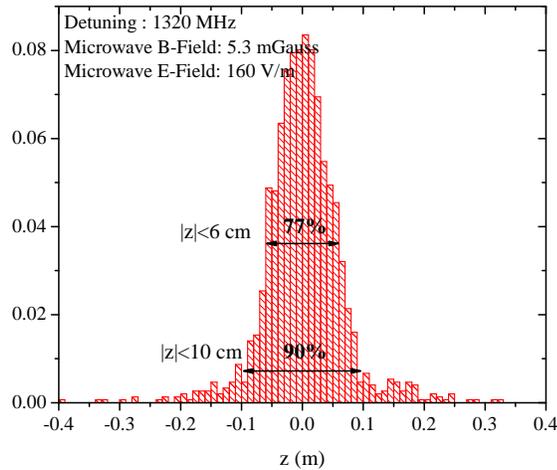
Figure 6.3: Axial distribution of annihilation event locations, as induced by spin flip transitions when the microwave frequency is detuned by 1320 MHz (corresponding to off-resonance experiments). Applying the requirement that $|z| < 6$ cm for an event to be considered valid (as was done in Ref. [1]) implies that 23% of the events caused by spin flip transitions are ignored. The looser constraint $|z| < 10$ cm is used in our conventional acceptance criteria. The detector spans the range -23 cm to +23 cm. The distribution has been normalized to 1.

If a simulated spin-flip transition occurs the atom trajectory is followed until it strikes the trap electrodes. In principle the distribution of simulated annihilation event locations should be identical to the distribution of reconstructed vertices for perfect detection. The radial distribution of simulated annihilation event locations is set by the inner radius (22.3 mm) of the electrode walls. For small detunings the axial distribution of annihilation events is concentrated near the midpoint of the trap because the surface over which resonant conditions are encountered is small and close to the center of the trap. For large detunings the axial distribution of annihilation events will be much more spread out, as is the case in Figs. 6.1, 6.2.g, and 6.2.h corresponding to a detuning of 1320 MHz (similar to experiments performed under off-resonance conditions).

As was mentioned in the previous chapter (Sec. 5.4.3) an axial cut $|z| < 6$ cm is applied to the location of detected annihilation event verticies. Using the simulation described here, we can examine the effect of this applied z-cut. Figure 6.3 shows that in the case of off-resonance experiments, 77% of antihydrogen annihilation events (caused by spin-flip transitions) fall inside the range $|z| < 6$ cm, and thus 23% of spin-flip transitions result in an annihilation event that would be rejected by a perfect detector.

6.3 Simulation of Experimental Results

Figure 6.4 shows a compilation of the time distribution of detected annihilation events for both appearance and disappearance modes. They are shown on the same plot despite the fact that the detection efficiencies for these two modes of operation are not the same. Appearance mode events during the 180 s microwave irradiation period are shown grouped in the usual 15 s bins. Disappearance mode events detected following the initiation of the magnet ramp down are shown in the single (and final) bin after t = 180 s. The width of this bin is drawn as if it is the same as the earlier bins, despite the fact that it is really only 30 ms in duration. In Fig. 6.4, data corresponding to experiments performed under on-resonance conditions are shown in red; data from off-resonance experiments are shown in blue. For on-resonance experiments the majority of trapped atoms are ejected in the first 30 s during which the $|c\rangle \rightarrow |b\rangle$ and $|d\rangle \rightarrow |a\rangle$ transitions are both probed. The number of counts in the rest of the bins are consistent with those expected given the cosmic ray background level (~ 2.7 counts/15-s bin). In contrast, for off-resonance experiments most of the trapped atoms survive the microwave irradiation period and are released and observed when the trapping fields are rapidly shut down.

Figures 6.5 and 6.6 show predictions of the simulation under conditions comparable to those employed in our experiments. Figure 6.5 shows the distribution of spin-flip probabilities for a single passage of an atom through the resonance region. This distribution peaks at around 0.002 which means that a trapped antihydrogen atom typically has to cross the resonance region several hundred times in order to have a decent chance of experiencing a spin-flip. Figure 6.6 shows the distribution of time delays between sequential crossings. It reveals that most atoms cycle through resonance at rates in excess of 50 Hz. These
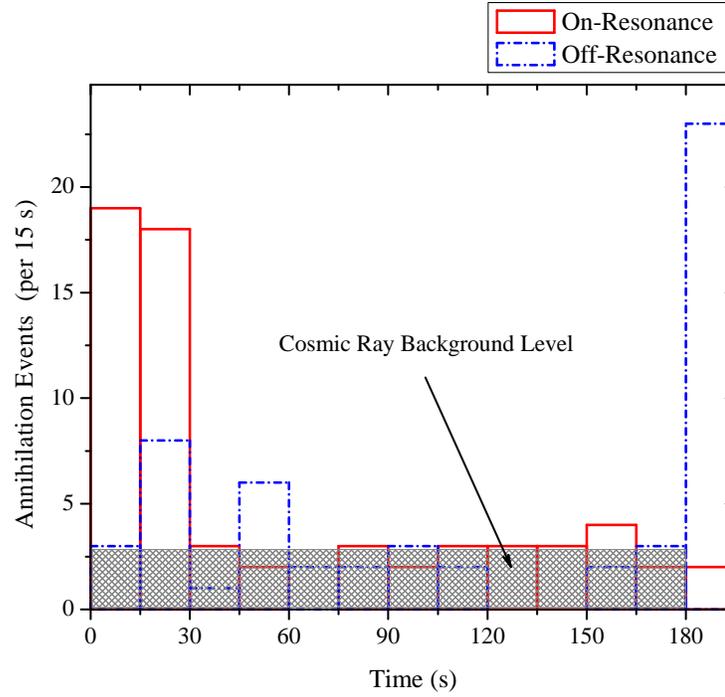
Figure 6.4: Compilation of experimental annihilation events in 'Appearance' and 'Disappearance' modes. The final bin corresponds to the annihilation of atoms that survived the 180 s microwave irradiation period. Data from experiments performed under on resonance and off resonance conditions are shown in red and blue, respectively. Most atoms survive the radiation applied in the off-resonance case (see the last bin). The gray band shows the anticipated cosmic background level during the microwave radiation window ($0 - 180$ s). The cosmic ray background level in the final bin (disappearance mode) for this bin is much smaller (0.16 ± 0.007 counts) because the observational window is far shorter (30 ms instead of 15 s). The event acceptance algorithm used in this mode is also different (see Secs. 5.4.2, 5.4.3, and App. A for details).

distributions form the main inputs to a second simulation, which attempts to model our experimental conditions in order to generate results like those shown in Fig. 6.4. The predictions shown in Figs. 6.5 and 6.6 will vary with respect to two important parameters: microwave power and microwave frequency. Increasing the microwave power will increase the spin-flip probability and increasing the detuning will tend to decrease the time between sequential passage through resonance (since the surface over which resonance conditions are encountered grows in size and is displaced outward from the trap centre, cf. Fig. 6.1).
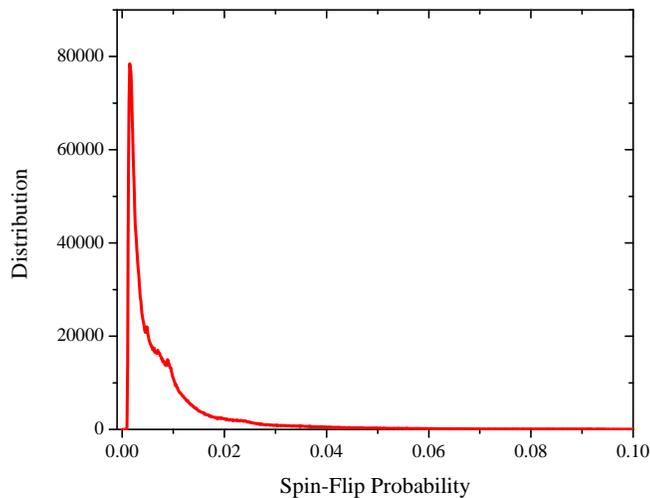
Figure 6.5: Distribution of spin-flip probabilities for a single passage of an antihydrogen atom through resonance. This distribution is the result of 1000 simulated antihydrogen trajectories. The magnitude of the microwave magnetic field is $B_{\mu W} = 5.3$ mGauss, and the detuning is set to 2.5 MHz. The surface over which the resonant conditions are encountered is therefore similar to the outer surface of a 4 mm thick, 1 cm diameter washer with an approximately 4 mm hole, located at the centre of the trap.

To generate simulation results analogous to Fig. 6.4 one needs to calculate the spin-flip probability as a function of time for an appropriate set of experimental conditions. The most direct way of doing this is to simply run the code used to model trajectories over and over using appropriate parameters. This procedure quickly becomes cumbersome when one is interested in exploring a large parameter space. We have chosen an alternate approach. From our numerical model, we have the spin-flip probability distribution for a single passage
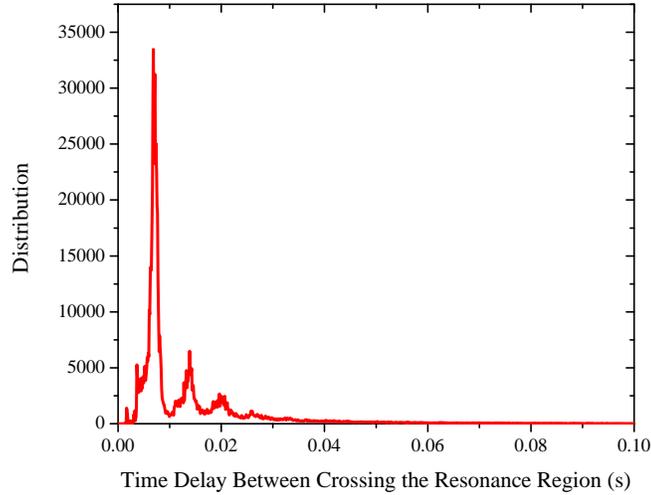
Figure 6.6: Distribution of time delays between sequential passages through resonance. The conditions here are identical to those outlined in the caption to Fig. 6.5.

through resonance for a given set of conditions (Fig. 6.5). Using this distribution and the distribution of times that it takes an atom to cycle back through resonance (Fig. 6.6), the cumulative spin-flip probability can be modeled as a function of time for those conditions:

$$P_{sf} = \sum_{i=1}^{n} P_i \prod_{m=1}^{i-1} (1 - P_m) \tag{6.2}$$

where P_{sf} is the cumulative spin-flip probability, P_i is the spin-flip probability for a single passage through resonance, and n represents the total number of passages through resonance. Alternately, the cumulative survival probability P_{su} (the probability that a spin flip transition is not induced after n passages through resonance) is:

$$P_{su} = \prod_{i=1}^{n} (1 - P_i) \tag{6.3}$$

and

$$P_{sf} + P_{su} = 1 \ . \tag{6.4}$$

Once the spin-flip probability for multiple crossings is known, one can numerically calculate the spin-flip probability as function of time. To do so a time loop is run, representing

the period of time during which the atom propagates inside the trap and every now and then passes through resonance. The algorithm to simulate this process is as follows: first a random spin-flip probability is selected from Fig. 6.5; this is the spin-flip probability when the atom passes through resonance for the first time. This value is recorded in the variable P_1. Then the time that it takes the atom to pass through resonance again is picked (randomly) from the distribution shown in Fig. 6.6 and the time loop steps forward. This process continues until the time loop reaches some pre-determined value. Then, using Eqs. 6.2, and 6.3 the spin-flip and survival probabilities are calculated. This algorithm is identical for other experimental conditions (different microwave powers and frequencies) except that the input distributions (Figs. 6.5, and 6.6) are changed accordingly. Figure 6.7 shows the cumulative spin-flip and survival probabilities for a microwave magnetic field $B_{\mu W} = 5.3$ mGauss and a detuning of 2.5 MHz. An effective spin-flip rate $\tau^{-1} = 0.38$ s^{-1} is obtained for these conditions by fitting exponential decay/buildup functions to the cumulative spin-flip/survival probabilities. As expected, this rate is consistent with the earlier observation that atoms cycle through resonance at rates > 50 Hz and that several hundred passes through resonance are required to obtain a significant spin-flip probability.

With an effective cumulative spin-flip probability function in hand, the expected time distribution of annihilation events (associated with PSR transitions) can be generated. A time loop is run; for each time step the cumulative spin-flip probability is calculated. A random number between 0 and 1 is generated. If the random number is smaller than the calculated spin-flip probability an annihilation event is assumed to occur and the corresponding time is recorded. The time loop runs for 180 s and if no spin-flips occur, an annihilation event is recorded to represent the annihilation of an atom at the end of the experiment when the trapping fields are ramped down. As the time loop progresses, the frequency at which the microwave field is applied is swept precisely as it is during our experiments (15 MHz sweeps over 15 s alternating between the two PSR lines).

