

On the Dynamics of Adiabatically Cooled Antihydrogen in an Octupole-Based Ioffe-Pritchard Magnetic Trap

A thesis submitted to The University of Manchester for the degree of Doctor of Philosophy
in the Faculty of Science and Engineering.



2022

Danielle L. Hodgkinson

School of Natural Sciences/Department of Physics and Astronomy

Contents

List of Tables	3
List of Figures	6
Abstract	22
Declaration	23
Copyright	24
Acknowledgements	25
1 Motivation and Background	28
1.1 Motivation for Antimatter Research	28
1.2 The ALPHA Apparatus	35
1.3 Antihydrogen Spectroscopy in ALPHA-2	45
1.4 Measuring Antihydrogen Gravity in ALPHA-g	50
1.5 Bounding the Charge of Antihydrogen in ALPHA-2	53
1.6 Antihydrogen Cooling Techniques	54
1.7 Outline of this Thesis	56
1.8 Author Contributions	57
2 Adiabatic Cooling Experimental Procedure	59
3 Theory	72
3.1 Antihydrogen Confinement in Octupole-Based Magnetic Minimum Traps	72
3.2 Principles of Adiabatic Expansion	79
3.3 Fundamentals of Energy Mixing in Magnetic Traps	83
3.4 Simulating Antihydrogen Trajectories	90
4 Simulation Results	101
4.1 Kolmogorov-Smirnov (KS) Statistics	102

4.2	Possible Explanations for Temporal Shifts between Experiment and Simulation	107
4.3	Energy of Adiabatically Cooled Antihydrogen	116
4.4	Estimating Energy Uncertainty from Temporal Offset	119
4.5	Reconstructing Energy Distribution from Annihilation-Time Distribution . . .	124
4.6	Final Energy Method Comparison	125
4.7	Convergence Studies	126
4.8	Octupole FRD as an Adiabatic Expansion Diagnostic	127
5	Adiabatic Cooling Models	136
5.1	Derivation of Adiabatic Cooling Models	136
5.2	Trap Volume and Axial Trap Length	139
5.3	Adiabatic Modelling Results	145
6	Optimisation and Applications of Adiabatic Cooling	148
6.1	Modifying the Expansion Rate	148
6.2	Charge Neutrality and Adiabatic Expansion	159
6.3	Compatibility of Adiabatic Cooling and 1S-2S Spectroscopy	173
7	Conclusion	181
7.1	Future Work	185
A	Infinite Octupole Model	188
B	Generic Adiabatic Expansion and Contraction	189
	REFERENCES	194

Word Count: 56211

List of Tables

2.1	Number of antihydrogen (denoted $\bar{\text{H}}$) annihilation events before FRD ($t < 0$ s), and during the ~ 1.5 s Octupole FRD detection window, N . Average detected annihilation time, $\langle t_a \rangle$, and ensemble-averaged trap depth at annihilation time, $\langle \Gamma(t_a) \rangle$, are also shown. The background rate is (53 ± 4) mHz or 0.239 ± 0.018 expected cosmic ray events during diagnostic ramp down. $\sigma_{\text{SE}} = \sigma/\sqrt{N}$ is the standard error on the mean annihilation time, $\langle t_a \rangle$, where σ is the sample standard deviation.	69
4.1	Typical simulation parameters for Adiabatic Expansion Trial A, Control Trial B and Control Trial C.	102
4.2	KS Test p value comparison for the three experimental trials. The “Linear ramp down” row indicates the KS p value result between experimental and simulated annihilation times based on a linear Octupole ramp down. The “DCCT Trace (A used for C)” row indicates the KS p value result between experimental annihilation times and simulated annihilation times, obtained using a modified ramp down in which the current as a function of time follows that recorded by a DCCT during the experiment. “(A used for C)” is used to remind the reader that I modify the Trial A and C simulations using the Trial A DCCT trace, whereas I modify the Trial B simulation using the Trial B trace, since the Trial C trace was not recorded. The results presented in this table can be interpreted as the p values at the intersection of the vertical black line with the three curves in Figure 4.7(a) and 4.7(b) for rows 1 and 2 respectively.	113
4.3	Ensemble averaged total energy, $\langle E(t) \rangle$, axial energy, $\langle E_{\parallel}(t) \rangle$, and transverse energy, $\langle E_{\perp}(t) \rangle$, for the three experimental trials at the start, t_i , and end, t_f , of adiabatic expansion (or corresponding control hold). $\sigma_{\text{SE}} = \sigma/N_{\text{sim}}$ is the standard error on the mean, where N_{sim} is the number of simulated trajectories (shown in Table 4.1) and σ is the sample standard deviation.	118

4.4	Ensemble-averaged trap depth at simulated annihilation time, $\langle \Gamma(t_{a,\text{sim}}) \rangle$, for Adiabatic Expansion Trial A, Control Trial B and Control Trial C. $\sigma_{\text{SE}} = \sigma/\sqrt{N_{\text{sim}}}$ is the standard error on the mean, where σ is the sample standard deviation and N_{sim} is the number of simulated antihydrogen atoms (see Table 4.1). I also show the range of temporal offsets, $\Delta\mathcal{T}$, of the experimental relative to simulated annihilation times for which the KS p value is > 0.05 (green region in Figure 4.7), which indicates reasonable agreement between simulated and experimental annihilation time distributions.	120
4.5	Range of ratios of ensemble-averaged final energy in Control Trial B and Control Trial C to Adiabatic Expansion Trial A, based on an energy correction applied using the method described in the text, which approximates the energy shift which leads to agreement between simulated and experimental annihilation times.	122
4.6	Reconstructed energy just prior to magnetic ramp down, $\langle \mathcal{E}(t_f) \rangle$, from experimental annihilation time and simulated map between annihilation time and energy just prior to magnetic ramp down, $E(t_f)$, using the inverse method described in [1] with a band-width of 0.01 s and sampling 100 samples per band. $\sigma_{\text{SE}} = \sigma/\sqrt{N}$ is the standard error on the mean, σ is the sample standard deviation and N is the total number of experimental antihydrogen counts during diagnostic ramp down (see Table 2.1).	124
4.7	Table comparing final energy predicted by naive method, $\langle \Gamma(t_a) \rangle$, final energy extracted from the simulation, $\langle E(t_f) \rangle$, and reconstructed energy, $\langle \mathcal{E}(t_f) \rangle$. The naive method involves calculating the trap depth at each experimental annihilation time, $\Gamma(t_a)$, and taking the ensemble average (denoted by an angled bracket). The energy extracted from the simulation involves determining the total energy (sum of Equation 3.33 and Equation 3.34) of an individual particle as it crosses the magnetic minimum. $\langle \mathcal{E}(t_f) \rangle$ is determined using the method described in Section 4.5. $\sigma_{\text{SE}} = \sigma/\sqrt{N}$ is the standard error on the mean, σ is the sample standard deviation and N is the total number of experimental antihydrogen counts during diagnostic ramp down (see Table 2.1).	125

5.1	Percentage change of ensemble averaged energies between the referenced simulation (<i>Simulation</i>) and the adiabatic model, Model 1 and Model 2) just before the dump ($t_f = -50$ ms). Quasi-trapped particles that are lost between t_i and t_f are excluded from the ensemble-average. Stated errors are propagated to the percentage ratios from the standard error on the sample means. No error is stated for the transverse ratio for Model 1 since the model assumes no energy mixing and axial expansion only and therefore the transverse energy is exactly constant.	145
5.2	Percentage change in ensemble averaged energy components in a simulation of adiabatic expansion Trial A and as predicted by adiabatic models, Model 1 and Model 2. Note that the first three rows are identical to that of Table 5.1; the third row gives the prediction of the modified adiabatic model that does not include radial compression, Model 2*. Stated errors are propagated to the percentage ratios from the standard error on the sample means. No error is stated for the transverse ratio for Model 1 since the model assumes no energy mixing and axial expansion only and therefore the transverse energy is exactly constant.	147
6.1	Comparison of simulation parameters for “No AE Charge Protocol” and “AE Charge Protocol” as defined in Figure 6.14. $I_{\text{mirr}}(t_1)$ and $I_{\text{oct}}(t_1)$ are the currents in Mirrors A and E at t_1 (as indicated by the dashed magenta line in Figure 6.14). Γ is the magnetic trap depth (defined as the lowest potential on the electrode surface). Γ_z is the minimum-maximum axial confining potential of the magnetic trap (the potential that must be overcome for an anti-atom to escape axially over the confining Mirror Coil). $\langle E(t_1) \rangle$ is the ensemble-averaged energy at t_1 , excluding particles that annihilate at $t < t_1$. $\sigma_{\text{SE}} = \frac{\sigma}{\sqrt{N}}$ is the standard error on the mean, where σ is the sample standard deviation and N is the number of simulated particles surviving until $t = t_1$ for $Q = 0$	171

List of Figures

- 1.1 Cross-sectional view of the ALPHA apparatus in 2018. Antiprotons enter the apparatus from CERN’s Antiproton Decelerator (as indicated by arrow), and positrons are obtained from a radioactive source of ^{22}Na (enter apparatus as indicated by arrow). The apparatus is under vacuum, and regions of different vacuum levels are indicated by the coloured regions specified in the legend. Antiproton and positron plasmas are guided by magnetic fields to the two Atom Traps (ALPHA-2 and ALPHA-g). Here, the two plasma species are merged to form neutral antihydrogen which is subsequently confined in an Ioffe-Pritchard magnetic trap. Note the direction of the gravitational force is indicated by an arrow labelled ‘g’. Figure from [2], courtesy of Mark Johnson. 35
- 1.2 (a) Diagram of the ALPHA-2 Atom Trap (AT), drawn to scale (axial scale lines up with that of (b) below) except from the annihilation detector radius. Antiprotons (denoted \bar{p}) enter from the left, whereas electrons (denoted e^-) and positrons (denoted e^+) enter from the right. (b) On-axis magnetic field strength as a function of axial position in the ALPHA-2 AT, where the axial position maps directly to the diagram in (a). Red solid curve gives the magnetic field strength with the flat trap energised, which is defined in Figure 2.3(c). Blue dashed curve gives the magnetic field strength during stacking where, in addition to the flat trap, Solenoid A and Solenoid B are energised to increase the magnetic field for enhanced capture, cyclotron cooling and rotating wall efficiency for charged species. Figure from [3]. 37

- 1.3 Diagram of ALPHA's positron source and Surko-style Positron Accumulator. Beneath the Positron Accumulator, the electrostatic potential generated by the Penning-Malmberg trap electrodes within the buffer gas volume is shown, where each of the three stages has a decreasing electrostatic potential, as a function of pressure, P . Positrons lose energy through collisions with a nitrogen buffer and accumulate in Stage 3 (lowest pressure stage). Note this is a diagram of ALPHA's positron source and accumulator from 2014, which differs from the current apparatus. The differences are local to the positron source (the Positron Accumulator end is the same): minor modifications were made to the coldhead (rotated), the guiding magnets and the vacuum chamber [4]. Figure from [5]. 40
- 1.4 Electrostatic potential manipulations for positron and antiproton mixing to form antihydrogen. In each subfigure, the on-axis electrostatic potential as a function of axial position is given, which can be shifted relative to the minimum of the antihydrogen trap and therefore the absolute value here is arbitrary but relative axial positions are relevant. (a) Before evaporative cooling. (b) Evaporative cooling. (c) Realignment such that the two potential wells almost overlap. (d) Antiproton and positron plasmas are mixed during 1 s. Dashed and solid curves represent electrostatic potentials before and after each step respectively. Purple and green regions indicate self-potentials and physical extents of positron and antiproton plasmas respectively. Figure from [3], details provided in [3]. 41
- 1.5 Radial profile diagram of the ALPHA-2 annihilation detector: the Silicon Vertex Detector (SVD). The letters label (a) Penning-Malmberg electrode wall, (b) Magnetic minimum trap superconducting magnet winding form, (c) Octupole windings, (d) liquid helium volume, (e) vacuum wall between helium volume and outer vacuum chamber (OVC), (f) vacuum wall between OVC and air, (g) three layers of silicon modules and (h) External Solenoid. An example antiproton annihilation vertex is indicated by a yellow star which results in tracks in the SVD of two charged pions, π^- and π^+ , and one neutral pion, π^0 , which quickly decays to form two photons, one of which is absorbed by the Octupole windings and the other undergoes electron-positron pair production. Figure from [6]. 43

1.6	Energy levels of the principal quantum number, $n = 1$ and $n = 2$, states in antihydrogen as a function of magnetic field. The first value in the ket notation represents the projection of the total angular momentum of the positron, $m_L + m_S$, where L is the orbital angular momentum (in S states $L = 0$ and in P states $L = 1$) and $S(= 1/2)$ is the spin. The second value in the ket notation represents the antiproton spin up (\uparrow) or down (\downarrow). Transitions stimulated in ALPHA experiments are indicated by solid arrows and the probable decays from those excited states are shown as dotted arrows. Energy levels at zero-field are shown in spectroscopic notation: nL_j , where j is the total angular momentum quantum number. At high field, $L = 0$ (S state) hyperfine levels are indicated by a subscript letter, whereas for $L = 1$ (P states), fine structure levels are indicated by a subscript letter and hyperfine levels are indicated by a further $+/-$ subscript. Figure adapted from [7].	46
1.7	ALPHA-g superconducting magnets, radial extent not to scale. The direction of the gravitational force is indicated by an arrow labelled ‘g’ and the trap z -axis (parallel to the gravitational force) is indicated by an arrow labelled ‘z’. Figure courtesy of Chukman So.	51
1.8	Illustration of gravity compensation in ALPHA-g. Shown is the total potential experienced by an antihydrogen atom (Equation 1.13) for an example value of F (exaggerated with respect to the hydrogen value $F = 1$). For this particular value of F in the <i>no compensation</i> potential, antihydrogen atoms would be more likely to escape at lower z . With correct <i>gravity compensation</i> , antihydrogen annihilations will be z -symmetric. Figure drawn as an indicator, not to scale. The ticks on the x -axis are labelled to indicate the separation between Mirror A and Mirror G in the up-down measurement trap.	52
2.1	ALPHA-2 neutral atom trap (AT). Penning-Malmberg trap electrodes (yellow) confine antiprotons and positrons axially. In Trials A and B, antihydrogen is formed at E11 (pink), whereas in E12 (orange) for Trial C. AT magnets are coaxial to the electrode stack: short solenoids, known as Mirrors A-E (blue), Octupole (red) and Solenoids A and B (green). The External Solenoid surrounds the entire AT and is not shown here. The z -axis is defined along the trap axis as shown.	59

2.2	On-axis magnetic potential, $U(r = 0, \phi, z)(\text{K})$ (red solid curve) and on-axis electrostatic potential, $U_E(r = 0, \phi, z)(\text{V})$, just prior to antihydrogen formation (black solid curve, corresponding to Figure 1.4(c)) and after mixing of positron (e^+) and antiproton (\bar{p}) plasmas for antihydrogen formation (black dotted curve, corresponding to Figure 1.4(d)). Magnetic potentials at the point of antihydrogen formation ((a) -23.55 s for Adiabatic Expansion Trial A and Control Trial B, (b) -24.523 s for Control Trial C). $U(r = 0, \phi, z)$ is relative to the magnetic minimum and is found using the method described in Section 3.4. $U_E(r = 0, \phi, z)$ is calculated using the COMSOL Multiphysics [®] potential solver. The green dotted line represents the center of the electrode in which positrons are held during mixing (E11 in (a) and E12 in (b)).	61
2.3	Typical on-axis confinement potentials, $U(r = 0, \phi, z)$, in ALPHA-2. In each case the Octupole is energised to ~ 900 A, which produces a mostly radial magnetic field, generating a magnetic potential of magnitude ~ 0.5 T at the trap wall, which provides radial confinement. (a) Standard trap configuration with Mirrors B-D at 0 A and Mirrors A and E at ~ 600 A, (b) Tight configuration with Mirrors A and E at ~ 600 A, Mirrors B and D at ~ 700 A and Mirror C at ~ -200 A, (c) Flat trap configuration with Mirrors A and E at ~ 600 A, Mirrors B and D at ~ -60 A and Mirror C at ~ -2 A.	62
2.4	Currents in each magnet of the neutral trap as a function of time during the adiabatic expansion experiment for experimental trials (a) Adiabatic Expansion Trial A, (b) Control Trial B and (c) Control Trial C. Times are given relative to the start of Octupole FRD and currents are shown until the simulation end time (around 7 s).	63
2.5	On-axis magnetic potential, $U(\mathbf{x}, t)$, as a function of axial position, z , at $x = y = 0$, for the three experimental trials. The y -axis label specifies the scale on each small graph.	65
2.6	Annihilation time data of the three experimental trials: (a) Adiabatic Expansion Trial A (98 counts), (b) Control Trial B (106 counts) and (c) Control Trial C (77 counts). Times are given relative to the start of Octupole FRD and green errorbars are given by Poissonian counting statistics (\sqrt{N} , where N is the number of detected counts). The average annihilation time, $\langle t_a \rangle \pm 1\sigma_{\text{SE}}$ is shown as a vertical magenta line, and the region within $1\sigma_{\text{SE}}$ is shown as a magenta region where $\sigma_{\text{SE}} = \sqrt{\sigma/N}$, is the standard error on the mean and σ is the sample standard deviation.	67

2.7	Trap depth, Γ , as a function of time, t (s), during Octupole FRD for (a) Trial A, (b) Trial B and (c) Trial C. Trap depth is defined as the minimum magnetic potential on the electrode inner surface relative to the minimum magnetic potential anywhere in the trap volume. Octupole FRD starts at 0 s in all three trials and the current decay is assumed to be perfectly linear. Trap depth is defined as the minimum magnetic potential on the inner electrode surface that is accessible from the minimum magnetic potential within the trapping volume.	69
2.8	Distribution of energies representing the trap depth (Figure 2.7) at the experimental annihilation time, $\Gamma(t_a)$. White error bars show Poissonian (\sqrt{N}) counting statistics, where N is the number of experimental annihilation counts (see Table 2.1). The ensemble-averaged energy at annihilation time, $\langle\Gamma(t_a)\rangle \pm \sigma_{SE}$, is shown as a vertical yellow line, where $\sigma_{SE} = \sigma/\sqrt{N}$ is the standard error on the mean and σ is the sample standard deviation. The yellow region indicates energies within $1\sigma_{SE}$ of $\langle\Gamma(t_a)\rangle$. (a) Adiabatic Expansion Trial A ($\langle\Gamma(t_a)\rangle = 0.081 \pm 0.007$ K), (b) Control Trial B ($\langle\Gamma(t_a)\rangle = 0.223 \pm 0.009$ K), and (c) Control Trial C ($\langle\Gamma(t_a)\rangle = 0.173 \pm 0.011$ K).	70
3.1	The effect of Zeeman splitting on the hyperfine structure of the antihydrogen ground state, $^1S_{1/2}$. The positron spin is indicated by \uparrow / \downarrow and the antiproton spin is indicated by $\uparrow\uparrow / \downarrow\downarrow$. The yellow region indicates the range of magnetic fields present in a typical ALPHA antihydrogen confinement potential. Figure from [8].	73
3.2	ALPHA-2 Octupole magnetic field components, (a) radial, B_r , and (b) azimuthal, B_ϕ , in the xy -plane at $z = 0$. The solid black line shows the inner radius at the electrode wall, $r_{\text{trap}} \approx 2.2$ cm. We plot the semi-analytic solution for the Octupole field described in Section 3.4.2.	74
3.3	Typical ALPHA-2 magnetic confining potential for low-field-seeking ground state ($n = 1$) anti-atoms. Mirrors A and E are energised to 600 A, Mirrors B... D 0 A, Octupole energised to 900 A and a 1 T axial background field is present. The contour plot of $U(r, \phi = 0, z, t)$ is given as a function of z and r , where z is defined along the trap axis and r is the radius defined in cylindrical coordinates (Figure 2.1). Negative r , although unphysical, is shown to demonstrate that the potential is effectively a magnetic bowl. $\phi = 0$ is given for the simplicity of the plot, although the trap is azimuthally asymmetric at all $z \neq 0$ as shown in Figure 3.6.	75

3.4	Diagram representing two extrema of antihydrogen motion in the trap, where the red dot indicates the instantaneous location of the anti-atom. We show the on-axis magnetic potential, $U(r = 0, \phi, z)$, as a function of axial position, z , for simplicity although the anti-atom moves in a 3D magnetic potential in reality. At <i>Position 1</i> the anti-atom has exchanged all of its magnetic potential energy for axial kinetic energy. At <i>Position 2</i> the anti-atom's total energy is entirely magnetic potential as it is at the turning point of its trajectory high on the Mirror Coil field.	77
3.5	A particle oscillating axially between <i>Wall 1</i> and <i>Wall 2</i> . Wall 2 moves away from Wall 1 at constant speed u , causing a slow increase in the trap length, L , as a function of time. The particle has speed v_z in the axial direction where $v_z \gg u$	81
3.6	(a) Diagram of magnetic trap: Octupole (red), Mirror Coils A-E (blue), Solenoids A and B (green) and electrode stack (yellow). The electrode in which antihydrogen formation takes place for adiabatic expansion Trials A and C is shown in pink, whereas Trial B is shown in orange. An xy -plane in close proximity to Mirror E is grey. (b) $ \mathbf{B}(\mathbf{x}) $ in the xy -plane close to Mirror E with Mirror E not energised (c) with Mirror E energised to 600A. All magnetic fields shown are modelled using the semi-analytic models presented in Section 3.4.2.	84
3.7	Cumulative distribution function (CDF) of ε_{\parallel} for a simulated distribution of 3757 trajectories as a function of energy at different sample times, $\bar{t} = 1s, 10s, 100s, 1000s$ [9].	87
3.8	Time averaged parallel energy, $\langle\langle E_{\parallel} \rangle\rangle$ (a) and time averaged normalised axial energy, $\langle\langle \varepsilon_{\parallel} \rangle\rangle$ (b) as a function of total anti-atom energy, E . Each trajectory was simulated for 1000s and all energies are in Kelvin. The red line (estimate of the location of minimum density between the two peaks in density) is drawn to separate two categories of trajectory which exhibit different energy mixing dynamics: those above the red line exchange axial and transverse energy quickly, whereas those below the red line do not. Figure from [9].	88
3.9	$\langle\langle \varepsilon_{\parallel} \rangle\rangle$ as a function of time, t , for 100 trajectories. Each coloured line is an individual simulated trajectory.	89

3.10	CDF of the time, t_{\downarrow} , taken for antihydrogen trajectories to change their initial normalised axial energy, $\varepsilon_{\parallel kj}$ by 0.2. The CDF is shown for different trap configurations, for which $ \mathbf{B} $ as a function of z ($r = 0$) is shown below the figure. The ‘quadrupole’ configuration represents the standard trap but with the Octupole magnet replaced by a quadrupole. Solid and dotted lines represent simulated trajectories without and with cascade initialisation respectively. Figure from [9].	90
3.11	Initialised antihydrogen (a) speed, $v(\text{m/s})$, (b) kinetic energy, $E_K(\text{K})$, distributions. In both subfigures the black curves represent the distribution given by the theory and the red histograms are the sampled distributions (10000 total samples) from the simulation.	93
3.12	Four points lie at time t in phase space (\mathbf{x}, \mathbf{p}) and are each mapped by the symplectic Leapfrog algorithm to a point in phase space $(\mathbf{x}', \mathbf{p}')$ at time t' . The quadrilaterals formed by connecting each of the points before and after application of the Leapfrog algorithm have area dA and dA' respectively. Figure from [10].	95
3.13	Diagram of the shape of the model for Solenoid A (can equally be used for Solenoid B). The model consists of a main Solenoid body and two end corrector loops that are inset by a distance, ΔZ_{loop} , from the ends of the main body. The pink region continues off to the right of the figure and indicates the approximate region in which we are required to calculate the magnetic field of Solenoid A (the antihydrogen trapping region).	100
4.1	Comparison of simulated (black line plot) and experimental annihilation time distributions (red, blue and green histograms) during Octupole FRD for the three experimental trials: (a) Adiabatic Expansion Trial A, (b) Control Trial B and (c) Control Trial C. The simulated annihilation distribution is normalised to the area of the detected annihilation histogram. Error bars (white) represent Poissonian (\sqrt{N} , where N is the number of experimental counts, see Table 2.1) counting statistics on experimental data. Similar error bars for the simulated annihilation time data are not visible on this scale. The average annihilation time, $\langle t_a \rangle \pm 1\sigma_{\text{SE}}$, for the simulated (vertical cyan line, $1\sigma_{\text{SE}}$ region not visible on this scale) and experimental (vertical magenta line, $1\sigma_{\text{SE}}$ magenta region) distributions are also shown, with uncertainty given by the standard error on the mean, $\sigma_{\text{SE}} = \sigma/\sqrt{N}$, where σ is the sample standard deviation. Simulation parameters are given in Table 4.1.	103

4.2	Comparison of simulated and experimental cumulative distribution function of annihilation times, $S_N(t)$, where I have simulated a linear magnetic ramp down with simulation parameters given in Table 4.1. Simulated cumulative distribution function of annihilation, $S_{N_{\text{sim}}}(t)$, is shown for (a) Trial A (red solid curve), (b) Trial B (blue dashed curve) and (c) Trial C (green dotted curve). In each subfigure, the experimental cumulative distribution function of annihilation time, $S_{N_{\text{exp}}}(t)$, is shown as a black solid curve. Above each subfigure I show the KS p value, which represents the confidence level to which I can conclude the simulated and experimental annihilation times are drawn from the same distribution. The maroon dashed line indicates the largest vertical deviation between the two cumulative distribution functions, or the position of the KS statistic, D	105
4.3	KS p value as a function of timing offset in seconds. Timing offset is defined as the offset made to each experimental data point with respect to the simulation i.e. a high KS p value with a negative offset suggests the experimental annihilation times tend to occur <i>later</i> than the simulated annihilation times. Trials A, B and C are displayed as red solid, blue dashed and green dotted lines respectively. The red region indicates the range of p values (< 0.05) at which the null hypothesis (simulated and experimental annihilation times are drawn from same distribution) is rejected, whereas the green region indicates the range of p values at which the null hypothesis is accepted.	106
4.4	Electrical circuit of the superconducting Octupole magnet, shown as an inductor labelled <i>Magnet</i> . The Silicon-Controlled Rectifier (labelled <i>SCR</i>), Isolated-Gate Bipolar Transistor (labelled <i>IGBT</i>) and <i>Diode</i> act as controllable switches. In normal operation, the Magnet is superconducting, the SCR is open and the IGBT and Diode are closed, allowing the power supply to drive current into the Magnet. In the event of a quench in the Magnet, the SCR is closed and the IGBT and Diode are closed, shorting the power supply out of the circuit and enabling the stored energy in the Magnet to dissipate in the 0.3Ω resistor.	108

- 4.5 (a) Trial A Octupole current as a function of time relative to the start of Octupole FRD, red solid curve shows perfectly linear current decay, whereas black solid curve shows current recorded by DCCT. (b) Trial B Octupole current as a function of time relative to the start of Octupole FRD, blue dashed curve shows perfectly linear current decay, whereas black solid curve shows current recorded by DCCT. (c) Temporal difference, Δt , at which the Trial A Linear (red solid curve in (a)) and Trial A DCCT (black solid curve in (a)) have the same Octupole current, as a function of Octupole Current. (d) Temporal difference, Δt , at which the Trial B Linear (blue dashed curve in (b)) and Trial B DCCT (black solid curve in (b)) have the same Octupole current, as a function of Octupole Current. Note that I do not show the equivalent for Trial C since the DCCT data was not recorded during the experiment. (e) Trial A experimental annihilation data (red histogram), white error bars show Poissonian counting statistics, magenta region is the one standard error region around the ensemble averaged annihilation time, $\langle t_a \rangle$. (f) Trial B experimental annihilation data (blue histogram), white error bars show Poissonian counting statistics, magenta region is the one standard error region around the ensemble averaged annihilation time, $\langle t_a \rangle$ 111
- 4.6 Cumulative distribution function of annihilation time, $S_N(t)$, as a function of time, t (s), relative to the start of Octupole ramp down (at $t = 0$ s) for (a) Adiabatic Expansion Trial A, (b) Control Trial B and (c) Control Trial C. $S_N(t)$ for the experimental data is labelled ‘experiment’ (black curve), $S_N(t)$ for the primary simulation in which the Octupole current ramps linearly from 900.1 A to 0 A in ~ 1.53 s is labelled ‘Linear FRD simulation’ (cyan curve), and $S_N(t)$ for the modified simulation in which the Octupole current during FRD follows the high resolution DCCT trace is labelled ‘DCCT Modified Simulation (magenta curve)’. 112

- 4.7 Kolmogorov-Smirnov (KS) p value as a function of temporal offset of experimental annihilation time data with respect to simulated annihilation times for Trial A (red solid curve), Trial B (blue dashed curve) and Trial C (green dotted curve) (a) with a linear Octupole ramp down (b) with a ramp down which follows the current measured by the DCCT during experimental ramp down. Note that in (b) Trials A and C use the DCCT trace obtained during Trial A (since the Trial C trace was not recorded) and Trial B uses the DCCT trace obtained during Trial B. The red region indicates the range of p values (< 0.05) at which the null hypothesis (simulated and experimental annihilation times are drawn from same distribution) is rejected, whereas the green region indicates the range of p values at which the null hypothesis is accepted. 115
- 4.8 Histograms of simulated antihydrogen energies at t_i (orange) and t_f (blue) for Trial A (a) E , (b) E_{\parallel} and (c) E_{\perp} , Trial B (d) E , (e) E_{\parallel} and (f) E_{\perp} , and Trial C (g) E , (h) E_{\parallel} and (i) E_{\perp} . Vertical dashed orange and blue lines show the trap depth, Γ , at t_i and t_f respectively. Quasi-trapped particles (particles that annihilate before $t = 0$ s) are excluded, but some anti-atoms that have energy greater than the trap depth survive until $t = 0$ s. Simulation parameters are shown in Table 4.1. 117
- 4.9 Ensemble-averaged (a) total energy, $\langle E \rangle$, (b) axial energy, $\langle E_{\parallel} \rangle$, and (c) transverse energy, $\langle E_{\perp} \rangle$, for simulations of Trial A (red solid), Trial B (blue dashed) and Trial C (green dotted) antihydrogen atoms as a function of time relative to the start of the Octupole ramp down. (d), (e) and (f) are energy-zoomed versions of (a), (b) and (c) respectively. Distribution averages are calculated every 250 ms. The energy of an individual particle is defined as its energy at the last z_{\min} crossing before each 250 ms threshold. Quasi-trapped particles are excluded. Standard error bars have been omitted since they are too small to see. Simulation parameters are shown in Table 4.1. 119
- 4.10 (a) Simulated ensemble-averaged final energy, $\langle E(t_f) \rangle$ for Adiabatic Expansion Trial A (red line), Control Trial B (blue line) and Control Trial C (green line). Standard error bars are not visible on this scale. (b) and (c) Range of energies, $\langle E(t_f) \rangle + \Delta \langle E(t_f) \rangle$ (see Equation 4.5), for which simulated and experimental annihilation time distributions have KS p values > 0.05 , are shown as coloured regions (red, blue and green for Trials A, B and C respectively) for (b) the primary linear FRD simulation and (c) the DCCT modified simulation. 121

4.11	Antihydrogen reconstructed energy distribution function, $f(\mathcal{E}(t_f))$ (black solid curve), as a function of anti-atom energy, $E(t_f)$, where the energy is reconstructed just prior to magnetic ramp down (at t_f) and is shown for (a) Adiabatic Expansion Trial A, (b) Control Trial B and (c) Control Trial C. The simulated energy distribution function, $f(E(t_f))$, is also shown for Trial A (red histogram), Trial B (blue histogram) and Trial C (green histogram). $f(\mathcal{E}(t_f))$ was determined using the inverse method described in [1] with a band-width of 0.01 s and sampling 100 samples per band. The average reconstructed energy, $\langle \mathcal{E}(t_f) \rangle \pm 1\sigma_{SE}$ is indicated by the vertical purple line, where $\sigma_{SE} = \sigma/\sqrt{N}$ is the standard error on the mean, σ is the sample standard deviation and N is the total number of experimental antihydrogen counts during diagnostic ramp down (see Table 2.1).	123
4.12	Absolute differences (%) relative to $dt = 3.5 \times 10^{-7}$ s, $\Delta\langle t_a \rangle$ (solid magenta), and $\Delta\langle E(t_f) \rangle$ (solid cyan), as a function of simulation timestep, dt . Each black cross represents a simulation of 200,000 Trial A trajectories with fixed timestep, dt	127
4.13	Energy components for mix (red) and no-mix (black) simulated antihydrogen trajectories, where trajectories are categorised based on their normalised axial energy 50 ms prior to the onset of Octupole FRD (mix have $\epsilon_{ } \geq 0.1$, and no-mix have $\epsilon_{ } < 0.1$). Axial, $E_{ }$ (a), (c) and (e), and transverse E_{\perp} (b), (d) and (f), energy components for Trials A, B and C respectively.	128
4.14	Axial annihilation location, z , (left) and annihilation time relative to the start of Octupole FRD, t , (right) for simulated Trial A [(a) and (b)], simulated Trial B [(e) and (f)] and simulated Trial C [(i) and (j)] respectively. Shown below each simulated z histogram is the predicted (labelled ‘‘Method Prediction’’) z annihilation distribution based on $E_{ }$ of the simulated trajectories 50 ms before the start of Octupole FRD and $U_z(z)$ (Figure 4.15) for Trial A (c), Trial B (g) and Trial C (k). Shown below each t histogram is the predicted (labelled ‘‘Method Prediction’’) t annihilation distribution based on E_{\perp} of the simulated trajectories 50 ms before the start of Octupole FRD and $\Gamma_r(\phi, z)$ (Figure 4.16) for Trial A (d), Trial B (h) and Trial C (l). ‘Mix’ and ‘no-mix’ categories are shown as red and black histograms respectively.	130
4.15	On-axis axial potential, $U_z(z)$, as a function of axial position, z , during Octupole FRD for Trial A (a), Trial B (b) and Trial C (c).	131

4.16	Radial trap depth $\Gamma_r(\phi, z, t)$ (K), as function of axial position, z , and azimuthal angle, ϕ , at different times (relative to start of Octupole FRD) (a) 0.0 s, (b) 0.5 s and (c) 1.0 s during Adiabatic Expansion Trial A, and equivalently during Control Trial C.	133
5.1	Potential, $U(r_{\text{trap}}, \phi, z)$ (K), on electrode wall ($r_{\text{trap}} = \sim 2.2$ cm) as a function of azimuthal angle, ϕ (rad), and axial position, z (mm). (a) Trial A $t = -50$ ms (before start of Octupole FRD) and (b) Trial A $t = 450$ ms (during Octupole FRD). Both subfigures are scaled to the same colour axis.	140
5.2	Ratios of trap volume and length at the end, t_f , to the start, t_i , of the three trials of adiabatic expansion as a function of anti-atom total energy, $E(K)$. (a) Trap volume ratio, $\frac{V(t_f)}{V(t_i)}$, and (b) axial trap length ratio, $\frac{L_z(t_f)}{L_z(t_i)}$, are shown for Trial A (red solid), Trial B (blue dashed) and Trial C (green dotted). The trap volume is calculated using the algorithm described in the text and the axial trap length ratio is obtained by averaging the length of this trap volume over radius. Cyan and magenta dotted lines in (a) are used to indicate the energies referenced in Figure 5.3.	142
5.3	Trappable volume in Cartesian coordinates, (x, y, z) , accessible by a particle of energy E at time t during adiabatic expansion Trial A, found by recursively searching a grid of 500 equally spaced points. Since the trap has larger axial extent than radial extent, the distance between grid points in the x and y directions is $dx = dy = 80 \mu\text{m}$ and in the z direction is $dz = 800 \mu\text{m}$. (a) $E = 0.01$ K, $t = t_i$ (b) $E = 0.01$ K, $t = t_f$, (c) $E = 0.4$ K, $t = t_i$, (d) $E = 0.4$ K, $t = t_f$	143
5.4	The <i>Mexican Hat Potential</i> in the ‘standard’ trap configuration (Mirrors A and E energised to 600 A, internal Mirror Coils at 0 A and the Octupole energised to 900 A) at $z = 0$. $ B(T) $ is shown on the vertical axis, with blue and red colours indicating smaller and larger relative values of $ B(T) $ respectively. The shape of the potential is dominated by the quadrature sum of the axial Mirror Coil field which is peaked at $r = 0$ and the radial Octupole field which has a minimum at $r = 0$	144
5.5	Ratio of area of trappable volume (averaged over axial position), $\frac{V(t_f)L_z(t_i)}{V(t_i)L_z(t_f)}$, at the end, t_f , to the start, t_i , of Trials A, B and C as a function of total anti-atom energy, $E(K)$. Note that $\frac{V(t_f)L_z(t_i)}{V(t_i)L_z(t_f)} < 1$ indicates radial compression over the course of the Trial.	146

6.1	(a) Superconducting magnet currents as a function of time during Adiabatic Expansion Trial A. (b) Superconducting magnet currents as a function of time during a 10.6s-long adiabatic expansion, where the timescale of the Solenoid A, Mirror B, Mirror C and Mirror D current ramps are scaled down by a factor of 2. Time, $t(s)$, is given relative to the onset of Octupole FRD.	149
6.2	Ensemble-averaged total energy, $\langle E(t) \rangle$ as a function of time, t , relative to the start of magnetic ramp-down for a 230.2s duration adiabatic expansion (red solid curve) and a 10.6 duration expansion followed by a hold in the standard trap configuration (defined in Figure 2.3(a)) (blue solid curve). Particles that annihilate before $t = 0s$ (quasi-trapped particles) have been excluded. 5000 trajectories were simulated in total, 4272 of which make it to $t = 0s$ for the 230.2s expansion, compared to 4576 for the 10.6s expansion.	150
6.3	Ensemble-averaged total energy, $\langle E(t) \rangle (K)$ as a function of scaled time, τ , in which the total expansion duration is scaled between 0 and 1. The initial energy of the 5000 simulated trajectories is truncated at 0.47 K (just below the magnetic trap depth) to prevent particle losses (quasi-trapped) before $\tau = 0$. .	151
6.4	Time-averaged normalised axial energy, $\langle \langle \epsilon_{\parallel} \rangle \rangle$ as a function of time, $t(s)$, during a 100s hold in the standard trap with mixing artificially turned off i.e. with an infinite length Octupole and the radial component of the Mirror magnetic field set to zero at all positions. Each line is an individual simulated trajectory. A non-constant $\langle \langle \epsilon_{\parallel} \rangle \rangle$ over time indicates energy mixing between axial and transverse dimensions.	153
6.5	(a) Ensemble-averaged total energy, $\langle E(t) \rangle (K)$, for simulated trajectories where mixing has been artificially turned off, as a function of scaled time τ , in which the total expansion duration is scaled between 0 and 1. The initial energy of the 5000 simulated trajectories is truncated at 0.47 K (just below the magnetic trap depth) to prevent particle losses (quasi-trapped) before $t = 0s$. Mixing is disabled by replacing the Octupole model with an infinitely long octupole and setting the radial field of the mirror coils to zero at all positions. (b) An energy-zoomed version of (a).	154

6.6	Diagram of a magnetic field sequence, in which antihydrogen trajectories are initialised in the short (tight) trap configuration at (1), before a slow adiabatic expansion takes place to arrive at the standard trap configuration at (2). After the expansion, the trap contracts back to the same tight configuration in (1) but the adiabatic contraction is faster than the slow expansion. More energy mixing is expected to occur during the slow expansion than the fast contraction. The length of the trap at each stage is indicated as a guide, although the true length of any one trajectory is energy and radius dependent.	155
6.7	Ratio between ensemble-averaged energies (green solid curve) at (3) and (1) in Figure 6.6, $\langle E_3/E_1 \rangle$, as a function of the ratio between trap lengths at (2) and (1), L_2/L_1 . I assume the trap expands only longitudinally and that energy equilibrates between all three dimensions between (2) and (3). Therefore, Equation 6.14 gives the functional form of the green curve. The blue region indicates an overall cooling achieved by the magnet manipulations shown in Figure 6.6, whereas the red region indicates overall heating. Note that $L_2/L_1 < 1$ is actually a slow contraction between (1) and (2) and a fast expansion between (2) and (3).	158
6.8	Ratio between ensemble-averaged energies at (3) and (1) in Figure 6.6, $\langle E_3/E_1 \rangle$, for a generic radial and axial trap expansion and an arbitrary degree of energy mixing between axial and transverse dimensions. Equation 6.16 governs the functional form of the contour plot, where I plot R and V_2/V_1 on the x and y axes respectively. The subplots represent different degrees of energy mixing where (a) $\langle f_2 \rangle = 0.1$, (b) $\langle f_2 \rangle = 1/3$ and (c) $\langle f_2 \rangle = 0.9$	160
6.9	Largest Lyapunov Exponent (LLE (s^{-1})) as a function of anti-atom total energy, $E(K)$, calculated for ~ 3000 trajectories simulated for 100 s in the standard trap configuration (Figure 2.3(a)). The points above the red line are chaotic since they have an LLE greater than 0. Figure from [9].	161
6.10	(a) A snapshot of the stochastic potential, $\phi(V)$, as a function of Cartesian position coordinate, $z(\text{mm})$, generated by the ALPHA-2 electrode stack on the inner surface of the electrode wall (green solid curve), on the trap axis (brown dashed curve) and at $r = 15.7 \text{ mm}$ (orange dashed curve). The sign of the potential shown in (a) is flipped in stochastic intervals as a function of time as shown in (b), where the sign of the electrostatic potential indicates the voltage on a single electrode. Note that (b) is just an example of the time dependence on the potential applied to a single electrode; for some types of electrode it is not possible to switch the sign of the potential as a square wave with a stochastic period (see [11] for details). Figure from [11].	163

6.11	Simulated survival probability, s , as a function of hypothetical antihydrogen charge, $ Q $, for the experimental protocol described in [11]. The orange band of varying thickness is the 1σ confidence region. Figure from [11].	165
6.12	Electrostatic potential at $r = 0$ (black solid line), at $r = 0.6R_{\text{wall}}$, where $R_{\text{wall}} \approx 2$ cm is the inner radius of the electrode wall. See [12] for a description of the orange dashed and orange dashed-dotted line as they are not relevant for the work discussed in this chapter. Figure adapted from [12].	166
6.13	Sketch of the configuration used to calculate the electrostatic potential generated by the n^{th} electrode raised to 1 V; with all other electrodes grounded. The inner radius of the electrode is r_w and the length of the n^{th} electrode is L_n . Figure from [13].	167
6.14	Definition of “No AE Charge Protocol” and “AE Charge Protocol”. Simulation time, $t(\text{s})$, is along the x -axis relative to the start of magnetic ramp down (at $t = 0$). A hashed region indicates that the simulation has not started. A pink region indicates a period of stochastic electrode switching with $\sigma/\bar{t} = 0.2$ and \bar{t} given by that in Table 6.1. The yellow region indicates a period of axial adiabatic expansion and trap depth reduction in which the Mirror Coil and Octupole currents are reduced to those given in Table 6.1. The on-axis magnetic potential as a function of time during the ‘AE + lower trap depth’ phase is shown in Figure 6.15. The dotted magenta line at t_1 indicates the start of the electrostatic field switching phase.	169
6.15	Snapshots of on-axis potential, U_z , as a function of axial location, z , at different times, t , during the AE Charge Protocol (see Figure 6.14). Time increments in 3 s intervals from (a)-(f). I_{oct} is the Octupole current at time t	170
6.16	(Total) Energy distributions at the start of the electrostatic field switching phase, $E(t_1)(\text{K})$, for the “No AE Charge Protocol” (black histogram) and “AE Charge Protocol” (blue histogram). The total histogram area is normalised to 1 in both cases.	172
6.17	Survival probability, s , (at $t = 0$) as a function of normalised anti-atom charge, $ Q $, (antihydrogen charge = Qe), relative to $Q = 0$. We compare the “No AE Charge Protocol” (black curve) to the “AE Charge Protocol” (blue curve). Note that a <i>lower</i> s at a given Q indicates an <i>increased</i> sensitivity to the hypothetical antihydrogen charge. The dotted red line indicates the observed survival fraction given that $Q = 0$	173

6.18	Diagram of the experimental setup for 1S-2S spectroscopy and laser cooling of antihydrogen. 121.6 nm laser light for antihydrogen laser cooling enters the trap at 2.3° to the trap axis, indicated by a purple beam. The 243.1 nm laser cavity for 1S-2S spectroscopy also at 2.3° to the trap axis is shown as a pink beam. Figure from [7], see [7] for a detailed description.	174
6.19	Definition of ‘No AE Protocol’ and ‘AE Protocol’. Each protocol refers to a specific set of magnetic field manipulations as indicated by the plots of magnetic potential, $U(K)$, at the bottom of the figure. The vertical dashed lines trace down to graphs of $U(K)$ which represent the status of the potential at that simulation time. Time, $t(s)$, is along the x -axis, relative to the start of magnetic ramp-down. A hatched region indicates that the simulation has not started at that time. The red region indicates a 60 s hold in the flat trap configuration, and the green region indicates a 23.5 s adiabatic expansion (labelled “AE”) of the same duration as Trial A, as defined in Chapter 2, but ending in the flat trap (Figure 2.3(c)) rather than the standard trap (Figure 2.3(a)).	175
6.20	Gaussian beam profile. ω_0 is the waist radius of the Gaussian beam, $\omega(z_{\text{laser}})$ is the radius of the beam at a given axial coordinate, z_{laser} , and z_R is the Rayleigh range. $\omega(z_{\text{laser}})$ is defined as the radial location at which the intensity of the beam drops to $1/e$ of the intensity at zero radius. Figure modified from [14]. .	176
6.21	Total time $\hat{t}(s)$ spent by $\sim 10,000$ simulated antihydrogen trajectories within the region of the 1S-2S laser beam during the 60 s hold of the (a) No AE Protocol and (b) AE Protocol. Both protocols are defined in Figure 6.19.	177
6.22	Time-summed xy -position, $\Sigma(s)$, shown as the greyscale axis is scaled to the minimum and maximum time spent in any single bin for No AE Protocol (a) and AE Protocol (b). (c) and (d) show the same data as (a) and (b) respectively but the greyscale axis is modified to make the features more prominent. A total of 10^6 square bins lie within a square grid of 3.0×3.0 mm. Σ is the total time spent by the $\sim 10,000$ trajectories in each square bin during the 60 s hold. No AE Protocol and AE Protocol are defined in Figure 6.19	178

Abstract

Antihydrogen is now routinely formed in ALPHA by combination of antiproton and positron plasmas. Formed anti-atoms with energy $< \sim 0.5$ K are trapped in an octupole-based Ioffe-Pritchard magnetic trap. Reducing trapped antihydrogen energy is expected to increase precision in experiments that measure fundamental antihydrogen properties for precise comparison to hydrogen. Cooling is expected to permit confinement in a shallower magnetic trap, hence reducing magnetic field errors in gravity experiments and increasing sensitivity in experiments that measure the hypothetical antihydrogen charge. In antihydrogen spectroscopy experiments, cooling is expected to increase laser interaction time, thus narrowing spectral linewidth.

A technique to adiabatically cool antihydrogen is presented, which involves trapping in a short magnetic well, before slowly (compared to the antihydrogen speed) expanding the trap volume along the trap axis. During magnetic release, lower energy anti-atoms are expected to escape and annihilate when the trap depth is *lower* which happens *later* in time. Adiabatically cooled (Trial A) annihilation times are compared to control samples held in the small (Trial B) and large (Trial C) magnetic volumes, showing Trial A anti-atoms are lost on average when the trap depth is 0.081 ± 0.007 K, compared to 0.223 ± 0.009 K and 0.173 ± 0.011 K for Trials B and C respectively.

Results of detailed Monte Carlo simulations are presented, which show that Trial A anti-atoms have average energies 1.616 ± 0.002 and 1.558 ± 0.002 times lower than Trials B and C respectively, corresponding to a $37.9 \pm 0.1\%$ energy decrease during expansion and achieving mean final energy 0.2226 ± 0.0002 K. The magnetic trap expands predominantly in the axial direction, giving rise to a $68.1 \pm 0.2\%$ axial and $14.2 \pm 0.2\%$ transverse energy decrease resulting from non-trivial energy mixing dynamics. Two adiabatic models that make limiting assumptions about the energy mixing dynamics are shown to bound the simulation results, confirming the cooling is consistent with an adiabatic process. A possible factor of ~ 10 increase in precision of antihydrogen charge measurements is presented and a potential issue for the use of axial adiabatic cooling in ALPHA antihydrogen spectroscopy experiments is raised, that a significant fraction of adiabatically cooled anti-atoms tend not to occupy the spatial region of the 1S-2S spectroscopy laser beam.

Declaration

No portion of the work referred to in the thesis has been submitted in support of an application for another degree or qualification of this or any other university or other institute of learning.

Copyright

The author of this thesis (including any appendices and/or schedules to this thesis) owns certain copyright or related rights in it (the “Copyright”) and she has given the University of Manchester certain rights to use such Copyright, including for administrative purposes.

Copies of this thesis, either in full or in extracts and whether in hard or electronic copy, may be made only in accordance with the Copyright, Designs and Patents Act 1988 (as amended) and regulations issued under it or, where appropriate, in accordance with licensing agreements which the University has from time to time. This page must form part of any such copies made.

The ownership of certain Copyright, patents, designs, trademarks and other intellectual property (the “Intellectual Property”) and any reproductions of copyright works in the thesis, for example graphs and tables (“Reproductions”), which may be described in this thesis, may not be owned by the author and may be owned by third parties. Such Intellectual Property and Reproductions cannot and must not be made available for use without the prior written permission of the owner(s) of the relevant Intellectual Property and/or Reproductions.

Further information on the conditions under which disclosure, publication and commercialisation of this thesis, the Copyright and any Intellectual Property and/or Reproductions described in it may take place is available in the University IP Policy (see <http://documents.manchester.ac.uk/DocuInfo.aspx?DocID=24420>), in any relevant Thesis restriction declarations deposited in the University Library, the University Library’s regulations (see <http://www.library.manchester.ac.uk/about/regulations/>) and in the University’s policy on Presentation of Theses.

Acknowledgements

Firstly, I'd like to thank my supervisor, Dr William Bertsche, for providing guidance, feedback and sharing his wealth of physics knowledge throughout. I have been incredibly fortunate to also have the guidance and support of Prof. Joel Fajans and Prof. Jonathan Wurtele, whose endless insight has been invaluable to my understanding and to the development of this research. Prof. Joel Fajans contributed substantially to the method of simulating the adiabatic cooling experiment and to the method of analysing the simulation results, Prof. Joel Fajans and Prof. Jonathan Wurtele contributed substantially to the method of simulating the charge neutrality experiment and to the method of analysing the simulation results. I would also like to thank Prof. Niels Madsen, Dr Chris Ørum Rasmussen and Dr Svante Jonsell for their guidance, as well as Prof. Francis Robicheaux whose work developing the principles of simulating antihydrogen is a critical aspect of this work.

I would like to thank every member of the ALPHA Collaboration for making my time at CERN incredibly valuable. In particular, thank you to my office mate Dr Steven Jones for making working much more enjoyable. I was fortunate to work closely with Peter Granum throughout, and have benefited from his amazing ability to solve problems, his incredible hard work and his valuable friendship. I would like to thank the members of the Magnet Team: Dr Daniel Maxwell, Dr Emma Anderson, Jaspal Singh, Dr Nathan Evetts, Peter Granum and Dr William Bertsche, for their hard work and dedication. The expertise of Dr Joseph McKenna and Dr Chukman So have been invaluable throughout. I have been fortunate to share my time at ALPHA with a number of talented physicists that I have discussed many aspects of my research with. In particular, discussions with Edward Thorpe-Woods and Andrew Christensen led to new understanding and perspective. I also had the opportunity to work with brilliant undergraduate students Evan Deddo and Nicolas Kalem, whose hard work simulating the charge neutrality experiment made parts of this research possible.

CERN wouldn't have been the same without the floats down the Rhône, games of Coup and R1 trips. Outside of ALPHA, my friend Luke Dyks and my housemates at 161 made my time at CERN. I'd also like to thank my family: thank you to my Mum for endless support during the long days of thesis writing, to my Dad for sparking my interest in science at a young age, to Ryan, spending time with you made writing up a lot easier. Also, thanks to

Pete for lots of meals on wheels. Finally, thank you to my outstanding physics and maths teacher Dr Gary Brawn, without you I wouldn't be studying physics.

This work was supported by an Engineering and Physical Sciences Research Council (EPSRC) Doctoral Training Award.

Dedicated to my Grandparents.

Motivation and Background

1.1 Motivation for Antimatter Research

1.1.1 A Brief Introduction to the History of Antimatter

It was in 1926 that Erwin Schrödinger published the linear partial differential equation that describes the non-relativistic spin-free mechanics of atoms and molecules, which became known as the Schrödinger equation [15]. For a system evolving with time, t , the functional form of the Schrödinger equation is

$$i\hbar \frac{d}{dt} |\Psi(t)\rangle = \hat{H} |\Psi(t)\rangle, \quad (1.1)$$

where $|\Psi(t)\rangle$ is the state vector of the quantum system, \hat{H} is the Hamiltonian operator and \hbar is the normalised Planck constant. Just a year later, Paul Dirac postulated the form of the equation that describes the wave function of a quantum mechanical system for particles with intrinsic spin and travelling at relativistic speeds,

$$i\hbar \gamma^\mu \partial_\mu \psi - mc\psi = 0, \quad (1.2)$$

where there is an implied summation over the index $\mu = [0, 1, 2, 3]$, ∂_μ is the four-gradient and ψ is the wave function. Dirac's key realisation was that the gamma matrices,

$$\gamma^0 = \begin{pmatrix} I & 0 \\ 0 & -I \end{pmatrix}, \quad \gamma^i = \begin{pmatrix} 0 & \sigma^i \\ -\sigma^i & 0 \end{pmatrix} \quad (1.3)$$

were 4×4 matrices rather than scalars [16], where $\sigma^i (i = [1, 2, 3])$ are the Pauli matrices, I indicates the 2×2 identity matrix and a 0 indicates a 2×2 matrix of zeros. His discovery gave rise to negative energy solutions which puzzled physicists around the world. Dirac imagined an unseen infinite “sea” of electrons that filled up the negative energy states and that, by the Pauli Exclusion Principle, electrons in the positive energy states were prevented from occupying the negative energy states. If this were true, however, an electron present in the

“sea”, given enough energy to be knocked out of the negative energy state, would appear as a positive charge. Dirac originally postulated that these “holes” could be protons but since the proton and electron do not have equal masses, he postulated the existence of the *anti-electron*: a particle with the same mass but opposite charge to the electron [16].

Such a particle was first observed by Carl Anderson in 1931, when he placed a sheet of lead in the center of a bubble chamber. The lead sheet changed the energy of the particle, which altered the radius of its track and allowed Anderson to determine the particle’s charge and to conclude that the particle had a mass much less than that of the proton [17]. Although Anderson had interpreted his experimental observation as the discovery of a new particle, the scientific community was reluctant to admit the existence of such a particle in a Universe that was thought to be made up of only protons and electrons [18]. In 1933, Blackett and Occhialini developed Anderson’s experimental work by sandwiching the cloud chamber between two Geiger-Müller counters [19]. By analysing tracks in the chamber, they came to the same conclusion as Anderson - that the tracks confirmed the existence of a new particle. It was the interpretation of their result that developed the discovery of this new particle, since they made the connection between Dirac’s negative energy solutions and Anderson’s original observations [18]. They postulated that the new particle, which later became known as the positron, could have been created during the process of a collision and that the conservation of charge would thus require electrons and positrons to be produced in equal numbers [18].

It is now known to be a profound universal feature of quantum field theory that every particle has an antiparticle, of which the positron was the first to be discovered. Antiparticles have the same mass and spin as their matter counterparts, but opposite charge. Some particles, such as the photon, are their own antiparticle [16]. When matter and antimatter meet, they annihilate and produce energy in the form of photons. Equivalently, a particle-antiparticle pair can be produced from a photon in the process of *pair production*. Other antiparticle discoveries followed those of the positron when, in 1955, the first observation of the antiproton was reported at the Berkeley Bevatron [20] and, just one year later, the same group reported the first observation of the antineutron [21].

1.1.2 The Baryon Asymmetry Problem

At the time of antimatter discovery, all physical processes were thought to be invariant under each independent inversion: C (charge conjugation), P (parity transformation) and T (time reversal). If such symmetries were held strictly in nature, equal amounts of matter and antimatter would be present in the Universe. Therefore, the fact that the Universe appears to be made almost entirely of matter, known as the Baryon Asymmetry Problem, was, and remains, a complete mystery.

The theory of the expanding Universe was first proposed by Georges Lemaître in 1927

[22] and backed up by Hubble’s observation that more distant galaxies travel away from earth faster [23]. This led to the theory that the Universe expanded from an extremely hot and dense initial state [24] [25], known as the Big Bang Theory, which was backed up by observation of the Cosmic Microwave Background (CMB) [26]. During the hot early stages of history, antimatter would have been present when pair production and annihilation were in thermal equilibrium. When the temperature became too small for pair production to occur, almost all particles and antiparticles would have annihilated to produce photons. However, the existence of the matter-dominated Universe in the modern day asserts that there must have been a small matter-antimatter asymmetry of one part in 10 billion to lead to the matter dominated Universe we inhabit today [27].

In 1967 Andrei Sakharov published his conditions for *baryogenesis* (the conditions for creation of a matter-antimatter asymmetric initial state): baryon number violation, C and CP (combined charge conjugation and parity transformation) violation, and deviation from thermal equilibrium. The Standard Model (SM) of Particle Physics is the theory describing three of the four fundamental forces of nature. Within the SM framework, physical processes have been known to fulfill all three Sakharov conditions. In 1956 the weak interaction was shown to violate independent C and P symmetries in studies of beta decays of ^{60}Co [28]. Later, CP violating processes were also discovered first in a system of neutral kaons [29] and then in B mesons [30]. Deviation from thermal equilibrium is fulfilled on account of the inflationary expansion of the Universe [27]. Despite the fact that the SM fulfills all three Sakharov conditions, the amount of CP violation observed to-date within the SM makes it extremely unlikely that baryogenesis is possible within the SM framework [27] [31]. Since the SM is unable to explain the observed matter-antimatter asymmetry in the Universe, it is natural to seek explanations beyond the SM.

Many have discussed the possibility that some regions of the Universe are matter-dominated and others are antimatter-dominated [32] [33]. The observation of positrons and antiprotons in cosmic rays may appear to back this up, but this is expected even in a Universe made up entirely of matter on account of their low pair production thresholds (compared to heavy nuclei) [27]. The lack of annihilation signals that would be present at the boundary between pockets of matter and antimatter has limited the size of the pockets to be greater than the observable Universe [34].

1.1.3 CPT symmetry

The CPT theorem is intrinsic to the SM and states that any Lorentz-invariant quantum field theory is invariant under the combined C, P and T operations [35]. As a consequence, CPT invariance asserts that particles have exactly the same charge, mass, lifetime and energy-level structure as their antiparticle counterpart. No process has ever been observed to break CPT

symmetry, but some beyond the Standard Model (BSM) theoretical approaches allow CPT symmetry violation [36] [37].

1.1.4 Antimatter Gravitational Acceleration

When it comes to the acceleration of an antiparticle under the force of gravity, the CPT theorem asserts that an *antiparticle* falls in the gravitational field of an *antimatter* body the same as a particle falls in the gravitational field of a matter body. The CPT theorem does not specify the behaviour of an *antiparticle* in the gravitational field of a *matter* body, such as the Earth [38]. Instead, it is Einstein’s Weak Equivalence Principle (WEP), within the framework of General Relativity, that asserts that the acceleration of an antiparticle in the Earth’s gravitational field be the same as its particle counterpart. There have been many indirect indications that there is no difference in the gravitational acceleration of matter and antimatter [39] [40], as well as many theoretical arguments against the existence of an asymmetry between the gravitational acceleration of matter and antimatter [41] [42]. However, in some cases it has been suggested that gravitational asymmetries between matter and antimatter could lead to baryogenesis [43] [44]. Further, the fact that the SM has not been unified with General Relativity and that our current understanding of physics does not explain the majority of mass or energy in the Universe, motivates experiments that directly test fundamental predictions of these models.

1.1.5 Antihydrogen as a Fundamental Symmetry Probe

Hydrogen is charge neutral and the simplest atom in existence. Many of its energy levels, described by current understanding of physics, have been measured experimentally [45] [46], with the highest precision being the 1S-2S energy level splitting measured at a relative precision of 4.2×10^{-15} [47]. Its antimatter equivalent is the antihydrogen atom: the bound state of an antiproton and a positron. Given that CPT invariance holds, antihydrogen is expected to have the same energy levels, (neutral) charge and (stable) lifetime as the hydrogen atom; making it an ideal candidate for precise fundamental matter-antimatter comparison.

Antihydrogen was first produced at the European Organization for Nuclear Research (CERN) in Geneva, Switzerland, in the Low Energy Antiproton Ring (LEAR) in 1995. LEAR bombarded xenon atoms with antiprotons, producing antihydrogen that travelled close to the speed of light, and detected the antihydrogen annihilation events [48]. This was the first bound system comprised entirely of antimatter ever observed.

Following the closure of CERN’s antiproton machines to free up resources for the Large Hadron Collider (LHC) in 1996, a community of physicists were interested in continuing low energy antiproton research, which led to approval of CERN’s Antiproton Decelerator (AD). The role of the AD is to provide a number of experiments with low energy antiprotons.

Antiprotons are provided to the AD by collision of the 26 GeV energy CERN Proton Synchrotron (PS) beam containing $\sim 1.5 \times 10^{13}$ protons with an iridium target. In some cases this leads to proton-antiproton pair production, of which antiprotons of energy ~ 3.6 GeV can be separated on account of their negative charge, focused into a parallel beam using a magnetic horn-type lens, and injected into the AD ring where they are decelerated using radiofrequency (RF) cavities. Electron and stochastic cooling techniques are also employed to reduce the bunch length, spatial extent and momentum spread of the antiproton pulse. Decelerated antiprotons are provided to one of five AD experiments every 100 s by the AD in bunches of 3×10^7 particles of length 100 – 200 ns, with energy 5.3 MeV [49].

By combining 5.3 MeV antiprotons with positrons produced in the beta-decay of ^{22}Na , the Antihydrogen Laser PHysics Apparatus (ALPHA) collaboration trapped antihydrogen for the first time in 2010 [50]. The antihydrogen atoms were confined in a magnetic minimum trap for 172 ms, before the magnetic trap was removed in 30 ms, during which time 38 antihydrogen annihilation events were detected [50].

In 2011, the ALPHA collaboration reported the confinement of antihydrogen for 1000 s [51], which made it possible to induce positron spin flips with resonant microwaves [52], to place limits on the ratio of gravitational to inertial mass of antihydrogen [53], to interrogate the charge neutrality of antihydrogen [54], and to measure the 1S-2S energy level splitting [55].

In more recent years, ALPHA has become capable of trapping 10-30 antihydrogen atoms every 4 minutes, accumulating around 1000 trapped antihydrogen atoms in several hours [7]. Improvements in trapping efficiency and refinement of experimental techniques led to a more precise bound on the antihydrogen charge [11], further characterisation of the 1S-2S transition [56], observation of the 1S-2P transition [57], and investigations of the fine structure of antihydrogen [58]. In addition, ALPHA has proposed an experiment to more precisely measure the effect of gravity on antihydrogen, known as ALPHA-g [59].

To date, no experiment has observed a CPT or WEP breaking asymmetry between matter and antimatter. As experiments place ever tighter bounds on the fundamental properties of antihydrogen, increases in experimental precision are required to search for ever more tightly bounded matter-antimatter asymmetries. A general way to increase precision of measurements of the fundamental properties of antihydrogen is to reduce the kinetic energy of trapped antihydrogen.

In the remainder of this chapter I will describe the process by which antihydrogen is routinely formed by merging positron and antiproton plasmas in ALPHA, how it can be trapped in a magnetic field, and how ALPHA detects its existence when antihydrogen atoms annihilate with internal trap structures. I will then introduce ALPHA's portfolio of antihydrogen spectroscopy, gravity and charge neutrality experiments, focusing on the

current limiting factors in experimental precision, which are often dependent on the trapped antihydrogen energy. Finally, I will introduce ALPHA's methods to reduce the energy of trapped antihydrogen atoms, including the adiabatic cooling technique which this thesis will investigate in detail.

1.2 The ALPHA Apparatus

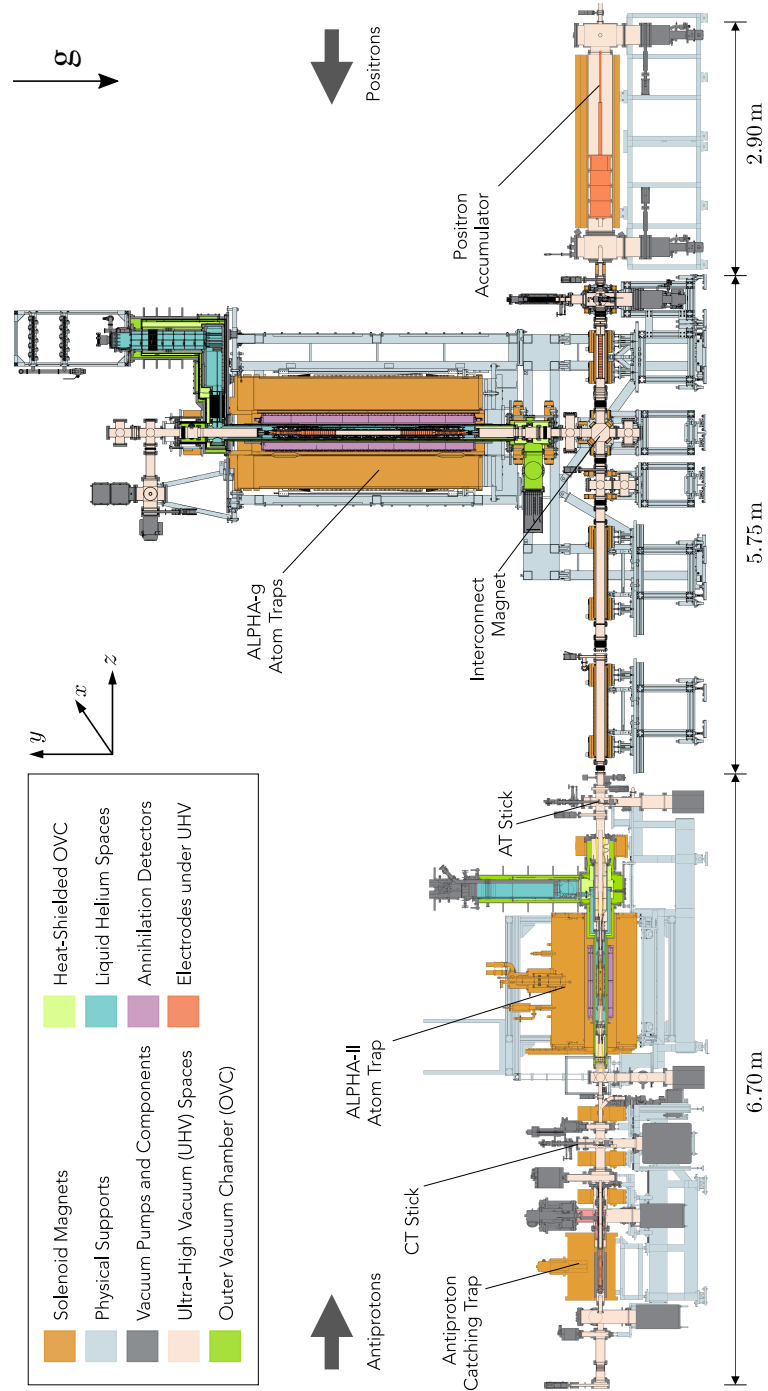


Figure 1.1: Cross-sectional view of the ALPHA apparatus in 2018. Antiprotons enter the apparatus from CERN’s Antiproton Decelerator (as indicated by arrow), and positrons are obtained from a radioactive source of ^{22}Na (enter apparatus as indicated by arrow). The apparatus is under vacuum, and regions of different vacuum levels are indicated by the coloured regions specified in the legend. Antiproton and positron plasmas are guided by magnetic fields to the two Atom Traps (ALPHA-2 and ALPHA-g). Here, the two plasma species are merged to form neutral antihydrogen which is subsequently confined in an Ioffe-Pritchard magnetic trap. Note the direction of the gravitational force is indicated by an arrow labelled ‘g’. Figure from [2], courtesy of Mark Johnson.

This section will begin with a broad overview of antihydrogen formation, trapping and detection in ALPHA, with more detail being provided in the subsections which follow. A complete description is given in [3] and [5].

A diagram of the ALPHA apparatus is shown in Figure 1.1. Spectroscopic measurements and experiments that place bounds on the antihydrogen charge take place in the ALPHA-2 antihydrogen trap, whereas proposed gravity measurements will take place in ALPHA-g, which is an antihydrogen trap with axis parallel to the gravitational force, to maximise measurement sensitivity (note that the force of gravity is indicated by an arrow in Figure 1.1).

A magnetic moment experiences a potential in a magnetic field. ALPHA exploits this by trapping antihydrogen atoms effectively in a magnetic bowl, known as an Ioffe-Pritchard magnetic minimum trap. The magnetic field (on-axis magnetic field shown in Figure 1.2(b)) is generated by a number of superconducting magnets: an *Octupole* magnet and five short solenoids known as *Mirror Coils* (Figure 1.2(a)). The magnetic field generated by the Octupole provides radial confinement of antihydrogen, and the magnetic field generated by the two outermost Mirror Coils (A and E) provides axial confinement. As will be explained in Section 3.4, ALPHA traps only antihydrogen atoms with positron spins aligned antiparallel to the magnetic field, known as *low-field-seeking* states, which experience a potential,

$$U(\mathbf{x}) = \mu_B(|\mathbf{B}(\mathbf{x})| - |\mathbf{B}(\mathbf{x}_{\min})|), \quad (1.4)$$

in the magnetic trap, where μ_B is the Bohr magneton, $\mathbf{B}(\mathbf{x})$ is the magnetic field strength at the coordinate vector \mathbf{x} , and \mathbf{x}_{\min} is the location of the minimum magnetic field strength, hence the magnetic potential is defined relative to the magnetic minimum. Typical trap depths are ~ 0.5 K, and effectively set an upper limit on the energy of confined antihydrogen. In Section 3.1 I give a detailed description of the theory of magnetic minimum trapping of neutral antihydrogen. Here, I focus on the procedures of antihydrogen formation and detection.

In ALPHA, antihydrogen is formed from antiprotons originating from CERN's Antiproton Decelerator (AD) (see Section 1.2.1) and positrons from radioactive beta decay (see Section 1.2.2). Constituent particle plasmas are confined in independent Penning-Malmberg traps (Figure 1.1, antiprotons in the *Antiproton Catching Trap (CT)* and positrons in the *Positron Accumulator*) in which an axial magnetic field confines charged particles radially and hollow cylindrical electrodes provide axial confinement. Constituent particles are in the plasma state, in which the potential of a typical particle due to its nearest neighbors is small compared to its kinetic energy. Plasmas are transported ballistically, guided by axial magnetic fields, to a Penning-Malmberg trap known as the Atom Trap (AT) (Figure 1.2).