As previously stated, the microwave magnetic field amplitude and the physical location of the surface on which atoms pass through resonance (set by the frequency) are the key parameters that govern the time distribution of annihilation events. From an experimental perspective our knowledge of these parameters is limited. We measure the on-axis static magnetic field and microwave electric field at the centre of the trap using electron cyclotron
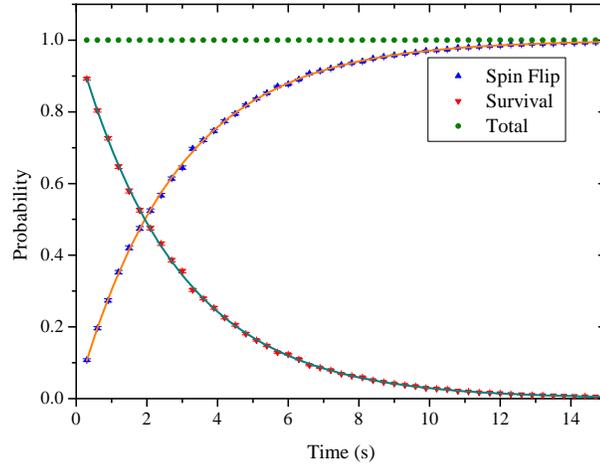
Figure 6.7: Cumulative spin-flip and survival probabilities. Exponentials of the form $p(t) = \alpha + \beta e^{-t/\tau}$ are fit to the calculated probabilities to extract τ, the effective spin-flip time constant. Here $\tau = 2.8$ s for $B_{\mu W} = 5.3$ mGauss, and a detuning of 2.5 MHz. Increasing the microwave field amplitude will result in a smaller value of τ. Note that sum of P_{sf} and P_{su} is 1, as expected.

resonance. More generally, one component of the microwave electric field can be mapped out along the axis of the trap. While this is informative, it does not give us the microwave magnetic field. In fact, it reveals that as expected, the microwave field pattern is a complex superposition of standing and traveling waves. The best we can do in our simulations to model this complex field is to treat it as an effective uniform field. As far as the absolute tuning of the microwave frequency with respect to the onset of the PSR transition frequencies at the trap centre is concerned, a useful observation comes from the appearance mode data acquired on resonance. We observe that annihilation events occur as soon as the microwave sweep begins (red data series in Fig. 6.4). This is true both at t = 0 s (start of the $|c\rangle \to |b\rangle$ transition) and at t = 15 s (start of the sweep covering the $|d\rangle \to |a\rangle$ transition). Recall that these sweeps nominally begin 5 MHz below the onset of the two transitions. This suggests that there is a systematic offset between the actual minimum magnetic field in the trap and the minimum magnetic field measured via ECR without the octupole field. This offset may be expressed in terms of an equivalent frequency offset $\Delta f \geq 0$ as illustrated in Fig. 6.8. The absolute value of this frequency offset is not known to us, but it cannot

be larger that 100 MHz; otherwise we would not have observed a decrease in off-resonance appearance mode annihilation events (and a corresponding increase in disappearance mode annihilation events).
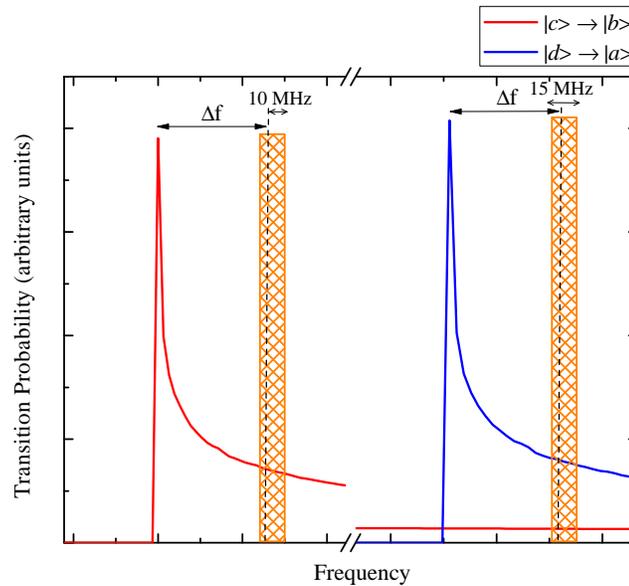
Figure 6.8: Probable systematic offset in the minimum magnetic field, with respect to the value inferred from ECR. Here the offset is shown as an equivalent frequency offset ($\Delta f \geq 0$).

My approach to exploring the consistency between our model and experimental data has been to generate simulated time distributions of annihilation events analogous to Fig. 6.4 for a broad range of 'effective' microwave magnetic field amplitudes and frequency offsets. I first calculated the distributions of spin-flip probabilities and times between sequential passages through resonance (analogous to Figs. 6.5 and 6.6) for a wide range of microwave magnetic field amplitudes and frequencies (Fig. 6.9). For each condition I generate cumulative spin-flip and survival probabilities that are then fed into a sequence that mimics the entire set of experimental runs. That is, a sequence that simulates the correct proportion of runs at each relative microwave power level, static field settings, microwave magnetic field amplitude variations for the lower and upper transitions, and microwave frequency sweep.
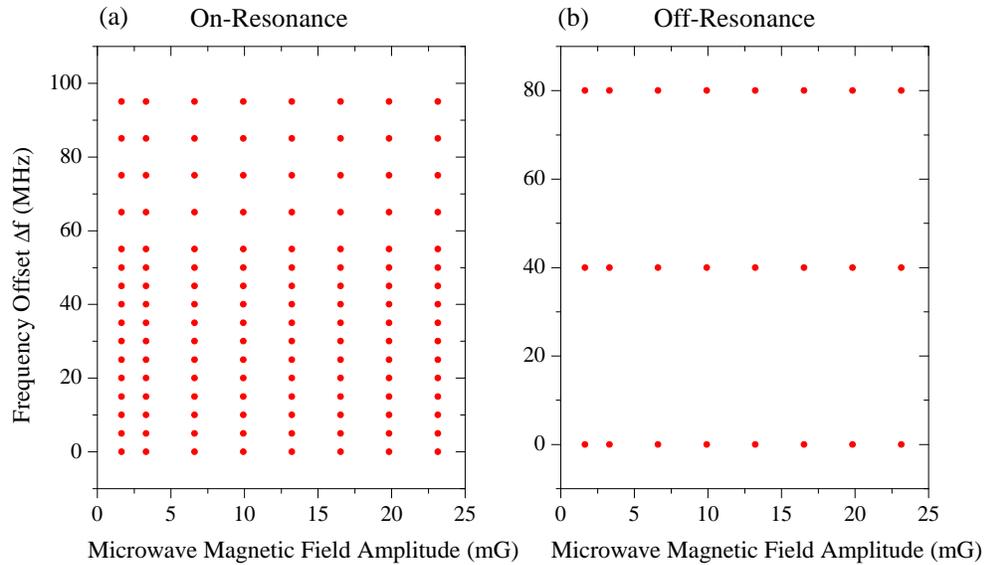
Figure 6.9: A grid showing the combinations of magnetic fields and frequency offsets Δf for which distributions of spin-flip probabilities and time between sequential passages through resonance (analogous to Figs. 6.5, and 6.6) are calculated. These distributions then form inputs to another simulation which generates the time distribution of annihilation events for both on-resonance and off-resonance cases. Depending on the target microwave power (or field) and frequency, the most relevant calculated input distributions are employed to generate the time distribution of annihilation events. As a point of reference, the amplitude of the microwave magnetic field (inferred from ECR) associated with our 'full power' data series (see Tab. 5.2) is 5.3 mGauss. The simulated off-resonance experiments are quite insensitive to frequency offset, and so only three different offsets (0 MHz, 40 MHz, and 80 MHz) are used to calculate the distributions of spin-flip probabilities and time between sequential passages through resonance.

It also accounts for the fact that two different PSR transitions can be induced when the microwave frequency is tuned to the upper transition. This sequence is then run many times, to generate an anticipated time distribution of annihilation events analogous to Fig. 6.4, for a given set of conditions.

The above procedure is then repeated for a range of microwave magnetic field amplitudes (0 to 25 mGauss) and microwave frequency offsets (0 MHz to 100 MHz for off-resonance experiments and -10 MHz to 100 MHz for on-resonance experiments). If the field amplitude and frequency offset happen to correspond to combination for which spin-flip probabilities and times between sequential passages through resonance (analogous to Figs. 6.5 and 6.6) have already been calculated (i.e. one of the combinations shown in Fig. 6.9), the appropriate distributions are used for the rest of the simulation. If on the other hand the field amplitude and frequency offset do not lie precisely on the grid shown in Fig. 6.9, the distributions are approximated by interpolating between the closest available distributions. The errors introduced by the interpolation process on the simulation results (Sec. 6.3.2) are marginal.

At this point the question becomes, over what range of conditions do the simulated time distributions of annihilation events mimic the experimental observations? I return to this question in Sec. 6.3.2, but first I introduce the 'likelihood ratio' as a statistical test for evaluating consistency between simulated and experimental results.

6.3.1   Likelihood Ratio

In experimental high-energy physics there are many situations in which one has two histograms in hand and is wondering if they are consistent. This is the case here. We generate histograms of simulated annihilation events as a function of time for various conditions and would like to know which ones are consistent with the data. From this comparison we hope to learn useful information about the range of simulation input parameters that are consistent with the data.

There are many statistical tests that address the question of consistency between a given dataset and a particular parent distribution. These tests may be adapted to address the

question of whether or not two datasets have been drawn from the same parent distribution ('two-sample' tests) [134]. These tests can be further customized to address the question of whether or not two histograms have the same shape. This is the case in which we are interested. In shape tests the relative normalization of the two histograms is not important; the densities (number of counts per bin) of the two histograms are examined bin-by-bin.

A variety of different test statistics can be used for shape comparison of histograms; examples include chi-squared, the Kolmogorov-Smirnov (KS) test, the Cramer-von-Mises test, and the likelihood ratio test [135, 136, 137]. I have used the likelihood ratio test because it is based on the Poisson distribution and hence is relatively robust against biases caused by low statistics. The predictions of this test are expected to asymptotically approach those obtained with a chi-squared (χ^2) distribution [137].

Consider two sample histograms. The bin contents of the first histogram are given by realization u of random variable U and the second by realization v of random variable V. The likelihood ratio test statistic, $-2 \ln\lambda$, is calculated from:

$$-2 \ln\lambda = -2 \sum_{i=1}^{k} \left[t_i \ln\left(\frac{1 + v_i/u_i}{1 + N_v/N_u} \right) + v_i \ln\left(\frac{N_v u_i}{N_u v_i} \right) \right] \tag{6.5}$$

where k is the number of bins, u_i and v_i are corresponding bin values, $t_i = u_i + v_i$, and N_u and N_v are the total counts for each histogram. This test statistic can be used to compare two histograms, neither of which has bins with zero counts. With some care, histograms with bins in which $u_i = 0$ or $v_i = 0$ can also be handled. Further discussion of the likelihood ratio test statistic can be found in Ref. [137].

Before using the likelihood ratio test, we first create binned time distributions from our experimental data and simulation results; the bins are 15 s wide. Then the expected cosmic ray events are added to the simulation results. Alternatively, one could subtract cosmic ray events from data. However, in our case where we have low statistics, subtracting the cosmic ray events results in bins with zero counts, which in turn can introduce biases in test results. Next, we consider the fact that the acceptance criteria used for our appearance mode experiment are 25% less efficient than those used for our disappearance mode experiments. Therefore, we reject 25% of the simulated disappearance events (last bin of the histograms).

Finally, Eq. 6.5 is used to quantify the shape-similarity of the binned time distribution of annihilation events (data) and the corresponding distribution predicted by simulation.

6.3.2 Simulation Results

Figures 6.10 and 6.11 show the likelihood test statistic as a function of simulation parameters for off-resonance and on-resonance conditions, respectively. In both cases, the vertical axis represents the frequency offset Δf (see Fig. 6.8) while the horizontal axis represents a scale factor for the microwave magnetic field amplitude. The scale factor is the ratio of the microwave magnetic field amplitude used in the simulation to that inferred from *in-situ* ECR-based measurements of the microwave electric field amplitude (Sec. 4.5). Thus a scale factor of 1 means that the calibration of the microwave magnetic field amplitudes used in the simulation is identical to that inferred from ECR and the plane wave approximation. A scale factor of 2 means that the calibration of the effective microwave magnetic field amplitudes used in the simulation are double those inferred from ECR, and so on.