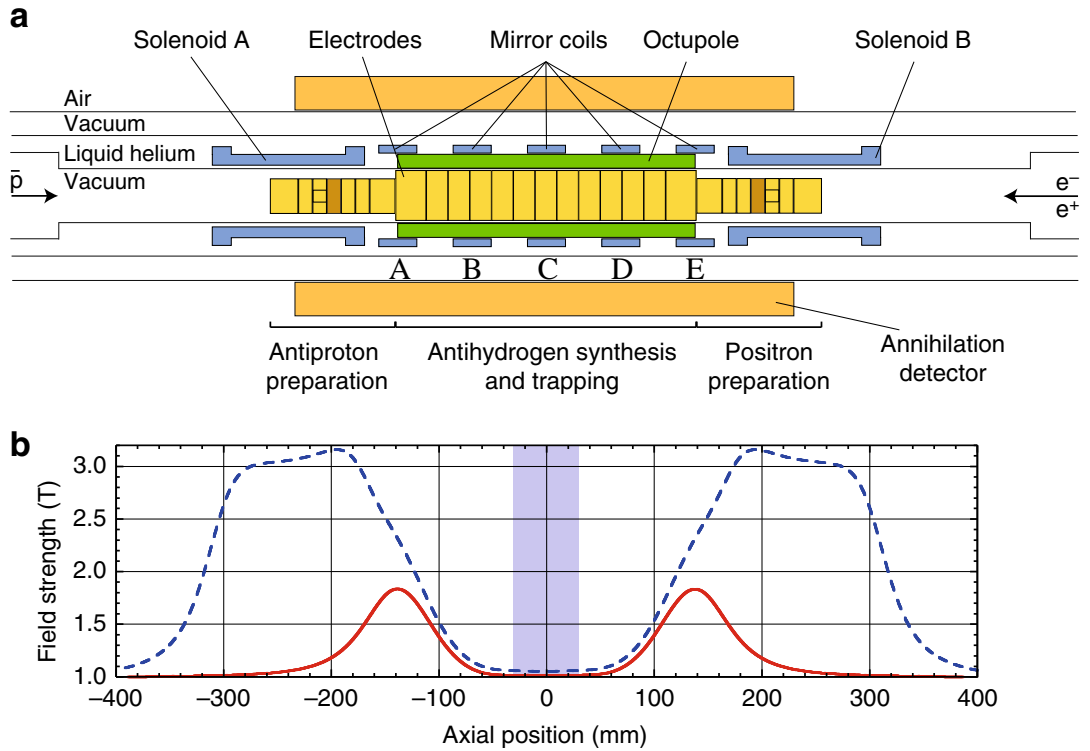


Figure 1.2: (a) Diagram of the ALPHA-2 Atom Trap (AT), drawn to scale (axial scale lines up with that of (b) below) except from the annihilation detector radius. Antiprotons (denoted \bar{p}) enter from the left, whereas electrons (denoted e^-) and positrons (denoted e^+) enter from the right. (b) On-axis magnetic field strength as a function of axial position in the ALPHA-2 AT, where the axial position maps directly to the diagram in (a). Red solid curve gives the magnetic field strength with the flat trap energised, which is defined in Figure 2.3(c). Blue dashed curve gives the magnetic field strength during stacking where, in addition to the flat trap, Solenoid A and Solenoid B are energised to increase the magnetic field for enhanced capture, cyclotron cooling and rotating wall efficiency for charged species. Figure from [3].

The AT is a hybrid trap, capable of trapping both charged constituent plasmas and neutral antihydrogen. To achieve this, a magnetic minimum antihydrogen trap is superimposed on the electric and magnetic fields of the Penning-Malmberg trap. In Figure 1.2(b) (red curve) the on-axis magnetic field strength is shown as a function of axial position with the antihydrogen trap energised. The *External Solenoid* (not shown in Figure 1.2(a) but sits radially outside of the AT magnets) provides a uniform 1 T axial magnetic field for radial confinement of charged particles, which raises the minimum magnetic field to ~ 1 T. The outer Mirror Coils (A and E) cause peaks in magnetic field strength at ± 137 mm, providing axial confinement of neutral antihydrogen.

To form antihydrogen, antiprotons and positrons are held in neighboring electrostatic potential wells with the magnetic minimum trap energised. The potential barrier between the plasmas is lowered, causing the two species to merge and form antihydrogen, which is

subsequently trapped by the magnetic field. The ALPHA-g AT will use the same principles as ALPHA-2 to form and trap antihydrogen, but in this thesis I focus on ALPHA-2 as the experimental and simulation results presented were obtained in ALPHA-2. Secondary pions, formed when antihydrogen atoms annihilate with trap structures, are detected in ALPHA-2 by a Silicon Vertex Detector (SVD) (see Section 1.2.4), and in ALPHA-g by a Radial Time Projection Chamber (rTPC). The detector is placed on the outside of the superconducting magnets of the AT to maximise the magnitude of the magnetic field at the inner trap radius (and therefore to maximise the trap depth), but inside the 1 T External Solenoid field to minimise scattering material between the trap and detector and to minimise detector cost.

The superconducting magnets sit in a liquid ^4He cryostat (Figure 1.2(a)) at 4.2 K, which is isolated from air by an outer vacuum chamber (OVC). The internal trapping region is under ultra-high (10^{-13} mbar) vacuum (UHV) [5], leading to a low rate of collision of trapped antihydrogen with background gases, which can result in antihydrogen loss via a number of rearrangement processes or by increases in energy from elastic scattering collisions [60].

In the following subsections, I refer to the number, radial extent and temperature of plasmas in the ALPHA experiment, which are determined experimentally. The number of particles in a plasma can be measured destructively by colliding the plasma with a capacitive plate, known as a Faraday cup. The plasma charges the Faraday cup; producing a voltage that decays exponentially with a ~ 1 ms time constant, which can be amplified and used to infer the amount of charge dissipated on the Faraday cup (and therefore to measure the number of charged particles in the plasma) [2]. The plasma radius can be determined destructively by colliding the plasma with a microchannel (MCP), consisting of many small channels with walls constructed from a high gain material. An incident charge causes a cascade of electrons which impact on a phosphor screen placed behind the MCP, causing fluorescence which can be imaged by a CCD camera, producing an image of the transverse plasma distribution [61]. An MCP also doubles as a diagnostic of the plasma temperature, by gradually lowering the potential at one end of an electrostatic well, the charged particles gradually escape and collide with the MCP. A fit to the charge dissipated on the MCP as a function of time enables the temperature of the plasma to be determined [62].

1.2.1 Antiprotons

As mentioned previously, the AD provides one of the five AD experiments with 3×10^7 antiprotons at 5.3 MeV every 100 s [49]. An additional deceleration stage has recently been added to the AD, known as the Extra Low Energy Antiproton Ring (ELENA), which reduces the energy of the antiprotons sent to the AD experiments to 100 keV [63]. In addition, ELENA delivers antiprotons to each of the AD experiments simultaneously, meaning experiments can obtain antiprotons 24 hours a day. In 2021 the ALPHA collaboration began using antiprotons

from ELENA. In this section I will describe parameters of antiprotons obtained prior to the ELENA era, as they are relevant in discussions of the adiabatic expansion experimental data (taken in 2016) presented in this thesis.

Antiprotons from the AD (or more recently, ELENA) enter the apparatus as indicated by the arrow in Figure 1.1. The beam is first passed through a series of thin degrader foils to slow the antiprotons to trappable energies ($< \sim 5$ keV) [64]. Then, antiprotons are caught between two high voltage (4 kV) electrodes in a pre-loaded electron plasma in the 3 T field of the CT Penning-Malmberg trap. One high voltage electrode is raised to 4 kV prior to capture, then on arrival of the antiproton pulse, the second high voltage electrode is raised to 4 kV to trap the antiprotons. The role of the CT is to produce an antiproton plasma that is sufficiently cold and has sufficiently small radius to produce large numbers of trappable antihydrogen. Specifically, the temperature of the antiproton plasma is ideally comparable to the AT trap depth (~ 0.5 K), although in practice the temperature is typically much higher. The plasma radius must be controlled for a number of reasons. Firstly, it is necessary to control the radius to ensure good overlap with the positron plasma during antihydrogen formation in order to maximise recombination probability. Secondly, maintaining a plasma with a small radius prevents the existence of antiprotons with high momentum due to rigid plasma rotation. Thirdly, plasma instabilities arise at radii $> 0.2r_{\text{trap}}$, where r_{trap} is the electrode inner radius, due the Octupole field (if a quadrupole were used, plasmas would become unstable at smaller radii) [64] [5].

1.2.2 Positrons

A positron beam enters the experimental apparatus as indicated by an arrow in Figure 1.1. The positrons are obtained from the beta decay of a ^{22}Na radioactive source with peak activity 2.8 GBq, and are formed into a beam using a solid neon moderator shown on the left of Figure 1.3 [5]. The energy of the beam is set using an electrical bias, and the beam is magnetically guided into the *Positron Accumulator*.

The Positron Accumulator is a three-stage Surko-style buffer gas accumulator [65], within which the positrons are confined using the Penning-Malmberg trapping technique: a 0.14 T magnetic field confines positrons radially and a number of hollow cylindrical electrodes provide axial confinement. One of these electrodes is azimuthally segmented into six sections, allowing use of the rotating wall technique to counteract plasma expansion. The three accumulation stages contain decreasing nitrogen buffer gas pressure and decreasing electrostatic potential (produced by applying decreasing voltage to the Penning-Malmberg electrodes), as shown in the graph of Figure 1.3. The positrons lose energy through collisions with the nitrogen gas and up to 250 million positrons can accumulate in the third stage in 3-4 minutes [5]. After accumulation, the nitrogen gas is pumped out of the system and the positrons are ready to




Figure 1.3: Diagram of ALPHA’s positron source and Surko-style Positron Accumulator. Beneath the Positron Accumulator, the electrostatic potential generated by the Penning-Malmberg trap electrodes within the buffer gas volume is shown, where each of the three stages has a decreasing electrostatic potential, as a function of pressure, P . Positrons lose energy through collisions with a nitrogen buffer and accumulate in Stage 3 (lowest pressure stage). Note this is a diagram of ALPHA’s positron source and accumulator from 2014, which differs from the current apparatus. The differences are local to the positron source (the Positron Accumulator end is the same): minor modifications were made to the coldhead (rotated), the guiding magnets and the vacuum chamber [4]. Figure from [5].

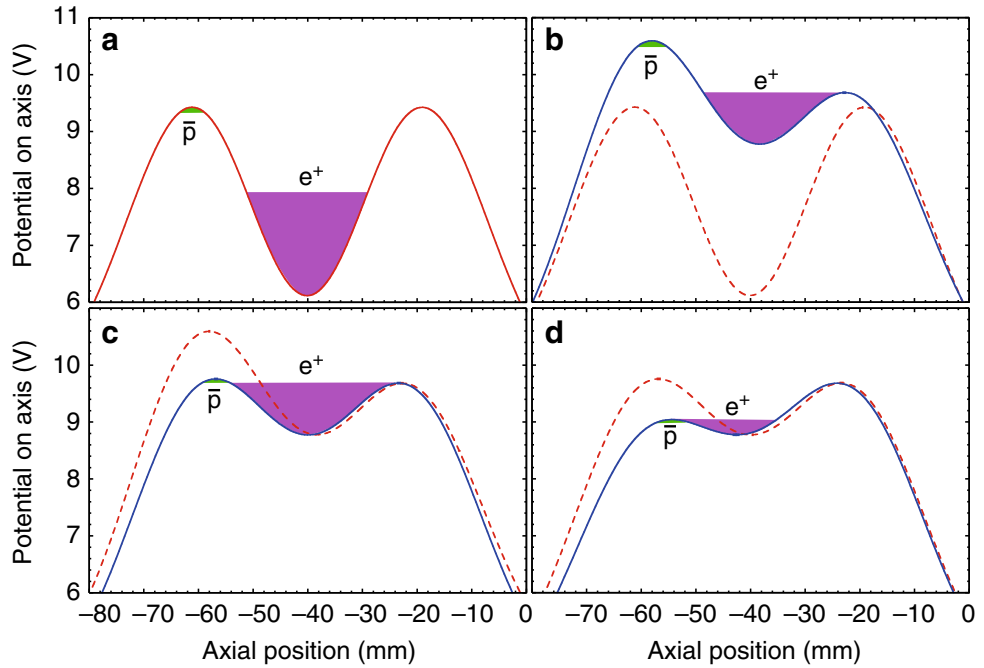


Figure 1.4: Electrostatic potential manipulations for positron and antiproton mixing to form antihydrogen. In each subfigure, the on-axis electrostatic potential as a function of axial position is given, which can be shifted relative to the minimum of the antihydrogen trap and therefore the absolute value here is arbitrary but relative axial positions are relevant. (a) Before evaporative cooling. (b) Evaporative cooling. (c) Realignment such that the two potential wells almost overlap. (d) Antiproton and positron plasmas are mixed during 1 s. Dashed and solid curves represent electrostatic potentials before and after each step respectively. Purple and green regions indicate self-potentials and physical extents of positron and antiproton plasmas respectively. Figure from [3], details provided in [3].

be transferred for use in other parts of the apparatus [5].

1.2.3 Antihydrogen Formation

Following antiproton and positron preparations in the CT and Positron Accumulator respectively, both species are moved to the AT, where they are independently prepared for *mixing* (combining antiprotons and positrons for antihydrogen formation) at opposite ends of the AT.

Antiprotons are caught in the AT in a pre-loaded electron well and undergo a similar set of manipulations as in the CT (compression, sympathetic cooling, electron ejection) [3]. Then the antiproton plasma is evaporatively cooled (EVC) to ~ 40 K, which involves lowering the electrostatic confining potential to allow the most highly energetic particles to escape, leaving the remaining charged particles to rethermalise at a lower temperature. At this stage, the plasma contains $\sim 9 \times 10^4$ antiprotons with plasma radius ~ 1 mm [3]. Meanwhile, on the other

side of the AT, the positron number and density is controlled using a combination of strong drive regime rotating wall compression and evaporative cooling known as (SDREVC) [66], intended to match the positron plasma radius to that of the antiproton plasma to maximise overlap during mixing. The positrons are then transferred to a shallower well (Figure 1.4(a)), causing them to cool adiabatically to ~ 30 K [3].

Once final plasma preparations are complete, the magnetic minimum antihydrogen trap is energised. Antiproton and positron plasmas sit in neighboring electrostatic wells (Figure 1.4(a)) that can both be shifted relative to the magnetic minimum trap so that antihydrogen formation takes place as close to the magnetic minimum as possible to produce the lowest energy antihydrogen. To mix the plasmas, the electrostatic potentials follow the sequence shown in Figure 1.4. The positrons are first evaporatively cooled (Figure 1.4(b)). After the evaporative cooling stage, the antiproton well is modified so that the two potentials almost overlap (Figure 1.4(c)) before the antiproton confining potential is gradually lowered and the plasmas merge (Figure 1.4(d)). Although this method causes loss of positrons prior to mixing, it is advantageous because the two plasmas come into contact without accelerating the antiprotons and the positrons are continuously cooled via evaporation during mixing [3]. Therefore, this technique is designed to optimise the production of low energy antihydrogen. Antihydrogen is expected to be formed via interaction of an antiproton with two positrons, in which one positron carries away the excess recombination energy and the other forms the antihydrogen bound state with the antiproton. Via this recombination method, antihydrogen is formed in high principal quantum number states (but the distribution of initial states is not accurately known), in a thermal distribution at the temperature of the positron plasma (~ 50 K). This process is known as *Three-Body Recombination* (TBR) [67] [68]. The fraction of antihydrogen atoms in this distribution with energy less than the depth of the magnetic minimum trap (~ 0.5 K) are subsequently confined.

The procedure described above (antiproton catching and preparation in the CT, positron accumulation, further preparations in the AT and mixing) takes 3-4 minutes in total and can be repeated (if the magnetic minimum trap remains energised) to accumulate ~ 150 antihydrogen atoms per hour (accumulation rate determined by removing the magnetic trap and detecting annihilations on internal trap structures) in a process known as *stacking* [3]. The trapped antihydrogen atoms have a finite lifetime which was recently shown to be >66 h due to collisions with background gas trap impurities [69].

1.2.4 Antihydrogen Detection

After formation, experiments can be carried out on the trapped antihydrogen. In some cases, for example in measurements of the 1S-2S energy level splitting of antihydrogen (more detail in Section 1.3), laser excitations lead to antihydrogen ionisation or transition

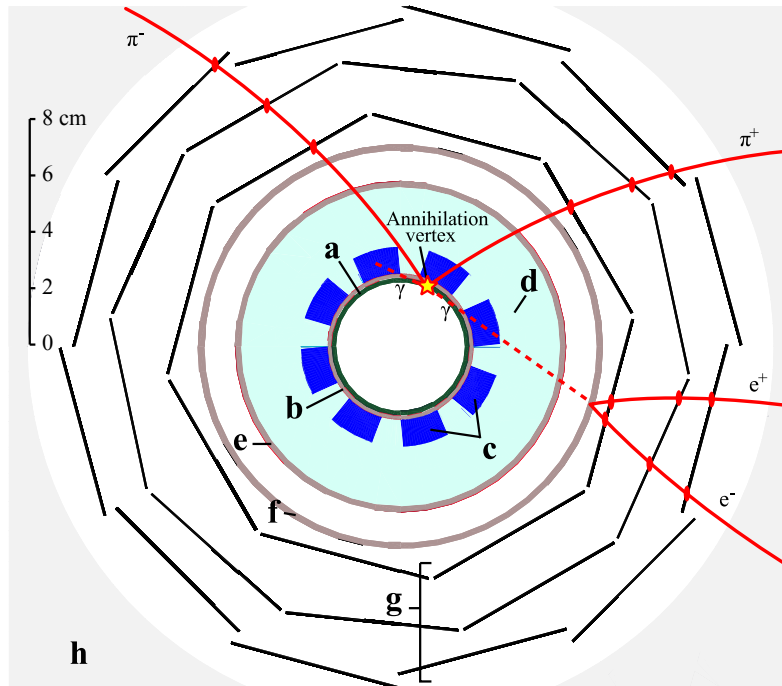


Figure 1.5: Radial profile diagram of the ALPHA-2 annihilation detector: the Silicon Vertex Detector (SVD). The letters label (a) Penning-Malmberg electrode wall, (b) Magnetic minimum trap superconducting magnet winding form, (c) Octupole windings, (d) liquid helium volume, (e) vacuum wall between helium volume and outer vacuum chamber (OVC), (f) vacuum wall between OVC and air, (g) three layers of silicon modules and (h) External Solenoid. An example antiproton annihilation vertex is indicated by a yellow star which results in tracks in the SVD of two charged pions, π^- and π^+ , and one neutral pion, π^0 , which quickly decays to form two photons, one of which is absorbed by the Octupole windings and the other undergoes electron-positron pair production. Figure from [6].

to untrappable hyperfine states [56], and in other cases the magnetic trap can be removed by ramping the current in the superconducting magnets to zero in a few seconds, known as Fast Ramp Down (FRD). Both antihydrogen loss mechanisms result in annihilations with internal trap structures, that are detected with spatial and temporal sensitivity as an experimental diagnostic.

As mentioned earlier in this chapter, antihydrogen annihilation events are detected by a Silicon Vertex Detector (SVD) in ALPHA-2 and a radial Time Projection Chamber (rTPC) in ALPHA-g. Here, I will focus on the SVD as the adiabatic expansion of antihydrogen experiment took place in ALPHA-2 and therefore the data presented in Chapter 2 was obtained by the SVD.

When an antihydrogen atom comes into contact with an internal trap structure, we detect the annihilation products of the antiproton annihilation; the positron is not detected, meaning antihydrogen and antiproton annihilations are indistinguishable. On average, one antiproton annihilates into three charged pions and two neutral pions [6]. The SVD is comprised of three concentric cylindrical layers containing a total of 60 double-sided microstrip modules covering a 90% solid angle at the axial center as shown in Figure 1.5 [6]. The neutral pions have a lifetime of only $\sim 10^{-16}$ s, decaying almost instantaneously and often producing electron-positron pairs as they traverse the apparatus (Figure 1.5). The charged pions have a lifetime of $\sim 10^{-8}$ s and move relativistically, meaning they are unlikely to decay as they pass through all layers of the SVD [6]. The charged pions can scatter off materials as they pass through the apparatus (particularly in the superconducting magnets), limiting detector resolution [70]. They may also liberate particles (mainly electrons) as they pass through, but the SVD is not particularly sensitive to these secondary particles [70].

As charged particles pass through the silicon modules they create electron-hole pairs, which produces an analog signal known as a “hit”. Hits in multiple layers are used to reconstruct charged particle tracks, which are helical due to the cyclotron motion of the antiproton in the External Solenoid field. The intersection of multiple tracks (e.g. as shown in Figure 1.5) is used to reconstruct an antiproton annihilation vertex using a machine learning based analysis [71] [72]. The analysis also distinguishes antiproton annihilation events from cosmic rays traversing the detector. Charged particles travel through a significant amount of material before reaching the detector, which results in deviations of a few mm in the charged tracks. Accounting for this, the spatial resolution of the detector at the electrode radius is 0.83 cm [6], and antiproton annihilation times can be determined with an uncertainty of $\sim 2 \mu\text{s}$ [70].

In the magnetic field of the AT (which varies spatially with z as shown in Figure 1.2), an antiproton can undergo cyclotron motion whilst also moving parallel to the magnetic field. If the magnetic field barely changes over a cyclotron period, one can show that

$$\mu = \frac{\frac{1}{2}mv_{\perp}^2}{B} \quad (1.5)$$

is a conserved quantity to very good approximation, where v_{\perp} is the velocity of the antiproton perpendicular to B . One can show that μ is a magnetic dipole moment and thus experiences a force $F = -\mu\nabla B$ in a magnetic field, known as the *mirror force* [73]. Antiprotons of energy

$$E < \mu B_{\max}, \quad (1.6)$$

where B_{\max} is the peak of the Mirror Coil field (at $z = 137$ mm), can be trapped by the mirror force akin to antihydrogen trapping via its magnetic moment as I will discuss in Section 3.1. For antiprotons exhibiting cyclotron motion at the magnetic field minimum, B_0 ,

$$\mu = \frac{\frac{1}{2}mv_{\perp 0}^2}{B_0}, \quad (1.7)$$

and

$$E = \frac{1}{2}mv_{\perp 0}^2 + \frac{1}{2}mv_{\parallel 0}^2 \quad (1.8)$$

where $v_{\perp 0}$ and $v_{\parallel 0}$ are the velocity components perpendicular and parallel to the magnetic field respectively at the location of the field minimum. Therefore, antiprotons can be mirror trapped given they have a $v_{\perp 0}$ that is sufficiently large compared to $v_{\parallel 0}$ [73].

The detection technique used by the SVD means antihydrogen annihilation events would be indistinguishable from mirror-trapped antiproton annihilations, but a method similar to that described in [1], which involves application of an electric field across the trap volume by applying a stair-step voltage ranging from ~ 5 V to ~ 140 V to the electrodes of the Penning-Malmberg trap, can clear mirror-trapped antiprotons whilst leaving antihydrogen atoms trapped, although ALPHA does not generally observe evidence for mirror trapped antiprotons under normal circumstances.

1.3 Antihydrogen Spectroscopy in ALPHA-2

With ALPHA's increased antihydrogen confinement lifetime [51] and the capability to trap ever increasing numbers of antihydrogen atoms [5], came the dawn of antihydrogen spectroscopy. Since CPT symmetry asserts that the energy levels of antihydrogen are identical to those of hydrogen, spectroscopic experiments can be compared to precise measurements in hydrogen as a stringent test of CPT invariance. In the following section I will briefly introduce the spectroscopic results achieved in the ALPHA-2 apparatus in recent years.

In the magnetic field of the ALPHA-2 AT, the antihydrogen energy levels are Zeeman shifted, as shown in Figure 1.6 (assuming CPT invariance holds) for the ground state (principal quantum number, $n = 1$) and first excited state ($n = 2$). The atomic transitions probed by ALPHA in the range of spectroscopic measurements are shown as solid coloured arrows in Figure 1.6, and the corresponding decays from the excited states are indicated by dashed arrows in corresponding colours.




Figure 1.6: Energy levels of the principal quantum number, $n = 1$ and $n = 2$, states in antihydrogen as a function of magnetic field. The first value in the ket notation represents the projection of the total angular momentum of the positron, $m_L + m_S$, where L is the orbital angular momentum (in S states $L = 0$ and in P states $L = 1$) and $S (= 1/2)$ is the spin. The second value in the ket notation represents the antiproton spin up (\uparrow) or down (\downarrow). Transitions stimulated in ALPHA experiments are indicated by solid arrows and the probable decays from those excited states are shown as dotted arrows. Energy levels at zero-field are shown in spectroscopic notation: nL_j , where j is the total angular momentum quantum number. At high field, $L = 0$ (S state) hyperfine levels are indicated by a subscript letter, whereas for $L = 1$ (P states), fine structure levels are indicated by a subscript letter and hyperfine levels are indicated by a further $+/-$ subscript. Figure adapted from [7].

As will be explained in Section 3.1, antihydrogen formed in the magnetic minimum trap cascades quickly from its initial high principal quantum number state to the ground state [74] [75]. Since $1S_d$ and $1S_c$ hyperfine states seek the low field, on account of their energy levels increasing with magnetic field strength, they are not confined by the magnetic minimum trap. On the other hand, the energy of the $1S_a$ and $1S_b$ hyperfine states decreases with increasing magnetic field strength, meaning the subset of antihydrogen atoms in these hyperfine states can be confined by the magnetic minimum trap. As a result, the distribution of confined antihydrogen contains only one positron spin-state, known as a singly spin-polarised distribution.

A tuned, oscillating magnetic field, on resonance with the ($1S_d \rightarrow 1S_a$ or $1S_c \rightarrow 1S_b$) transition frequency, can be used to stimulate a positron spin-flip to an untrappable ($1S_a$ or $1S_b$) hyperfine state. In ALPHA, such transitions are driven using resonant microwaves at about 29 GHz. Doubly spin-polarised (containing only one hyperfine state) samples are often desirable in antihydrogen spectroscopy experiments as they simplify the experimental procedure. Such samples are often obtained by stimulating a positron spin-flip using resonant microwaves of the $1S_c$ state anti-atoms into the untrappable $1S_b$ hyperfine state (grey arrow, Figure 1.6), leaving only $1S_d$ state trapped antihydrogen atoms [52]. This was the first atomic transition to be driven in antihydrogen, and this technique is used in some of the spectroscopic studies described below.

1.3.1 Observation of the Hyperfine Spectrum of Antihydrogen

As mentioned above, resonant microwaves at about 29 GHz can be used to stimulate the $1S_c \rightarrow 1S_b$ (grey arrow, Figure 1.6) or $1S_d \rightarrow 1S_a$ (brown arrow, Figure 1.6) transition in antihydrogen. The untrappable $1S_b$ or $1S_a$ state anti-atoms then annihilate with internal trap structures, producing a detectable signal. The frequency of both transitions can be measured by sweeping the microwave frequency across the resonance of the transition, where a peak in annihilation signal indicates resonant microwaves. The difference in the two transition frequencies represents the ground-state field-independent hyperfine splitting which was measured by ALPHA in 2017 as $1,420.4 \pm 0.5$ MHz using this technique [8]. The measurement uncertainty comes from magnetic field drifts during the scan (0.3 MHz), contributions from the procedure used for combining the data from separate days (0.3 MHz) and from the uncertainty in the determination of the onset frequencies (reflecting the precision of the frequency range over which the peak in annihilation signal could be determined) of the two transitions (0.3 MHz) [8]. Precision may be improved by reducing the discrete microwave frequency steps, by balancing the microwave power at the two transitions or by measuring or controlling the stability of the magnetic field [8]. A full description of the technique to determine the ground state hyperfine splitting in antihydrogen is given in [8].

1.3.2 Observation of the 1S-2P Transition in Trapped Antihydrogen

The purple line in Figure 1.6 indicates the transition from the $1S_d$ to $2P_c$ state that can be excited by radiation of wavelength 121.6 nm. Linearly polarised pulsed radiation at 121.6 nm was produced in a krypton/argon gas medium via third-harmonic generation of 364.7 nm radiation, which was generated by frequency doubling pulse-amplified continuous-wave-seeded 729.4 nm radiation in a solid-state laser system [7]. Excited $2P_c$ state anti-atoms can decay back to either the trappable $1S_d$ or untrappable $1S_a$ states [57]. In the latter case the annihilations can be detected by the SVD and provide the experimental diagnostic. When sweeping the frequency of the ultraviolet laser through resonance with the transition, the peak detected annihilation signal indicates the transition frequency, and the width of the lineshape traced out indicates the precision to which the transition frequency can be determined.

In 2018 the ALPHA collaboration used this technique to measure the frequency of the $1S_d$ to $2P_c$ transition at 1.033 T to be $2,466,051.7 \pm 0.12$ GHz [57]. The broad natural linewidth of the transition of $2\pi \times 99$ MHz makes it undesirable as a precision CPT probe, but observation of the transition was a step towards laser cooling of antihydrogen [57] [7].

The experimental linewidth of the transition is dominated by Doppler broadening [57]. On account of the average speed of the trapped antihydrogen atom parallel to the laser beam, v_L , in the reference frame of the antihydrogen atom, the frequency of the incoming photon is Doppler shifted by,

$$\Delta f_{\text{FWHM}} = \frac{2\sqrt{\ln 2}v_L}{\lambda}, \quad (1.9)$$

where λ is the rest wavelength of the incoming photon and Δf_{FWHM} is the change in frequency spread of the full-width at half-maximum (FWHM) [76]. Since Δf_{FWHM} is proportional to v_L , the precision of the measurement can be increased by decreasing the energy of the antihydrogen atoms parallel to the laser beam. A full description of the technique to measure the 1S-2P transition in antihydrogen is given in [57].

1.3.3 1S-2S Transition in Trapped Antihydrogen

The $1S_d \rightarrow 2S_d$ transition can be excited by two counter-propagating 243 nm photons on a doubly spin-polarised sample of $1S_d$ state antihydrogen atoms, as indicated by the two stacked black arrows in Figure 1.6. The transition has a long lifetime of 0.125 s and therefore a narrow natural linewidth, making it an ideal transition as a precision CPT probe in comparison with hydrogen. Excited $2S_d$ state anti-atoms can be lost either by ionisation by a third 243 nm photon, or by a positron spin-flip transition to the untrappable $1S_a$ state via coupling to the 2P state [77], as indicated by the black dotted arrow in Figure 1.6. Annihilations via the two loss mechanisms are detected by the SVD and the transition lineshape can be traced by sweeping the laser through resonance.

In 2017 the ALPHA collaboration used this technique to measure the 1S-2S transition frequency to a precision of about 5 kHz at 2.5×10^{15} Hz [56], showing consistency with CPT at a relative precision of 2×10^{-12} . Since the transition is excited by two counter-propagating photons, it is free of the first-order Doppler shift and the dominating effect that broadens the linewidth is transit-time broadening which arises from the Heisenberg Uncertainty Principle on account of the finite interaction time of the antihydrogen atom with the laser beam [76], Transit-time broadens the FWHM of the lineshape by

$$\Delta f_{\text{FWHM}} = 2\sqrt{\ln(2)} \frac{v_T}{\pi w_0}, \quad (1.10)$$

where w_0 is the waist of the Gaussian laser beam and v_T is the velocity of the antihydrogen atom perpendicular to the laser beam [77]. Hence, increasing the precision of this measurement requires decreasing the velocity of trapped antihydrogen perpendicular to the laser beam. More detail is provided in Chapter 6.3, and a full description is given in [55], [56] and [77].

1.3.4 Investigations of the Fine Structure of Antihydrogen

The zero-field fine structure splitting of the $2P_{3/2}$ and $2P_{1/2}$ levels can be found by exciting the purple, green, pink and orange transitions shown in Figure 1.6, and extrapolating to zero-field [58]. Experiments can also first doubly spin polarise the samples and excite $1S_d$ state purple and green transitions only [58]. The excited $2P$ state anti-atoms can decay back to the original trappable hyperfine state or undergo a positron spin flip into an untrappable hyperfine state, in which case annihilations can be detected by the SVD. Again, frequency sweeps of the laser through the transition frequency allow the transitions to be measured and the weighted average of transition frequencies obtained with singly and doubly spin polarised samples is taken [58].

Using this method, in 2018 the ALPHA collaboration inferred the zero-field fine structure splitting of the $2P_{3/2}$ and $2P_{1/2}$ levels to be 10.88 ± 0.19 GHz [58]. The result can be combined with the most accurate measurement of the 1S-2S transition (see previous subsection), to determine the splitting between the $2P_{1/2}$ and $2S_{1/2}$ states at zero field, known as the Lamb shift which can be explained by effects of Quantum Electrodynamics (QED) [78] [79] [80]. This method gives a Lamb shift of 0.99 ± 0.11 GHz [58], which is consistent with the most precise measurement of the hydrogen Lamb shift, 1.0578298 ± 0.0000032 GHz [45].

The uncertainty of this measurement is dominated by Doppler broadening which is proportional to the antihydrogen velocity parallel to the laser beam (Equation 1.9). A full description of the technique to determine the zero-field fine structure splitting in antihydrogen is given in [58].

1.4 Measuring Antihydrogen Gravity in ALPHA-g

Charge neutral antihydrogen trapped in a magnetic minimum trap can be used to probe the gravitational interaction of an antimatter particle and the matter Earth, as a fundamental test of Einstein’s Weak Equivalence Principle (WEP), which asserts that antihydrogen should fall exactly as hydrogen falls. Assuming antihydrogen experiences a modified gravitational force,

$$F_{\bar{H}} = F \cdot Mg, \quad (1.11)$$

where M is the antihydrogen mass and g is the acceleration of hydrogen due to the Earth’s gravitational field, gravity experiments in ALPHA measure F which can be interpreted as antihydrogen’s ratio of gravitational to inertial mass. The first such experiment of its kind took place in ALPHA-2 in 2013, and ruled out $F > 75$ at a statistical significance level of 5 % [53]. The pioneering experiment acted as a proof-of-principle which led to the proposal of ALPHA’s new antihydrogen trap: ALPHA-g [59], which was first constructed in 2018 and is yet to confine antihydrogen.

In 2013 ALPHA could only confine small numbers of antihydrogen atoms, which were released from the trap during the near-exponential decay of the magnetic fields with a time constant ~ 9 ms to ensure the detected annihilation signal was above the background [53]. This is in contrast to the modern Fast Ramp Down (FRD) technique in which the trap is removed linearly in a few seconds. During removal of the magnetic trap, the antihydrogen trajectories became increasingly sensitive to the force of gravity which is perpendicular to the axis of the ALPHA-2 AT (indicated by an arrow in Figure 1.1). According to the WEP, antihydrogen atoms experience a downward force due to the Earth’s gravitational field. Effects other than the gravitational force can cause the anti-atoms to escape the trap at a given height, for example the finite energy of trapped antihydrogen, but the gravitational force gives rise to a vertical asymmetry in annihilation location. The time and location in the plane transverse to the z -axis (defined in Figure 2.1) of detected annihilations were compared with Monte Carlo simulations in which F is a free parameter, allowing ranges of F (> 75 at a statistical significance level of 5 %) to be ruled out [81].

Assuming antihydrogen experiences the modified gravitational force given by Equation 1.11, the modified gravitational potential is then,

$$U_{\bar{g}} = Fmgh \quad (1.12)$$

where h is the vertical height of the antihydrogen atom relative to the bottom of the trap. In ALPHA-2 h is perpendicular to the trap (z) axis, which limits the sensitivity of the measurement. The ALPHA-g AT is rotated relative to the ALPHA-2 AT such that h is parallel to the trap (z) axis (see Figure 1.1).




Figure 1.7: ALPHA-g superconducting magnets, radial extent not to scale. The direction of the gravitational force is indicated by an arrow labelled ‘g’ and the trap z -axis (parallel to the gravitational force) is indicated by an arrow labelled ‘z’. Figure courtesy of Chukman So.

ALPHA-g is designed to use techniques developed in ALPHA-2 to confine antihydrogen. The ALPHA-g AT is also a hybrid Penning-Malmberg and magnetic minimum trap capable of confining both charged constituent plasmas and neutral antihydrogen. Antiprotons and positrons are magnetically guided by beamlines into the ALPHA-g AT, where the *Interconnect Magnet* is responsible for bending the charged trajectories around the corner (Figure 1.1). The ALPHA-g AT consists of a large number of superconducting magnets as shown in Figure 1.7. As in ALPHA-2, a position-sensitive detector sits on the outside of the magnetic minimum trap magnets and the vacuum wall, but on the inside of a uniform axial 1 T magnetic field.

Two measurements are intended to be made in ALPHA-g. The goal of the first is to determine the sign of the gravitational force on antihydrogen, known as the *up-down measurement*. In Figure 1.8 the total on-axis potential,

$$U(r = 0, \phi, z) = \mu_B(|\mathbf{B}(r, \phi, z)| - |\mathbf{B}(\mathbf{x}_{\min})|) + Fmgz, \quad (1.13)$$

is shown for an example value of F , where the two peaks in $U(r = 0, \phi, z)$ at high $|z|$ are caused by the confining Mirror Coils and (r, ϕ, z) are cylindrical coordinates. The principle of the up-down measurement is to compensate to some extent for the gravitational potential by applying a slightly different current to the upper/lower Mirror coil, and then to gradually ramp down the Mirror Coils retaining the magnetic potential compensation during the ramp-down. If the gravitational force is perfectly compensated (green dashed line, Figure 1.8) antihydrogen annihilations will be symmetric about the z -axis, whereas asymmetric annihilations are an indicator of overcompensation (red dashed line) or undercompensation (blue dashed line).

For the up-down measurement, antihydrogen will be formed in the up-down measurement magnetic minimum trap shown in Figure 1.7 where strong radial confinement is provided by both the *Short Octupole* and *Long Octupole*. During formation, *Mirrors A, B, F* and *G* are all energised, trapping antihydrogen between B and F, which are both slowly (compared to the antihydrogen speed) ramped to 0 A, resulting in confinement between A and G. During trap expansion, anti-atoms are expected to undergo adiabatic cooling, losing energy through collisions with the receding magnetic walls, a mechanism that will be explored in detail in this

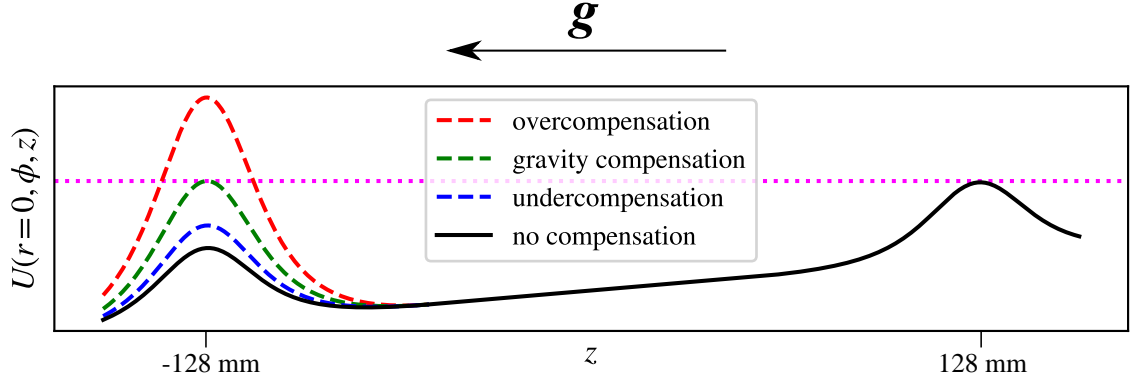


Figure 1.8: Illustration of gravity compensation in ALPHA-g. Shown is the total potential experienced by an antihydrogen atom (Equation 1.13) for an example value of F (exaggerated with respect to the hydrogen value $F = 1$). For this particular value of F in the *no compensation* potential, antihydrogen atoms would be more likely to escape at lower z . With correct *gravity compensation*, antihydrogen annihilations will be z -symmetric. Figure drawn as an indicator, not to scale. The ticks on the x -axis are labelled to indicate the separation between Mirror A and Mirror G in the up-down measurement trap.

thesis. The slower antihydrogen atoms become more sensitive to the gravitational potential.

Prior to antihydrogen release, the Long Octupole is ramped to 0 A since its magnetic windings are asymmetric about the trapping region. This process evaporates fewer antihydrogen atoms on account of the aforementioned adiabatic cooling. The experimental procedure of trapping and magnetic release with gravity compensation will be repeated for different degrees of compensation (as shown in Figure 1.8). Symmetric annihilations about the z -axis indicate correct gravity compensation, from which the direction of the gravitational force can be inferred.

The second measurement in ALPHA-g intends to increase the precision of the measurement as a test of the WEP at the 1% level. The same process of trapping and adiabatic expansion is followed by a transfer to the *precision measurement* magnetic minimum trap in which the Long Octupole provides radial confinement and the *Analysis Coils* (A_n) confine axially. The Long Octupole provides a weaker confining potential that is more accurate since it has fewer windings and the fact that the analysis trapping region is further from its less uniform wire end turns (the connections between the eight poles). At the 1% level, the measurement is sensitive to the persistent magnetic fields of the superconducting coils, which is mitigated by placing the precision trap between two copies of the up-down trap that are powered in series (so that persistence fields are largely symmetric about the precision region).

For the precision ALPHA-g measurement, adiabatic expansion will be critical for the confinement of large numbers of antihydrogen in the shallow precision measurement trap. In addition, sensitivity to the gravitational force is increased by reducing antihydrogen velocity

parallel to the gravitational force.

1.5 Bounding the Charge of Antihydrogen in ALPHA-2

CPT symmetry asserts that the charge of an antimatter particle be equal and opposite to that of its matter counterpart. Therefore, according to CPT symmetry the charge of antihydrogen is the same as that of hydrogen, which is zero. In 2014 the ALPHA collaboration set the first bound on the antihydrogen charge of $q/e < (-13 \pm 11 \pm 4)$ ppb, where the stated errors are from statistics and known systematic errors respectively (as described in [54]), q is the antihydrogen charge and e is the unit charge [54]. The principle of the measurement is that a hypothetically charged antihydrogen atom would experience a force, \mathbf{F} , in an electric field, \mathbf{E} , that is proportional to its charge on account of the Lorentz force law, $\mathbf{F} = q(\mathbf{E} + \mathbf{v} \times \mathbf{B})$. The Penning-Malmberg trap electrodes were used to generate a gradient in electrostatic potential, ϕ , across the trap, which gives rise to an axial electric field via $\mathbf{E} = -\nabla\phi$. During the near-exponential release of the magnetic minimum trap, hypothetically charged antihydrogen atoms would be deflected to one side of the trap (dependent on the sign of the applied electric field) and the axial annihilation locations were indicative of the magnitude and sign of the antihydrogen charge. The charge bound was set by comparison of experimental annihilation locations with those extracted from Monte Carlo simulations of antihydrogen atoms with charge as a free parameter. Measurement precision was limited by the achievable magnitude of the applied electric field and the spatial detector resolution [53].

In 2016 the ALPHA collaboration set a new bound on the antihydrogen charge of $|q/e| < 0.71$ ppb, using modified method in which the antihydrogen atoms are subjected to kicks from an electric field that varies stochastically [11]. Hypothetically charged anti-atoms undergo a random walk in energy that leads them to escape the trap. By interleaving experimental trials with stochastic electric fields on and off, asymmetries between the number of trapped antihydrogen atoms surviving until magnetic ramp-down between stochastic and null trials are an indicator of non-zero charge. Comparison with Monte Carlo simulation provides the bound [11]. A detailed description of this experiment is provided in Section 6.2.1 and in [11] and [12].

Reducing the energy of the trapped antihydrogen makes it possible to reduce the depth of the magnetic confining potential whilst retaining large numbers of trapped anti-atoms. Performing the charge neutrality experiment using the stochastic acceleration method with a lower trap depth would mean the random walk would have to transfer less energy to the anti-atoms to cause them to escape. Therefore, the experiment would become more sensitive to a smaller antihydrogen charge.

1.6 Antihydrogen Cooling Techniques

ALPHA intends to employ two main techniques to reduce the energy of a distribution of trapped antihydrogen: laser cooling and adiabatic expansion cooling. In addition, the mean energy of a distribution of trapped antihydrogen can be reduced by lowering the depth of the magnetic trap to allow the highest energy antihydrogen atoms to be released, with the downside of reducing the number of trapped anti-atoms. Note that this is not typical atomic evaporative cooling (EVC, introduced in Section 1.2.3) where the remaining trapped particles will re-thermalise, since collisions between trapped antihydrogen atoms in ALPHA are negligible.

As well as techniques to cool trapped antihydrogen, a number of techniques are used to produce colder antihydrogen. Since, as introduced in Section 1.2.3, the antiprotons are expected to thermalise at the temperature of the positron plasma prior to mixing, the temperature of the formed antihydrogen distribution is critically dependent on the temperature of the positron plasma. Techniques to achieve lower positron temperatures involve EVC [82], adiabatic expansion and more recently sympathetic cooling with laser cooled Be^+ ions [83].

Because the trapped antihydrogen distribution is truncated at ~ 0.5 K from a relatively hot distribution, the main advantage of a colder initial positron plasma is the expected increase in the trapping efficiency, though it is also expected to cause only a relatively small shift in the energy distribution of the trapped population, due to current plasma temperatures relative to typical trap depths. Since achievable precision of spectroscopy (Section 1.3), gravity (Section 1.4) and charge neutrality (Section 1.5) measurements is limited by the energy of the trapped antihydrogen distribution, techniques to cool the anti-atoms after trapping are of great importance.

1.6.1 Laser Cooled Antihydrogen

The observation of the 1S-2P transition [57] motivated the possibility to laser cool trapped antihydrogen. The ALPHA collaboration recently reported Doppler cooled antihydrogen [7] on a singly spin polarised sample of $1S_d$ state trapped antihydrogen. The ALPHA-2 experiment used a red detuned laser almost parallel to the trap axis to excite a transition to the $2P_a$ state (red arrow, Figure 1.6). Since the cooling laser beam is red detuned such that the photon frequency is slightly below resonance of the cooling transition for the anti-atom at rest, the anti-atom will have a higher probability of absorption of the photon (which is massless, but carries momentum) when its velocity is opposite to that of the photon. The photon momentum at 121.6 nm leads to a velocity change to the antihydrogen atom of 3.3 m/s. After photon absorption, the anti-atom decays back to the $1S_d$ state via photon

emission (excited state lifetime is ~ 1.6 ns) in a random direction. The maximum energy change is achieved when the photon is reemitted in the opposite direction to absorption. This technique primarily cools the component of antihydrogen velocity parallel to the cooling laser. Doppler cooling can achieve antihydrogen temperatures ~ 20 mK and therefore has the capability to greatly improve precision antihydrogen experiments, as was highlighted by ALPHA's observed narrowing of the 1S-2S transition linewidth following laser cooling [7]. Since (as mentioned in Section 1.3) the 1S-2S transition linewidth is dominated by transit-time-broadening, narrowing in the measured 1S-2S linewidth of a population of laser cooled antihydrogen suggests the population was cooled in the direction perpendicular to the laser beam, which is expected to result from the non-trivial orbit mixing dynamics described in Section 3.3 [7] [9]. A more detailed description of the technique to laser cool antihydrogen is provided in [7].

1.6.2 Adiabatic Expansion Cooling of Antihydrogen

As raised in discussions of proposed gravity measurements in Section 1.4, antihydrogen can be initially confined within a small magnetic volume (between two confining Mirror Coils that are close together), before the trap volume is expanded slowly (compared to the antihydrogen speed) to a larger trapping volume (where axial confinement is provided by Mirror Coils that are separated by a larger distance). During the magnetic trap expansion, antihydrogen atoms are expected to lose energy via collisions with the receding magnetic barrier in a process known as adiabatic cooling.

The adiabatic cooling technique will be critical for measurements of antihydrogen gravity at the 1% level and could be used in conjunction with laser cooling in ALPHA-2 to further decrease antihydrogen energy for spectroscopy measurements. For example, antihydrogen could be confined within a small volume trap, where it is laser cooled to the laser cooling limit (which is above the Doppler limit in practice on account of the field not being completely uniform), before the trap is adiabatically expanded, reducing the energy of the trapped antihydrogen even further. In addition, adiabatic expansion is a relatively simple technique (requires no laser interactions) that could gain increases in precision in measurements of the antihydrogen charge introduced in Section 1.5 in combination with increases in trapping efficiency.

In 2016, the ALPHA collaboration trapped antihydrogen atoms in a small magnetic trapping volume in the ALPHA-2 trap, before expanding the volume of the magnetic trap in ~ 23 s in a proof-of-principle experiment. The magnetic trap was then linearly ramped down in ~ 2 s, during which the time of annihilation of the antihydrogen atoms on internal trap structures was recorded. This annihilation time distribution was compared to two experimental control trials in which antihydrogen was held in a static magnetic confining

potential for ~ 23 s, before also being released from the trap during magnetic ramp down. The antihydrogen atoms in the expanded trap volume were observed to annihilate later in time than the control populations, suggesting the anti-atoms that were confined in the slowly expanding magnetic volume had a lower energy than the control populations at the time of magnetic ramp-down. This observation suggests the energy of the antihydrogen atoms was adiabatically cooled, and that the energy of the population was reduced below typical energies of trapped antihydrogen in ALPHA-2, where antihydrogen is normally formed and held in the larger magnetic volume.

At present, Monte Carlo simulations of the experimental procedure are required to confirm that the observed tendency for the adiabatically expanded populations to annihilate later in time than control populations is indeed an indicator of energy reduction via adiabatic cooling since we had no means of measuring the energy of the trapped antihydrogen atoms directly. Further, simulations can provide insight into the magnitude of achievable cooling, including the magnitude of cooling in the direction transverse to the expansion which we expect would be primarily caused by non-trivial orbit mixing dynamics [9]. Ultimately, simulations will enable optimisation of the adiabatic cooling technique for application to precision measurements of the fundamental properties of antihydrogen.

To simulate adiabatic expansion cooling of antihydrogen, we are required to track the motion of a magnetic moment in a time-dependent magnetic field. Mixing of antiproton and positron plasmas for antihydrogen formation has been simulated [84], but since the simulations are computationally extensive, we opt for Monte Carlo methods which involve pseudo-random initialisation based on informed knowledge of antihydrogen formation. Such simulations have been used by ALPHA for many years and have been validated by comparison against experimental data [51]. Once initialised, the equations of motion of the anti-atom must be solved in a way that reproduces physical characteristics of the trajectories. For example, in a static field the total energy of an antihydrogen atom must be conserved within bounds on long timescales. Details of the methods used to simulate adiabatic expansion cooling of antihydrogen are provided in Section 3.4 and the results of such simulations are presented throughout the entire thesis.

1.7 Outline of this Thesis

In this thesis, through detailed Monte Carlo simulations of the experimental procedure of adiabatic expansion that took place in 2016, I will test the hypothesis that the later detected annihilation times of antihydrogen atoms, subject to a slowly expanding trap volume, indicate a reduction in energy. Additionally, I will quantify the magnitude of the change in energy that was achieved in the 2016 experiment, question whether the energy decrease during magnetic expansion is consistent with adiabatic cooling and investigate whether I can gain further

insight into the dynamics of antihydrogen during adiabatic expansion, with the ultimate goal of exploiting adiabatic expansion to improve precision measurements on antihydrogen.

In Chapter 2 the reader will be introduced to the adiabatic expansion experimental procedure that took place in 2016 and I will compare the experimental annihilation time data of adiabatically cooled populations to control samples held in static wells. I will determine the magnitude of the trap depth at each experimental annihilation time, to determine whether the experimental data indicates that adiabatically cooled anti-atoms tend to escape the magnetic trap when the energy required to escape is lower.

In Chapter 3 I will present the theoretical background surrounding dynamic expansion of a magnetic trap, in which neutral anti-atoms are confined by their magnetic moment. In addition, I will present in detail a description of the simulation which is used throughout this thesis to investigate the dynamics of adiabatically cooled antihydrogen.

I will first present simulation results in Chapter 4, where I will compare outputs of detailed simulations of the experimental procedure of adiabatic expansion to experimental data, and quantify agreement using Kolmogorov Smirnov statistical analysis. Given that the simulated and experimental annihilation-time data are in sufficient agreement, I will extract energy distributions from the simulation. I will use this to quantify the energy decrease achieved during the adiabatic expansion experiment, as well as making a comparison to the energy distributions of control populations held in a static well. Further, I will use an established technique to deduce the energy of the experimental antihydrogen atoms using the simulated annihilation times and energies.

In Chapter 5 I will determine whether the simulated decrease in energy presented in Chapter 4 is consistent with adiabatic cooling by comparing to two semi-analytic models.

In Chapter 6 I will begin to introduce techniques to optimise adiabatic cooling by investigating the effect of changing the expansion duration. Also in this chapter, I will begin to investigate the applications of adiabatic expansion by presenting simulations of the charge neutrality experiment using adiabatically cooled antihydrogen atoms. I will then provide insight into orbital dynamics of adiabatically cooled antihydrogen atoms that could hinder their use in precision spectroscopy experiments, such as those raised in Section 1.3.

The work presented in this thesis is a detailed review of the author's understanding of adiabatic expansion of antihydrogen atoms. The technique is yet to be fully optimised and protocols of its application to precision antimatter experiments continue to be developed.

1.8 Author Contributions

The adiabatic expansion experiment (described in Chapter 2) was proposed by Prof. Joel Fajans and performed by the ALPHA Collaboration in 2016. The author was not a member of the ALPHA collaboration at the time and therefore did not contribute to the experimental

procedure.

The principles of simulating antihydrogen in ALPHA's magnetic minimum trap have been developed over many years, by Prof. Francis Robicheaux, Prof. Joel Fajans and other members of the ALPHA Collaboration. The author made large contributions to the development of an object-oriented C++ antihydrogen simulation, based on established principles of simulating antihydrogen alongside Dr Joseph McKenna, Peter Granum and (also under the supervision of) Dr William Bertsche. Many sections of the simulation rely on code originally written by Prof. Francis Robicheaux. The studies presented in this thesis use the same simulation diagnostics of antihydrogen axial to transverse energy mixing as used to present results in [9], written by the original authors (Adrienne Zhong, Prof. Joel Fajans and Abe Zukor).

The author ran simulations of the adiabatic expansion experiment performed in 2016 and wrote simulation diagnostics to extract the energy of the trapped antihydrogen distributions. The author performed the statistical comparisons of experimental and simulated data (Section 4.1), determined the energy of the trapped antihydrogen distributions in the experimental adiabatic expansion trials via simulation (Section 4.3), reconstructed the energy of the experimental distribution of antihydrogen using an established technique (Section 4.5), performed convergence studies (Section 4.7), and analysed and interpreted annihilation patterns during Octupole Fast Ramp Down (Section 4.8).

The author derived the adiabatic cooling models presented in Chapter 5 and wrote the algorithm used to determine the volume and axial length of the magnetic trap which enter the adiabatic models.

The author ran and analysed simulations of different expansion durations presented in Chapter 6. Simulations of adiabatic cooling prior to performing the charge neutrality experiment, also presented in Chapter 6, were performed originally by undergraduate student Evan Deddo, and reproduced and optimised by undergraduate student Nicolas Kalem. The author supervised the work of both students alongside Prof. Joel Fajans and Prof. Jonathan Wurtele. Simulation and analysis of the dynamics of adiabatically cooled antihydrogen for application to 1S-2S spectroscopy of antihydrogen presented in Chapter 6 are the author's work.

The author originally wrote, ran and analysed results of the simple laser cooling simulation discussed in the conclusion, but the simulation complexity was developed in collaboration with Edward Thorpe-Woods, guided by the work of Prof. Francis Robicheaux. Following these developments, the author re-ran and analysed simulations of antihydrogen laser cooling.

Adiabatic Cooling Experimental Procedure




Figure 2.1: ALPHA-2 neutral atom trap (AT). Penning-Malmberg trap electrodes (yellow) confine antiprotons and positrons axially. In Trials A and B, antihydrogen is formed at E11 (pink), whereas in E12 (orange) for Trial C. AT magnets are coaxial to the electrode stack: short solenoids, known as Mirrors A-E (blue), Octupole (red) and Solenoids A and B (green). The External Solenoid surrounds the entire AT and is not shown here. The z -axis is defined along the trap axis as shown.

The first demonstration of adiabatic cooling of antihydrogen atoms took place in a short experimental campaign in 2016. This following chapter will detail the experimental procedure and present the experimental data.

The goal of the proof-of-principle adiabatic expansion experiment was to demonstrate that antihydrogen atoms, initially confined in a relatively small magnetic trap that is gradually expanded to a larger magnetic trapping volume, lose kinetic energy during the expansion. When the magnetic trap expands sufficiently slowly such that

$$\tau \frac{dU}{dt} \ll U, \quad (2.1)$$

where τ is the period of a particle's oscillation in the trap potential and U is the magnetic potential in the trap relative to the magnetic minimum (given by Equation 1.4), changes to the particle's Hamiltonian are sufficiently slow that the phase space volume of the trajectory is conserved. As a result, a particle loses kinetic energy as the length of its trajectory increases. After the trap expansion is finished, the magnetic trap is then gradually removed and the annihilation time of the anti-atoms provides the experimental energy diagnostic, which relies on the hypothesis that *slower* antihydrogen atoms tend to escape and annihilate when the magnetic trap is *shallower*, which happens *later* in a monotonic magnetic ramp down. Therefore, comparing annihilation times of an adiabatically cooled population to control samples held in a static well allows us to determine whether the magnetic trap expansion leads to a reduction in total (kinetic plus magnetic potential) energy as a result of adiabatic cooling.

In this experiment, adiabatic cooling was achieved by forming antihydrogen in a short magnetic trap, where axial confinement is provided by Mirror Coils (see Figure 2.1) that are close together, before the Mirror Coil currents were manipulated to cause an axial trap expansion. The experiment consisted of three experimental trials which we will refer to as *Trial A*, *Trial B* and *Trial C*. Trial A is the adiabatic cooling trial in which the trap volume expands axially. Trial B is a control trial which enables comparison to a distribution of antihydrogen which was initialised in a well identical to that of Trial A but which did not undergo an expansion. Trial C is another control trial which was used as a reference to typical antihydrogen confinement used for antihydrogen spectroscopy experiments in 2016.

Trials A and B require antihydrogen formation in a magnetic trap with a smaller volume relative to that of Trial C. This smaller initial magnetic trap is generated by energising Mirrors A, C, D and E to allow trapping between Mirror A and Mirror C. In Figure 2.2(a) we show the initial (at the time of antihydrogen formation) on-axis magnetic potential, $U(r = 0, \phi, z)$, for Trials A and B. By contrast, the Trial C initial magnetic trapping potential (Figure 2.2(b)) is generated by energising Mirrors A and E, to allow trapping in a larger magnetic volume.

In the experiment, Trials A, B and C were interleaved in 13 independent runs per trial. An individual run involved forming antihydrogen by merging a single antiproton shot from the

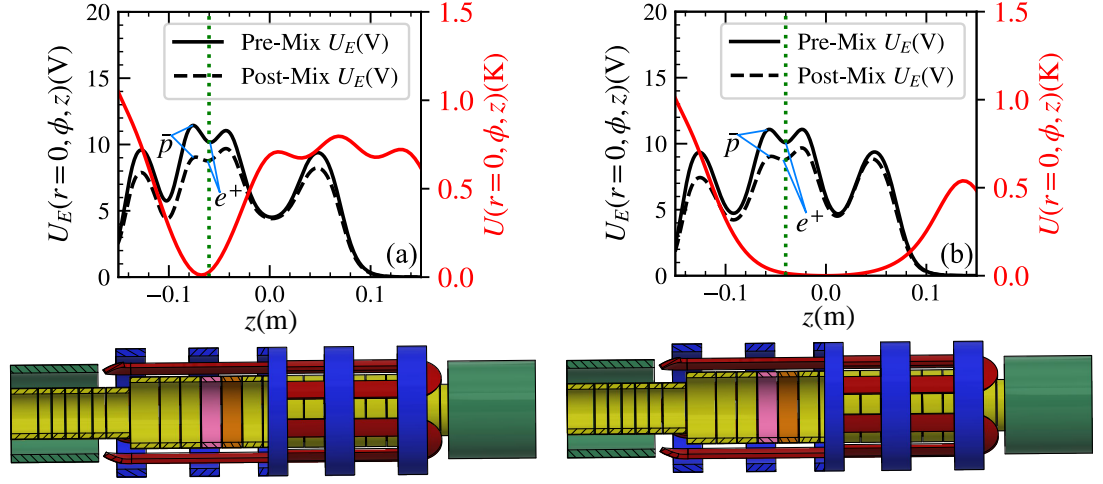


Figure 2.2: On-axis magnetic potential, $U(r = 0, \phi, z)(\text{K})$ (red solid curve) and on-axis electrostatic potential, $U_E(r = 0, \phi, z)(\text{V})$, just prior to antihydrogen formation (black solid curve, corresponding to Figure 1.4(c)) and after mixing of positron (e^+) and antiproton (\bar{p}) plasmas for antihydrogen formation (black dotted curve, corresponding to Figure 1.4(d)). Magnetic potentials at the point of antihydrogen formation ((a)–23.55s for Adiabatic Expansion Trial A and Control Trial B, (b) –24.523s for Control Trial C). $U(r = 0, \phi, z)$ is relative to the magnetic minimum and is found using the method described in Section 3.4. $U_E(r = 0, \phi, z)$ is calculated using the COMSOL Multiphysics[®] potential solver. The green dotted line represents the center of the electrode in which positrons are held during mixing (E11 in (a) and E12 in (b)).




Figure 2.3: Typical on-axis confinement potentials, $U(r = 0, \phi, z)$, in ALPHA-2. In each case the Octupole is energised to ~ 900 A, which produces a mostly radial magnetic field, generating a magnetic potential of magnitude ~ 0.5 T at the trap wall, which provides radial confinement. (a) Standard trap configuration with Mirrors B-D at 0 A and Mirrors A and E at ~ 600 A, (b) Tight configuration with Mirrors A and E at ~ 600 A, Mirrors B and D at ~ 700 A and Mirror C at ~ -200 A, (c) Flat trap configuration with Mirrors A and E at ~ 600 A, Mirrors B and D at ~ -60 A and Mirror C at ~ -2 A.

AD with positrons prepared by the Positron Accumulator. Typically, this led to formation of 5-10 antihydrogen atoms which were then either subjected to a magnetic trap expansion or held in a static magnetic trap, depending on the experimental trial.

Antihydrogen is formed in ALPHA when antiproton and positron plasmas, confined in neighboring electrostatic wells, are mixed together (see Section 1.2.3). The technique involves minimising the potential difference between the two electrostatic wells without accelerating the antiprotons and whilst continuously cooling the positrons via evaporation [3]. The merging process results in low energy antiprotons traversing the positron plasma, and hence antihydrogen is formed within the volume of the positron plasma. This means antihydrogen formation is centered around the minimum of the positron confining well prior to mixing. Therefore, the axial center of antihydrogen formation, z_{form} , can be chosen from discrete axial locations corresponding to the center of each electrode in the stack of Penning-Malmberg trap electrodes (Figure 2.1).

The magnetic potential at the center of antihydrogen formation, U_{init} , and the initial depth of the trapping potential, $\Gamma(t_{\text{init}})$, were important considerations for the adiabatic expansion experiment, as they effect the initial energy distribution of the trapped antihydrogen. They can be adjusted by changing z_{form} (by adjusting the electrode voltages), and by adjusting the currents in the magnetic minimum trap magnets. In addition, one must ensure antiproton and positron plasmas are not mixed in a large magnetic field gradient, as this can lead to expansion of the plasma, subsequent heating, and hence an inability to trap in large numbers. In Figure 2.2 we overlay the *Pre-Mix* (black solid curve) and *Post-Mix* (black dashed curve) on-axis electrostatic potentials, $U_E(r = 0, \phi, z)$, with the on-axis magnetic potential during




Figure 2.4: Currents in each magnet of the neutral trap as a function of time during the adiabatic expansion experiment for experimental trials (a) Adiabatic Expansion Trial A, (b) Control Trial B and (c) Control Trial C. Times are given relative to the start of Octupole FRD and currents are shown until the simulation end time (around 7 s).

mixing for Trials A and B (Figure 2.2(a)) and Trial C (Figure 2.2(b)).

Spectroscopy experiments in ALPHA benefit from maximising the trap volume at a uniform magnetic field (around 1 T) because the energy levels of atoms within this volume will have energy levels Zeeman shifted to the same extent. Therefore, spectroscopy experiments typically use a *flat* trap configuration, in which relatively small negative currents on the internal Mirror Coils (B, C, D) are used to flatten out the magnetic field, as shown in Figure 2.3(c). Recall that Trial C was intended to mimic antihydrogen formation in such experiments and hence, for Trial C, $\Gamma(t_{\text{init}})$ and U_{init} were chosen to match that of typical spectroscopy experiments in ALPHA as closely as possible. Unfortunately, at the time of the experiment, some of the central-most Penning-Malmberg trap electrodes were shorted together which ruled out antihydrogen formation in certain axial locations. As a consequence, the Trial C energy distribution was matched as closely as possible to spectroscopy trapping by mixing with

positrons held in E12 (orange electrode, Figure 2.1 and green dotted line in Figure 2.2(b)) with Mirror B energised to 33.0 ± 0.2 A, acting to reduce the local magnetic field. At the origin of antihydrogen formation in Trial C, $\Gamma(t_{\text{init}}) \approx 0.493$ K and $U_{\text{init}} \approx 0.014$ K.

For Trials A and B, if we were to hold positrons in E12 prior to mixing (as in Trial C), this would set $U_{\text{init}} \approx 0.21$ K and hence the energy distribution would be significantly higher than for Trial C. For this reason, positrons are held in E11 (pink electrode, Figure 2.1 and green dotted line in Figure 2.2(a)) during mixing and the initial magnetic trapping potential is produced by energising Mirrors A, C, D and E to 539.4 ± 1.5 A (currents measured experimentally, as described later in this section), the Octupole to 899.4 ± 1.0 A and Solenoid A to 249.5 ± 0.5 A. Mirror B was initialised with a current of -252.6 ± 1.0 A which, since it sets the minimum magnetic field, sets $\Gamma(t_{\text{init}}) \approx 0.490$ K and $U_{\text{init}} \approx 0.028$ K. The difference in U_{init} between Trial A/B and Trial C of $\Delta U_{\text{init}} \approx 0.014$ K, and the difference in $\Gamma(t_{\text{init}})$, $\Delta \Gamma(t_{\text{init}}) \approx 0.003$ K, will lead to some difference between the initial energy distributions which turns out to be insignificant and will be presented in Chapter 4. One could equivalently perform an adiabatic expansion experiment with initial confinement provided by Mirrors A, B, D and E, such that antihydrogen is confined between Mirrors B and D. Then the magnetic trap can be expanded to confinement between Mirrors A and E, the symmetry of which seems a natural choice. However, the shorts between central neighboring electrodes prevented this.

Positrons and antiprotons were independently prepared for mixing using methods described in Section 1.2.1. Antiprotons were caught and prepared in the CT, transferred to the AT where they underwent further preparation and evaporative cooling. Solenoid A was energised to 249.5 ± 0.5 A which raised the magnetic field in the AT antiproton preparation region to ~ 3 T which increases cyclotron cooling and is required since the plasma preparation routine is optimised to this high magnetic field. Antiprotons were then held in E10 for Trials A and B and E11 for Trial C. Positrons were simultaneously accumulated for 170 s in the Positron Accumulator, before being transferred to the AT for further preparation and then held in E11 for Trials A and B and E12 for Trial C. At this stage, positrons and antiprotons were confined in neighboring electrostatic wells (equivalent to Figure 1.4(a)). Then the magnetic minimum trap was energised (red curve in Figure 2.2). In preparation for mixing, the positrons were evaporatively cooled (Figure 1.4(a) \rightarrow 1.4(b)), and then the difference in electrostatic potential was reduced to the Pre-Mix potential (black solid curve, Figure 2.2), so that the two species almost overlap. This preparation sequence resulted in $\sim 4 \times 10^4$ antiprotons with a temperature of ~ 31.9 K and $\sim 30 \times 10^6$ positrons at ~ 50 K (measured as described in Section 1.2). In 1 s the electrostatic potential was morphed from Pre-Mix (black solid curve, Figure 2.2) to Post-mix (black dotted curve, Figure 2.2), enabling formation of a relatively hot distribution of antihydrogen via three-body recombination (as described in Section 1.2.3), which was then effectively truncated around magnetic trap depth ($\Gamma(t_{\text{init}}) \approx 0.490$ K in Trials A and B and

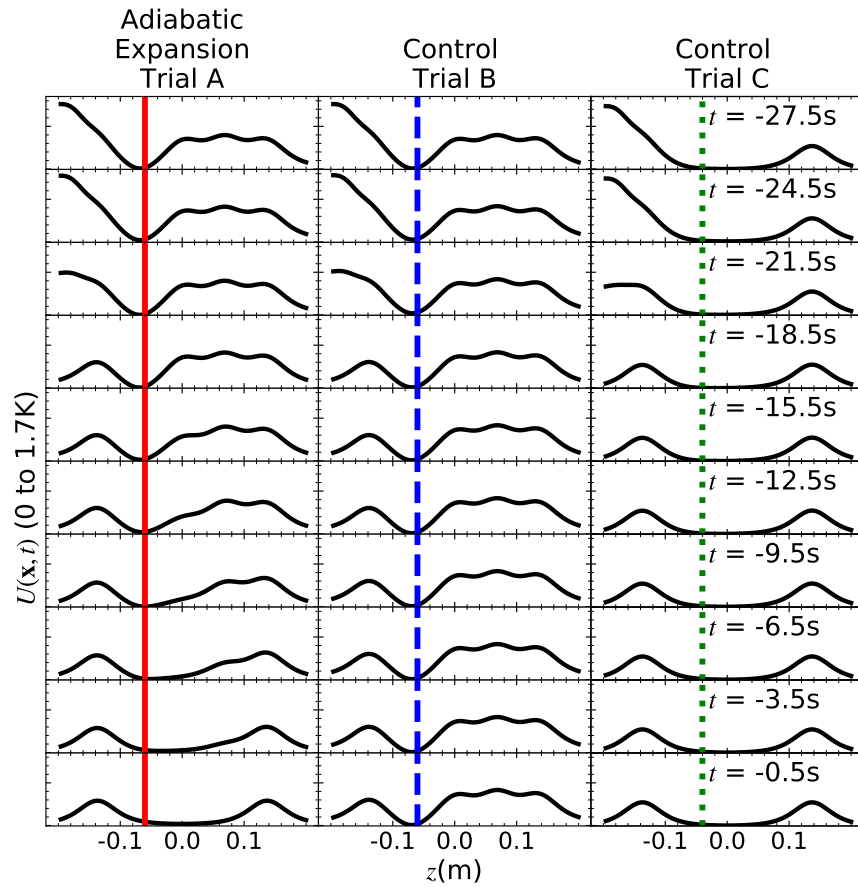


Figure 2.5: On-axis magnetic potential, $U(\mathbf{x}, t)$, as a function of axial position, z , at $x = y = 0$, for the three experimental trials. The y -axis label specifies the scale on each small graph.

$\Gamma(t_{\text{init}}) \approx 0.493 \text{ K}$ in Trial C).

After antihydrogen formation and trapping, the experiment entered the ~ 23 second-long magnetic trap manipulation phase; during which the currents in the superconducting magnets were manipulated to expand the Trial A magnetic volume. During each experimental run, the currents in the superconducting magnets were recorded at high precision but low (1 s) resolution by a Direct Current Current Transducer (DCCT). These DCCTs are coaxial to the leads carrying the current to the superconducting magnets, and are designed to output a voltage proportional to the current in the lead. The current in the magnet lead passes through three magnetisable toroidal cores, inducing a flux. A secondary current-carrying wire is wound many times around the cores. The principle of the DCCT measurement is to drive a current in the secondary current-carrying wire that cancels the flux in the cores. Since the secondary winding is wound multiple times around the cores, the current is much less than that in the magnet lead. Therefore, a measurement of the voltage across an accurately calibrated resistor provides an accurate measure of the current in the magnet lead, avoiding issues related to excessive heating of the resistor (and therefore changes in resistance) [85].

The high precision of the DCCT current measurement made it possible to estimate the errors on the superconducting magnets stated in the previous paragraphs by the variation between runs. The low resolution of this reading meant it was not possible to accurately measure the relative timing of the current ramps. However, a few days after the experiment, five runs of each experimental trial were repeated without particles and the currents in the superconducting magnets were recorded by the DCCT at 250 ms resolution. In Figure 2.4 we show the average current behaviour as a function of time for the three experimental trials, which is based on determining the average start and end time of each current ramp from the 250 ms DCCT reading, and assuming perfectly linear current ramps.

At the start of the magnet manipulation phase ($t \approx -23 \text{ s}$) Solenoid A was ramped to 0 A in 5 s in all three experimental Trials. Then, during Trial A, in which initial confinement was provided by Mirrors A and C (shorter trap), the magnetic trap was gradually expanded by first starting to ramp down the current in Mirror C and then (before the Mirror C ramp was finished) starting to ramp down the current in Mirror D, ending with confinement between Mirror A and Mirror E, as indicated by the on-axis magnetic potential as a function of time shown in Figure 2.5. The expansion is slow (Equation 2.1) compared to the axial bounce time of the anti-atoms which is $\sim 10 \text{ ms}$. During Trial A, the barrier recedes axially which is expected to lead predominantly to a reduction in axial energy. Simulations are required to determine the extent to which the transverse energy component was reduced during Trial A as a result of energy mixing between axial and transverse dimensions, as will be discussed in Section 3.3. For nontrivial reasons, there is also a small amount of radial compression which leads to minor radial heating as will be discussed in Section 5.