Figure 6.10 shows the likelihood ratio test statistic comparing the off-resonance simulations and data. The color represents the value of the test statistic ($-2 \ln\lambda$ in Eq. 6.5). Note that the global minimum value of the test statistic has been subtracted from all other values so that the values of the test statistic start from 0. Smaller values indicate better consistency between simulation and data. An increment of 1 in the value of the test statistic represents a one-standard deviation decrease in consistency between the simulated binned time distribution of annihilation events, and the data. The colored bands shown in Figs. 6.10 and 6.11 are correspondingly drawn to represent regions over which the test statistic increases by one.

The off-resonance likelihood test result (Fig. 6.10) reveals almost no sensitivity to changes in frequency offset; this is expected. During off-resonance experiments the spin-flip probability for the $|c\rangle \rightarrow |b\rangle$ transition changes very little over a 100 MHz band while microwaves are injected just below the $|d\rangle \rightarrow |a\rangle$ transition (see Fig. 5.9). However, this test is sensitive to the microwave magnetic field amplitude. In effect, the field amplitude controls the balance between the number of annihilation events that are detected during the appearance and disappearance mode phases of the experiment. As is evident in Fig. 6.10, the simulation generates results that are most similar to the experimental observations if the effective microwave field amplitudes are approximately double (1.2 to 2.5) the expectations inferred
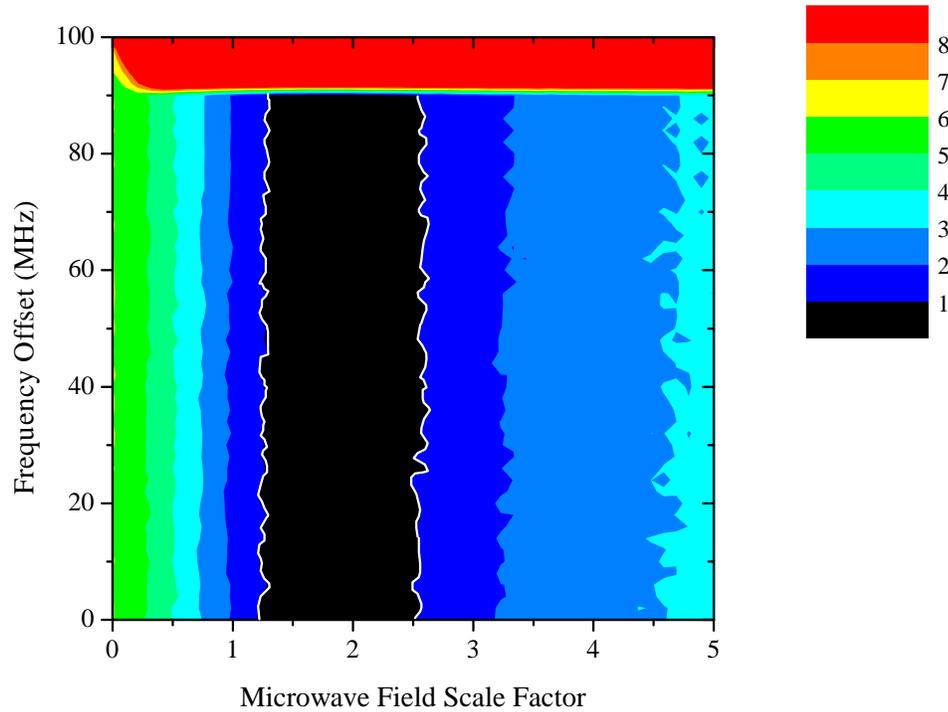
Figure 6.10: Likelihood ratio test statistic (cf. Eq. 6.5) comparing simulated binned time distributions of annihilation events with data (cf. Fig. 6.4) during the off-resonance portion of the experiment. The smaller the test statistic the better the agreement between simulation and experiment. The global minimum value of the test statistic has been subtracted from all other values so that the plotted values start from 0. The region over which the two are most consistent forms a nearly vertical band where the microwave field scale factor spans the range 1.2 to 2.5. Notice that the fine (ragged) structure on the vertical contour lines are due to the interpolation between the points and do not carry significant information (one can simply replace them with straight lines). Colored bands represent changes in the value of the test statistic by 1. When the frequency offset is larger than 90 MHz, the upper range of both 15 MHz frequency sweeps overlaps with the PSR lines (see Fig. 5.9). This means the simulation is no longer 'off resonance' and provides time distributions of annihilation events that are highly inconsistent with the data (red band on the top of the figure).
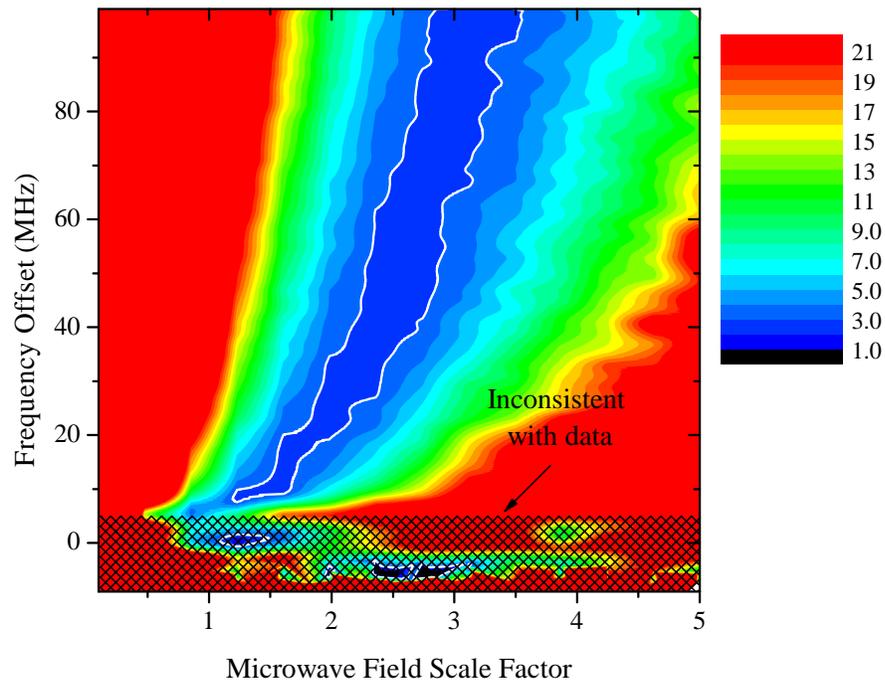
Figure 6.11: Likelihood ratio test statistic (cf. Eq. 6.5) comparing simulated binned time distributions of annihilation events with data (cf. Fig. 6.4) during the on-resonance portion of the experiment. The color codes can be interpreted in the same manner as those in Fig. 6.10. From the on-resonance event data, we know that annihilation events are observed as soon as the microwave sweeps start. This implies that the bands over which the frequencies are swept overlap the lineshape completely. This sets the minimum frequency offset to 5 MHz, and excludes the region for which $\Delta f \leq 5$ MHz (shaded).

from ECR (the central dark blue band). If the frequency offset is set to values larger than 90 MHz, the upper ends of the frequency sweeps start to overlap the two PSR lines, generating binned time distributions of annihilation events that resemble those observed during on-resonance experiments. The likelihood test predicts extremely poor consistency between simulation and data under these conditions (red band at the top of Fig. 6.10).

Figure 6.11 shows the likelihood test statistics comparing the on-resonance simulations and data. The color scheme should be interpreted in the same way as for the off-resonance case. Unlike the off-resonance case, the on-resonance test statistic depends on the frequency offset (and the microwave field amplitude). In this case we also examine negative frequency offsets. Frequency offsets $\Delta f < -10$ MHz correspond to no overlap between the microwave sweeps and the PSR lines, giving rise to simulation results that are highly inconsistent with the data.
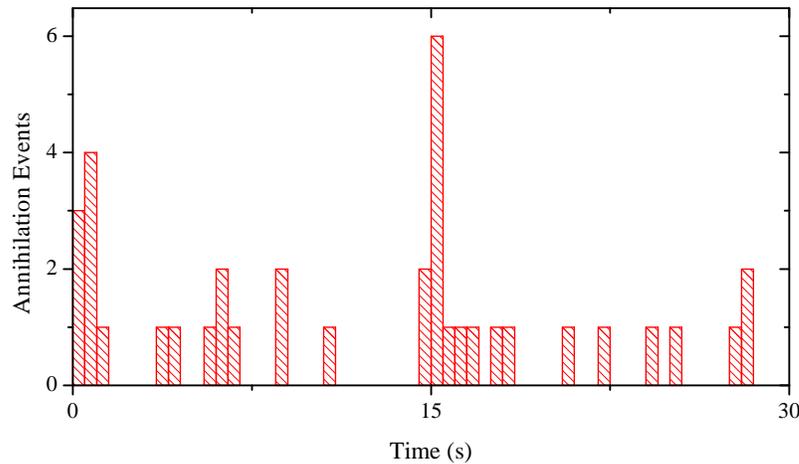
Figure 6.12: Time distribution of events during the first 30 s of experiments performed on-resonance, collected into 0.5 s bins. During the first 15 s the $|c\rangle \to |b\rangle$ transition is irradiated. During the second 15 s the $|d\rangle \to |a\rangle$ transition is irradiated.

The region for which the frequency offset in Fig. 6.11 is less than +5 MHz has been shaded. This has been done to highlight an inconsistency between the simulated binned time distribution of annihilation events and the data, that becomes apparent when the 15 s

bins used up to this point are broken down into smaller bins. Figure 6.12 shows the observed distribution of events for the first 30 s of the on-resonance experiments with 0.5 s binning. The first 15 s interval is associated with the period during which the $|c\rangle \to |b\rangle$ transition is irradiated and the second 15 s interval is associated with the period in which the $|d\rangle \to |a\rangle$ transition is irradiated. As is evident, many events are observed right at the start of the frequency sweep, where the microwave frequency is 5 MHz below the nominal target transition frequency (see Fig. 6.8). This suggests that the sweep starts at a frequency that already overlaps the PSR line, restricting the minimum frequency offset to 5 MHz. Statistically, the events observed in the first bin after the start of the sweep and in the first bin after the change in frequency at 15 s are highly unlikely to be associated with the cosmic ray background. The probability that the first bin (3 events) is all due to cosmic ray events is 1×10^{-4} and the probability that the first bin after 15 s (6 events) is all due to cosmic ray events is 6×10^{-10}. On the basis of this observation, we exclude the hatched region for which $\Delta f < 5$ MHz in Fig. 6.11.

The regions of parameter space for which the time distribution of simulated events best reproduces the data is clearly different for experiments performed on resonance and off resonance. One may overlay the two likelihood ratio test statistic plots (Figs. 6.10 and 6.11) and look for the intersection of the regions of best agreement. Figure 6.11 has been reproduced in Fig. 6.13. The vertical band shows the region of best agreement between simulations and data for the off resonance experiments (see Fig. 6.10). The striped region bounded by the vertical lines and the on resonance contour (dark blue contour showing the best consistency between on resonance simulations and data) shows the intersection of the region of best agreement between both sets of simulations and data. The minimum and maximum microwave field scale factor and frequency offsets consistent with this overlap are 1.2 to 2.5 and 10 MHz to 70 MHz, respectively.