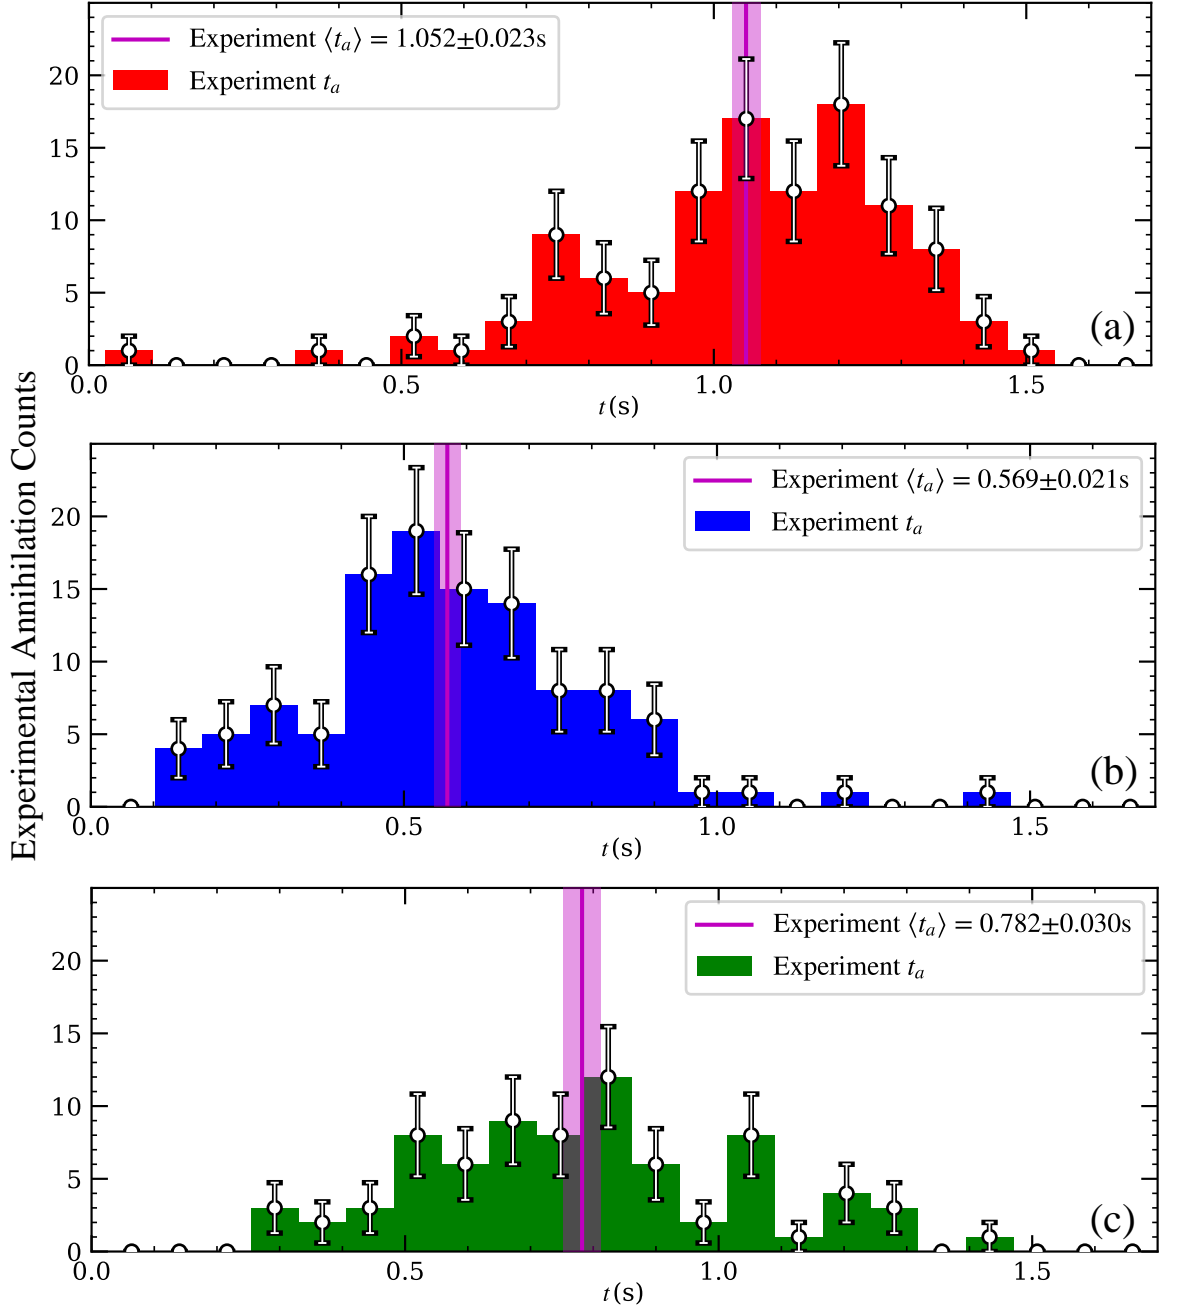


Figure 2.6: Annihilation time data of the three experimental trials: (a) Adiabatic Expansion Trial A (98 counts), (b) Control Trial B (106 counts) and (c) Control Trial C (77 counts). Times are given relative to the start of Octupole FRD and green errorbars are given by Poissonian counting statistics (\sqrt{N} , where N is the number of detected counts). The average annihilation time, $\langle t_a \rangle \pm 1\sigma_{SE}$ is shown as a vertical magenta line, and the region within $1\sigma_{SE}$ is shown as a magenta region where $\sigma_{SE} = \sqrt{\sigma^2/N}$, is the standard error on the mean and σ is the sample standard deviation.

By contrast, control Trials B and C involved keeping the magnetic trap largely static during the magnetic trap manipulation phase. Recall that a 33.0 ± 0.2 A current in Mirror B was used, in Trial C, to reduce the magnetic field in the region of antihydrogen formation. To avoid releasing the anti-atoms from an atypical magnetic configuration, Mirror B was gradually ramped to 0.0 ± 0.3 A over the course of the hold, with Trial C ending in a magnetic configuration resembling the standard trap (Figure 2.3(a)), with Mirrors A and E energised to 539.4 ± 1.5 A, the Octupole at 899.4 ± 1.0 A and the internal Mirror Coils at 0.0 ± 0.3 A. This gradual ramping of Mirror B is expected to cause minor adiabatic heating, the extent of which can be determined via simulation as will be presented in Chapter 4.

After the adiabatic expansion (or corresponding control hold) is finished, magnetic ramp down begins (at $t = 0$ s). For each of the three trials, magnetic ramp down involves linearly ramping the Octupole magnet to 0 A in ~ 1.5 s, during which the Mirror Coil currents remain static. This is a specific choice of FRD known as *Octupole FRD*. The reason for this choice is that the simultaneous (to the Octupole ramp) ramping of the Mirror Coils during regular FRD causes axial adiabatic cooling. This cooling would be more significant in Trial B than in Trial A or Trial C. As a result, the contrast between annihilation times in the Trial A (or C) and Trial B diagnostic ramps would be reduced. As the Octupole current was reduced, the annihilations of the anti-atoms on internal trap structures were detected by the SVD, the times of which were detected with an uncertainty of ~ 2 μ s. During Octupole FRD, it is mainly the radial magnetic potential that reduces, which either means particles annihilating later in time had lower radial energy or the axial energy of the trapped anti-atoms was somehow coupled to the annihilation time during Octupole FRD. Both turn out to be true to some extent as will be discussed in Chapter 4.

As mentioned earlier, the lower limit on the expansion time is set by the axial bounce time of the anti-atoms (~ 10 ms) but in reality the speed at which the superconducting Mirror Coils can be ramped from full to zero current is ~ 2 s. This limit is set by the rate at which the power supply can sink the energy stored in the magnet. The current can be removed from the magnet much quicker than this (in ~ 10 ms) if the Mirror Coils are quenched and the power supply is shorted out of the circuit. The upper limit on the expansion time is set by the finite lifetime of the trapped anti-atoms as a result of elastic and inelastic scattering collisions on background gas, which was measured to be around 1000 s around the time of this experiment. A reasonable and somewhat arbitrary expansion time of ~ 23 s was chosen for Trial A. Trial B involved holding the antihydrogen atoms in the static (short) potential for the same duration (~ 23 s) and the Trial C hold (in the larger magnetic volume) lasted ~ 24 s.

Figure 2.6 shows the experimental annihilation times during Octupole FRD for the three experimental trials. The events shown are those that pass the criteria of the machine learning based analysis and have not been adjusted according to detector efficiency ($67.6 \pm 0.3\%$) or




Figure 2.7: Trap depth, Γ , as a function of time, t (s), during Octupole FRD for (a) Trial A, (b) Trial B and (c) Trial C. Trap depth is defined as the minimum magnetic potential on the electrode inner surface relative to the minimum magnetic potential anywhere in the trap volume. Octupole FRD starts at 0s in all three trials and the current decay is assumed to be perfectly linear. Trap depth is defined as the minimum magnetic potential on the inner electrode surface that is accessible from the minimum magnetic potential within the trapping volume.

expected background rate (53 ± 4 mHz or 0.239 ± 0.018 expected cosmic ray events during diagnostic ramp down). The number of annihilation events during the detection window for each of the three trials is shown in Table 2.1, which is similar for each trial and is significantly above the expected number of background annihilations. Also shown, is the average annihilation time of the detected annihilations, $\langle t_a \rangle$, which, for Trial A, is 1.85 ± 0.08 times higher than Trial B and 1.35 ± 0.06 times higher than Trial C. Crucially, the adiabatically cooled Trial A population tends to annihilate *later* in time than the control populations of Trial B and C.

The main principle of the adiabatic expansion experimental protocol is the assumption

Trial	\bar{H} counts ($t < 0$ s)	\bar{H} counts ($t > 0$ s), N	$\langle t_a \rangle \pm 1\sigma_{SE}$ (s)	$\langle \Gamma(t_a) \rangle \pm 1\sigma_{SE}$ (K)
A	8	110	1.052 ± 0.023	0.081 ± 0.007
B	13	111	0.569 ± 0.021	0.223 ± 0.009
C	23	76	0.782 ± 0.030	0.173 ± 0.011

Table 2.1: Number of antihydrogen (denoted \bar{H}) annihilation events before FRD ($t < 0$ s), and during the ~ 1.5 s Octupole FRD detection window, N . Average detected annihilation time, $\langle t_a \rangle$, and ensemble-averaged trap depth at annihilation time, $\langle \Gamma(t_a) \rangle$, are also shown. The background rate is (53 ± 4) mHz or 0.239 ± 0.018 expected cosmic ray events during diagnostic ramp down. $\sigma_{SE} = \sigma/\sqrt{N}$ is the standard error on the mean annihilation time, $\langle t_a \rangle$, where σ is the sample standard deviation.




Figure 2.8: Distribution of energies representing the trap depth (Figure 2.7) at the experimental annihilation time, $\Gamma(t_a)$. White error bars show Poissonian (\sqrt{N}) counting statistics, where N is the number of experimental annihilation counts (see Table 2.1). The ensemble-averaged energy at annihilation time, $\langle \Gamma(t_a) \rangle \pm \sigma_{\text{SE}}$, is shown as a vertical yellow line, where $\sigma_{\text{SE}} = \sigma/\sqrt{N}$ is the standard error on the mean and σ is the sample standard deviation. The yellow region indicates energies within $1\sigma_{\text{SE}}$ of $\langle \Gamma(t_a) \rangle$. (a) Adiabatic Expansion Trial A ($\langle \Gamma(t_a) \rangle = 0.081 \pm 0.007 \text{ K}$), (b) Control Trial B ($\langle \Gamma(t_a) \rangle = 0.223 \pm 0.009 \text{ K}$), and (c) Control Trial C ($\langle \Gamma(t_a) \rangle = 0.173 \pm 0.011 \text{ K}$).

that *slower* antihydrogen atoms annihilate when the magnetic trap is *shallower* which occurs *later* in time during Octupole FRD. It is possible to approximately quantify this principle with a simple analysis: we calculate the magnetic trap depth at the experimental annihilation time of each particle. This can be interpreted as a lower bound on the energy of the anti-atom at the time of escape from the trap because the trajectory may be trapped on an orbit that does not sample the entire trapping volume as will be discussed in Chapters 3 and 5. Note that this does not represent the energy of the anti-atom at the start of Octupole FRD ($t = 0$ s), since there will be some radial adiabatic cooling during the dump on account of the radial trap expansion from the decreasing Octupole current. However, the magnitude of radial cooling should be similar between trials. The magnetic trap depth, Γ , is shown in Figure 2.7 as a function of time during Octupole FRD. A histogram of the trap depth at each experimental annihilation time, $\Gamma(t_a)$, is shown in Figure 2.8 for the three experimental trials, and the ensemble-averaged trap depth at the experimental annihilation time, $\langle \Gamma(t_a) \rangle$, is shown in Table 2.1. This naive analysis suggests the Trial A anti-atoms had a mean energy distribution that was 2.75 ± 0.72 times lower than that of Trial B and 2.14 ± 0.49 times lower than that of Trial C. Detailed simulations are required to determine the decrease in energy during adiabatic expansion more accurately and are presented in Chapter 4.

Theory

In this chapter I will discuss the range of theoretical knowledge that is required to understand the dynamics of adiabatically cooled antihydrogen atoms. I will start by describing how antihydrogen is routinely trapped in a static magnetic potential generated by a number of superconducting magnets, before deriving an expression relating to the antihydrogen motion in the trap that is invariant to very good approximation when the magnetic potential is varied slowly (compared to the particle speed). The principle of adiabatic cooling is the conservation of such a pseudo-invariant, since it states that the energy of the trapped antihydrogen atoms can be reduced as the length of an oscillating trajectory increases. I will then introduce the concept of energy mixing between axial and transverse dimensions, as it controls the extent to which an axial adiabatic expansion can cool the transverse energy (by transfer of energy from the transverse to the axial dimension) of the trapped anti-atoms. Finally, I will detail the method used to simulate antihydrogen atoms confined within the ALPHA-2 magnetic minimum trap. In this simulation, the magnetic fields can be time-varied to simulate adiabatic cooling of antihydrogen.

3.1 Antihydrogen Confinement in Octupole-Based Magnetic Minimum Traps

As discussed in Section 1.2.3, antihydrogen is formed in ALPHA when antiproton and positron plasmas, confined in neighbouring electrostatic wells, are merged together. Neutral antihydrogen is expected to be formed during the process of three-body recombination (TBR) by the interaction of two positrons with an antiproton, where one positron forms a bound state with the antiproton and the other carries away the excess recombination energy [67] [68]. The resulting antihydrogen is formed in a thermal distribution approximately at the temperature of the positron plasma (~ 50 K) and in high principal quantum number states (although the distribution of initial principal quantum number states is not accurately known). In the




Figure 3.1: The effect of Zeeman splitting on the hyperfine structure of the antihydrogen ground state, $^1S_{1/2}$. The positron spin is indicated by \uparrow / \downarrow and the antiproton spin is indicated by \uparrow / \downarrow . The yellow region indicates the range of magnetic fields present in a typical ALPHA antihydrogen confinement potential. Figure from [8].

following section I will describe how low energy antihydrogen, formed through TBR, can be confined in an octupole-based Ioffe-Pritchard magnetic minimum trap.

In ALPHA-2, the magnetic minimum trap is generated by a number of superconducting magnets. Trapping is based on the principle that a magnetic moment, $\boldsymbol{\mu}$, in a magnetic field, $\mathbf{B}(\mathbf{x})$ gives the Zeeman energy,

$$U(\mathbf{x}, t) = -\boldsymbol{\mu} \cdot \mathbf{B}(\mathbf{x}), \quad (3.1)$$

where \mathbf{x} is the position vector of the antihydrogen atom. The spin of the bound positron largely determines whether the magnetic moment is aligned parallel or anti-parallel to the magnetic field. The fraction of anti-atoms formed with positron spin anti-parallel to the magnetic field lines lie in energy levels that increase in energy as a function of magnetic field strength due to the Zeeman effect. In Figure 3.1 we show the energy of these trappable *low-field-seeking* states (in green) as a function of magnetic field strength for the ground state (principal quantum number $n = 1$). Antihydrogen is formed in high n states and can lose energy as it radiatively cascades into the ground state in a process known as *radiative cascade cooling* as will be discussed later in this section.

The magnetic potential of an anti-atom in the trap magnetic field in a low-field-seeking $n = 1$ state simplifies to,

$$U(\mathbf{x}) = \mu_B (|\mathbf{B}(\mathbf{x})| - |\mathbf{B}(\mathbf{x}_{\min})|), \quad (3.2)$$

where \mathbf{x}_{\min} is the location of the magnetic field minimum, μ_B is the Bohr magneton and we have defined the magnetic potential minimum as the minimum magnetic field strength within




Figure 3.2: ALPHA-2 Octupole magnetic field components, (a) radial, B_r , and (b) azimuthal, B_ϕ , in the xy -plane at $z = 0$. The solid black line shows the inner radius at the electrode wall, $r_{\text{trap}} \approx 2.2$ cm. We plot the semi-analytic solution for the Octupole field described in Section 3.4.2.

the trap. A magnetic minimum trap is effectively a magnetic bowl in which the modulus of the magnetic field strength increases in all directions away from the minimum magnetic field. A low-field-seeking anti-atom will be confined in the magnetic minimum trap if it has energy (sum of kinetic energy and magnetic potential energy) less than the depth of the magnetic bowl.

The other fraction of antihydrogen atoms are formed in *high-field-seeking* states which decrease in energy as magnetic field strength increases and are not trappable in our apparatus since they experience a force directing them towards regions of high magnetic field strength, which means they are directed towards the trap wall where they annihilate. As shown in Figure 3.1, the low-field-seeking (green) and high-field-seeking (blue) states are each split into two hyperfine states, due to the additional interaction of the antiproton magnetic moment with the magnetic flux density of the positron. Both low-field-seeking hyperfine states are trapped in the ALPHA magnetic minimum trap, creating a singly spin polarised (containing only positron spin state) antihydrogen population. As described in Section 1.3, resonant microwaves can be used to excite a transition from a trapped hyperfine state into an untrapped state, leaving a doubly spin polarised trapped antihydrogen population [58].

In ALPHA, the magnetic minimum trap is generated by the sum of the magnetic field strengths of the Octupole and Mirror Coils (see Figure 2.1 for diagram of trap magnets). In Figure 3.2 we show the radial, B_r , and azimuthal, B_ϕ , components of the Octupole magnetic field where (r, ϕ, z) are cylindrical coordinates. The quadrature sum of these field components gives rise to a magnetic field modulus $|\mathbf{B}_{\text{oct}}| \propto r^3$; providing a radial restoring force toward the trap center and hence radial confinement of antihydrogen. Two outer mirror coils (for example, Mirror A and Mirror E) give rise to an increasing magnetic field at the axial ends of the trap;




Figure 3.3: Typical ALPHA-2 magnetic confining potential for low-field-seeking ground state ($n = 1$) anti-atoms. Mirrors A and E are energised to 600 A, Mirrors B... D 0 A, Octupole energised to 900 A and a 1 T axial background field is present. The contour plot of $U(r, \phi = 0, z, t)$ is given as a function of z and r , where z is defined along the trap axis and r is the radius defined in cylindrical coordinates (Figure 2.1). Negative r , although unphysical, is shown to demonstrate that the potential is effectively a magnetic bowl. $\phi = 0$ is given for the simplicity of the plot, although the trap is azimuthally asymmetric at all $z \neq 0$ as shown in Figure 3.6.

providing axial confinement. The magnetic field strength of the Octupole and Mirror Coils is added to the 1 T axial External Solenoid field that is used for radial confinement of charged particles to generate the total magnetic field strength, $|\mathbf{B}(\mathbf{x})|$, of the magnetic minimum trap.

A typical trap configuration, often referred to as the *standard* trap, produces a magnetic potential distributed as shown in Figure 3.3. Outer Mirror Coils A and E are energised to ~ 600 A, generating a peak magnetic field strength of ~ 1 T, and are responsible for the two local maxima in potential at $z = \pm 137$ mm. The Octupole is energised to ~ 900 A, generating a peak magnetic field strength of ~ 0.85 T at the trap radius (~ 22 cm). The standard trap configuration generates a potential with trap depth ~ 0.5 K (with the limiting factor being the magnitude of the Octupole magnetic field at the trap wall), meaning the distribution of formed antihydrogen (which is much hotter than this at the temperature of the positron plasma, ~ 50 K) is truncated at a relatively low energy. As a result, around 0.1% of the total distribution of formed antihydrogen is can be confined in the magnetic minimum trap.

Given that the potential of a low-field-seeking ground state antihydrogen atom is given by Equation 3.2, the Hamiltonian of such an anti-atom confined within the magnetic trap is then,

$$H(\mathbf{x}, \mathbf{p}) = \frac{|\mathbf{p}|^2}{2M} + \mu_B |\mathbf{B}(\mathbf{x})|, \quad (3.3)$$

where \mathbf{p} is the cartesian momentum vector and M is the antihydrogen mass, which we assume to be equal to the hydrogen mass. We have neglected the gravitational potential ($\sim 10^{-4}$ K

over the diameter of the trap) as it is small compared to the magnetic potential ($\sim 10^{-1}$ K).

Antihydrogen confinement in a magnetic minimum trap relies on the assumption that the magnetic moment of a low-field-seeking anti-atom remains anti-aligned to the magnetic field. The magnetic moment is adiabatically locked anti-parallel to the magnetic field lines by the concept of *Rapid Adiabatic Passage* [86], which results from the fast precession of the magnetic moment about the magnetic field. A classical magnetic moment precesses around a magnetic field line at the Larmor frequency, Ω_L . The magnetic moment of an antihydrogen atom arises from the quantum mechanical orbit of a positron about an antiproton. Using the Hamiltonian of a stationary antihydrogen atom in a magnetic field, $B = B_0 \hat{z}$, I will briefly show that the expectation value of the magnetic moment of the antihydrogen atom is analogous to the classical precession of a magnetic moment in a magnetic field. The Hamiltonian of the stationary anti-atom is given by $H = \gamma \hat{L}_z B_0$, where γ is the gyromagnetic ratio and \hat{L}_z is the quantum mechanical angular momentum operator. The components (i, j, k) of the magnetic moment operator, $\hat{\mu}$, each obey,

$$\frac{d}{dt} \hat{\mu}_i = \frac{i\gamma}{\hbar} [H, \hat{L}_i], \quad (3.4)$$

where \hbar is the normalised Planck constant. Using the commutation relations for the angular momentum operators,

$$[\hat{L}_i, \hat{L}_j] = i\hbar \varepsilon_{ijk} \hat{L}_k, \quad (3.5)$$

where ε_{ijk} is the three-dimensional Levi-Civita symbol, and taking the expectation value,

$$\frac{d}{dt} \langle \hat{\mu} \rangle = \gamma \langle \hat{\mu} \rangle \times \mathbf{B}. \quad (3.6)$$

Therefore, the expectation value of the quantum mechanical magnetic moment is exactly analogous to the classical precession of a magnetic moment about the axis of the magnetic field [86], and hence we are free to use the analogy of the antihydrogen magnetic moment precessing classically about the magnetic field lines in the magnetic minimum trap.

In the frame of reference of the antihydrogen atom, a magnetic field that is static in the lab frame is time-varying due to the anti-atom motion. If this magnetic field varies slowly compared to Ω_L , then the magnetic moment is adiabatically locked to the magnetic field line. In the case where the magnetic field varies on the order of Ω_L , there is a finite probability of inducing a positron spin-flip into an untrappable state. These are known as *Majorana Spin-Flips* and occur in the ALPHA trap in atypical cases when the 1 T background field is lowered, giving rise to regions of zero magnetic field [87]. A potential motivation to reduce the background field is its effect on mixing between dimensions of antihydrogen kinetic energies. As will be discussed in Section 3.3, the axial background field adds in quadrature to the radial Octupole field which increases as r^3 . The combined magnetic potential increases like r^6 for small r , tending towards r^3 for larger r . In other words, the background field steepens




Figure 3.4: Diagram representing two extrema of antihydrogen motion in the trap, where the red dot indicates the instantaneous location of the anti-atom. We show the on-axis magnetic potential, $U(r = 0, \phi, z)$, as a function of axial position, z , for simplicity although the anti-atom moves in a 3D magnetic potential in reality. At *Position 1* the anti-atom has exchanged all of its magnetic potential energy for axial kinetic energy. At *Position 2* the anti-atom's total energy is entirely magnetic potential as it is at the turning point of its trajectory high on the Mirror Coil field.

the walls of the radial confining potential which hinders energy mixing into (and out of) the radial dimension. The motivation to gain control over mixing between axial and transverse energy components is discussed in Section 3.3. In adiabatic expansion Trials A, B and C the background field remains static at 1 T.

3.1.0.1 Radiative Cascade Cooling

Three-body-recombination leads to antihydrogen formation in high principal quantum number, n , states with a range of allowed orbital angular momentum quantum numbers, $0 \leq l \leq n - 1$. These antihydrogen atoms decay to the ground state relatively quickly (in a few milliseconds for anti-atoms initially formed in $n = 25$ to $n = 30$ states [75]) via photon emission. Since the magnitude of the antihydrogen magnetic moment decreases as the anti-atom decays to the ground state (expression below), the effective volume of the magnetic trap increases as a function of time, causing the anti-atoms to lose energy akin to adiabatic cooling. The initial distribution of n is not accurately known, but assuming the initial energy of the antihydrogen distribution is high compared to the trap depth (~ 0.5 K) and that after several decays the energy of the distribution will still be high compared to the trap depth, simulations of the cascade will give similar energy distributions when the anti-atoms reach the ground state for any $n \gtrsim 25$ [75]. We are required to include the radiative cascade in simulations of trapped antihydrogen to obtain an accurate energy distribution. In this section we will briefly describe radiative cascade cooling of antihydrogen, although a full description is given in [75].

Investigations described in [75] have focused on two cases. The first is decay via states with $l = |m| = n - 1$, where m is the azimuthal quantum number. The second case involves

decay via states with $m < l$. Here, we focus on decay via circular states as approximations described in [75] hold best for this state, Monte Carlo simulations described in [75] have shown both cases give similar results, and cascades from high n tend to populate the circular states [75].

The modified magnetic potential experienced by an antihydrogen atom in a low-field-seeking circular state is

$$U'(\mathbf{x}) = n\mu_B|\mathbf{B}(\mathbf{x})|. \quad (3.7)$$

Since the initial state is circular ($l = |m| = n - 1$), the photon spin sets $\Delta l = \pm 1$, $\Delta m = 0$ or ± 1 , $-l \leq m \leq l$ and $l < n$, the anti-atom must remain in a circular state by decreasing each of l , $|m|$ and n by one. Therefore, a single photon emission leads to a shift in magnetic potential of,

$$\Delta U'(\mathbf{x}) = -\mu_B|\mathbf{B}(\mathbf{x})|, \quad (3.8)$$

resulting from the anomalous Zeeman effect. Figure 3.4 depicts an antihydrogen atom at two extrema of its axial oscillatory motion in the magnetic trap. Equation 3.8 tells us that an antihydrogen atom undergoing a circular state decay at *Position 2* (high magnetic field) loses more energy than if it were at *Position 1* (low magnetic field). If this change in magnetic potential is slow compared to the antihydrogen speed, the energy loss will be adiabatic (the phase space volume of the trajectory will be conserved). On the contrary, if changes in magnetic potential are fast compared to the antihydrogen speed the energy loss will be diabatic (the phase space volume of the trajectory will not be conserved). The relative change in total magnetic potential due to an individual circular state decay is inversely proportional to n . For example, a circular decay from $n = 50$ to $n = 49$ results in a fractional change in magnetic potential of $1/50$, whereas a circular decay from $n = 2$ to $n = 1$ results in a corresponding fractional change of $1/2$. Therefore, high n circular decays are more likely to lead to adiabatic energy changes.

The circular state decay rate is given to very good approximation by [75]

$$\xi(n) = \frac{2}{3} \frac{1.61 \times 10^{10} \text{ Hz}}{(n - 1/2)^2 n^3}, \quad (3.9)$$

meaning the radiative cascade is simple to simulate using Monte Carlo methods. A random number, r , uniformly distributed between 0 and 1 is generated. If

$$r < \xi(n)dt \quad \{r \in \mathbb{R} \mid 0 \leq r \leq 1\}, \quad (3.10)$$

where dt is the simulation timestep, the simulated antihydrogen trajectory successfully decays to the lower circular state [75]. This method is based on an exponential decay approximation and therefore requires $\xi(n)dt \ll 1$ to be valid. Therefore, with the timestep used in the simulation (3.5×10^{-6} s), this approximation becomes inaccurate for $n \sim 13$. As a result, the n of a simulated anti-atom undergoes a single decay per timestep, whereas in reality an

anti-atom would undergo multiple decays within a 3.5×10^{-6} s window. However, this has no practical effect because the sudden cascade happens at a random location and therefore the effect of the bad approximation is that the sudden cascade occurs at a slightly different random location [88].

The decay rate (Equation 3.9) is lower for high n states. The average time taken to decay from $n = 25$ to $n = 20$ is ~ 3 ms, whereas the average time taken to decay from $n = 20$ to $n = 1$ is ~ 0.1 ms. Since a typical radial bounce time of a trapped anti-atom is ~ 1 ms, antihydrogen atoms tend to have sufficient time to move around in the magnetic trap before undergoing consecutive circular decays during decays between states with $n \gtrsim 20$. This effect contributes to the adiabatic nature of energy loss at high n and means that the overall energy loss at high n is small since photon loss happens at a range of trap locations (and hence magnetic field strengths, Equation 3.8). Specifically, $n > \sim 20$ is known as the *adiabatic regime*. At $n \lesssim 20$, the antihydrogen atom does not have sufficient time to move in the trap before cascading to the ground state. In this *sudden regime*, the increase in decay rate contributes to the diabatic nature of energy loss at low n .

Energy loss in the sudden regime depends on the particle location (which is effectively fixed during the sudden cascade). If the particle occupies Position 1 (low magnetic field strength), the energy loss will be small compared to the total energy of the particle; whereas at Position 2 (high magnetic field strength) the energy loss is will be a significant fraction of the total energy of the particle (if the particle was stationary at Position 2, the energy loss would be equal to the change in magnetic potential at Position 2, specified by Equation 3.8). Therefore, the cascade effectively selects antihydrogen atoms present at high field during the onset of the sudden regime. Due to the minimal energy changes at high n , simulation results do not change significantly when antihydrogen atoms are initialised in states with $n > 25$, compared to those initialised in $n = 25$. For this reason, to save computation time simulations of the radiative cascade are initialised with $n = 25$. The fraction of cooled anti-atoms (and therefore the trapping efficiency from TBR) depends on the spatial distribution of the magnetic field strength, but the radiative cascade can have a significant effect [75].

3.2 Principles of Adiabatic Expansion

In this section I will determine changes to a particle's Hamiltonian given that the confining magnetic field is non-constant in time, by following the derivation given in [89]. The derivation will assume changes to the magnetic field can change the length of the antihydrogen trajectory in an arbitrary direction, q ($q = x, y$ or z), and that the particle oscillates with momentum, p , in the q direction (for example, an antihydrogen atom bouncing between the axial confining magnetic field of the Mirror Coils quasi-periodically). The derivation will assume the magnetic

field varies slowly in time and specifically that,

$$\tau \frac{dU}{dt} \ll U, \quad (3.11)$$

where τ is the period of the anti-atom motion in the q direction and U is the magnetic potential as defined by Equation 3.2. If the magnetic field is *stationary*, the antihydrogen atom constantly exchanges kinetic and magnetic potential energy as it oscillates in the magnetic trap, but its total energy, E (sum of kinetic and potential), is constant over time. If $\frac{dU}{dt} \neq 0$, the total energy of the particle is no longer a conserved quantity. Since the magnetic field varies slowly (Equation 3.11), the rate of change in energy will also be slow. Here, I will derive a combination of time-varying quantities that is constant in time to very good approximation, known as an *adiabatic invariant*.

The rate of change of energy in the system is given by the time derivative of the Hamiltonian,

$$\frac{dH}{dt} = \frac{\partial H}{\partial q} \dot{q} + \frac{\partial H}{\partial p} \dot{p} + \frac{\partial H}{\partial t}. \quad (3.12)$$

Using Hamilton's equations,

$$\frac{\partial H}{\partial p} = \dot{q}, \quad (3.13)$$

$$\frac{\partial H}{\partial q} = -\dot{p}, \quad (3.14)$$

Equation 3.12 simplifies to,

$$\frac{dH}{dt} = \frac{\partial H}{\partial t} = \frac{\partial H}{\partial U} \frac{dU}{dt}. \quad (3.15)$$

Since this depends on the rapidly varying q and p as well as the slow varying U , the average the energy change must be averaged over a period of oscillation,

$$\frac{\overline{dH}}{dt} = \frac{\overline{\partial H}}{\partial U} \frac{dU}{dt}, \quad (3.16)$$

where dU/dt is excluded from the average since the variation in U and therefore in dU/dt is small. In addition, the averaged quantity can be considered not to depend on U . Specifically, the average is defined as,

$$\frac{\overline{\partial H}}{\partial U} = \frac{1}{\tau} \int_0^\tau \frac{\partial H}{\partial U} dt = \frac{1}{\tau} \oint \frac{\partial H}{\partial U} \frac{dq}{\partial H / \partial p}, \quad (3.17)$$

where the period, τ , is given by,

$$\tau = \int_0^\tau dt = \oint \frac{dq}{\partial H / \partial p}, \quad (3.18)$$

and in both cases dt has been replaced by that given by Equation 3.13. It is important to note that the closed loop integral is performed over a closed path in q assuming a constant U . Using Equations 3.17 and 3.18, the expression for the averaged energy change becomes,

$$\frac{\overline{dH}}{dt} = \frac{dU}{dt} \frac{\oint \frac{\partial H}{\partial U} \frac{dq}{\partial H / \partial p}}{\oint \frac{dq}{\partial H / \partial p}}. \quad (3.19)$$




Figure 3.5: A particle oscillating axially between *Wall 1* and *Wall 2*. Wall 2 moves away from Wall 1 at constant speed u , causing a slow increase in the trap length, L , as a function of time. The particle has speed v_z in the axial direction where $v_z \gg u$.

The total derivative of the Hamiltonian with respect to U is given by,

$$\frac{dH}{dU} = \frac{\partial H}{\partial U} + \frac{\partial H}{\partial p} \frac{\partial p}{\partial U}, \quad (3.20)$$

and since the integrals in Equation 3.19 are constant in U because changes in U are small during one period,

$$\frac{dH}{dU} = \frac{\partial H}{\partial U} + \frac{\partial H}{\partial p} \frac{\partial p}{\partial U} = 0 \implies \frac{\partial H / \partial U}{\partial H / \partial p} = -\frac{\partial p}{\partial U}. \quad (3.21)$$

Substituting this into Equation 3.19 leads to,

$$\frac{d\overline{H}}{dt} = -\frac{dU}{dt} \frac{\oint (\partial p / \partial U) dq}{\oint (\partial p / \partial H) dq}. \quad (3.22)$$

which can alternatively be expressed as

$$\oint \left(\frac{\partial p}{\partial H} \frac{d\overline{H}}{dt} + \frac{\partial p}{\partial U} \frac{dU}{dt} \right) dq = 0. \quad (3.23)$$

This expression allows definition of the adiabatic invariant, I , where

$$\overline{dI/dt} = 0, \quad (3.24)$$

and

$$I \equiv \oint p dq, \quad (3.25)$$

where we have used the expression for the total derivative of p with respect to t ,

$$\frac{dp}{dt} = \frac{\partial p}{\partial H} \frac{dH}{dt} + \frac{\partial p}{\partial U} \frac{dU}{dt}, \quad (3.26)$$

and we have excluded terms from the average that do not vary significantly during a period.

As explained in Chapter 2, the adiabatic expansion experiment in the ALPHA-2 apparatus involved expanding the trap mainly axially (although there is also a small amount of radial compression). Relating to the axial oscillatory motion of such an anti-atom, I is expressed as

$$I = \oint p_z dz, \quad (3.27)$$

where p_z is the axial momentum component. The closed loop integral is simple to solve since, as mentioned earlier, the integral is performed over one period assuming a constant magnetic field. If the particle were bouncing in a hard-walled box, expanding slowly compared to the particle motion as shown in Figure 3.5, the adiabatic invariant would be

$$v_z L = \text{constant}, \quad (3.28)$$

where L is the axial length of the box. Finding an expression for the adiabatic invariant is useful as it makes it possible to estimate changes in particle speed (and therefore kinetic energy) corresponding to changes in trap length. Since the adiabatic invariant is a constant throughout an adiabatic expansion, its value before and after the expansion can be equated, leading to an expression for the expected kinetic energy (after expansion) of the anti-atom in terms of the kinetic energy (before expansion) and the change in trap length. In Chapter 5 I will use this method to model the expected change in energy during adiabatic expansion.

However, the reality of antihydrogen trajectories is not as simple as the particle in a box model and hence the adiabatic invariant is not as simple as Equation 3.28. Firstly, trajectories oscillate in the transverse plane as well as axially. This, alone, would not complicate the adiabatic invariant, but azimuthal asymmetries in $|\mathbf{B}(\mathbf{x})|$ are known to mix axial and transverse energies as will be discussed in Section 3.3. Energy mixing can alter v_z during a period of oscillation, which changes the power law of the adiabatic invariant, as we will show in Chapter 5. Secondly, a single trap length L cannot be defined for a particle oscillating in a magnetic minimum trap given that it also oscillates in the transverse plane and L is a function of radius in the ALPHA-2 trap. In Chapter 5 we will demonstrate a technique to approximate the axial length of the trap by calculating the volume accessible by a trapped antihydrogen atom, and averaging the length of this volume over radius. Further, this volume is a function of the particle energy (because higher energy anti-atoms can climb higher up the walls of the magnetic potential), which means the volume changes as the particle energy is adiabatically reduced with all other things being equal.

It is noteworthy that we have used only arguments of classical mechanics to derive Equation 3.28. Often for this problem, thermodynamics are used to derive the adiabatic formula for a large number of non-interacting particles in a box with temperature $T \sim \langle v^2 \rangle$, where the angled brackets denote the ensemble average. In the thermodynamics context, an adiabatic process is defined as one which is thermally isolated from its surroundings and is subject to external conditions which vary sufficiently slowly [89]. Generalising Equation 3.28 to the other two spatial dimensions, assuming the box expands self-similarly (same length in all three dimensions), and squaring to bring the adiabatic invariant in terms of energy,

$$v^2 L^2 = (v_x^2 + v_y^2 + v_z^2) L^2 = \text{constant}, \quad (3.29)$$

where v is the total particle speed and we have assumed the box is a cube with sides of length

L expanding self-similarly. Replacing the expression for the temperature of the gas and using $V = L^3$ where V is the volume of the box,

$$TV^{2/3} = \text{constant}. \quad (3.30)$$

In combination with the ideal gas law, $PV = Nk_B T$, where P is the gas pressure, N is the number of particles and k_B is Boltzmann's constant, we arrive at the familiar adiabatic invariant,

$$PV^{5/3} = \text{constant}. \quad (3.31)$$

Hence, the classical mechanics derivation of adiabatic expansion of a particle in a box that expands self-similarly is consistent with the thermodynamics derivation of the adiabatic expansion of an ideal gas, which is typically derived from the first law of thermodynamics with infinitesimal changes in heat, $\delta Q = 0$ [90].

3.3 Fundamentals of Energy Mixing in Magnetic Traps

Energy mixing between degrees of freedom in a magnetic antihydrogen trap can be beneficial in some cases and undesired in others. For example, for antihydrogen cooling techniques such as laser cooling and adiabatic expansion cooling, energy mixing between axial and transverse degrees of freedom during cooling increases the total energy loss when the cooling is primarily occurring in a single dimension. However, total energy loss is not always the priority. In some applications, it is more important for the energy in a given dimension to be minimised. For example, for measuring antihydrogen gravity in ALPHA-g, the priority is cooling in the direction parallel to the force of gravity. Similarly, for reducing the effect of Doppler-broadening on the lineshape of an atomic transition, the goal is to minimise the energy component parallel to the spectroscopy laser beam. On the other hand, if the dominant error is transit-time broadening, such as in the case of the 1S-2S lineshape [56], it is reduction of the energy in the direction transverse to the spectroscopy laser that matters. Knowledge of the mixing timescale informs the experimental protocol for adiabatic expansion because it allows us to estimate the rate at which we should expand the trap volume to achieve a particular goal.

The following section follows studies presented in [9], in which ensembles of antihydrogen atoms are simulated in the ALPHA-2 magnetic minimum trap using a similar simulation method presented in this thesis (Section 3.4). I have reproduced many of the simulated results using the simulation presented in this thesis. All presented figures are from [9], except Figures 3.6 and 3.9 which were produced by the author using the simulation presented in this thesis. In this section I will describe the method used to determine the axial and transverse energy components of a simulated antihydrogen atom as a function of time, which I use to determine the degree of axial and transverse cooling achieved during adiabatic expansion Trial




Figure 3.6: (a) Diagram of magnetic trap: Octupole (red), Mirror Coils A-E (blue), Solenoids A and B (green) and electrode stack (yellow). The electrode in which antihydrogen formation takes place for adiabatic expansion Trials A and C is shown in pink, whereas Trial B is shown in orange. An xy -plane in close proximity to Mirror E is grey. (b) $|\mathbf{B}(\mathbf{x})|$ in the xy -plane close to Mirror E with Mirror E not energised (c) with Mirror E energised to 600A. All magnetic fields shown are modelled using the semi-analytic models presented in Section 3.4.2.

A (described in Chapter 2) as will be presented in Chapter 4. I will use results presented in [9] to describe how individual antihydrogen atoms can exchange axial and transverse energy as a function of time whilst they are trapped in the static standard trap. These results enable two categories of trapped antihydrogen atoms that exhibit significantly different energy mixing dynamics to be identified. In Section 4.8, I will use simulation techniques presented in this section to explain the simulated distributions of annihilation time and axial annihilation location of antihydrogen atoms during diagnostic ramp down. I will present the method used to determine the time taken for an individual trajectory to mix a significant proportion of axial and transverse energy. This diagnostic of the mixing timescale allows the fraction of an ensemble of antihydrogen trajectories that have mixed energy to be determined as a function of time. I present results from [9] which determine this mixing timescale in trap configurations comparable to those at the start and end of adiabatic expansion Trial A, allowing me to estimate the extent to which antihydrogen atoms are expected to mix energy components during adiabatic expansion of the trap volume.

For the following section, we will work in cylindrical coordinates, (r, ϕ, z) , where the z -axis is parallel to the trap axis as shown in Figure 3.6. Energy exchange between trap degrees of freedom requires a magnetic field gradient in one dimension that changes as a function of another. In other words,

$$\frac{\partial^2 |\mathbf{B}(r, \phi, z)|}{\partial a_i \partial b_j} \neq 0, \quad (3.32)$$

where $a_i = r, \phi, z$ for $i = 1, 2, 3$, $b_j = r, \phi, z$ for $j = 1, 2, 3$, and $a_i \neq b_j$. A ϕ -asymmetry arises in $|\mathbf{B}(r, \phi, z)|$ as a result of the vector sum of the radial Mirror Coil field and the radial components of the Octupole field entering $|\mathbf{B}(r, \phi, z)|$. Figure 3.6(a) defines an xy -plane in close proximity to Mirror E. Figure 3.6(b) shows $|\mathbf{B}(r, \phi, z)|$ in this plane given that Mirror E is not energised. By contrast, Figure 3.6(c) also shows $|\mathbf{B}(r, \phi, z)|$ in this plane, but this time Mirror E is energised to 600A. The resulting field is azimuthally asymmetric. For a standard confinement trap, both Mirrors A and E are energised. In order for both of their axial fields to add to the axial 1 T background field, their inward facing radial fields point in opposite directions. As a result, their radial fields exactly cancel at $z = 0$ and the azimuthal asymmetry induced by the radial field of Mirror A is rotated by 45 degrees relative to that of Mirror E. The result is that the trap ‘twists’ as a function of z ; resulting in, for example, $\frac{\partial}{\partial z} \frac{\partial |\mathbf{B}(\mathbf{r}, \theta, \mathbf{z})|}{\partial \theta} \neq 0$ and hence it is possible to mix energy components. Other field characteristics also mix energy components, for example, the azimuthally asymmetric axial field produced by the Octupole wire end turns (connections between the eight poles of the Octupole).

We define the axial energy, E_{\parallel} , as the energy parallel to the z -axis and the transverse energy, E_{\perp} , as the energy perpendicular to this i.e. the sum of radial and azimuthal energy components. For the following discussion we will use Cartesian coordinates, where a particle has position (x, y, z) and velocity (v_x, v_y, v_z) . The fact that both the Mirror Coil and Octupole

fields have radial and axial components, means the Hamiltonian (Equation 3.3) does not split into independent axial and transverse components. As a consequence, if we determine the axial kinetic energy ($\frac{Mv_z^2}{2}$) and transverse kinetic energy ($\frac{Mv_x^2}{2} + \frac{Mv_y^2}{2}$) of a simulated antihydrogen trajectory oscillating in the magnetic trap at a random sample time, the particle is likely to have climbed the confining magnetic walls to some extent. Since the magnetic field modulus, $|\mathbf{B}(\mathbf{x})|$, does not split into independent axial and transverse components, it is not possible to know to what extent this magnetic potential can be exchanged for axial or transverse kinetic energy as the particle continues its oscillation.

To evade this issue, we use the method described in [9] which involves sampling E_{\parallel} and E_{\perp} of a simulated anti-atom at every trajectory crossing of the axial trap minimum, z_{\min} , at which time the anti-atom's axial energy is entirely kinetic and hence,

$$E_{\parallel} = \frac{Mv_z^2}{2}. \quad (3.33)$$

An antihydrogen atom crosses z_{\min} at a position (x_0, y_0, z_{\min}) , where x_0 and y_0 are not necessarily 0. Therefore, the transverse energy has a potential component, U_{\perp} ,

$$E_{\perp} = \frac{Mv_x^2}{2} + \frac{Mv_y^2}{2} + U_{\perp}(x_0, y_0, z = z_{\min}), \quad (3.34)$$

Thus, the method allows us to isolate the components of total (kinetic plus magnetic potential energy) in the axial and transverse degrees of freedom (E_{\parallel} and E_{\perp}).

To quantify the relative components of axial and transverse energy for a given simulated anti-atom, we define the normalised axial energy,

$$\varepsilon_{\parallel} = \frac{E_{\parallel}}{E_{\parallel} + E_{\perp}}, \quad (3.35)$$

which is a measure of the fraction of the total anti-atom energy that is in the axial degree of freedom. For example, an anti-atom which bounces back and forth between the confining Mirror Coils and does not move in the transverse plane has $\varepsilon_{\parallel} = 1$, whereas an anti-atom which moves only in the transverse plane has $\varepsilon_{\parallel} = 0$.

Figure 3.7 shows the result of a simulation of an ensemble of 3757 antihydrogen trajectories held for 1000 s in the standard trap configuration (as defined in Figure 2.3(a)) as presented in [9]. At different sample times in this simulation, $\bar{t} = 1\text{s}, 10\text{s}, 100\text{s}$ and 1000s , the cumulative distribution function (CDF) of ε_{\parallel} is determined. The CDF of ε_{\parallel} can be interpreted as the fraction of simulated antihydrogen atoms with energy at or below a given value of ε_{\parallel} . For example, a sharp increase in the CDF of ε_{\parallel} over a narrow range of ε_{\parallel} would tell us that a large fraction of trajectories occupy that narrow range of ε_{\parallel} . On the other hand, if the CDF of ε_{\parallel} were flat over a range of ε_{\parallel} , we could conclude that no antihydrogen trajectories occupy that band of ε_{\parallel} . Comparing the CDF of ε_{\parallel} at the presented sample times allows us to conclude that the overall distribution of ε_{\parallel} does not change significantly between 1 s and 1000 s.




Figure 3.7: Cumulative distribution function (CDF) of ε_{\parallel} for a simulated distribution of 3757 trajectories as a function of energy at different sample times, $\bar{t} = 1s, 10s, 100s, 1000s$ [9].

An individual anti-atom that exchanges axial and transverse energy will have an ε_{\parallel} and E_{\parallel} that fluctuates as a function of time. Taking the time average of these quantities provides us with information about the extent to which axial and transverse energies are exchanged. For example, if energy is shared on-average equally between axial and transverse degrees of freedom, the time averaged normalised axial energy, which we will denote $\langle\langle\varepsilon_{\parallel}\rangle\rangle$, would tend to 0.5 over long time averages. Alternatively, if a particle does not exchange axial and transverse energy on long timescales, the time averages would be close to the initial value. Since simulations are required to determine E_{\parallel} and E_{\perp} at discrete time intervals (when the anti-atom crosses the magnetic minimum), the time average is given by the average of the values at consecutive minimum crossings. Specifically, for an individual anti-atom denoted j ,

$$\langle\langle E_{\parallel j} \rangle\rangle = [1/(t_{N_j,j})] \sum_{i=1}^{N_j} E_{\parallel ij} \Delta t_{ij}, \quad (3.36)$$

where i indexes each crossing of z_{\min} made by the particle and the particle makes a total of N_j crossings, $\Delta t_{ij} = t_{ij} - t_{(i-1)j}$ is the time between consecutive crossings and $t_{N_j,j}$ is the total time taken for the particle to undergo N_j crossings [9]. Note that the time-averaged normalised axial energy, $\langle\langle\varepsilon_{\parallel}\rangle\rangle$, has an equivalent definition.

Figure 3.8 shows the result presented in [9] of $\langle\langle E_{\parallel}\rangle\rangle$ and $\langle\langle\varepsilon_{\parallel}\rangle\rangle$ for an ensemble of antihydrogen trajectories held in the static standard trap, where the time-average is taken over 1000s. It is apparent that the simulated antihydrogen trajectories form two bands in $\langle\langle E_{\parallel}\rangle\rangle$ and $\langle\langle\varepsilon_{\parallel}\rangle\rangle$. Those in the band with higher $\langle\langle E_{\parallel}\rangle\rangle$ have $\langle\langle\varepsilon_{\parallel}\rangle\rangle \sim 0.5$, whereas those in the band with lower $\langle\langle E_{\parallel}\rangle\rangle$ have $\langle\langle\varepsilon_{\parallel}\rangle\rangle < \sim 0.1$. The red lines in Figure 3.8 are drawn to separate the two populations.

The information in Figure 3.8 alone is not sufficient to deduce whether the antihydrogen atoms mix axial and transverse energy components. Take the category with $\langle\langle\varepsilon_{\parallel}\rangle\rangle \sim 0.5$ for example: either this population is exchanging axial and transverse energies on-average equally




Figure 3.8: Time averaged parallel energy, $\langle\langle E_{\parallel}\rangle\rangle$ (a) and time averaged normalised axial energy, $\langle\langle \varepsilon_{\parallel}\rangle\rangle$ (b) as a function of total anti-atom energy, E . Each trajectory was simulated for 1000s and all energies are in Kelvin. The red line (estimate of the location of minimum density between the two peaks in density) is drawn to separate two categories of trajectory which exhibit different energy mixing dynamics: those above the red line exchange axial and transverse energy quickly, whereas those below the red line do not. Figure from [9].

over time, or this population was initialised with $\varepsilon_{\parallel} \sim 0.5$ and did not exchange axial and transverse energies. However, if the latter case were true the CDF of ε_{\parallel} (Figure 3.7) would indicate a preferential sampling of ε_{\parallel} around 0.5 i.e. the CDF would increase sharply around $\varepsilon_{\parallel} = 0.5$. Since this is not observed in Figure 3.7, the category with $\langle\langle \varepsilon_{\parallel}\rangle\rangle \sim 0.5$ must be exchanging axial and transverse energies equally over time.

By looking at $\langle\langle \varepsilon_{\parallel}\rangle\rangle$ as a function of time for a set of individual simulated trajectories, we can observe whether or not these categories of trajectory exchange axial and transverse energy components. In Figure 3.9 we show $\langle\langle \varepsilon_{\parallel}\rangle\rangle$ as a function of time for 100 trajectories, simulated using the method described in Section 3.4. If an anti-atom has a non-constant $\langle\langle \varepsilon_{\parallel}\rangle\rangle$ over time, it exchanges axial and transverse energy components. Firstly, note that at 1000s we see the equivalent behaviour shown in Figure 3.8(b). Secondly, the category with $\langle\langle \varepsilon_{\parallel}\rangle\rangle < \sim 0.1$ at 1000s, have an $\langle\langle \varepsilon_{\parallel}\rangle\rangle$ that does not change significantly during the period from 0s to 1000s. Thirdly, the category with $\langle\langle \varepsilon_{\parallel}\rangle\rangle \sim 0.5$ at 1000s tend to exchange significant fractions of axial and transverse energy. Therefore, we will refer to the $\sim 1/3$ of trajectories with $\langle\langle \varepsilon_{\parallel}\rangle\rangle < \sim 0.1$ as the *no-mix* category and the $\sim 2/3$ of trajectories with $\langle\langle \varepsilon_{\parallel}\rangle\rangle \sim 0.5$ as the *mix* category.

The two categories are thought to arise because only the *mix* category have sufficient axial energy to sample regions of the trap where they are able to mix a significant proportion of their axial and transverse energies. As explained earlier, the azimuthal field asymmetries become more pronounced closer to the outer Mirror Coils, and can only be sampled by antihydrogen atoms with sufficiently high E_{\parallel} . However, we should note that qualifying for the *mix* category requires a high ε_{\parallel} , rather than E_{\parallel} . This can be seen in Figure 3.8(a), since those above the red line can have a very small $\langle\langle E_{\parallel}\rangle\rangle$ at 1000s, but have a relatively large $\langle\langle \varepsilon_{\parallel}\rangle\rangle$. The reason for this is not known.




Figure 3.9: $\langle\langle \varepsilon_{\parallel} \rangle\rangle$ as a function of time, t , for 100 trajectories. Each coloured line is an individual simulated trajectory.

One might wonder why, for the mix category, $\langle\langle \varepsilon_{\parallel} \rangle\rangle \rightarrow \frac{1}{2}$, rather than $\frac{1}{3}$, as is expected from three-dimensional energy mixing. Since the radial field of the Octupole adds in quadrature to the 1 T axial background field, $U(r)$ is flat bottomed with steep confining walls. Trajectories that approach these steep walls tend to bounce back without mixing z and r energy components, which means energy does not share equally between all three dimensions.

The energy mixing time is quantified by taking $\varepsilon_{\parallel kj}$ where k indexes a random z_{\min} crossing between 1 s and 2 s and j indicates an individual trajectory. This provides a starting energy which is free from initial condition effects at the very beginning of the simulation. Then, the time taken for this ratio to change by 0.2 is computed. In other words, when the threshold $\varepsilon_{\parallel kj} \pm 0.2$ is crossed, the particle is considered to have mixed [9]. This threshold is somewhat arbitrary but is chosen such that it ensures a significant proportion of energy has mixed before the threshold is crossed but also that it is not so high that it fails to count particles that significantly exchange their energies. The time at which the threshold is first crossed is defined as $t_{\uparrow j}$. The CDF of t_{\uparrow} then provides an indication of the percentage of trajectories have significantly mixed energies at a given time. Figure 3.10 shows the CDF of t_{\uparrow} as a function of time for three different trap configurations: standard, tight and flat (as defined by Figure 2.3). In addition, labelled as *quadrupole* is the standard trap configuration but with the Octupole replaced by a quadrupole magnet. The different trap configurations lead to significantly different extents of axial to transverse energy mixing. The mixing timescale is shortest in the tight trap, where the radial Mirror Coil field is strong closer to the trap center. In the flat trap, where a slight negative current on the internal Mirror Coils (B... D) is used to flatten the magnetic field as a function of z , the mixing timescale is also faster than in the

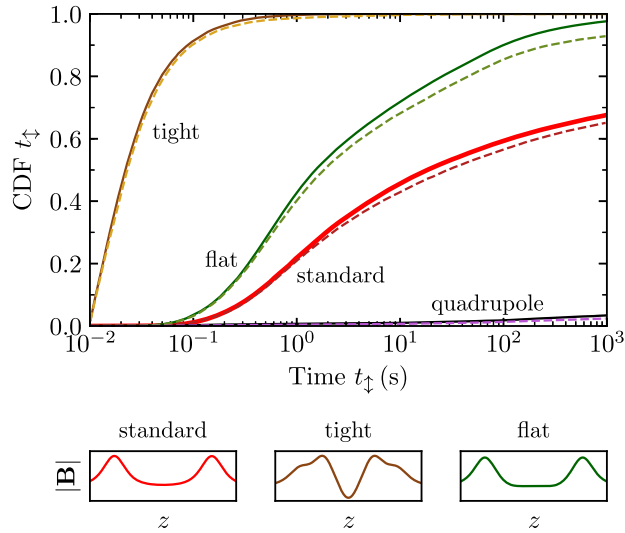


Figure 3.10: CDF of the time, t_{\downarrow} , taken for antihydrogen trajectories to change their initial normalised axial energy, $\varepsilon_{\parallel kj}$ by 0.2. The CDF is shown for different trap configurations, for which $|\mathbf{B}|$ as a function of z ($r = 0$) is shown below the figure. The ‘quadrupole’ configuration represents the standard trap but with the Octupole magnet replaced by a quadrupole. Solid and dotted lines represent simulated trajectories without and with cascade initialisation respectively. Figure from [9].

standard trap (due to the radial field of Mirror Coils B... D).

Based on the mixing timescale in the tight trap, we estimate that mixing is initially fast during Adiabatic Expansion Trial A (start of Trial A is comparable to the tight well), with approximately 90% of trajectories mixing in 0.1s and almost all trajectories mixing in 1s. As the trap expands, the magnetic configuration in Adiabatic Expansion Trial A gradually morphs into a trap comparable to the standard trap. When this happens, the rate of energy mixing decreases, and in the final trap configuration, approximately 50% of trajectories mix in 10s and approximately 33% of trajectories do not mix in 1000s.

However, the energy mixing analysis described in this section has studied simulated antihydrogen ensembles that were not adiabatically cooled. Since the antihydrogen atoms are expected to lose energy during the adiabatic expansion of the magnetic trap, the energy mixing dynamics may differ from that of the non-adiabatically-cooled population.

3.4 Simulating Antihydrogen Trajectories

In this section we will describe the Monte Carlo simulation used to produce the results presented in Chapters 4, 5 and 6. Simulated anti-atoms are initialised with a pseudo-random initial position and velocity, sampled from distributions that represent our best knowledge of the spatial and energy distribution of antihydrogen formed through three-body-recombination

(see Section 1.2.3). The force on the anti-atom is proportional to the gradient of the magnetic field at the particle location, which leads to a pair of coupled differential equations that specify the equations of motion of the simulated anti-atom. We will describe the symplectic Leapfrog algorithm used to solve these equations of motion for the coordinates of the particle in six dimensional phase space, (x, y, z, v_x, v_y, v_z) , as a function of time.

Since the magnetic potential experienced by a trapped ground-state antihydrogen atom in the magnetic trap is given by Equation 3.2, the force on the anti-atom leads to a pair of coupled differential equations,

$$\frac{d\mathbf{x}}{dt} = \mathbf{v}, \quad (3.37)$$

and

$$M \frac{d\mathbf{v}}{dt} = -\mu_B \nabla |\mathbf{B}(\mathbf{x})|, \quad (3.38)$$

where M is the antihydrogen mass and $\mathbf{v} = (v_x, v_y, v_z)$ is the Cartesian velocity vector. The Leapfrog algorithm provides a second order symplectic solution to the trajectory at fixed time intervals, dt , by solving for the position at the midpoint of the time interval,

$$\begin{aligned} x_{n+1/2}^i &= x_{n-1/2}^i + v_n^i dt & \{i \in x, y, z\}, \\ v_{n+1}^i &= v_n^i - \frac{\mu_B}{M} \nabla |B^i(\mathbf{x}_{n+1/2})| dt & \{i \in x, y, z\}, \end{aligned} \quad (3.39)$$

where a superscript i denotes the spatial dimension, a subscript n denotes iteration order, and a half interval n denotes the solution at half time intervals e.g. $x_{n+1/2}$ is the position at $t + dt/2$ [10]. A more detailed description of this symplectic Leapfrog integration method is provided later in this section. We are required to know the vector sum of the magnetic fields from the Octupole magnet, Mirror Coils, Solenoids A and B and the background 1 T field at every timestep at the position of the particle, and to calculate its gradient. At a particular position, $\mathbf{x} = (x, y, z)$, the \hat{x} component of the magnetic field gradient is approximated as,

$$(\nabla |\mathbf{B}(x, y, z)|)_{\hat{x}} = \frac{|\mathbf{B}(x + \frac{\delta x}{2}, y, z)| - |\mathbf{B}(x - \frac{\delta x}{2}, y, z)|}{\delta x}, \quad (3.40)$$

where $\delta x = 1.0 \times 10^{-5} r_{\text{trap}}$, where $r_{\text{trap}} \approx 2.2$ cm is the inner radius of the electrode stack (on which antihydrogen atoms escaping the magnetic confining potential annihilate). Similar expressions are used for the \hat{y} and \hat{z} components of $\nabla |\mathbf{B}(x, y, z)|$, where the same displacement is made in the y or z direction respectively. $\mathbf{B}(x, y, z)$ of each magnet is calculated using the semi-analytic solutions described later in this section.

To initialise the symplectic Leapfrog algorithm, we require the anti-atom initial conditions: v_0 and $x_{1/2}$. The velocity in each dimension is sampled from an independent normal distribution,

$$f(v_i) = \sqrt{\frac{M}{2\pi k_B T}} \exp\left(-\frac{M(v_i)^2}{2k_B T}\right) \quad \{i \in x, y, z\}, \quad (3.41)$$

where M is the antihydrogen mass, k_B is Boltzmann's constant and T is the temperature of the distribution of formed antihydrogen. To sample this distribution, for each dimension we

generate a random angle [91],

$$\theta_i = 2\pi\alpha_i \quad \{\alpha_i \in \mathbb{R} \mid 0 \leq \alpha_i \leq 1\}, \quad (3.42)$$

and a random amplitude, R_i ,

$$R_i = \sqrt{\frac{2k_B T \ln \beta_i}{M}} \quad \{\beta_i \in \mathbb{R} \mid 0 \leq \beta_i \leq 1\}. \quad (3.43)$$

The speed is then given by,

$$v_i = R_i \cos \theta_i \quad \{i \in x, y, z\}. \quad (3.44)$$

Since the particle speed is given by the quadrature sum of each component,

$$v = \sqrt{v_x^2 + v_y^2 + v_z^2}, \quad (3.45)$$

the speed is distributed according to

$$f(v) = \left(\frac{M}{2\pi k_B T}\right)^{3/2} 4\pi v^2 \exp\left(-\frac{Mv^2}{2k_B T}\right). \quad (3.46)$$

Rearranging in terms of the kinetic energy, $E_K = \frac{1}{2}Mv^2$, we arrive at the Maxwell-Boltzmann distribution,

$$f_E(E_K) = 2\sqrt{\frac{E_K}{\pi}} \left(\frac{1}{k_B T}\right)^{3/2} \exp\left(\frac{-E_K}{k_B T}\right). \quad (3.47)$$

We assume a formed antihydrogen temperature, $T = 50$ K, which is approximately the temperature of the positron plasma. This is based on the assumption that antiprotons thermalise at the temperature of the positron plasma prior to antihydrogen formation via three-body recombination (see Section 1.2.3). This choice has been motivated by comparing annihilation times and axial annihilation locations between experiment and simulation [51].

Figures 3.11(a) and 3.11(b) compare the distribution functions given by Equations 3.46 and 3.47 to the initial speed and kinetic energy of 10000 simulated anti-atoms respectively. We conclude that the method samples the energy distribution effectively. Antihydrogen atoms with energy much greater than the depth of the magnetic trap will quickly escape the trap and annihilate. Therefore, to save simulation time we propagate only simulated anti-atoms with energy greater than 0.75 K, which is several hundreds of millikelvin above a typical trap depth (~ 0.5 K). We observe that anti-atoms with energy greater than the trap depth, known as *quasi-trapped* particles, can remain confined on long timescales because they do not sample the entire trap volume. The majority of quasi-trapped particles are quickly lost from the trap which effectively truncates the Maxwell Boltzmann distribution at the trap depth and thus $f_E(E_K) \propto \sqrt{E_K}$. As a result of this truncation, the simulated distribution is fairly insensitive to T .

The initial position of the trajectory is sampled from a uniform distribution throughout the volume of the positron plasma since, as mentioned in Section 1.2.3, antihydrogen is thought

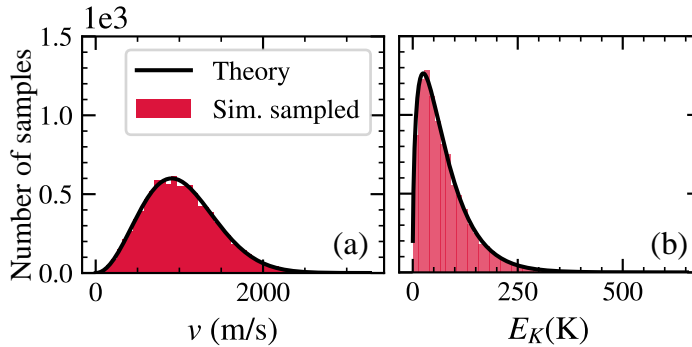


Figure 3.11: Initialised antihydrogen (a) speed, $v(\text{m/s})$, (b) kinetic energy, $E_K(\text{K})$, distributions. In both subfigures the black curves represent the distribution given by the theory and the red histograms are the sampled distributions (10000 total samples) from the simulation.

to be formed through three-body-recombination when antiprotons transit this volume [9]. We assume the positrons occupy an ellipsoid of length 10 mm and radius 0.8 mm (based on knowledge of the positron plasma prior to antihydrogen formation via simulation and experimental Micro-channel Plate (MCP) measurements). We must note that, before arriving at our set of initial conditions, we need to step our position by half a timestep to convert x_0 into $x_{1/2}$ to begin the leapfrog algorithm. The Taylor expansion,

$$x\left(t + \frac{dt}{2}\right) \approx x(t) + x'(t)\frac{dt}{2} + \frac{1}{2}x''(t)\left(\frac{dt}{2}\right)^2 + \dots, \quad (3.48)$$

allows us to step to second order, given that the first time derivative of the position, $x'(t) = v_0$ and the second, $x''(t)$, is the acceleration which we can deduce from the force, $F = -\mu_B \nabla |\mathbf{B}(\mathbf{x}(t))|$.

The set of pseudo-random numbers used to generate v_0 and x_0 for a given particle is specified by a random ‘seed’ which allows us to reproduce an exact initial position and velocity, whilst retaining the random nature of initialisation of each particle. This allows us to, for example, run a particle with identical initial conditions through simulations of different adiabatic expansion trials.

The total energy of an anti-atom in the simulation is the sum of the magnetic potential and kinetic energy components. The magnetic potential is defined relative to the minimum magnetic field, $\mathbf{B}(\mathbf{x}_{\min})$ (Equation 3.2). Therefore, we find the location of the magnetic minimum, \mathbf{x}_{\min} , at each timestep by searching the magnetic potential using ROOT’s *Minuit2 Minimization Library* [92]. The minimisation is too slow to compute at runtime and we therefore calculate the magnetic minimum at discrete time intervals in advance of running the simulation. The magnetic minimum is then found at an arbitrary time using linear interpolation.

We initialise the particles in high principal quantum number, n , states and simulate the

radiative cascade as described in Section 3.1.0.1, along with the particle trajectory. Once a particle is initialised, the radiative cascade forms part of what is known as the ‘pre-simulation’; the purpose of which is to propagate a v_0 and $x_{1/2}$ through the Leapfrog algorithm for 1 s to determine whether the particle is sufficiently confined to graduate to the ‘main simulation’. The steps of the pre-simulation are as follows. E is used to denote the anti-atom total energy.

1. Randomly sample initial velocity, v_0 , in each spatial dimension (sample from normal distribution, Equation 3.41) and initial position, $x_{1/2}$ (from uniform distribution over volume of positron plasma, and apply Taylor expansion, Equation 3.48). Initialise $n = 25$.
2. If($n \neq 1$) {Model radiative cascade: roll r (Equation 3.10) and calculate $\xi(n)$ (Equation 3.9). If($r < \xi(n)dt$) { $n = n - 1$ }}
3. Evolve position, velocity and time (symplectic stepper, Equation 3.39 and $t = t + dt$).
4. If($E > 0.75K$) {Go to 1.}
5. If($x^2 + y^2 > r_{\text{trap}}$) {Go to 1.}
6. If($t > 1\text{s}$ and $n = 1$) {Enter main simulation.} Else {Go to 2.}

The main simulation then repeats step 3 and the check in step 5 until either the simulation end time is reached or the anti-atom fails the check in step 5 (indicating annihilation with the trap wall). A number of diagnostics can also be called on each timestep to extract and save information about the trajectory. One such diagnostic is used to determine the axial (Equation 3.33) and transverse (Equation 3.34) energy components of a simulated anti-atom. Since the simulation has finite timestep, the energy components at the crossing of the magnetic minimum are determined using quadratic interpolation [9].

3.4.1 Symplectic Leapfrog Algorithm

The simplest way to solve the coupled differential equations (Equation 3.37 and 3.38) is to use the Euler method,

$$x_{n+1} = x_n + v_0 dt, \quad (3.49)$$

where n denotes the iteration index. The accuracy of the solution is improved by taking the value of the velocity at the midpoint of the interval, otherwise known as the Leapfrog method of integration,

$$x_{n+1} = x_n + v_{n+1/2} dt, \quad (3.50)$$

where the velocity evolves according to,

$$v_{n+3/2} = v_{n+1/2} - \frac{\mu_B}{M} \nabla |\mathbf{B}(\mathbf{x}_{n+1/2})|, \quad (3.51)$$




Figure 3.12: Four points lie at time t in phase space (\mathbf{x}, \mathbf{p}) and are each mapped by the symplectic Leapfrog algorithm to a point in phase space $(\mathbf{x}', \mathbf{p}')$ at time t' . The quadrilaterals formed by connecting each of the points before and after application of the Leapfrog algorithm have area dA and dA' respectively. Figure from [10].

for an individual spatial dimension [10]. It is trivial to generalise this to three dimensions and we could equivalently evaluate the position (rather than the velocity) at half timestep intervals as in Equation 3.39. It is important to note that the velocity step in Equation 3.51 requires evaluation of the vector field, $\mathbf{B}(\mathbf{x}_{m+1/2})$, and therefore Equation 3.50 must be performed in each dimension prior to the velocity step [10]. The Leapfrog method has an overall error on the order of dt^2 , making it a second order solution. For comparison, the Euler method is first order, making it less accurate for the same dt . Since the Leapfrog and Euler algorithms require similar computational time to complete, the Leapfrog algorithm is more accurate when computational time is matched. In addition to this, the Leapfrog algorithm has a number of desirable characteristics:

1. It retains the time-reversal-invariance of Newton's laws.
2. Its equations conserve angular momentum exactly.
3. The algorithm ensures the total energy of a trajectory remains fixed within bounds for a static magnetic field. It is conservation of symplectic invariance that leads to this quality; namely that the errors in the energy do not multiply over time and cause the energy to diverge.

Another name for a symplectic integrator is an 'area preserving' integrator. Let us use Figure 3.12 to depict the meaning of this name by taking four points (labelled 1... 4) in phase space (\mathbf{x}, \mathbf{p}) at time t . The four points have coordinates (x, p) , $(x, p + dp)$, $(x + dx, p)$ and $(x + dx, p + dp)$ and enclose an area dA . Then the system evolves to time t' ; forming four new

coordinates in phase space which now enclose an area dA' . Liouville's theorem asserts

$$dA' = dA. \quad (3.52)$$

Without initially asserting the theorem, the two areas are related by,

$$dA' = \det J dA, \quad (3.53)$$

where the Jacobian,

$$J = \begin{pmatrix} \frac{\partial x'}{\partial x} & \frac{\partial x'}{\partial p} \\ \frac{\partial p'}{\partial x} & \frac{\partial p'}{\partial p} \end{pmatrix} \quad (3.54)$$

Therefore $\det J = 1$ for a symplectic integrator [10].