Figures 6.14 and 6.15 show examples of binned time distributions of simulated annihilation events overlayed on top of the data for on resonance and off resonance experiments, respectively. In both cases the simulated effective microwave magnetic field is 1.8 times larger than that inferred from ECR experiments, and the frequency offset is 20 MHz above the target value. These parameters lie in the region corresponding to the intersection of the regions of best correspondence between the simulation and the data. In both cases (on
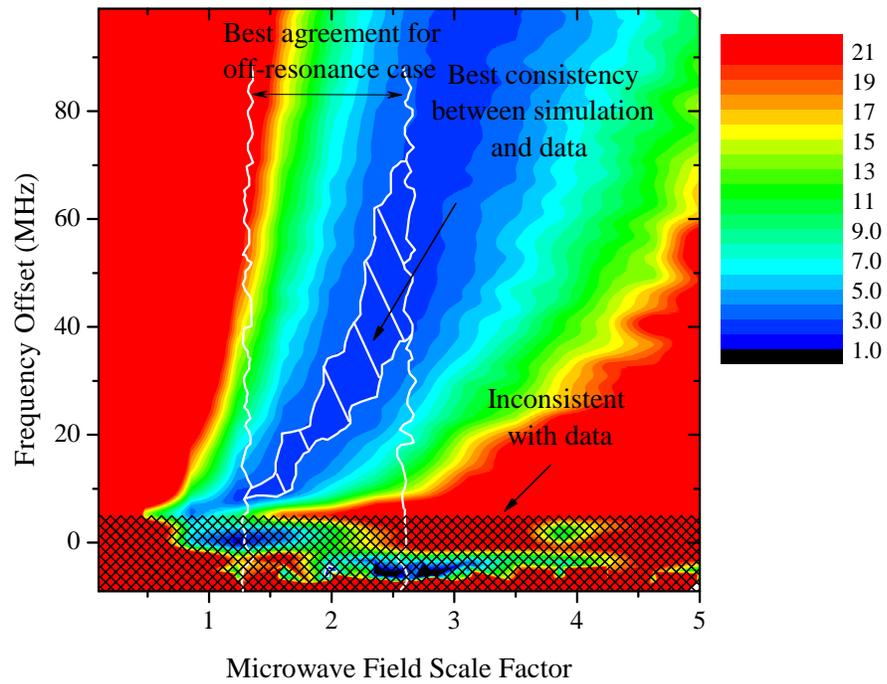
Figure 6.13: Intersection of the regions for which the simulated time distribution of events best matches the on-resonance data (contours) and off-resonance data (vertical white lines). The best consistency between simulations and data is observed when the microwave field scale factor ranges from 1.2 to 2.5, and the frequency offset ranges from 10 MHz to 70 MHz.
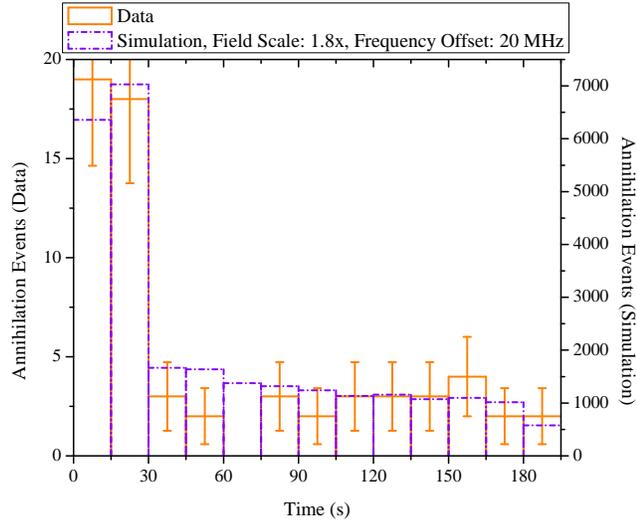
Figure 6.14: Comparison of a simulated time distribution of annihilation events with experimental observations for the case where the microwaves are tuned on-resonance. Here the microwave field amplitude is 1.8 times larger than that inferred from ECR. A systematic 20 MHz offset in frequency has been also applied. This corresponds to a 7 Gauss deficit in static magnetic field at the centre of the trap. The error bars are due to counting statistics.

resonance and off resonance), the simulated binned time distribution of annihilation events closely resemble the corresponding data. Similar levels of agreement are observed throughout the striped region marked on Fig. 6.13.

Before leaving this topic, note that both of the input parameters used to generate the simulations shown in Figs. 6.14 and 6.15 are reasonable and are consistent with our understanding of the experiment. A frequency offset $\Delta f = 20$ MHz implies that the minimum magnetic field with all trap magnets energized is approximately 7 Gauss less than the value anticipated, based on ECR measurements with the static and mirror fields on (octupole off) and with the static and octupole fields on (mirrors off). It is reasonable that interactions between the various magnets (mechanical force, flux redistribution, and induced persistent currents) could result in a 7 Gauss deficit at the trap centre. Similarly, a microwave field scale factor of 1.8 implies that the effective uniform magnetic field used in the simulation is a bit less than double the value inferred from ECR-based measurements of one component of the non-uniform microwave electric field. Again, this level of agreement is reasonable.
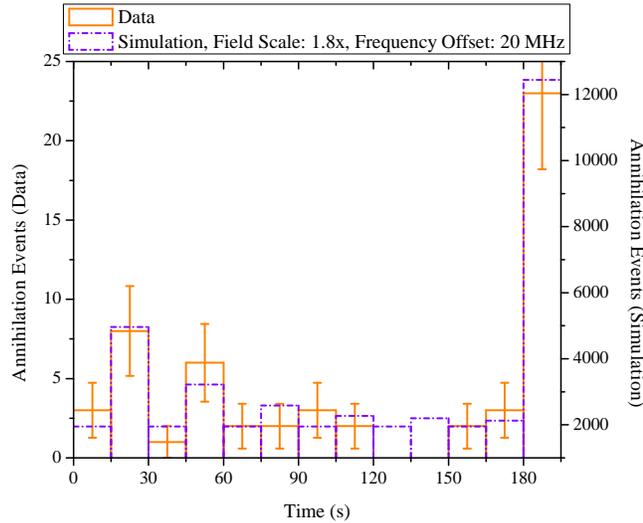
Figure 6.15: Comparison of a simulated time distribution of annihilation events with experimental observations for the case where the microwaves are tuned off-resonance. Here the microwave field amplitude is 1.8 times larger than that inferred from ECR. A systematic 20 MHz offset in frequency has been also applied. This corresponds to a 7 Gauss deficit in static magnetic field at the centre of the trap. The error bars are due to counting statistics.

6.4 Experimental Bound on the Hyperfine Splitting of the Antihydrogen Atom

Collectively, our appearance and disappearance mode data demonstrate strong evidence for the annihilation of antihydrogen atoms associated with intentionally induced PSR transitions. These annihilations vanish when the microwave frequency is detuned by 100 MHz, so that it lies below the anticipated onset of the transition [1]. Close examination of the annihilation data reveals that immediately after the injection of microwave radiation begins, annihilation events are observed (Fig. 6.12). This is true both for the first time microwaves are applied at the $|c\rangle \to |b\rangle$ transition and again for the $|d\rangle \to |a\rangle$ transition. This implies that the entire microwave sweep overlaps with the PSR lines. This observation can be used to constrain the hyperfine splitting of the antihydrogen atom $\Delta\nu_{HFS}$.

To quantify the experimental bound on the hyperfine splitting of the antihydrogen atom we seek the maximum and minimum values of $\Delta\nu_{HFS}$ that are consistent with our observations. In Fig. 6.16 we hypothesize that $\Delta\nu_{HFS}$ is larger for the antihydrogen atom than it is for the hydrogen atom (1420 MHz). The maximum hypothetical splitting such that the on-resonance experiments remain on resonance (Fig. 6.16.a) and the off-resonance experiments remain off resonance (Fig. 6.16.b) is 1505 MHz. In Fig. 6.17 we hypothesize that $\Delta\nu_{HFS}$ is smaller for the antihydrogen atom than it is for the hydrogen atom. The minimum hypothetical splitting such that the on-resonance experiments remain on resonance (Fig. 6.17.a) and the off-resonance experiments remain off resonance (Fig. 6.17.b) is 1335 MHz.
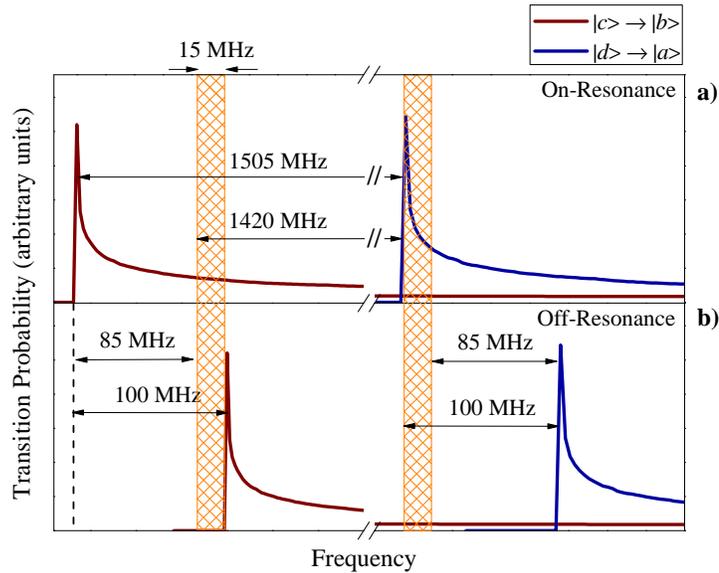
Figure 6.16: Maximum antihydrogen hyperfine splitting that is consistent with our data: $\Delta\nu_{HFS}(\bar{H}) = \Delta\nu_{HFS}(H) + 85$ MHz. A larger splitting would still be consistent with experiments performed on resonance (panel a), but would be inconsistent with experiments performed off resonance (panel b). This figure is schematic; it is not drawn to scale.

Combining these results we conclude that the hyperfine splitting of the antihydrogen atom is consistent with that of the hydrogen atom to within 85 MHz. That is:
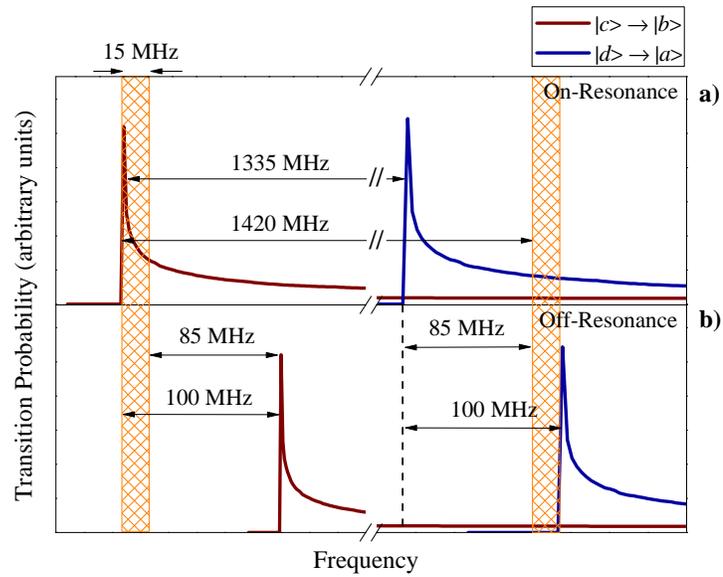
Figure 6.17: Minimum antihydrogen hyperfine splitting that is consistent with our data: $\Delta\nu_{HFS}(\bar{\text{H}}) = \Delta\nu_{HFS}(\text{H}) - 85$ MHz. A smaller splitting would still be consistent with experiments performed off resonance (panel b), but would be inconsistent with experiments performed on resonance (panel a). This figure is schematic; it is not drawn to scale.

$$\Delta\nu_{HFS} = 1420 \pm 85 \text{ MHz}. \tag{6.6}$$

for antihydrogen. One could certainly try to refine these bounds by examining the relative number of events observed when microwaves are applied on resonance for the two lines, but this would require making an assumption about the relative number of atoms trapped in the two low-field seeking states.

6.5 Spin-Flip: an Alternative for Detection of Trapped Antihydrogen

Our conventional approach to antihydrogen detection (disappearance mode) has been to turn off the trapping magnets and let the atoms strike the electrode walls and annihilate. Rapid discharge of the magnets (with a time constant of 9 ms) requires sophisticated electronics to avoid potential safety hazards and damage. More importantly, the mechanical impulses associated with the discharge are incompatible with a precision measurement environment. As an alternative, we envision driving PSR transitions to induce spin flip transitions, ultimately leading to annihilation of trapped atoms.

It has already been demonstrated that most atoms are ejected from the trap during the first sweep (of six) of the microwave frequency over ranges that overlap the two PSR lines (eg. Figs. 6.4 or 6.14, starting at $t = 0$ s and $t = 15$ s). Increasing the microwave field amplitude increases the probability of a spin flip each time an atom passes through resonance (Eq. 6.1) and thus decreases the amount of time required to effect a spin flip transition. This behaviour is expected to continue until the spin flip probability for a single pass through resonance approaches one, beyond which point the time required to empty the trap is limited by atom dynamics (eg. Fig. 6.6).