All higher order symplectic integrators require more magnetic field evaluations per timestep, such as the 'Position Extended Forest-Ruth Like' (PERFL) fourth order time reversal invariant algorithm [93][10], which requires four magnetic field evaluations per timestep. A fair comparison between the two is hence a comparison of the accuracies of the two methods; with the timestep used by PERFL a factor of four greater than that used by Leapfrog. With the fair comparison, PERFL is still about 340 times more accurate [10]. However, the simplicity of the Leapfrog algorithm makes it favourable when second order accuracy is sufficient. Studies in ALPHA have compared outputs of simulations using 2nd and 4th order integrators and found that our simulations converge sufficiently using the 2nd order Leapfrog algorithm, though this could be a topic for future study [94].

3.4.2 Magnetic Field Modelling

Since the second step of the symplectic Leapfrog integrator (Equation 3.39) requires knowledge of $|\mathbf{B}(\mathbf{x})|$ at every timestep, we calculate the vector field $\mathbf{B}(\mathbf{x})$ from each magnet in the trap (Octupole, Mirror Coils A... E and Solenoid A), and take the modulus of their vector sum every timestep. The method used to determine $\mathbf{B}(\mathbf{x})$ for each of these magnets is described below. The methods are optimised by adjusting parameters in the solution by comparison with an accurate Biot Savart model. To generate this model, the most accurate knowledge of the placement of the individual wires on the magnet form is used. Each wire is then split into N straight line segments, where each line segment has length $d\mathbf{l}$ and points in the direction of the current flow. The Biot Savart law then states that this segment contributes

$$d(\mathbf{B}(\mathbf{x})) = \frac{\mu_0 I}{4\pi} \frac{d\mathbf{l} \times \mathbf{r}}{|\mathbf{r}|^3}, \quad (3.55)$$

where \mathbf{r} is the coordinate vector from the line segment to the position where the field is sampled, \mathbf{x} , and I is the current in the magnet [95]. The magnetic field at \mathbf{x} is then approximated by the sum over line segments, i ,

$$\mathbf{B}(\mathbf{x}) = \frac{\mu_0 I}{4\pi} \sum_i^N \frac{d\mathbf{l}_i \times \mathbf{r}_i}{|\mathbf{r}_i|^3}. \quad (3.56)$$

The Biot Savart model is particularly good at accurately solving for complicated details of the magnetic field, such as the magnetic field contribution from the end turns of the Octupole that connect the poles together, as it can model any shape of current flow. In addition, it is capable of accurately solving for the magnetic field at locations close to the current carrying wires, where other models may be inaccurate. The number of line segments can be arbitrarily increased within machine precision but computation time increases linearly with N [95]. Unfortunately, obtaining an accurate model for $\mathbf{B}(\mathbf{x})$ requires increasing N so much that the method is very slow compared to the semi-analytic models discussed below for the same levels of accuracy [95]. It is possible, however, to calculate the Biot Savart solution once with a large N and to then find the parameters of a semi-analytic model that minimise the difference between the solutions of the semi-analytic model and the accurate Biot Savart model, as we do in the methods described below.

3.4.2.1 Octupole Model

We use the approximate solution for the magnetic field of a finite-length octupole (with end turns that connect the currents in the eight poles), symmetric about its axial center, derived in [1]. The method assumes a general expression for the magnetic vector potential of the octupole, involving arbitrary functions of z . By applying Maxwell's laws, these arbitrary functions can be expressed in terms of derivatives of a single function of z . We then assume an expression for this function, which contains free parameters that are determined by fitting to a Biot Savart model, with derivatives that are simple to compute.

In cylindrical coordinates, (r, ϕ, z) , the components of the magnetic vector potential are approximated as

$$\begin{aligned} A_s &= [G_5(z)r^5 + G_7(z)r^7 + \dots] \cos(4\phi), \\ A_\phi &= [H_5(z)r^5 + H_7(z)r^7 + \dots] \sin(4\phi), \\ A_z &= [F_4(z)r^4 + F_6(z)r^6 + F_8(z)r^8 + \dots] \cos(4\phi), \end{aligned} \quad (3.57)$$

where $r^2 = x^2 + y^2$ and F , G and H are functions of z to be determined [1]. Using the condition

$$\nabla^2 A_z = 0, \quad (3.58)$$

we find an expression for each $F_i(z)$ ($i = 6, 8, 10\dots$) in terms of $F_4(z)$ and derivatives,

$$F_i(z) = F_4^{i-4}(z) \prod_{j=6,8,10\dots}^i \left(\frac{-1}{j^2 - 16} \right) \quad \{i = 6, 8, 10\dots\}, \quad (3.59)$$

where a superscript integer denotes the derivative order; for example, $F_4^1(z) = \frac{dF_4(z)}{dz}$. Using the condition

$$\nabla^2 A_x = \nabla^2 A_y = 0, \quad (3.60)$$

we find

$$H_k(z) = G_k(z) \quad \{k = 5, 7, 9, \dots\}. \quad (3.61)$$

In addition, the condition

$$\nabla \cdot \mathbf{A} = 0 \quad (3.62)$$

allows us to find an expression for $G_k(z)$ ($k = 5, 7, 9, \dots$),

$$G_k(z) = -\frac{F_{k-1}^1(z)}{k+5} \quad \{k = 5, 7, 9, \dots\}, \quad (3.63)$$

which, in combination with Equation 3.59, allows us to express $G_k(z)$ in terms of $F_4(z)$ and derivatives [1]. Thus, we have expressed each of $F_i(z)$ ($i = 6, 8, 10, \dots$), $G_k(z)$ and $H_k(z)$ ($k = 5, 7, 9, \dots$) in terms of $F_4(z)$ and derivatives. We assume [1],

$$F_4(z) = D [\operatorname{erfc}((z - z_f)/\Delta z) - \operatorname{erfc}((z + z_f)/\Delta z)] \quad (3.64)$$

where D is a constant, $\pm z_f$ are the approximate ends of the Octupole, Δz is the axial distance over which the octupole field drops approximately from full to zero field and $\operatorname{erfc}(z)$ is the complementary error function,

$$\operatorname{erfc}(z) = \frac{2}{\sqrt{\pi}} \int_z^\infty e^{-z'^2} dz', \quad (3.65)$$

which can be numerically evaluated with a fractional error everywhere less than 1.2×10^{-7} using a routine based on Chebyshev fitting given in [96]. The magnetic field is then calculated from the magnetic vector potential using

$$\mathbf{B} = \nabla \times \mathbf{A}. \quad (3.66)$$

The parameters that enter the magnetic vector potential are found by fitting to a Biot Savart model to be $z_f = 129.46$ mm, $D = 19.665$ and $\Delta z = 16.449$ mm, which gives the difference between the Biot Savart model and the Octupole model to be less than 0.02 T at all positions within the trap when the Octupole field was ~ 1.5 T at the electrode inner surface [1].

3.4.2.2 Mirror Coil Model

The Mirror Coils are modelled as two current loops separated by an axial distance, z_{sep} . Although an exact expression for the magnetic vector potential from a loop of current exists [97], numerical evaluation of the expression converges slowly within the trap near to the Mirror Coil windings [1]. We therefore use the approximate form,

$$A_\phi = C \frac{1}{2a\lambda} \left[(a^2 + s^2 - 2a\lambda r)^{-1/2} - (a^2 + s^2 + 2a\lambda r)^{-1/2} \right] \quad (3.67)$$

where $s^2 = r^2 + z^2$, $C = I\mu_0 a^2/4$, a is the loop radius, μ_0 is the vacuum permeability and λ is a dimensionless fit parameter. When $\lambda = \sqrt{3}/2$, this approximation is exactly equal

to the first two terms of the exact expression for A_ϕ from a current loop. The parameters are again found by fitting to a Biot Savart model. When $a = 45.238$ mm, $\lambda = 0.9019$ and $z_{\text{sep}} = 8.251$ mm, differences between the analytic expression given above and the Biot Savart model are less than 0.02 T when the peak Mirror Coil field is ~ 1 T [1].

This solution for the magnetic field of a current loop has become known as the *Truncated Approximate Vector Potential Model* (TAVP) and is much less accurate off-axis than other models for the magnetic field of a circular current loop [95]. The McDonald model has been proposed as a faster and more accurate alternative to TAVP [95]. However, for unknown reasons, the McDonald model is significantly slower than TAVP when used within the framework of the simulation outlined in this chapter. Therefore, we proceed with TAVP for this application (which is the same model used in [9]).

3.4.2.3 Solenoid Model

We use an off-axis expansion method to solve for the magnetic field of Solenoids A and B. The solenoids are modelled as a thin shell solenoid extending axially from Z_L to Z_R and with radius R , for which the on-axis magnetic field is given in [95]. As shown in Figure 3.13, there are also two end corrector current loops which sit at an axial inset, ΔZ_{loop} , relative to the main body of the solenoid and each have radius r_{loop} . We are required to calculate the magnetic field of Solenoid A in the antihydrogen trapping region, which is shown approximately as the pink region in Figure 3.13 and extends axially off to the left of the figure (see Figure 2.1 for a diagram of the ALPHA-2 neutral trap magnets).

The analytic expression for the on-axis components of magnetic field strength from the thin shell solenoid and corrector loops is

$$\begin{aligned}
B_z(0, 0, z) &= IK \left\{ \frac{Z_R - z}{\sqrt{(Z_R - z)^2 + R^2}} - \frac{Z_L - z}{\sqrt{(Z_L - z)^2 + R^2}} \right. \\
&\quad \left. + \frac{K_{\text{loop}} r_{\text{loop}}^3}{\left(r_{\text{loop}}^2 + (Z_L + \Delta Z_{\text{loop}} - z)^2\right)^{\frac{2}{3}}} + \frac{K_{\text{loop}} r_{\text{loop}}^3}{\left(r_{\text{loop}}^2 + (Z_R - \Delta Z_{\text{loop}} - z)^2\right)^{\frac{2}{3}}} \right\}, \\
B_x(0, 0, z) &= 0 \\
B_y(0, 0, z) &= 0
\end{aligned} \tag{3.68}$$

where I is the current and K and K_{loop} are constant scaling factors. The derivatives of the magnetic field with respect to z are approximated using the finite difference method, by calculating the on-axis field at positions shifted by a small axial displacement, ϵ . The first derivative is given by

$$\begin{aligned}
\frac{dB(0, 0, z)}{dz} &= \frac{1}{12\epsilon} [B(0, 0, z - 2\epsilon) - 8B(0, 0, z - \epsilon) + 8B(0, 0, z + \epsilon) \\
&\quad - B(0, 0, z + 2\epsilon)],
\end{aligned} \tag{3.69}$$




Figure 3.13: Diagram of the shape of the model for Solenoid A (can equally be used for Solenoid B). The model consists of a main Solenoid body and two end corrector loops that are inset by a distance, ΔZ_{loop} , from the ends of the main body. The pink region continues off to the right of the figure and indicates the approximate region in which we are required to calculate the magnetic field of Solenoid A (the antihydrogen trapping region).

and similar expressions for the higher order derivatives are given in [98]. At high values of ϵ , the derivative becomes inaccurate, whereas at very low values the result fluctuates largely with ϵ . An optimal value of $\epsilon = 5 \times 10^{-4}$ m was found that lies just above the low ϵ fluctuations. We then use the McDonald method [99] to 4th order to generate the off-axis magnetic field components,

$$\begin{aligned}
 B_r(x, y, z) &= \frac{r^6}{2} \frac{d^3 B_z(0, 0, z)}{dz^3} - r^2 \frac{dB_z(0, 0, z)}{dz}, \\
 B_x(x, y, z) &= B_r(x, y, z) \frac{x}{r^2}, \\
 B_y(x, y, z) &= B_r(x, y, z) \frac{y}{r^2}, \\
 B_z(x, y, z) &= B_z(0, 0, z) + \frac{r^8}{4} \frac{d^4 B_z(0, 0, z)}{dz^4} - r^4 \frac{d^2 B_z(0, 0, z)}{dz^2}.
 \end{aligned} \tag{3.70}$$

We again find the parameters, $Z_R = 166.74$ mm, $Z_L = 308.43$ mm, $R = 32.608$ mm, $K = 0.0043284$, $K_{\text{loop}} = 0.32619$, $r_{\text{loop}} = 42.744$ mm and $\Delta Z_{\text{loop}} = 42.744$ mm, fitting to a Biot Savart model for Solenoid A. The least accurate field in the relevant region is the radial field component off-axis closest to the Solenoid, where the field error is ~ 0.05 T at $I = 217$ A.

Simulation Results

In this chapter, I present results of detailed Monte Carlo simulations of the proof-of-principle adiabatic expansion experiment described in Chapter 2. I compare experimental and simulated annihilation times during magnetic ramp down, and use the Two-Sample Kolmogorov-Smirnov (KS) test to quantify the agreement. Reasonable agreement between simulated and experimental annihilation times provides confidence in the simulation to accurately represent the experimental trapped antihydrogen population. Since we have no means of measuring the energy distribution of the trapped antihydrogen population directly, I extract the energy of the simulated population. I then employ an established technique which uses the experimental annihilation times and the simulated energy distribution, to reconstruct the (experimental) energy of the trapped antihydrogen population. Finally, I answer an important question raised in Chapter 2: if the trap expansion is predominantly in the axial direction and therefore reduces the axial energy of the trapped anti-atoms, why do we observe a significant difference in annihilation time of adiabatically cooled anti-atoms during Octupole FRD, in which the anti-atoms are released radially? Answering this question is vital to understanding what the experimental diagnostic measures: does it measure only the small degree of transverse cooling resulting from axial to transverse energy mixing? Or is there an effect coupling the axial energy of the trapped anti-atoms to the annihilation time during Octupole FRD?

Simulation Parameter	Trial A	Trial B	Trial C
Total number of simulated anti-atoms	200000	200000	200000
Simulated Duration	~ 24 s	~ 24 s	~ 24 s
Simulation Timestep, dt	3.5×10^{-6} s	3.5×10^{-6} s	3.5×10^{-6} s
Computation Time per Anti-atom	~ 40 s	~ 40 s	~ 40 s
Number of particles surviving until FRD ($t = 0$ s)	179631	180553	175399

Table 4.1: Typical simulation parameters for Adiabatic Expansion Trial A, Control Trial B and Control Trial C.

The simulations presented in this chapter involve initialisation techniques described in Section 3.4, with the positron plasma ellipsoid centered at the electrode in which positrons are held during antihydrogen formation (see Section 1.2.3 for a review of antihydrogen formation). Trajectory propagation and magnetic field determination follow protocols outlined in Section 3.4. Until stated otherwise, presented results refer to the standard simulation with simulation parameters shown in Table 4.1.

For adiabatic expansion, I am required to simulate around 200,000 anti-atoms. At numbers much less than this, the ensemble-averaged axial (and transverse) energy components (which I will present in Section 4.3) become macroscopically uneven as a function of time. Since each adiabatic expansion trial is required to run for ~ 24 s of real time for each particle which requires ~ 40 s of computation time per particle (see Table 4.1), I am required to parallelise individual trajectory simulations by running batch processes on independent computers (this is feasible since the low numbers of neutral antihydrogen atoms confined in the trap means the interaction between anti-atoms is negligible). Running 100 parallel batch processes of 2000 series particles for each adiabatic expansion trial results in a simulation duration of ~ 1 day.

4.1 Kolmogorov-Smirnov (KS) Statistics

In Figure 4.1, I compare the experimental annihilation time data (presented in Chapter 2), determined by reconstructing antihydrogen annihilation vertices from tracks of secondary charged particles in the ALPHA-2 SVD (see Section 1.2.4), to simulated annihilation times for Trials A, B and C assuming a linear Octupole ramp down. For all three trials, the simulated distribution appears broadly similar to that of the experimental data. Crucially, the simulated data agrees with the experimental data that adiabatically cooled (Trial A) anti-atoms tend to annihilate later in time than control samples held in a static magnetic trap (Trials B and C). In addition, the simulated annihilation time distributions have the same general shape as the distributions of experimental data. More specifically, the tails of the distributions are

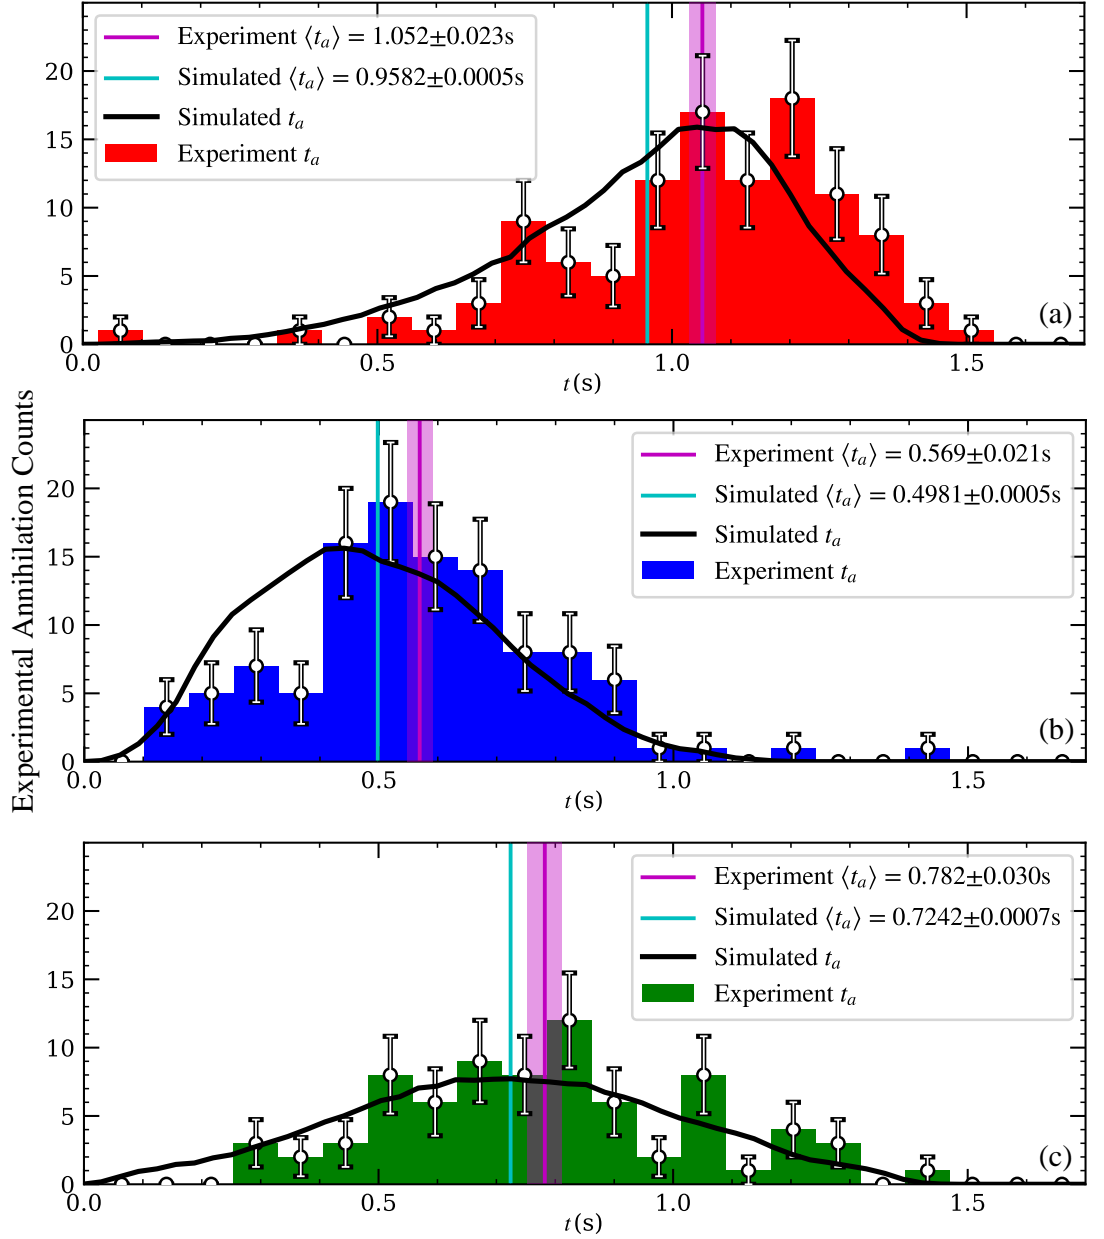


Figure 4.1: Comparison of simulated (black line plot) and experimental annihilation time distributions (red, blue and green histograms) during Octupole FRD for the three experimental trials: (a) Adiabatic Expansion Trial A, (b) Control Trial B and (c) Control Trial C. The simulated annihilation distribution is normalised to the area of the detected annihilation histogram. Error bars (white) represent Poissonian (\sqrt{N} , where N is the number of experimental counts, see Table 2.1) counting statistics on experimental data. Similar error bars for the simulated annihilation time data are not visible on this scale. The average annihilation time, $\langle t_a \rangle \pm 1\sigma_{SE}$, for the simulated (vertical cyan line, $1\sigma_{SE}$ region not visible on this scale) and experimental (vertical magenta line, $1\sigma_{SE}$ magenta region) distributions are also shown, with uncertainty given by the standard error on the mean, $\sigma_{SE} = \sigma/\sqrt{N}$, where σ is the sample standard deviation. Simulation parameters are given in Table 4.1.

in agreement i.e. the tails of the Trial A and B distributions are directed towards earlier and later times respectively, whereas the Trial C annihilation time distribution is roughly symmetric about the central annihilation time.

To quantitatively determine whether the simulated and experimental annihilation times are in agreement, I use the two-sample Kolmogorov-Smirnov (KS) test, which is a statistical measure of the probability that two independent discrete samples are drawn from the same distribution [100]. For a sample of N events, I generate the cumulative distribution function of the discrete annihilation times, $S_N(t)$, for both simulation, $S_{N_{\text{sim}}}(t)$, and experiment, $S_{N_{\text{exp}}}(t)$. The KS statistic, D , is the largest vertical separation (greatest probabilistic difference) between $S_{N_{\text{sim}}}(t)$ and $S_{N_{\text{exp}}}(t)$,

$$D = \max_{-\infty < t < \infty} |S_{N_{\text{sim}}}(t) - S_{N_{\text{exp}}}(t)|. \quad (4.1)$$

Therefore, a small D indicates a small difference between two samples. The p value for an observed value of D is given approximately by [100],

$$p = \text{Probability}(D > \text{observed}) \approx Q_{KS}([\sqrt{N_e} + 0.12 + \frac{0.11}{\sqrt{N_e}}]D), \quad (4.2)$$

where

$$Q_{KS}(\lambda) = 2 \sum_{j=1}^{\infty} (-1)^{j-1} e^{-2j^2\lambda^2}, \quad (4.3)$$

and the effective number of data points,

$$N_e = \frac{N_{\text{sim}}N_{\text{exp}}}{N_{\text{sim}} + N_{\text{exp}}}. \quad (4.4)$$

If the p value is greater than the level of significance, α , I accept the null hypothesis (that the two samples are drawn from the same distribution). In this case, the difference between the two samples is not significant enough to say that they have different distributions. The default level of significance is $\alpha = 0.05$ [101].

The KS test is effective at finding shifts between distributions, which makes it a good candidate for our comparison because temporal shifts of experimentally-measured annihilation times relative to simulated annihilation times could indicate a difference between experimental and simulated energy distributions. The reason the KS test is good at finding shifts is that the variance of a cumulative distribution function, $S_N(t)$, is proportional to $S_N(t)[1 - S_N(t)]$ and is therefore largest at $S_N(t) = 0.5$ i.e. the median value. The KS statistic, being the maximum absolute difference between the two cumulative distributions, is therefore weighted more significantly toward the median value. This makes it less sensitive at the tails of the annihilation distributions where the number of data points are low; qualifying it further for our comparison [96].

In Figure 4.2 I compare $S_N(t)$ between simulated and experimental data based on a linear Octupole ramp down. For Trials A and B, the p values are $\sim 10^{-4}$ and $\sim 10^{-3}$ respectively;




Figure 4.2: Comparison of simulated and experimental cumulative distribution function of annihilation times, $S_N(t)$, where I have simulated a linear magnetic ramp down with simulation parameters given in Table 4.1. Simulated cumulative distribution function of annihilation, $S_{N_{\text{sim}}}(t)$, is shown for (a) Trial A (red solid curve), (b) Trial B (blue dashed curve) and (c) Trial C (green dotted curve). In each subfigure, the experimental cumulative distribution function of annihilation time, $S_{N_{\text{exp}}}(t)$, is shown as a black solid curve. Above each subfigure I show the KS p value, which represents the confidence level to which I can conclude the simulated and experimental annihilation times are drawn from the same distribution. The maroon dashed line indicates the largest vertical deviation between the two cumulative distribution functions, or the position of the KS statistic, D .

indicating a low probability that the simulated and experimental annihilation times were sampled from the same distribution. For Trial C, however, the p value is 0.137 (> 0.05) which suggests the difference between the simulated and experimental annihilation time distributions is not significant enough to reject the null hypothesis.

To determine whether the low p values for Trial A and Trial B are consistent with an offset between simulated and experimental annihilation times, I can shift each experimental data point by the same offset and then recalculate the KS p value. Figure 4.3 shows the KS p value as a function of this timing offset. Trials A and B peak to very high agreement between simulation and data at offsets of approximately -90 ms and -70 ms respectively. For Trial C, although the KS p value is consistent with an offset of 0 ms, there is a higher probability that the experimental data is offset by -45 ms. In addition, I accept the null hypothesis in all three experimental trials with the same offset ranging from approximately -50 ms to -100 ms. Peak agreement for the three experimental trials is found with an offset of ~ -70 ms. Note that a high KS p value with a negative offset indicates that the experimental data is annihilating *later* in time than that predicted by the simulation.




Figure 4.3: KS p value as a function of timing offset in seconds. Timing offset is defined as the offset made to each experimental data point with respect to the simulation i.e. a high KS p value with a negative offset suggests the experimental annihilation times tend to occur *later* than the simulated annihilation times. Trials A, B and C are displayed as red solid, blue dashed and green dotted lines respectively. The red region indicates the range of p values (< 0.05) at which the null hypothesis (simulated and experimental annihilation times are drawn from same distribution) is rejected, whereas the green region indicates the range of p values at which the null hypothesis is accepted.

The aim of the KS analysis described in this section was to determine whether the simulation (described in Section 3.4) is capable of accurately reproducing features of the

experimental distribution of antihydrogen. Since both simulation and experiment produce an annihilation time distribution, a comparison of these distributions (using KS analysis) provides a test of the simulation accuracy. In this section, I identified a range of possible temporal offsets of the experimental annihilation times relative to the simulated annihilation times. The outstanding question is whether the offsets I have identified are within the experimental uncertainty on the measured annihilation times. In the following section, I will discuss the range of uncertainty contributions to the experimentally-measured annihilation time distribution.

4.2 Possible Explanations for Temporal Shifts between Experiment and Simulation

The adiabatic expansion experiment, as presented in Chapter 2, was conducted in a single experimental run in 2016. For this reason, there was insufficient opportunity to perfect the experimental protocol. With this in mind, I consider effects leading to an offset between experimental and simulated annihilation time data ~ 70 ms ($\sim 4.6\%$).

In ALPHA, a National Instruments (NITM) LabVIEW program known as the ‘sequencer’ has central control over the experiment. A user can program a ‘sequence’ which specifies the specific experimental procedure - for example - it specifies the voltages on the Penning-Malmberg trap electrodes as a function of time. The sequencer is then responsible for communicating these commands with the relevant hardware or more specific control software. Another experimental parameter that must be controlled is the set of currents in the superconducting magnets of the magnetic minimum trap as a function of time. To achieve control, the sequencer communicates with an independent LabVIEW program, known as the ‘magnet controller’, which is responsible for executing the sequencer commands by communicating with the magnet power supplies.

To begin diagnostic ramp down, the sequencer sends a command, known as the *FRD trigger*, to the magnet controller to ramp the Octupole magnet linearly to 0 A in 1.5 s. The time of the FRD trigger is recorded by a 10 MHz clock that is accurately referenced to the SI standard second. Antihydrogen annihilation times, recorded by the SVD, are referenced to the same 10 MHz clock, resulting in annihilation times measured relative to the FRD trigger ($t = 0$ s) with an uncertainty of ~ 2 μ s, which is very small compared to the ~ 70 ms predicted offset between experimental and simulated annihilation times.

A lack of knowledge of the Octupole current as a function of time could be a source of uncertainty between experiment and simulation. For example, if the simulation models the Octupole current as a faster linear ramp than was achieved in the experiment, with all other parameters being equal, I would observe simulated annihilation times appearing earlier in time

than experimentally detected annihilation times. As mentioned in Chapter 2, currents in the superconducting magnets were recorded at high precision throughout the experiment using a Direct Current Current Transducer (DCCT) with 1 s timing resolution. A few days after the experiment, each of the three experimental trials were repeated five times without particles, during which time the magnet controller recorded the DCCT currents with 250 ms timing resolution. In addition, upon receiving the FRD trigger, it was intended that the magnet controller record the DCCT current at 20 μ s resolution relative to the trigger. Unfortunately, this high resolution DCCT reading was not recorded for the Trial C experimental protocol.

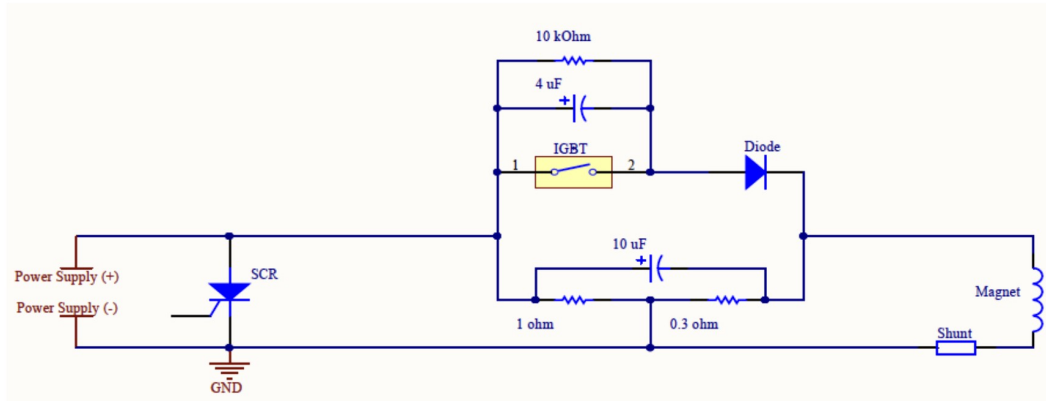


Figure 4.4: Electrical circuit of the superconducting Octupole magnet, shown as an inductor labelled *Magnet*. The Silicon-Controlled Rectifier (labelled *SCR*), Isolated-Gate Bipolar Transistor (labelled *IGBT*) and *Diode* act as controllable switches. In normal operation, the Magnet is superconducting, the SCR is open and the IGBT and Diode are closed, allowing the power supply to drive current into the Magnet. In the event of a quench in the Magnet, the SCR is closed and the IGBT and Diode are closed, shorting the power supply out of the circuit and enabling the stored energy in the Magnet to dissipate in the 0.3 Ω resistor.

The 250 ms resolution DCCT readings suggested the Octupole current did not ramp to 0 A in the requested 1.5 s; instead it suggested the ramp down took on-average ~ 1.53 s. For this reason, the primary simulation data presented in this thesis has assumed a linear Octupole ramp down in ~ 1.53 s, starting immediately when the sequencer triggers the magnet controller to ramp down the Octupole. It is possible that there was a delay between the time at which the power supply received the FRD command from the magnet controller and the time at which the current began to decay, that is not resolved with 250 ms resolution. In addition, the current decay may have been nonlinear in some way.

Nonlinearities and delays may have resulted from the temperature-dependent behaviour of electrical components in the circuit of the superconducting Octupole, a diagram of which is shown in Figure 4.4. The circuit is designed to protect the power supply and superconducting magnet in the event of a quench. The Silicon Controlled Rectifier (SCR), Isolated-Gate Bipolar Transistor (IGBT) and Diode act as switches that can be controlled. In normal

operation (the magnet has not quenched), the SCR is open and the IGBT and Diode are closed, allowing current to flow through the superconducting Octupole. In the event of a quench, the SCR is closed and the IGBT is opened, shorting the power supply out of the circuit and allowing the stored energy in the Octupole to dissipate in the $0.3\ \Omega$ resistor. During Octupole FRD, the circuit is constantly in normal operation mode. The IGBT and Diode have a finite $1 - 2\ \text{V}$ voltage drop, meaning the $1\ \Omega$ resistor constantly dissipates power. This $1\ \Omega$ resistor forms a non-trivial inductor-resistor (LR) circuit, acting as a low-pass filter. This effect, in combination with the voltage drop across the $1\ \Omega$ resistor means the current will tend to lag linear control from the power supply. Since the resistance of the $1\ \Omega$ resistor and the voltage across the IGBT and Diode are temperature dependent, the current in the Octupole as a function of time during FRD can be subject to dependence on fluctuations in laboratory temperature.

Since the $250\ \text{ms}$ resolution DCCT trace and the high ($20\ \mu\text{s}$) DCCT trace during FRD were recorded a few days after the experiment, we cannot rule out differences in laboratory temperature leading to differences in the measured current behaviour compared to the experiment itself. To characterise the Octupole current as a function of time during ramp down, in Section 4.2.1, I will use the high resolution DCCT readouts for Trials A and B to estimate the effect of nonlinearities and delays in the Octupole current ramp, assuming the Trial C ramp down is identical to that of Trial A.

Another possible source of temporal offset is a delay between the time at which the sequencer triggers the magnet controller and the time at which the magnet controller commands the power supply to ramp down. Such offsets are expected to be small ($\sim 1\ \text{ms}$), but they have not been measured and therefore an offset of this kind cannot be ruled out.

I should also consider simulation inaccuracies as a possible source of timing offset between experimental and simulated annihilation times. For example, the axial center of the electrode in which positrons are held during mixing is expected to have an uncertainty of $\sim 1\ \text{mm}$. Axial shifts in the center of antihydrogen initialisation of $1\ \text{mm}$ in Trials A and B (same starting magnetic configuration) result in shifts in magnetic potential at the center of initialisation of $\sim 5\ \text{mK}$. For Trial C the equivalent shift in initial magnetic potential is $\sim 1\ \text{mK}$ since initialisation in Trial C takes place in a smaller magnetic field gradient. I can use the trap depth as a function of time (Figure 2.7) to estimate the temporal offset corresponding to these possible shifts in magnetic potential at the initialisation location. Around the mean time of annihilation in each trial (Table 2.1), the energy shifts of $5\ \text{mK}$, $5\ \text{mK}$ and $1\ \text{mK}$ in Trials A, B and C respectively correspond to estimated temporal shifts of $\sim 15\ \text{ms}$, $\sim 10\ \text{ms}$ and $\sim 2\ \text{ms}$.

Uncertainties in the initial temperature of the total distribution of formed antihydrogen were expected to be large ($\sim 10\ \text{K}$), since it is difficult to accurately measure the temperature of the positron plasma during mixing, due to constant evaporative cooling and interaction

with the antiproton plasma. However, uncertainties in the positron plasma temperature are expected to have a small effect on the initial energy distribution of formed antihydrogen as the magnetic trap truncates the relatively hot distribution at the trap depth (~ 0.5 K).

It is also possible that inaccuracies in simulated magnetic field models (Section 3.4.2) relative to the magnetic field of the physical magnets are responsible for the observed temporal offset. In Section 4.8 I will investigate the timing offset corresponding to known magnetic field modelling inaccuracies. It is also likely that other differences exist between the simulated and physical magnets that we are not aware of. This is backed up by the observation that antihydrogen annihilations in the ALPHA-2 trap are slightly asymmetric in z , suggesting an error in the magnet windings.