A simple illustration of the effect of increasing the microwave field amplitude (albeit with low statistics) is shown in Fig. 6.18. In this figure a subset (corresponding to the 'full power' series) of the on resonance data is shown. These data have been grouped into 3 s bins, and as was the case for Fig. 6.18, include the contribution of cosmic rays.
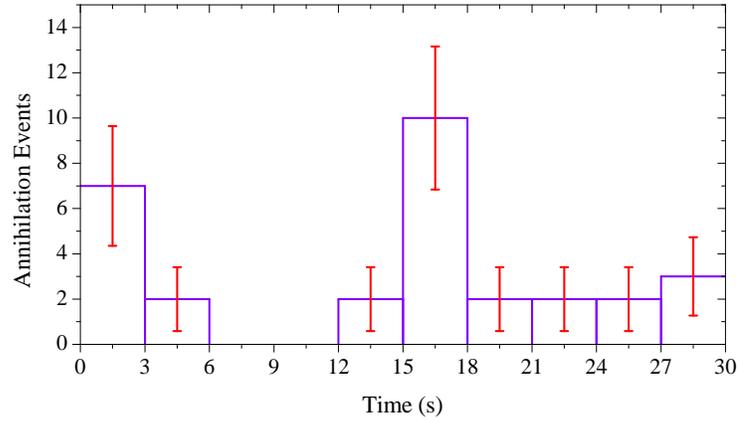
Figure 6.18: Time distribution of annihilation events when the microwaves are set to 'full-power' ($\sim$ 700 mW at the injection port; see Fig. 4.7 and Sec. 5.4.1).
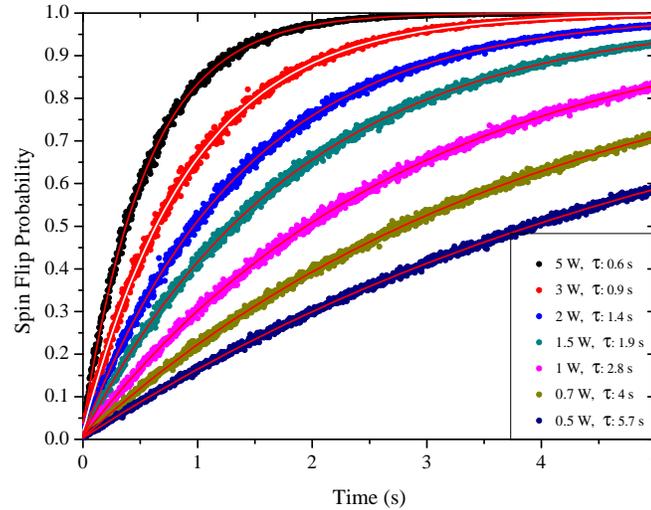
Figure 6.19: Simulated spin-flip probability as function of time and microwave power level. The curve labeled 0.7 W nominally corresponds to our 'full power' spin-flip experiments. The input parameters to the simulation have been described in the text. An exponential recovery function $P = P_0 \left(1 - e^{-t/\tau} \right)$ has been fit to each curve to extract the spin flip rate $1/\tau$.

Figure 6.19 shows the spin flip probability as a function of time at various injected power levels, as calculated using the simulations described in Sec. 6.3. The frequency offset is $\Delta f = 0$ in all of the calculations presented here, and the microwave frequency is swept over a 15 MHz band around the $|c\rangle \rightarrow |b\rangle$ transition frequency (as in the spin flip experiments). At each power level, an exponential recovery function has been fit to the simulation results to extract a time constant τ. As expected from Eq. 6.1, in the limit where the coupling matrix element V (proportional to microwave magnetic field $B_{\mu W}$) is small, the spin flip probability is proportional to the microwave power level (hence to $B_{\mu W}^2$). This is shown explicitly in Fig. 6.20 where the spin flip rate $1/\tau$ is plotted as a function of injected power. At the upper end of this plot (5 W) the spin flip rate (~ 2 s^{-1}) is still small compared to the rate at which atoms are cycled through resonance (> 50 Hz, cf. Fig. 6.6).
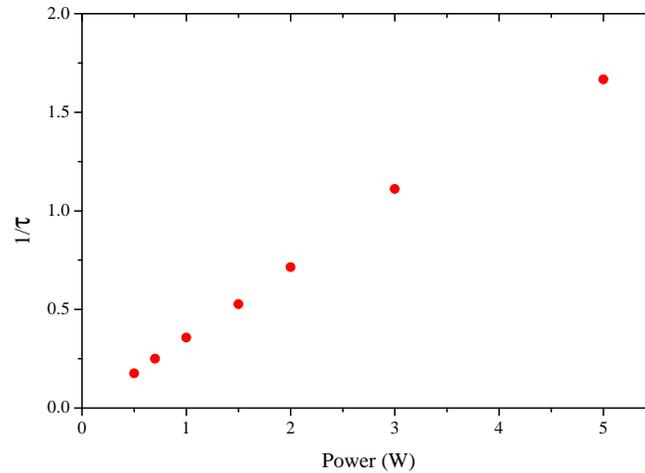
Figure 6.20: Spin flip rate $1/\tau$ as a function of injected power.

6.6 Summary

The experiments described in Chapter 5 and first reported in Ref. [1] are important, in the sense that they demonstrate the feasibility of inducing positron spin flip transitions in trapped antihydrogen atoms, and subsequently detecting the by-products of annihilation events as the atoms encounter the walls of the apparatus. The analysis presented in Chapter 6 goes beyond that reported to date. The simulations described in Secs. 6.2 and 6.3

provide insight regarding the rate at which spin flip transitions are induced, and the physical locations at which these transitions occur. The latter is closely tied to the distribution of annihilation event locations observed by the silicon vertex detector. The simulations also provide a compelling consistency check, in the sense that the time distribution of annihilation events can be modeled using parameters that closely reflect the environment that we believe exists deep within the trapping apparatus, as inferred from independent measurements. We find, for example, that a uniform microwave magnetic field with an amplitude that is 1.8 times that inferred from ECR experiments will produce the observed time series of events, as long as a systematic offset in the static magnetic field of 7 Gauss is introduced (equivalent to a +20 MHz shift in electron cyclotron resonance frequency). Equivalent results are obtained for field scale factors spanning the range 1.2 to 2.5 and frequency offsets spanning the range 10 MHz to 70 MHz (see Fig. 6.13 for the actual bounds on this parameter space). Perhaps most importantly, using the analysis presented in this chapter, the first experimental bound is determined for the zero-field hyperfine splitting of the antihydrogen atom: $\Delta\nu_{HFS} = 1420 \pm 85$ MHz. This bound is based on the first proof-of-principle spectroscopy measurements reported in Ref. [1] and in Chapter 5 of this thesis. Finally, we have used simulations to help extrapolate from the experiments that were performed in Ref. [1] to high power levels. This simple study provides support for the viability of using microwave induced spin flip transitions as a state-selective mechanism for antihydrogen detection.

Chapter 7

Future Prospects, Summary, and Conclusion

The production and confinement of cold antihydrogen atoms have paved the way for precision studies of matter-antimatter asymmetry in the atomic sector. There are also a number of other fundamental and intriguing questions that could be addressed through experiments on trapped antihydrogen atoms, including the study of gravitational interactions involving antimatter [139]. The next strategic move of the ALPHA collaboration after the work reported here was to design a new apparatus, tailor-made for precision spectroscopic measurements (both microwaves and laser). The design of the new apparatus began in 2012 and the machine was commissioned in late 2013. It is referred to as ALPHA-II, to distinguish it from the original apparatus.

The limiting factor for microwave spectroscopy experiments in the original ALPHA apparatus is associated with the magnetic field homogeneity. The precision of microwave spectroscopy experiments can be significantly increased by increasing the field uniformity. This may be accomplished by introducing additional shim coils, which effectively increase the multipolar order of the mirror coils. This capability was included in the design of ALPHA-II. With a more uniform magnetic field, the microwave spectroscopy program will be continued. The PSR measurements reported in [1] and in this thesis could certainly be repeated with higher resolution. However, there are other complementary microwave experiments with greater physics reach that are likely to be carried out in the near future.

Below I briefly describe one of these experiments. I then give a summary of the contents of this thesis.

7.1 Measurement of NMR and PSR Transitions

A microwave spectroscopy experiment that will almost certainly be performed on trapped ground state antihydrogen in the future involves inducing both NMR and PSR transitions. In this scenario the first step is to eject all of the atoms in the $|c\rangle$ state by driving the $|c\rangle \rightarrow |b\rangle$ transition, leaving only $|d\rangle$ state atoms in the trap (see Fig. 5.1). Then the $|d\rangle \rightarrow |c\rangle$ transition (NMR) and the $|c\rangle \rightarrow |b\rangle$ transition (PSR) are driven, either sequentially or simultaneously. This leads to atoms being ejected from the trap. By varying the frequencies at which the various transitions are driven one can measure f_{cd}, f_{bc} and f_{ad}.

There is an important feature of this experiment that can be used to advantage for spectroscopy. The transition frequency f_{cd} passes through a maximum at a magnetic field $B' = 0.65$ Tesla (Fig. 7.1). In other words, to first order f_{cd} is independent of B at B'. The absolute value of the transition frequency at this particular field is a characteristic of the antihydrogen atom that is closely linked to the zero-field splitting, and could be compared to a corresponding measurement for hydrogen. The fact that it is associated with an extremum suggests that in principle one could measure f_{cd} at $B = B'$ more accurately than at other fields. In this context one can show that

$$f'_{cd} = \frac{a}{h} \left[\frac{1}{2} - \frac{\sqrt{\eta}}{1 + \eta} \right] \tag{7.1}$$

where

$$\eta = \frac{\gamma_{\bar{p}}}{\gamma_{e^+}} \ll 1 \tag{7.2}$$

and f'_{cd} denotes f_{cd} at $B = B'$. Thus, measurement of f'_{cd} gives the zero-field hyperfine splitting of ground state antihydrogen a without requiring an independent measurement of B (the ratio $\gamma_{e^-}/\gamma_{e^+}$ is known to 8 parts in 10^{13} [138] and the antiproton magnetic moment is known to 4.4 parts in 10^6 [21]). As with a measurement of PSR transition frequencies alone, there is significant spectroscopic advantage to performing this measurement on atoms as they pass through the turning point in the field near the centre of the trap [11]. To put this statement in perspective, the influence of a 7 Gauss effect in B (relative to B') only

changes f_{cd} by 32 Hz with respect to f'_{cd} (Fig. 7.1). This corresponds to a relative precision of 2×10^{-8} relative to a/h.
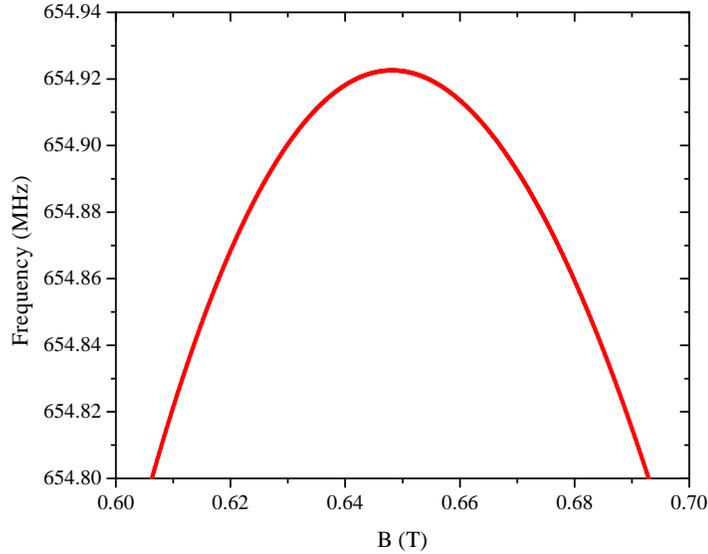
Figure 7.1: Frequency of the $|c\rangle \rightarrow |d\rangle$ transition as a function of magnetic field in the vicinity of the maximum that occurs at $B = B' \approx 0.65$ Tesla.

The frequency of the $|c\rangle \rightarrow |d\rangle$ NMR transition in a magnetic field of B=0.65 T is approximately 655 MHz (Fig. 7.1; notice how one can see this from expression 7.1, since $\eta \ll 1$, and $\frac{a}{h} \approx 1420$ MHz). The free space wavelength of electromagnetic radiation at this frequency is an order of magnitude larger than the electrode stack diameter. This precludes simply injecting microwaves from one end of the electrode stack and having them propagate to the trap centre. In order to drive this transition some sort of built-in resonator tuned to the appropriate frequency and compatible with the Penning trap is required. However, once installed, experiments analogous to those reported in this thesis (in scope) are expected to yield measurements of a/h to perhaps a few kHz (10^{-6} level of precision). A measurement at this level will have already suppressed the benchmarks for precision (of order 100 kHz) at which antihydrogen spectroscopy experiments are expected to enter uncharted territory. Further cooling of trapped antihydrogen atoms to mK temperatures (and lowering the trap depth by a corresponding amount) should allow for spectroscopic resolution of a/h at the

level of a few hundred Hz (10^{-7} level of precision). These types of experiments are envisioned for the future.

7.2 Summary and Conclusions

This dissertation reviews aspects of nearly four years of the Project ALPHA experimental program, from 2009 to 2012. During this period, the ALPHA collaboration made significant progress towards one of its ultimate goals: a precision study of the internal quantum states of atomic antihydrogen. A summary of the scope of this is given below. Key conclusions are also highlighted.

Chapter 1 presents the reader with a general picture of the ALPHA experiment. It also gives a brief historical overview of measurements of the ground state hyperfine splitting of atomic hydrogen.

Chapter 2 discusses the ALPHA apparatus. The ALPHA apparatus involves two superimposed traps: a combined magneto-electrostatic trap for charged plasma confinement and a superconducting octupole-based magnetic minimum trap for atom confinement. Antihydrogen atoms are synthesized by mixing antiproton and positron plasmas in an Ultra-High Vacuum (UHV) environment. Some of these atoms are formed with low enough momentum that they are trapped in the magnetic potential well, which is approximately 0.8 T or 0.5 K (in temperature units) deep for ground state antihydrogen atoms. The apparatus benefits from a number of particle detectors and plasma diagnostics. The particle detectors (such as plastic scintillator paddles and a silicon-vertex detector) function as counting and/or position-location devices to identify the annihilation of antiparticles. The plasma diagnostic devices (such as a Faraday cup, a micro-channel plate, and a phosphor screen) are used to characterize the properties of ensembles of (anti)particles. Antihydrogen atoms are detected when they strike the surrounding material walls of the apparatus and annihilate, which occurs when the magnetic trap walls are lowered. The location of these annihilation events are then determined by a silicon vertex detector which consists of 60 double-sided silicon wafer modules. These modules surround the ALPHA trap in a three-concentric-layer arrangement [43, 100].

Chapter 3 reviews experimental methods for production and confinement of antihydrogen in ALPHA [72]. These procedures set the stage for the central subject of the thesis: the first microwave spectroscopic measurement of an antiatomic energy interval. The antihydrogen production procedure begins with the extraction of a bunch of antiprotons (approximately 4×10^7 particles at an energy of order 5.3 MeV) from the CERN Antiproton Decelerator (AD) ring. For successful confinement of antihydrogen atoms, the energy of the extracted antiprotons needs to be reduced by many orders of magnitude. A number of energy reduction and cooling procedures, such as electron cooling and evaporative cooling, are therefore applied [79]. After these procedures a plasma of $(30 \pm 5) \times 10^3$ antiprotons with a density of 6.5×10^6 cm^{-3} at a temperature 200 ± 40 K is prepared for antihydrogen synthesis. In parallel, a positron plasma with a density of 5.5×10^7 cm^{-3} and a temperature 40 ± 15 K is prepared and confined in a nested-well, adjacent to the prepared antiproton plasma. Prior to mixing the two plasmas, the magnetic trap is energised. Then the antiprotons are autoresonantly injected into the positron plasma leading to production of antihydrogen atoms [96]. Most of the synthesized antihydrogen atoms are too energetic to be confined in the magnetic potential well, and thus escape and annihilate when they strike the surrounding material walls. The two particle species are allowed to interact for 1 s during which time we observe 5000 ± 400 annihilation events in the silicon detector. The remaining charged particles are then cleared using a sequence of electric field pulses (up to 500 V/m). Once all charged particles have been removed, the magnetic trap fields are rapidly ramped down, and we look for annihilation events associated with any atoms that might have been trapped. In our initial attempts at antihydrogen trapping, we repeated the above sequence 335 times (in three different variations) and observed 38 trapped atoms [72]. We also tried to increase the storage time of the confined antihydrogen atoms. We observed that even at a 1000 s confinement time there is strong evidence for survival of trapped atoms. The probability that the events observed after a 1000 s confinement time are caused by cosmic rays is 10^{-5}, corresponding to a statistical significance of 8σ [74]. In practical terms this confinement time is essentially infinite, as far as spectroscopy experiments are concerned, and thus opens the door to a variety of potential experiments.

Chapter 4 discusses two schemes (external and internal) for the injection of microwaves into the UHV environment of the ALPHA apparatus. The external injection scheme is suitable for injecting low power microwaves into the apparatus for electron cyclotron resonance

experiments. It involves two metallic mirrors or reflectors. One of these reflectors is a simple stainless steel plate mounted at the bottom end of the vertical translator (inside the vacuum system). The second reflector consists of a 5×10^2 cm^2 section of the inner surface of a prolate aluminum ellipsoid of revolution, located outside of the apparatus and positioned in front of a 4″ diameter glass view port. It is positioned such that one of its focal points is located outside of the apparatus while the other sits on or near the surface of the internal reflector. The throat of a 20 dB horn antenna is positioned at the external focal point of the ellipsoidal mirror and is oriented so that microwave radiation is focused onto the internal reflector and down the bore of the electrode stack. The attenuation in microwave power associated with this scheme between the output of the microwave synthesizer and the centre of the ALPHA trap, is of order 30 dB in the frequency range 28-29 GHz.

The internal injection scheme provides higher injection efficiency, and is better adapted to the antihydrogen spin-flip experiments reported in Chapter 5. It involves injection of microwaves via a horn antenna mounted on the vertical translator (inside the vacuum system). A vertical length of WR28 waveguide connects a waveguide window (on the top of the vertical translator) to the microwave horn. Provision is made outside the apparatus to amplify the microwaves (using a broad-band Miteq AMF-4B amplifier) prior to injection. The attenuation in microwave power associated with this scheme, between the output of the microwave synthesizer and the centre of the ALPHA trap, is of order 13.6 dB in the frequency range 28-29 GHz.

Chapter 4 also introduces a novel diagnostic tool for in-situ measurement of the static magnetic field on the axis of the ALPHA trap. It involves measuring the cyclotron resonance frequency of electrons from which the local magnetic field can be inferred. Typically an electron plasma comprising approximately 7×10^7 particles (density of 6.5×10^{14} m^{-3}) is first prepared at the center of the trapping apparatus and then a series of 4 μs duration microwave pulses is injected, at a rate of one pulse every 30 s. The microwave frequency is incremented every time a pulse is injected, in order to scan through the cyclotron resonance. While the microwave pulses are being injected, the plasma quadrupole frequency is monitored. Under appropriate conditions changes in this frequency are proportional to the plasma temperature. The increase in plasma temperature that occurs when microwaves are injected is maximum at the cyclotron resonance frequency. The magnetic field is then inferred based on the measured cyclotron resonance frequency. This diagnostic tool enables us to determine the electron cyclotron resonance frequency at the minimum magnetic

field of the ALPHA apparatus with an uncertainty of 10 MHz, corresponding to a relative magnetic field measurement of $\Delta B/B = 3.4 \times 10^{-4}$ [109]. Using this diagnostic tool we are also able to extract information about the amplitude of the local microwave electric field that is responsible for heating the electron plasma. This connection relies on a single particle theory that relates the electron plasma temperature increase to the amplitude of the component of the microwave electric field that co-rotates with respect to the cyclotron motion. Once the local microwave electric field is known, reasonable estimates of the local microwave magnetic field can be made. This parameter is then crucial to the understanding of magnetic resonance experiments on trapped antihydrogen.

Chapter 5 is dedicated to microwave spectroscopy experiments. Ground state antihydrogen atoms trapped in the ALPHA apparatus must be in one of the two trappable low-field seeking states, $|c\rangle$ or $|d\rangle$. A transverse microwave magnetic field $B_{\mu\text{W}}$ resonant with the $|c\rangle \rightarrow |b\rangle$ transition is first applied to convert $|c\rangle$ state atoms into $|b\rangle$ state atoms. These $|b\rangle$ state atoms are high-field seeking and are thus ejected from the trap and annihilate when they strike the electrode walls. Once the $|c\rangle$ state atoms have been removed from the trap the process can be repeated with the microwave radiation tuned to drive the $|d\rangle \rightarrow |a\rangle$ transition. In practice, during our experiments the two positron spin resonance (PSR) transitions are alternately driven, in a pattern that repeats several times. If both PSR transition frequencies (f_{bc} and f_{ad}) are measured under precisely the same conditions, their difference yields the zero-field hyperfine splitting $\Delta\nu_{HFS} = a/h$ independent of magnetic field.

After the microwave irradiation period the microwaves are turned off and the trapping field currents are rapidly ramped down. Annihilation events are identified throughout the entire experiment, and are analyzed in two distinct modes. If trapped antihydrogen atoms survive the microwave radiation, their annihilations are classified as disappearance mode events. If instead they are ejected during the irradiation period, their annihilations are classified as appearance mode events. During the experiments in which the frequency of microwave fields are detuned by 100 MHz from the anticipated resonance, annihilation events during the microwave radiation period largely vanish (off-resonance experiments). The small signal that does remain can be understood in terms of much weaker interactions with $|c\rangle$ state atoms as they pass through resonance 'far' from the trap centre. From a spectroscopic perspective, the dominant uncertainty in this series of measurements are associated with the uncertainty in the measurement of the magnetic field at the centre of the ALPHA trap.

We can infer the minimum magnetic field in the trap to 15 G or an equivalent electron cyclotron frequency of 41 MHz. This uncertainty is dominated by a systematic effect. We were not able to directly measure the field with full magnetic trap energised, and were forced to extrapolate from conditions under which direct measurements were possible. Although these experiments were only intended as a proof-of-principle demonstration of antihydrogen spectroscopy, they represent a significant achievement because they demonstrate that it is possible work with very small numbers of atoms. Much higher levels of precision are anticipated for experiments that will be performed using the second generation apparatus: ALPHA-II.

Chapter 6 takes the experimental results reported in Chapter 5 and presents an analysis that yields further insight. A model that accounts for the probability of spin-flip transitions occurring is introduced and used to reproduce the time distribution of annihilation events. Our simulations use a mixture of classical and quantum mechanics. We use classical mechanics to calculate the motion of atoms in the ALPHA trap and we use quantum mechanics to compute the transition probability each time an atom encounters a resonance condition. Any time that a resonance condition is met, the location and momentum of the atom is recorded and the spin-flip probability is estimated using the Lanadu-Zener approximation. This information is then used to calculate the effective cumulative spin-flip probability as a function of time and the expected time distribution of annihilation events associated with PSR transitions. The microwave magnetic field amplitude and the frequency offset (of the applied microwave radiation with respect to the onset of the PSR lineshapes) are treated as parameters that govern the time distribution of annihilation events. Consistency between our numerical model and experimental data is explored by generating simulated time distributions of annihilation events for a broad range of effective microwave magnetic field amplitudes and frequency offsets. The simulated time distributions of annihilation events are then compared with data using the likelihood ratio test statistic, and we look for the range of parameters (microwave field amplitudes and frequency offsets) over which the simulated time distributions of annihilation events are consistent with the experimental observations. The calculations suggest that the highest degree of consistency is obtained for frequency offsets spanning the range 10 MHz to 70 MHz and microwave field amplitudes 1.2 to 2.5 times larger than that inferred directly from ECR experiments. These values are in turn consistent with our understanding of the experiment.

In conclusion, this thesis reviews the first proof-of-principle microwave spectroscopy experiments on the antihydrogen atom [1]. We observe strong evidence for intentionally-induced PSR spin-flip transitions between the hyperfine levels of ground state atoms when microwave radiation is injected into the ALPHA apparatus at appropriate frequencies, but not when the frequency is detuned. No attempt was made to accurately localize a resonance or determine a lineshape, but the size of the step taken to detune the frequency (or field) sets a loose bound of 4 parts in 10^3 on the registration of the observed resonance with respect to those expected for hydrogen. Our data also constrain the hyperfine splitting of the antihydrogen atom to 1420 ± 85 MHz. These bounds are too crude to be useful as tests of fundamental symmetries, but they do mark the advent of antiatomic spectroscopy and demonstrate the viability of performing fundamental measurements on small numbers of trapped antihydrogen atom.

Over the last several years, the antihydrogen physics program at CERN has gained significant momentum, and now appears to be on the verge of pushing the boundaries of modern physics. In light of the achievements of the ALPHA collaboration reported in this thesis, as well as those of other experiments such as ATRAP and ASACUSA, it seems likely that the next few years will be an exceptional period in the history of antimatter. These efforts are expected to lead to valuable new insight and a better understanding of the fundamental laws that govern our universe.

"Nature uses only the longest threads to weave her patterns, so that each small piece of her fabric reveals the organization of the entire tapestry."

— Richard Feynman

Bibliography

[1] C. Amole et al. (ALPHA collaboration), Nature **483**, 439 (2012). xvii, 1, 86, 101, 102, 103, 104, 106, 111, 117, 135, 140, 141, 142, 150

[2] J. H. Christenson, J. W. Cronin, V. L. Fitch, and R. Turlay, Phys. Rev. Lett. **13**, 138 (1964). 2

[3] B. Feng, M. Li, J. Q. Xia, X. Chen, X. Zhang, Phys. Rev. Lett. **96**, 221302 (2006). 2

[4] B. Aubert et al., hep-ex/0607103 (2006). 2

[5] A. Matveev et al., Phys. Rev. Lett. **110**, 230801 (2013). 2

[6] C. G. Parthey et al., Phys. Rev. Lett. **107**, 203001 (2011). 2

[7] M. Fischer et al., Phys. Rev. Lett. **92**, 230802 (2004). 2

[8] S. G. Karshenboim, Phys. Rep. **422**, 1-63 (2005). 3

[9] H. Hellwig et al., IEEE Trans. Inst. Meas. **19**, 200 (1970). 3

[10] N. F. Ramsey, World Scientific Publishing Co., *Nobel Lecture* (1989). 3

[11] M. D. Ashkezari et al. (ALPHA collaboration), Hyperfine Interact. **212**, 81-90 (2011). 3, 67, 79, 80, 93, 143

[12] B. Juhász, E. Widmann, Hyperfine Interact. **193**, 305 (2009). 3, 89

[13] Particle Data Group, J. J. Hernández et al., Phys. Lett. B **239**, 1 (1990). 3

[14] M. Kobayashi, A. I. Sanda, Phys. Rev. Lett. **69**, 3139 (1992). 3

[15] I. I. Bigi, Nucl. Phys. A **692**, 227c (2001). 3

[16] C. E. Carlson, V. Nazryan, K. Griffioen, Phys. Rev. A **78**, 022517 (2008). 4

[17] R. S. Hayano, Hyperfine Interact. **175**, 53 (2006). 3

[18] D. Kleppner, Experiments with Atomic Hydrogen in Atomic Physics and Astrophysics, Vol. I, eds. M Chrétien and E. Lipworth, Gordon and Breach (1971). 4, 5

[19] A. Kreissl et al., Z. Phys. C **37**, 557 (1988). 4

[20] T. Pask et al., Phys. Lett. B **678**, 55 (2009). 4

[21] J. DiSciacca et al. (ATRAP Collaboration), Phys. Rev. Lett. **110**, 130801 (2013). 5, 143

[22] J. E. Nafe, E. B. Nelson, and I. I. Rabi, Phys. Rev. **71**, 914 (1947). 5

[23] D. E. Nagle, R. S. Julian, and J. R. Zacharias, Phys. Rev. **72**, 971 (1947) 5

[24] P. Kusch and H. M. Foley, Phys. Rev. **72**, 1256 (1947). 5

[25] S. D. Drell, J. D. Sullivan, Phys. Rev. **154**, 1477 (1967). 5

[26] R. Pohl et al., Nature **466**, 213 (2010). 5

[27] R. Pohl et al., Ann. Rev. Nucl. Part. Sci. **63**, 175-204 (2013) 5

[28] D. Colladay, V. A. Kostelecký, Phys. Rev. D **55**, 6760 (1997). 6

[29] R. Bluhm, V. A. Kostelecký, N. Russell, Phys. Rev. Lett. **79**, 1432 (1997). 6

[30] R. Bluhm, V. A. Kostelecký, N. Russell, Phys. Rev. D **57**, 3932 (1998). 6

[31] V. A. Kostelecký, Phys. Rev. Lett. **80**, 1818 (1998). 6

[32] R. Bluhm, V. A. Kostelecký, N. Russell, Phys. Rev. Lett. **82**, 2254 (1999). 6

[33] J. M. Brown, S. J. Smullin, T. W. Kornack, and M. V. Romalis, Phys. Rev. Lett. **105**, 151604 (2010). 6

[34] F. G. Major, V. N. Gheorghe, and G. Werth, Charged Particle Traps, Springer (2005). 11

[35] L. S. Brown, and G. Gabrielse, Rev. Mod. Phys. **58**, 233 (1986). 11

[36] G. B. Andresen et al. (ALPHA collaboration), Phys. Rev. Lett. **100**, 203401 (2008). 12, 23

[37] D. E. Pritchard, Phys. Rev. Lett. **51**, 1336 (1983). 14

[38] W. Bertsche et al. (ALPHA collaboration), Nucl. Instrum. Methods in Phys. Res. A **566**, 746-756 (2006). 14, 15, 16, 51

[39] G. B. Andresen et al. (ALPHA collaboration), Phys. Rev. Lett. **98**, 023402 (2007). 14, 51

[40] D. L. Eggleston and T. M. O'Neil, Phys. Plasmas **6**, 2699 (1999). 14

[41] J. Fajans, W. Bertsche, K. Burke, S. F. Chapman, and D. P. van der Werf, Phys. Rev. Lett. **95**, 155001 (2005). 14

[42] G. Ganetis, Private communication, Brookhaven National Laboratory (2003). 16

[43] C. Amole et al. (ALPHA collaboration), Nucl. Instrum. Methods in Phys. Res. A **735**, 319-340 (2014). 8, 13, 16, 22, 24, 64, 145

[44] X. Fei, PhD Thesis, Harvard University (1990). 19

[45] D. Möhl, CERN-PS-96-034-DI (1996). 20

[46] P. Belochitskii, T. Eriksson, and S. Maury, Nucl. Instrum. Methods in Phys. Res. B **214**, 176 (2004). 20

[47] P. Belochitskii (AD team), AIP Conf. Proc. **821**, 48 (2006). 20

[48] S. van der Meer, Rev. Mod. Phys. **57**, 689 (1985). 21

[49] G. I. Budker, Sov. J. At. Energ. **22**, 438 (1967).

[50] T. J. Murphy and C. M. Surko, Phys. Rev. A **46**, 5696 (1992). 22

[51] M. Amoretti et al. (ATHENA Collaboration), Nucl. Instrum. Methods in Phys. Res. A **518** (2004). 22

[52] L. V. Jørgensen et al. (ATHENA Collaboration), Phys. Rev. Lett. **95**, 025002 (2005). 22

[53] M. Charlton and J. Humberston, Positron Physics, Cambridge University Press, (2001). 23

[54] A. P. Mills, E. M. Gullikson, Appl. Phys. Lett. **49** 1121, (1986). 23

[55] X. P. Huang, F. Anderegg, E. M. Hollmann, C. F. Driscoll, and T. M. O'Neil, Phys. Rev. Lett. **78**, 875 (1997). 23

[56] R. G. Greaves and C. M. Surko, Phys. Rev. Lett. **85**, 1883 (2000). 23

[57] E. Klempt, C. Batty, and J. M. Richard, Phys. Rep. **413**, 197 (2005). 47

[58] J. L. Wiza, Nucl. Instrum. Methods **162**, 587601 (1979). 26

[59] http://www.elmul.com 26

[60] G. B. Andresen et al. (ALPHA collaboration), Rev. Sci. Instrum. **80**, 123701 (2009). 26, 27

[61] M. C. Fujiwara et al. (ALPHA collaboration), AIP Conf. Proc. **1037**, 208 (2008). 28

[62] G. B. Andresen et al. (ALPHA collaboration), J. Instrum. **7**, C01051 (2012). 28

[63] C. Regenfus, Nucl. Instrum. Methods in Phys. Res. A **501**, 65 (2003). 28

[64] M. C. Fujiwara et al. (ALPHA collaboration), Phys. Rev. Lett. **92**, 065005 (2004). 28

[65] M. Amoretti et al. (ATHENA collaboration), Nucl. Instrum. Methods in Phys. Res. A **518**, 679 (2004). 28

[66] G. B. Andresen et al. (ALPHA collaboration), Nucl. Instrum. Methods in Phys. Res. A **684**, 73-81 (2012). 28, 29, 30, 51

[67] R. A. Hydomako, Ph.D. Thesis, University of Calgary (2011). 28, 47, 48

[68] G. B. Andresen et al. (ALPHA collaboration), Phys. Lett. B **695**, 95104 (2011). 29, 30, 48, 51, 53, 54

[69] G. Baur et al. (PS210), Phys. Lett. B **368**, 251 (1996). 30

[70] G. Blanford et al., Phys. Rev. Lett. **80**, 3037 (1998). 30

[71] M. Amoretti et al. (ATHENA collaboration), Nature **419**, 456459 (2002). 30, 43, 44

[72] G. Gabrielse et al. (ATRAP collaboration), Phys. Rev. Lett. **89**, 213401 (2002). 30, 146

[73] G. B. Andresen et al. (ALPHA collaboration), Nature **468**, 673 (2010). 30, 48, 51, 56, 57, 58, 95, 160

[74] G. B. Andresen et al. (ALPHA collaboration), Nature Physics **7**, 558 (2011). 30, 60, 90, 146

[75] J. F. Ziegler, U. Littmark, J. P. Biersack, The stopping and ranges of ions in matter, Pergamon, New York, NY, 1985. 31

[76] G. Gabrielse et al., Phys. Rev. Lett. **57**, 2504 (1986). 31

[77] J. D. Jackson, Classical Electrodynamics, 3^{rd} ed. Wiley, New York, NY, 1999. 33, 68

[78] M. H. Holzscheiter et al., Phys. Lett. A **214**, 279 (1996). 33

[79] G. B. Andresen et al. (ALPHA collaboration), Phys. Rev. Lett. **105**, 013003 (2010). 30, 36, 38, 39, 146

[80] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell, Science **269** 198 (1995). 36

[81] H. F. Hess, Phys. Rev. B **34**, 3476 (1986). 36

[82] W. Ketterle and N. J. van Druten, Adv. At. Mol. Opt. Phys. **37** (1995). 37

[83] G. B. Andresen, Ph.D. Thesis, Aarhus University (2010). 38, 49

[84] E. Butler, Ph.D. Thesis, Swansea University (2011). 38, 40, 44, 48

[85] G. Gabrielse, S. L. Rolston, L. Haarsma, and W. Kells, Phys. Lett. A **129**, 38 (1988). 40

[86] M. Charlton, S. Jonsell, L.V. Jørgensen, N. Madsen and D.P. van der Werf, Contemporary Physics **49**, 29-41 (2008). 41, 44

[87] H. A. Bethe and E. E. Salpeter, Quantum Mechanics of One- and Two-Electron Atoms, Plenum, New York, NY, 1977. 41

[88] F. Robicheaux, J. Phys. B: At. Mol. Opt. Phys. **41**, 192001 (2008). 42

[89] M. Amoretti et al. (ATHENA collaboration), Phys. Rev. Lett. **97**, 213401 (2006). 42

[90] M. Charlton, Phys. Lett. A **143**, 143146 (1990). 42

[91] B. I. Deutch et al., Hyperfine Interact. **76**, 151161, (1993). 42

[92] C. H. Storry et al. (ATRAP collaboration), Phys. Rev. Lett. **93**, 263401, (2004). 42

[93] E. A. Hessels, D. M. Homan, and M. J. Cavagnero, Phys. Rev. A **57**, 16681671 (1998). 42

[94] G. Gabrielse, S. L. Rolston, L. Haarsma, and W. Kells, Phys. Lett. A **129**, 38 (1988). 42

[95] M. Amoretti et al. (ATHENA collaboration), Phys. Lett. B **590**, 133 (2004). 44

[96] G. B. Andresen et al. (ALPHA collaboration), Phys. Rev. Lett. **106**, 025002 (2011). 30, 45, 146

[97] W. Bertsche, J. Fajans, and L. Friedland, Phys. Rev. Lett. **91**, 265003 (2003). 45

[98] J. Fajans and L. Friedland, Am. J. Phys. **69**, 1096 (2001). 45

[99] J. Beringer et al. (Particle Data Group), Phys. Rev. D **86**, 010001 (2012). 47, 59

[100] G. B. Andresen et al. (ALPHA collaboration), Nucl. Instrum. Methods in Phys. Res. A **684**, 73-81 (2012). 47, 50, 145

[101] C. Amole et al. (ALPHA collaboration), New J. Phys. **14**, 105010 (2012). 30, 51, 53, 55, 58

[102] F. Chen. Introduction to Plasma Physics, Plenum Press, New York, NY, 1974. 53

[103] H. Hess et al., Phys. Rev. Lett. **59**, 672-675 (1987). 59

[104] P. A. Willems, and K. G. Libbrecht, Phys. Rev. A **51**, 1403-1406 (1995). 59

[105] D. Griffiths, Introduction to Electodynamics, 3rd ed. Prentice Hall, Upper Saddle River, New Jersey (1999). 68

[106] D. H. E. Dubin, Phys. Rev. Lett. **66**, 2076 (1991). 74, 75

[107] M. D. Tinkle, R. G. Greaves, C. M. Surko, R. L. Spencer, and G. W. Mason, Phys. Rev. Lett. **72**, 352 (1994). 74, 75

[108] M. D. Tinkle, R. G. Greaves, and C. M. Surko, Phys. Plasmas **2**, 2880 (1995). 76, 79

[109] C. Amole et al. (ALPHA collaboration), New J. Phys. **16**, 013037 (2014). 76, 77, 80, 81, 83, 109, 148

[110] M. E. Hayden, F. Robicheaux, W. N. Hardy, ALPHA internal communications. 83

[111] R. E. Collin, Foundations for Microwave Engineering, 2^{nd} ed. McGraw Hill, New York (1992). 84

[112] N. Kuroda, et al., (ASACUSA collaboration), Nature Comm. **5**, 3089 (2014) DOI: 10.1038/ncomms4089 86, 89, 90

[113] ASACUSA proposal addendum, CERN/SPSC 2005-002, SPSC P-307 Add.1 (2005). 89

[114] A. Mohri, Y. Yamazaki, Europhys. Lett. **63**, 207 (2003). 89

[115] R. S. Dyck, P. B. Schwinberg, H. G. Dehmelt, Phys. Rev. Lett. **59**, 26 (1987).

[116] J. Przyborowski, H. Wilenski, Biometrika **31**, 313323 (1939). 101, 160

[117] L. Breiman, Mach. Learn. **45**, 532 (2001). 102, 159

[118] I. Narsky, Preprint at http://arxiv.org/abs/physics/0507143 (2005). 102, 159

[119] I. Narsky, Preprint at http://arxiv.org/abs/physics/0507157 (2005). 102, 159

[120] M. Inokuti, B. Bederson, H. Walther, Advances in Atomic, Molecular, and Optical Physics, Volume 33: Cross Section Data 105

[121] G. Punzi, Proceedings of PHYSTAT2003: Statistical Problems in Particle Physics, Astrophysics, and Cosmology (Menlo Park, California, 2003), 79-83, http://inspirebeta.net/record/634798. 160

[122] H. O. Lancaster. Significance Tests in Discrete Distributions. *J. Amer. Stat. Assoc*, **56**, 223 (1961). 160

[123] T. Friesen, Ph.D. Thesis, University of Calgary (2014). 109

[124] P. H. Donnan, M. C. Fujiwara, and F. Robicheaux, J. Phys. B **46**, 025302 (2013). 111

[125] C. Amole, et al. (ALPHA collaboration), New J. Phys. **14**, 015010 (2012). 112

[126] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. Numerical Recipes 2^{nd} ed. Cambridge University Press, New York, 2007. 112

[127] J. L. Hurt, et al., J Phys. B: At. Mol. Opt. Phys. **41**, 165206 (2008). 112

[128] J. M. Thijssen, Computational Physics, Cambridge University Press, Cambridge (2000). 112

[129] L. Landau, Physikalische Zeitschrift der Sowjetunion **2**: 4651 (1932). 113

[130] C. Zener, Proc. Royal Soc. London A **137** (6) 696702 (1932). 113

[131] E. C. G. Stueckelberg, Helvetica Physica Acta **5**: 369 (1932). 113

[132] E. Majorana, Il Nuovo Cimento **9** (2): 4350 (1932). 113

[133] Jan R. Rubbmark, Michael M. Kash, Michael G. Littman, and Daniel Kleppner, Phys. Lett. A **23**, 3107 (1981). 113

[134] T. W. Anderson, Annals Math. Stat., **33**, 1148 (1962). 127

[135] A. Kolmogorov, G. Ist. Ital. Attuari, **4** 8391 (1933). 127

[136] T. W. Anderson, Annals Math. Stat., **33** 11481159 (1962). 127

[137] F. Porter, arXiv:0804.0380v1 (2008). 127

[138] R.S. Dyck, P.B. Schwinberg, H.G. Dehmelt, Phys. Rev. Lett., **59**, 26 (1987). 143

[139] C. Amole et al. (ALPHA collaboration), Nature Comm. **4**, 1785 (2013). 142

Appendix A

Multivariate Analysis

While the magnetic trap is energized and microwave radiation is applied, the time and axial position distributions of antihydrogen annihilation vertices are sensitive to the mechanism responsible for their release from the trap. We therefore monitor annihilation events throughout the entire hold time.

We base the event selection criteria on an alternative classifier (to that described in Chapter 4), optimized for this extended time window, and use a bagged decision tree, in the random forest approach [117, 118, 119], to separate the signal (annihilations on the trap walls) from the background (cosmic ray events). The random forest classifier (RF) is chosen for its stability with higher dimensionality, training stability, and insensitivity to input variables with weak discriminating power.

Nine variables are used for event classification: the (I) radial and (II) azimuthal coordinate of the reconstructed annihilation vertex, if present, (III) the total number of hits, (IV) the number of 3-hit combinations used as track candidates, (V) the number of reconstructed tracks, (VI) the sum of the squared residual distances of hits from a best fit straight line, and three topological variables. The topological variables comprise (VII) a sphericity variable, (VIII) the cosine of the angle between the event axis and the detector axis, and (IX) the angle between the event axis and the vertical direction in the X-Y plane. The sphericity variable is defined as the quantity $\frac{3}{2}(\lambda_2 + \lambda_3)$. Here $\lambda_1 \geq \lambda_2 \geq \lambda_3$ are the eigenvalues of the tensor $S^{\alpha\beta} = (\sum_i^N p_i^\alpha p_i^\beta |p_i|^{-2})/N$, where p_i^α is the component α ($\alpha = x, y, z$)

of the momentum associated with the i^{th} track. The event axis is defined as the line passing through the center of the detector and oriented along the eigenvector associated with λ_1.

We avoid overtraining by separating the control sample events equally between three sets. We train the classifier on one set, and use the other to determine if the classifier is overtrained. An unbiased estimate of the signal and background efficiencies when cutting on the classifier output is obtained from the third set.

We optimize the selection on the RF (Random Forest) output, by optimizing the Punzi figure of merit [121], $s/(N_\sigma + \sqrt{N_B})$ where N_σ is the sigma level of significance, taken here to be 3, and s and N_B are the signal efficiency and the expected number of background events, obtained from antiproton annihilations recorded during the mixing phase of trapping experiments and dedicated cosmic ray data-taking periods, respectively. The event-selection criteria have been determined without direct reference to the neutral trapping data used in the experiments described in Chapter 5, to avoid experimenter bias.

An annihilation candidate is retained if RF > 0.90. Compared to the selection applied to events occurring during the magnet quenches [73], the RF selection is about ten times more effective in rejecting cosmic background, while retaining about 75% of the signal. Based on simulations of the axial position z of the annihilations of spin-flipped antihydrogen with respect to the center of the trap, we furthermore require $|z| < 6$ cm. This requirement sacrifices approximately 23% of the signal and further suppresses the background by a factor of 3, resulting in an expected background rate of $(1.7 \pm 0.3) \times 10^{-3}$ s^{-1}.

The time and z distributions of the selected events are shown in Fig. 5.10 and Fig. 5.11. In the 30 s observation window we record a significant excess of counts in on-resonance data compared to off-resonance data, corresponding to a p-value $p = 2.8 \times 10^{-5}$ [116, 122].

Appendix B

Microwave Electric Field

This appendix summarizes key elements of a simple single particle theory that describes interaction of an electron plasma confined in a magnetic field $\mathbf{B}$ with a square amplitude-modulated pulse of microwaves tuned to the vicinity of the cyclotron resonance. It enables one to extract information about the amplitude of the microwave electric field from the measurement of the extent to which the plasma is heated by the pulse.

In a magnetic field $\mathbf{B}$ electrons undergo cyclotron motion at an angular frequency $\omega_c = qB/m$ where q and m refer to the magnitude of the electron charge and mass, respectively. When a microwave field is applied at an appropriate frequency it can drive cyclotron motion and store kinetic energy in the transverse motion of the electrons. To quantify the work that is done, we start from the equation of motion for a single charged particle in the presence of external electric $\mathbf{E}$ and magnetic $\mathbf{B}$ fields:

$$m\frac{d\mathbf{v}}{dt} = q\mathbf{E} + q\mathbf{v} \times \mathbf{B}. \tag{B.1}$$

One can decompose the transverse (x and y) components of the microwave electric field and the particle velocity into components that co- and counter-rotate with respect to the cyclotron motion. That is:

$$E_{\pm}(t) = E_x(t) \pm iE_y(t), \tag{B.2}$$

and

$$v_\pm(t) = v_x(t) \pm iv_y(t). \tag{B.3}$$

The particle equation of motion (Eq. B.1) can then be rewritten:

$$\frac{dv_\pm(t)}{dt} = \frac{q}{m}E_\pm(t) \mp i\omega v_\pm(t). \tag{B.4}$$

If the microwave pulse duration is short compared to damping and collisional timescales of the plasma, the solution to Eq. B.4 is of the form:

$$v_\pm(t) = \left[v_\pm(t_0)e^{\pm i\omega t_0} + \frac{q}{m}\int_{-\infty}^{t} e^{\pm i\omega t'} E_\pm(t')dt' \right] e^{\mp i\omega t}. \tag{B.5}$$

where t_0 marks the time at which the interaction with microwaves begins and $E_\pm = 0$ for $t < t_0$. The average change in transverse kinetic energy $\langle KE_\perp \rangle = \frac{1}{2}m\langle v_+ v_- \rangle$ for an ensemble of particles is then

$$\langle KE_\perp \rangle = \frac{q^2}{2m} \left| \int_{-\infty}^{\infty} E_+(t')e^{i\omega t'} dt' \right|^2 \tag{B.6}$$

where $E_+ = E_x(t) + iE_y(t)$ is the co-rotating component of the microwave electric field. This deposited energy result in a temperature change of

$$\Delta T = \frac{2}{3k_B} \langle KE_\perp \rangle. \tag{B.7}$$

For an amplitude-modulated microwave pulse with a square envelope, the components of the co-rotating component of the microwave electric field E_+ can be expressed:

$$E_x(t) = E_{x,0} \cos(\omega_0 t) \left[H(t + \tau/2) - H(t - \tau/2) \right], \tag{B.8}$$

$$E_y(t) = E_{y,0} \cos(\omega_0 t + \delta_y) \left[H(t + \tau/2) - H(t - \tau/2) \right], \tag{B.9}$$

where ω_0 is the microwave frequency, τ is the pulse width, and H is the Heaviside step function. The integral appearing in Eq. B.6 will thus have the form

$$\int_{-\infty}^{\infty} E_+(t')e^{i\omega t'} dt' = \left(\frac{\sin\left[(\omega_0 - \omega)\tau/2\right]}{\omega_0 - \omega} + \frac{\sin\left[(\omega_0 + \omega)\tau/2\right]}{\omega_0 + \omega} \right) E_0, \tag{B.10}$$

where $E_0 = E_{x,0} + iE_{y,0}e^{i\delta_y}$ is the amplitude of the co-rotating component of the microwave electric field. Near resonance $\omega \simeq \omega_0$, and so $\Delta\omega = \omega_0 - \omega \ll \omega_0 + \omega$. In this case Eq. B.10 can be approximated as

$$\int_{-\infty}^{\infty} E_+(t')e^{i\omega t'}\, dt' = \frac{\tau}{2}\text{sinc}\left(\frac{\Delta\omega\tau}{2}\right) E_0, \tag{B.11}$$

and the change in transverse kinetic energy of the particles (Eq. B.6) can be rewritten as

$$\langle KE_\perp \rangle = \frac{q^2\tau^2}{8m}\text{sinc}^2\left(\frac{\Delta\omega\tau}{2}\right)|E_0|^2. \tag{B.12}$$

Combining Eqs. 4.7, B.7, and B.12 one can solve for the amplitude of the co-rotating component of the microwave electric field in terms of quadrupole frequency change Δf_2 of the plasma for a square pulse of microwaves applied at the cyclotron resonance frequency:

$$|E_0| = \frac{2\sqrt{3mk_B\Delta f_2/\beta}}{q\tau} \tag{B.13}$$