4.2.1 Current Decay Temporal Shifts

As described above, a lack of knowledge of the Octupole current during FRD is a possible source of temporal offset between experimental and simulated annihilation times during the adiabatic expansion experiment. To investigate this issue, I will re-simulate Trials A, B and C with a modified Octupole current during FRD. Rather than modelling the Octupole current ramp as a perfectly linear ramp from 0 s to ~ 1.53 s, in the modified simulation the Octupole current will follow the 20 μ s resolution DCCT trace. As stated above, these high resolution DCCT traces were obtained a few days after the experiment, and therefore the current decay may differ from that in the experiment itself due to temperature fluctuations in the electrical circuit of the Octupole magnet. In addition, since the Trial C high resolution DCCT trace is missing, I will assume the current decay in Trial C is identical to that of Trial A. I choose Trial A (rather than Trial B) since Trials A and C ramp down from the same magnetic configuration (see Figure 2.5).

In Figure 4.5(a) and 4.5(b), I compare perfectly linear current ramps to those measured by the DCCT at 20 μ s resolution during current ramps that mimic Trial A and Trial B respectively (Trial C is missing since the data was not recorded). In Figure 4.5(c) and 4.5(d), I plot the temporal difference, Δt , between the linear current ramp and the DCCT trace for Trials A and B respectively. Δt is around 100 ms for the majority of the ramp down in both Trial A and Trial B, but Δt is about 10 ms larger in Trial B than in Trial A. Beneath Δt , I plot the simulated and experimental annihilation time data (the same data as plotted in Figure 4.1) to give the reader an idea of the temporal region in which Δt would have the largest impact. For example, the Trial A experimental annihilation times are peaked at around 1.1 s, at which time $\Delta t \sim 110$ ms. I can infer that there are subtle differences between the measured DCCT currents of Trial A and Trial B, but in general the offsets between the linear and DCCT current ramps are large enough, if not too large, to correct for the observed shift between experimental and simulated data.

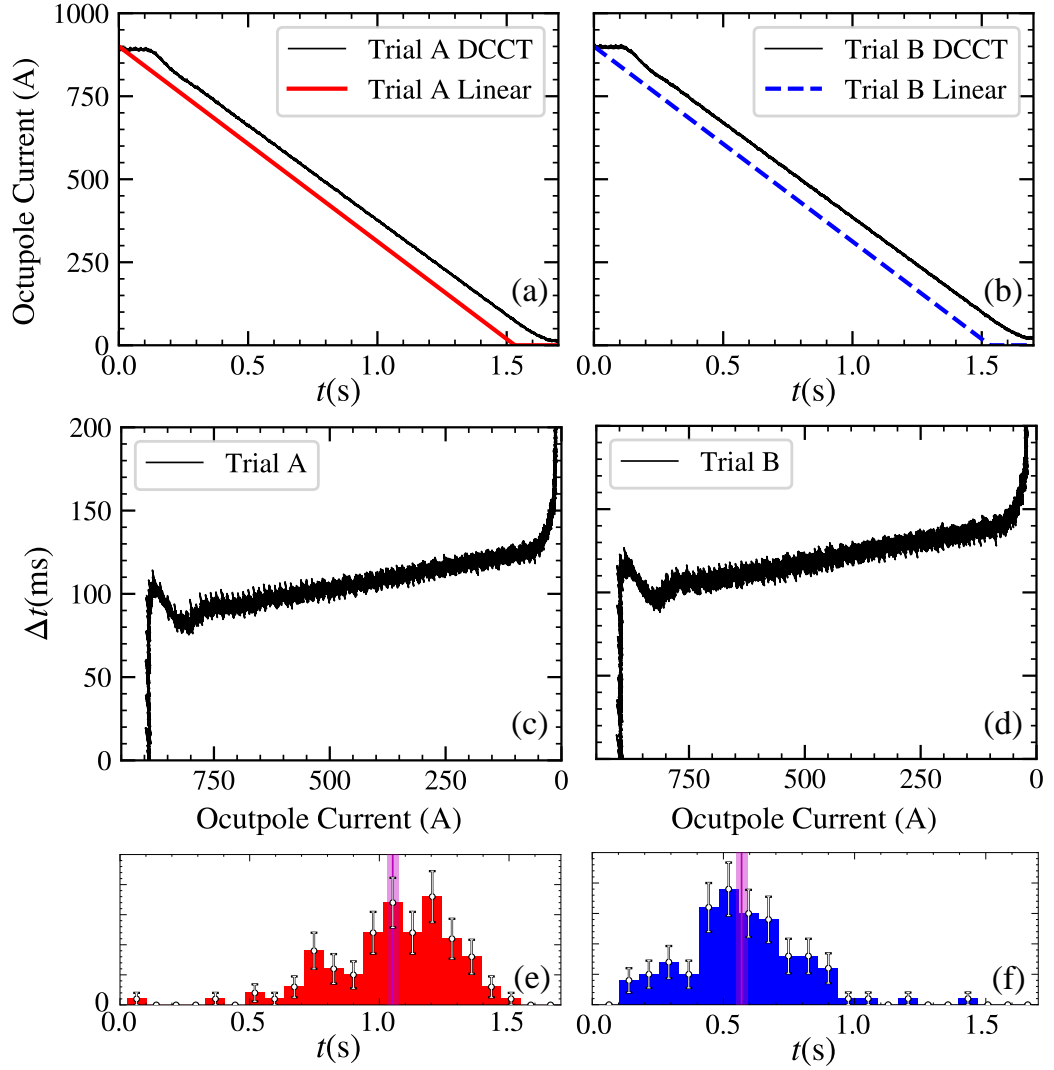


Figure 4.5: (a) Trial A Octupole current as a function of time relative to the start of Octupole FRD, red solid curve shows perfectly linear current decay, whereas black solid curve shows current recorded by DCCT. (b) Trial B Octupole current as a function of time relative to the start of Octupole FRD, blue dashed curve shows perfectly linear current decay, whereas black solid curve shows current recorded by DCCT. (c) Temporal difference, Δt , at which the Trial A Linear (red solid curve in (a)) and Trial A DCCT (black solid curve in (a)) have the same Octupole current, as a function of Octupole Current. (d) Temporal difference, Δt , at which the Trial B Linear (blue dashed curve in (b)) and Trial B DCCT (black solid curve in (b)) have the same Octupole current, as a function of Octupole Current. Note that I do not show the equivalent for Trial C since the DCCT data was not recorded during the experiment. (e) Trial A experimental annihilation data (red histogram), white error bars show Poissonian counting statistics, magenta region is the one standard error region around the ensemble averaged annihilation time, $\langle t_a \rangle$. (f) Trial B experimental annihilation data (blue histogram), white error bars show Poissonian counting statistics, magenta region is the one standard error region around the ensemble averaged annihilation time, $\langle t_a \rangle$.




Figure 4.6: Cumulative distribution function of annihilation time, $S_N(t)$, as a function of time, $t(s)$, relative to the start of Octupole ramp down (at $t = 0$ s) for (a) Adiabatic Expansion Trial A, (b) Control Trial B and (c) Control Trial C. $S_N(t)$ for the experimental data is labelled ‘experiment’ (black curve), $S_N(t)$ for the primary simulation in which the Octupole current ramps linearly from 900.1 A to 0 A in ~ 1.53 s is labelled ‘Linear FRD simulation’ (cyan curve), and $S_N(t)$ for the modified simulation in which the Octupole current during FRD follows the high resolution DCCT trace is labelled ‘DCCT Modified Simulation (magenta curve)’.

In Figure 4.6, I compare the cumulative distribution functions of annihilation time, $S_N(t)$, between the primary simulation (linear current ramp) and the DCCT modified simulation. The modified simulation is significantly different from the primary simulation in all three experimental trials. As predicted by the magnitude of Δt in Figure 4.5(c) and 4.5(d), antihydrogen atoms in the DCCT modified simulation tend to annihilate later than the experimental data, since the DCCT modification appears to overcorrect for the observed offset between the primary simulation annihilation times and the experimental data.

Based on annihilation times of trajectories in the DCCT modified simulation, I recalculate the agreement between simulation and experiment using the KS test described in Section 4.1. The result is shown in Table 4.2, where the first row duplicates the results of the KS test with the primary linear ramp down (presented in Section 4.1) and the second row indicates the KS test result for the DCCT modified simulation. The p values for Trials A and B are significantly increased (good agreement between simulation and experiment), whereas for Trial C the p value decreases. However, in all three experimental trials, the difference between the experimental and DCCT modified simulation data is not significant enough to reject the null hypothesis. In other words, the DCCT modified simulation is in reasonable agreement with the experimental data in all three experimental trials.

Trial	Trial A	Trial B	Trial C
Linear ramp down	1.3×10^{-4}	2.1×10^{-3}	0.137
DCCT Trace (A used for C)	0.709	0.181	0.095

Table 4.2: KS Test p value comparison for the three experimental trials. The “Linear ramp down” row indicates the KS p value result between experimental and simulated annihilation times based on a linear Octupole ramp down. The “DCCT Trace (A used for C)” row indicates the KS p value result between experimental annihilation times and simulated annihilation times, obtained using a modified ramp down in which the current as a function of time follows that recorded by a DCCT during the experiment. “(A used for C)” is used to remind the reader that I modify the Trial A and C simulations using the Trial A DCCT trace, whereas I modify the Trial B simulation using the Trial B trace, since the Trial C trace was not recorded. The results presented in this table can be interpreted as the p values at the intersection of the vertical black line with the three curves in Figure 4.7(a) and 4.7(b) for rows 1 and 2 respectively.

In order to determine whether a temporal offset between the simulation and experiment still exists, with the Octupole current following the DCCT trace during ramp down, I can offset the experimental data with respect to the DCCT modified simulation data and calculate the KS p value as a function of this offset. In Figure 4.7(a) I re-plot the KS p value as a function

of experimental data offset for the linear ramp down simulation and in Figure 4.7(b) I plot the equivalent for the DCCT modified simulation. By specifying the simulation to follow the DCCT trace, I overshoot, leading to good agreement between experimental and modified simulation annihilation times with an offset somewhere between 0 ms and $\sim +60$ ms (optimal agreement at $\sim +40$ ms). Note that a positive offset indicates that the experimentally trapped antihydrogen atoms annihilate *earlier* than that predicted by the DCCT modified simulation. I can conclude that nonlinearities in the Octupole current ramp do cause a significant shift in annihilation time. The shift leads to agreement between simulation and data in all three experimental trials, but suggests optimal agreement with a smaller temporal shift in the opposite direction.

4.2.2 Establishing Reasonable Agreement

In the previous subsection I found that the current, as measured by the DCCT during experimental Trials A and B, was significantly different from the linear ramp down modelled in the primary simulation. A simulation of the as-measured Octupole current during magnetic ramp down showed reasonable agreement between simulated and experimental annihilation times in all three trials, but predicted experimental annihilation times were likely to be on-average around 40 ms earlier than simulated annihilation times. A shift in this direction cannot be attributed to an unmeasured software delay, but may be a result of changes in laboratory temperature in the few days between the adiabatic expansion experiment and the time when the high resolution DCCT trace was recorded. The missing Trial C high resolution DCCT trace means I am required to assume the Trial C trace is identical to that of Trial A. For this reason, I choose to retain the linear DCCT ramp down model as the primary simulation since this linear ramp relies on DCCT readings (250 ms resolution) that were obtained during repeats of the respective specific trial.

Since the superconducting magnet currents were measured to high accuracy and low (1 s) resolution by DCCTs during the experiment itself, and the steady state currents fluctuate between experimental runs at the $<1\%$ level (see errors on currents in Chapter 2), the experimental uncertainty is expected to be dominated by the timing of the current ramps. As I showed in the previous section, the specifics of the Octupole current decay during FRD have a relatively large impact on the distribution of annihilation times. In addition, I cannot rule out uncertainties in the electrode positions relative to the Mirror Coils, or inaccuracies in our models of the magnetic fields as an additional correction to the observed offset. In combination with the missing data during Octupole FRD of Trial C, a careful experimental campaign would be required to further diagnose the offset. Given the experimental uncertainties, the simulation and experiment agree well, which provides us with confidence in the simulation to reproduce features of the experimental trapped antihydrogen distribution, such as the mean




Figure 4.7: Kolmogorov-Smirnov (KS) p value as a function of temporal offset of experimental annihilation time data with respect to simulated annihilation times for Trial A (red solid curve), Trial B (blue dashed curve) and Trial C (green dotted curve) (a) with a linear Octupole ramp down (b) with a ramp down which follows the current measured by the DCCT during experimental ramp down. Note that in (b) Trials A and C use the DCCT trace obtained during Trial A (since the Trial C trace was not recorded) and Trial B uses the DCCT trace obtained during Trial B. The red region indicates the range of p values (< 0.05) at which the null hypothesis (simulated and experimental annihilation times are drawn from same distribution) is rejected, whereas the green region indicates the range of p values at which the null hypothesis is accepted.

energy.

It is likely that the observed temporal shift is a diagnostic shift i.e. that the magnetic ramp down is not accurately modelled by the simulation, which is backed up by increases in agreement when simulating ramp down which follows the DCCT current. A diagnostic shift does not indicate a difference in energy as determined by the simulation relative to the experiment. Despite this, in Section 4.4 I will estimate an uncertainty on the energy of the distribution of trapped antihydrogen based on the worst case scenario that the observed temporal shift indicates a difference in energy between experiment and simulation. I should also note that the KS results show that all three trials are consistent with the same temporal shift, meaning all three trials would tend to have energy distributions shifted in the same direction by a similar amount in the event the temporal shift is indicative of an energy difference between simulation and experiment.

4.3 Energy of Adiabatically Cooled Antihydrogen

During the adiabatic expansion experiment, we had no means of measuring the energy of the trapped antihydrogen atoms directly and hence we rely on extracting the energy distributions from simulations of the experimental procedure. In this section, I will determine whether the experimental observation that anti-atoms, subject to a gradually expanding magnetic trap, lose energy during the expansion. I will extract the energy of the simulated anti-atoms before and after the expansion to determine the energy lost in the process. I will compare the final energy achieved via adiabatic expansion to the energy of antihydrogen atoms held in the small and large magnetic volumes, to determine whether adiabatic cooling achieves lower energy antihydrogen than a typical antihydrogen trapping sequence. In the presented adiabatic expansion protocol, the magnetic trap expands predominantly in the axial direction. I will determine the extent to which the axial energy of the trapped antihydrogen population is reduced in the expansion, as well as the extent to which the transverse energy is reduced by energy mixing between axial and transverse degrees of freedom (see Section 3.3).

I determine the energy of the simulated anti-atoms at the start, t_i , and end, t_f , of adiabatic expansion Trial A. The axial and transverse energy components (Equation 3.33 and 3.34) can be determined when a simulated anti-atom crosses the magnetic minimum, using the method described in Section 3.3. The total energy of an anti-atom is given by the sum of the axial and transverse energy components when the particle crosses the magnetic minimum. $t_i \sim -23$ s is chosen to be 500 ms after the simulation starts to ensure all particles cross the magnetic minimum before the energy sample is taken, and $t_f \sim -50$ ms is just prior to magnetic ramp down.

In Figure 4.8(a), I plot the total energy of the simulated anti-atoms at the start, $E(t_i)$, and end, $E(t_f)$, of adiabatic expansion Trial A. I only show energies of antihydrogen atoms that

survive until $t = 0$ s i.e. I ignore quasi-trapped particles that are lost before the diagnostic ramp, meaning the total energy decrease observed in Figure 4.8(a) is not a result of preferential loss of high energy anti-atoms. Instead, the energy of the anti-atoms appears to have been reduced via adiabatic cooling. In Figures 4.8(d) and 4.8(g), I show that there is no significant change in total energy between t_i and t_f during control Trials B and C respectively. Note that the magnetic fields are almost completely static during this interval in Trials B and C (see Figure 2.4 for the currents as a function of time).




Figure 4.8: Histograms of simulated antihydrogen energies at t_i (orange) and t_f (blue) for Trial A (a) E , (b) E_{\parallel} and (c) E_{\perp} , Trial B (d) E , (e) E_{\parallel} and (f) E_{\perp} , and Trial C (g) E , (h) E_{\parallel} and (i) E_{\perp} . Vertical dashed orange and blue lines show the trap depth, Γ , at t_i and t_f respectively. Quasi-trapped particles (particles that annihilate before $t = 0$ s) are excluded, but some anti-atoms that have energy greater than the trap depth survive until $t = 0$ s. Simulation parameters are shown in Table 4.1.

In Table 4.3, I present the ensemble-averaged energy at t_i and t_f , $\langle E(t) \rangle$, for the three trials, excluding quasi-trapped particles (simply the mean of histograms shown in Figure 4.8) as well as the corresponding standard error on the mean, σ_{SE} . The simulations show a $37.9 \pm 0.1\%$ decrease in mean total energy over the course of adiabatic expansion Trial A. In contrast, for Trials B and C, changes in $\langle E \rangle$ are less than 1%. The simulations show that the final energy of the adiabatically cooled antihydrogen atoms is on-average 1.616 ± 0.002 times less energetic than Trial B anti-atoms and 1.558 ± 0.002 times less energetic than Trial

C anti-atoms.

Trial, t	$\langle E(t) \rangle \pm 1\sigma_{\text{SE}}(\text{K})$	$\langle E_{\parallel}(t) \rangle \pm 1\sigma_{\text{SE}}(\text{K})$	$\langle E_{\perp}(t) \rangle \pm 1\sigma_{\text{SE}}(\text{K})$
A, t_i	0.3586 ± 0.0003	0.1578 ± 0.0003	0.2007 ± 0.0003
A, t_f	0.2226 ± 0.0002	0.0504 ± 0.0001	0.1722 ± 0.0003
B, t_i	0.3607 ± 0.0003	0.1571 ± 0.0003	0.2036 ± 0.0003
B, t_f	0.3598 ± 0.0003	0.1555 ± 0.0003	0.2043 ± 0.0003
C, t_i	0.3471 ± 0.0003	0.1535 ± 0.0002	0.1936 ± 0.0003
C, t_f	0.3467 ± 0.0003	0.1479 ± 0.0002	0.1987 ± 0.0003

Table 4.3: Ensemble averaged total energy, $\langle E(t) \rangle$, axial energy, $\langle E_{\parallel}(t) \rangle$, and transverse energy, $\langle E_{\perp}(t) \rangle$, for the three experimental trials at the start, t_i , and end, t_f , of adiabatic expansion (or corresponding control hold). $\sigma_{\text{SE}} = \sigma/N_{\text{sim}}$ is the standard error on the mean, where N_{sim} is the number of simulated trajectories (shown in Table 4.1) and σ is the sample standard deviation.

In addition to the total energy, in Figure 4.8 I also show histograms of the axial, E_{\parallel} (Equation 3.33), and transverse, E_{\perp} (Equation 3.34) energy components at t_i and t_f . For Trial A, the axial trap expansion leads mainly to a reduction in $\langle E_{\parallel} \rangle$ as expected, although there is also a smaller reduction in $\langle E_{\perp} \rangle$, resulting from energy mixing between axial and transverse degrees of freedom (see Section 3.3). As will be shown in Chapter 5, the axial trap expansion actually results in a small radial compression, which increases $\langle E_{\perp} \rangle$ adiabatically. However, this effect is counteracted (to a larger extent) by energy mixing, resulting in an overall decrease in $\langle E_{\perp} \rangle$. In Table 4.3 I present corresponding ensemble-averaged axial and transverse energy components at t_i and t_f , which show that the energy loss during adiabatic expansion Trial A can be broken down into a $68.1 \pm 0.2\%$ decrease in $\langle E_{\parallel} \rangle$ and a $14.2 \pm 0.2\%$ decrease in $\langle E_{\perp} \rangle$.

In Chapter 2, I raised the point that the Trial A (or equivalently B) and Trial C experimental protocols involve formation of antihydrogen at slightly different magnetic potentials and in traps with slightly different trap depths. I stated that simulations would be required to show that the initial distributions of trapped antihydrogen are similar in Trials A, B and C; something I have shown in Figure 4.8.

In Figure 4.9 I show the ensemble-averaged energies as a function of time (at 250 ms intervals) prior to magnetic ramp down during Trials A, B and C. A 250 ms interval duration is chosen since it allows sufficient time for every particle to cross the magnetic minimum during an interval, but retains resolution on the ensemble-average. The ensemble average as a function of time makes it possible to identify the effect of individual magnet ramps on the mean energy of the trapped antihydrogen population. For example, during the first 5 s of all

three trials, Solenoid A ramps from full to zero current (currents are graphed in Figure 2.4) which leads to a small amount of axial adiabatic cooling, as shown in zoomed Figures 4.9(d) and 4.9(e). In contrast, the Solenoid A ramp leads to a radial compression and increases $\langle E_{\perp} \rangle$ as shown in zoomed Figure 4.9(f). In addition, the change in $\langle E \rangle$ during Trial C between -20 s and -5 s is a result of Mirror B ramping from 33.0 ± 0.2 A to zero current during the hold; no such ramp occurs for Trial B.




Figure 4.9: Ensemble-averaged (a) total energy, $\langle E \rangle$, (b) axial energy, $\langle E_{\parallel} \rangle$, and (c) transverse energy, $\langle E_{\perp} \rangle$, for simulations of Trial A (red solid), Trial B (blue dashed) and Trial C (green dotted) antihydrogen atoms as a function of time relative to the start of the Octupole ramp down. (d), (e) and (f) are energy-zoomed versions of (a), (b) and (c) respectively. Distribution averages are calculated every 250 ms. The energy of an individual particle is defined as its energy at the last z_{\min} crossing before each 250 ms threshold. Quasi-trapped particles are excluded. Standard error bars have been omitted since they are too small to see. Simulation parameters are shown in Table 4.1.

4.4 Estimating Energy Uncertainty from Temporal Offset

In Section 4.1, I used the Two-Sample Kolmogorov-Smirnov (KS) Test to predict the temporal offset between simulated and experimental annihilation time distributions for the three experimental trials, by offsetting the experimental annihilation times with respect to the simulated annihilation times and calculating the KS p value. Figure 4.7 showed the KS p value as a function of this temporal offset for the primary linear FRD simulation (Figure 4.7(a)) and for the DCCT modified simulation (Figure 4.7(b)), in which the simulated Octupole current follows the high resolution DCCT trace during FRD. For the range of temporal offsets

with KS p value > 0.05 , the difference between the simulated and (offset) experimental data is insufficient to conclude the distributions are different, and hence I bracket the range of temporal offsets for which simulation and data are in agreement.

In Section 4.2.2 I established reasonable agreement between simulated and experimental annihilation time distributions within experimental uncertainty and remarked that I expect the observed temporal offset between experiment and data is a diagnostic shift i.e. that it indicates inaccuracy in our modelling of magnetic ramp down, rather than indicating a shift in energy between simulation and experiment. For these reasons, in the previous section I determined the energy of the anti-atoms in each experimental trial by extracting the energy from the simulation. In this section I will estimate the uncertainty in the simulated ensemble-averaged energies in the worst case scenario: that the observed temporal shift between experimental and simulated annihilation times is indicative of a shift in energy between simulation and experiment.

Trial	$\langle \Gamma(t_{a,\text{sim}}) \rangle \pm 1\sigma_{\text{SE}}(\text{K})$	Range of $\Delta\mathcal{T}$ (ms)
A	0.1078 ± 0.0002	-48 to -148
B	0.2542 ± 0.0002	-28 to -120
C	0.1971 ± 0.0003	16 to -108

Table 4.4: Ensemble-averaged trap depth at simulated annihilation time, $\langle \Gamma(t_{a,\text{sim}}) \rangle$, for Adiabatic Expansion Trial A, Control Trial B and Control Trial C. $\sigma_{\text{SE}} = \sigma/\sqrt{N_{\text{sim}}}$ is the standard error on the mean, where σ is the sample standard deviation and N_{sim} is the number of simulated antihydrogen atoms (see Table 4.1). I also show the range of temporal offsets, $\Delta\mathcal{T}$, of the experimental relative to simulated annihilation times for which the KS p value is > 0.05 (green region in Figure 4.7), which indicates reasonable agreement between simulated and experimental annihilation time distributions.

Recall that in Chapter 2, I made a naive estimation of the energy distribution of the trapped antihydrogen atoms in the three experimental trials by calculating the trap depth at the experimental annihilation time, $\Gamma(t_a)$. To estimate the energy uncertainty corresponding to the temporal shift between simulation and experiment, I will calculate the equivalent parameter for the simulated annihilation times, $\Gamma(t_{a,\text{sim}})$, then shift the simulated annihilation times such that the simulation and experiment are in agreement, before recalculating $\Gamma(t_{a,\text{sim}})$. The difference between $\Gamma(t_{a,\text{sim}})$ with zero offset and with an offset that leads to agreement between simulation and experiment is our estimate of the energy uncertainty. Since the KS analysis (Figure 4.7) provides a range of possible temporal offsets for each experimental trial, I will calculate a range of energy shifts which correspond to agreement between simulation and data.




Figure 4.10: (a) Simulated ensemble-averaged final energy, $\langle E(t_f) \rangle$ for Adiabatic Expansion Trial A (red line), Control Trial B (blue line) and Control Trial C (green line). Standard error bars are not visible on this scale. (b) and (c) Range of energies, $\langle E(t_f) \rangle + \Delta \langle E(t_f) \rangle$ (see Equation 4.5), for which simulated and experimental annihilation time distributions have KS p values > 0.05 , are shown as coloured regions (red, blue and green for Trials A, B and C respectively) for (b) the primary linear FRD simulation and (c) the DCCT modified simulation.

In Table 4.4, I show the ensemble-averaged trap depth at each simulated annihilation time, $\langle \Gamma(t_{a,\text{sim}}) \rangle$, with no applied temporal offset for the three experimental trials. In addition, I show the range of offsets, $\Delta \mathcal{T}$, that I will apply to the simulated annihilation times, before calculating $\langle \Gamma(t_{a,\text{sim}} - \Delta \mathcal{T}) \rangle$. I will calculate the energy uncertainty corresponding to the range of valid offsets between experimental and simulated annihilation times for the primary (linear FRD) simulation (Figure 4.7(a)) and for the DCCT modified simulation (Figure 4.7(b)). In each case, I approximate the uncertainty in the ensemble-averaged final simulated energy, $\langle E(t_f) \rangle$, as

$$\Delta \langle E(t_f) \rangle = \langle \Gamma(t_{a,\text{sim}}) \rangle - \langle \Gamma(t_{a,\text{sim}} - \Delta \mathcal{T}) \rangle, \quad (4.5)$$

which can be interpreted as the difference in trap depth corresponding to the temporal shift between experimental data and simulation. In Figure 4.10(a) I show the simulated ensemble-averaged final energy, $\langle E(t_f) \rangle$, for the three experimental trials, where the standard error is not visible on the scale of the figure. In Figure 4.10, I show the range of energies corresponding to $\langle E(t_f) \rangle + \Delta \langle E(t_f) \rangle$, which is an estimate of the range of final simulated energies with

annihilation time distributions consistent with the experimental annihilation times, for the primary linear FRD simulation (Figure 4.10(b)) and the DCCT modified simulation (Figure 4.10(c)).

Since the KS test predicts all three trials generally have offsets in the same direction, the range of energies corresponding to agreement between simulated and experimental annihilation times tends to shift in the same direction for all trials. In Table 4.5 I plot the range of possible ratios of ensemble-averaged energies between control Trials B and C and adiabatic expansion Trial A. These ratios allow lower limits on the final energy of the control trials relative to the adiabatic expansion trial to be placed, given that the temporal offset between simulation and experiment can take any value with KS p value > 0.05 in all three trials, and on the assumption that the observed temporal offset between simulation and experiment is indicative of a shift in energy between simulation and experiment.

Recall from the previous section that the simulations showed that the Trial A ensemble-averaged final energy was 1.616 ± 0.002 times lower than Trial B and 1.558 ± 0.002 times lower than Trial C. These values lie close to the lower limits of the ranges of possible ratios stated in Table 4.5. This is because the range of allowed offsets is different for the three trials and therefore, when the offsets are applied, they tend to estimate larger differences between the final energy in the three trials.

Trial	Linear FRD	DCCT Modified
Ratio	$\langle E(t_f) \rangle + \Delta \langle E(t_f) \rangle (K)$	$\langle E(t_f) \rangle + \Delta \langle E(t_f) \rangle (K)$
B/A	1.47–1.88	1.47–1.83
C/A	1.47–1.92	1.43–1.80

Table 4.5: Range of ratios of ensemble-averaged final energy in Control Trial B and Control Trial C to Adiabatic Expansion Trial A, based on an energy correction applied using the method described in the text, which approximates the energy shift which leads to agreement between simulated and experimental annihilation times.

The result shows that, even in the worst case scenario that the temporal offsets between simulation and data indicate a difference in energy distribution, there is still a significant difference between the energy of the adiabatically cooled population and that of the control samples. This is the case whether the Octupole current ramp down is modelled as a perfectly linear ramp or whether it is modelled on the high resolution current trace. This analysis allows us to place a lower bound on the ratio of the final mean energy in Trial B to Trial A of 1.47 and in Trial C to Trial A of 1.43.

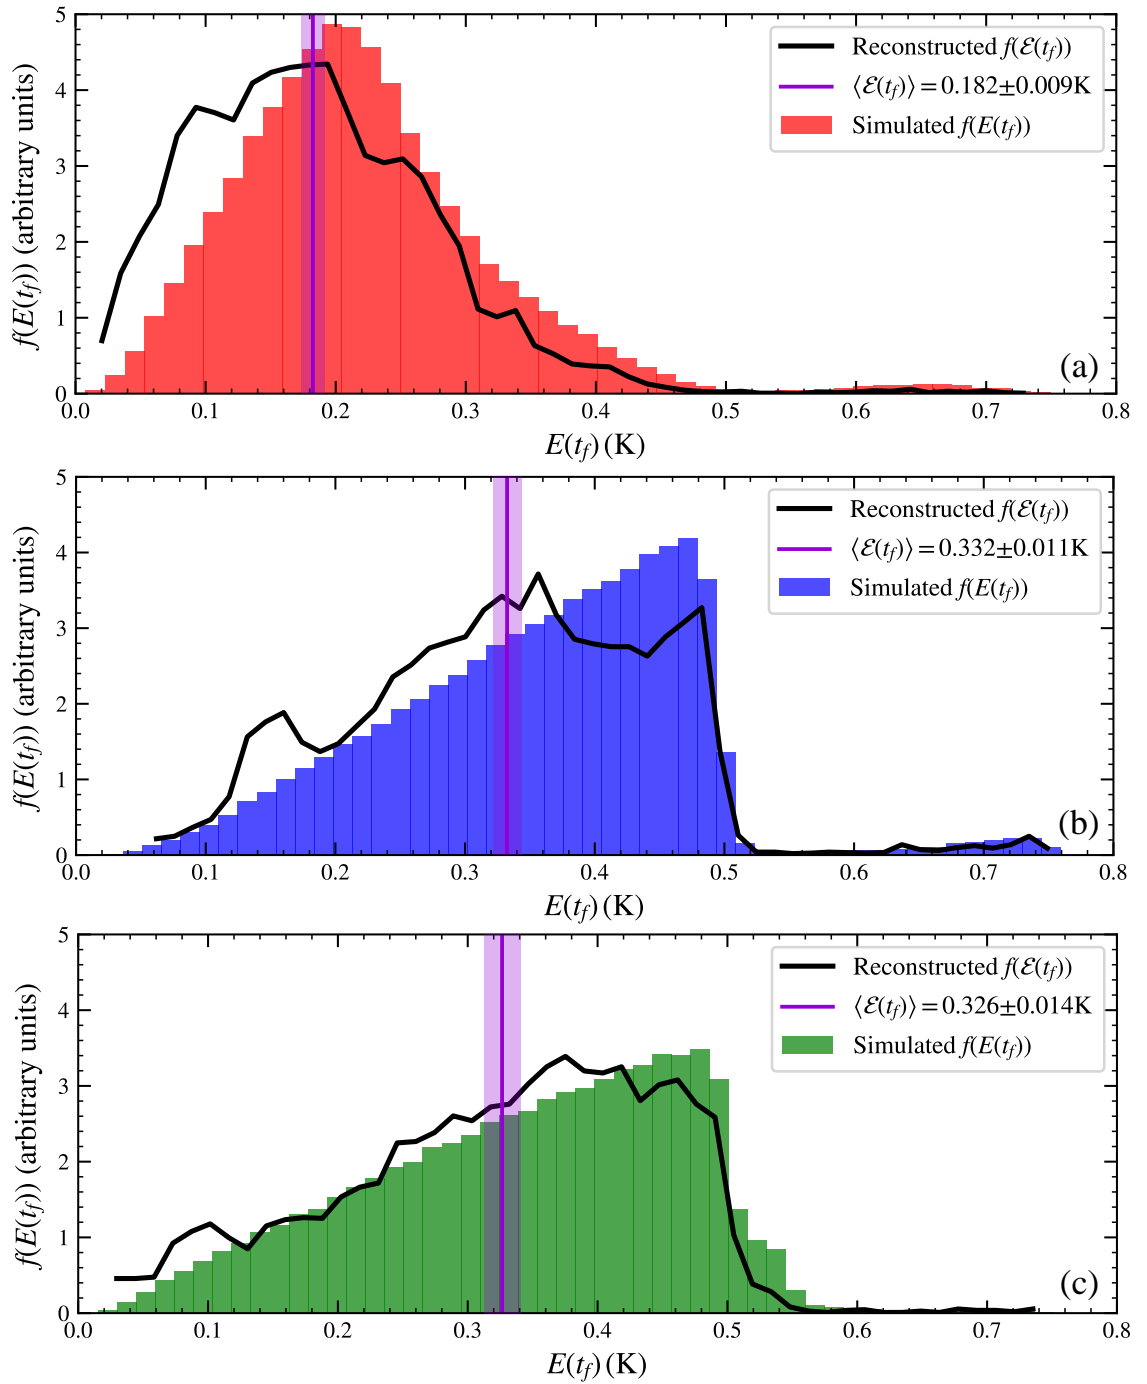


Figure 4.11: Antihydrogen reconstructed energy distribution function, $f(\mathcal{E}(t_f))$ (black solid curve), as a function of anti-atom energy, $E(t_f)$, where the energy is reconstructed just prior to magnetic ramp down (at t_f) and is shown for (a) Adiabatic Expansion Trial A, (b) Control Trial B and (c) Control Trial C. The simulated energy distribution function, $f(E(t_f))$, is also shown for Trial A (red histogram), Trial B (blue histogram) and Trial C (green histogram). $f(\mathcal{E}(t_f))$ was determined using the inverse method described in [1] with a band-width of 0.01 s and sampling 100 samples per band. The average reconstructed energy, $\langle \mathcal{E}(t_f) \rangle \pm 1\sigma_{SE}$ is indicated by the vertical purple line, where $\sigma_{SE} = \sigma/\sqrt{N}$ is the standard error on the mean, σ is the sample standard deviation and N is the total number of experimental antihydrogen counts during diagnostic ramp down (see Table 2.1).

4.5 Reconstructing Energy Distribution from Annihilation-Time Distribution

In Section 4.3, I presented the energy of simulated antihydrogen trajectories at t_f (~ 50 ms before the onset of Octupole FRD). In combination with the simulated annihilation times, I have effectively mapped energy prior to ramp down, $E(t_f)$, to annihilation time, t_a . In this section I present results of an established technique which uses the experimental annihilation times in combination with the simulated results to infer the energy of the experimental antihydrogen distribution. I intend to determine whether the experimental annihilation events form a cooled population based on the simulated map of energy to annihilation time.

Trial	$(\langle \mathcal{E}(t_f) \rangle \pm 1\sigma_{SE})(\text{K})$
A	0.183 ± 0.009
B	0.331 ± 0.011
C	0.326 ± 0.014

Table 4.6: Reconstructed energy just prior to magnetic ramp down, $\langle \mathcal{E}(t_f) \rangle$, from experimental annihilation time and simulated map between annihilation time and energy just prior to magnetic ramp down, $E(t_f)$, using the inverse method described in [1] with a band-width of 0.01 s and sampling 100 samples per band. $\sigma_{SE} = \sigma/\sqrt{N}$ is the standard error on the mean, σ is the sample standard deviation and N is the total number of experimental antihydrogen counts during diagnostic ramp down (see Table 2.1).

To calculate this measure of the experimental energy distribution function after adiabatic expansion, $f(\mathcal{E}(t_f))$, I employ the ‘inverse’ method described in [1]. Annihilations occur at time $t_a > t_f$ during diagnostic ramp down. I define $P(\mathcal{E}(t_f) | t_a)$ as the probability of annihilation at time t_a of anti-atoms with specific energy $\mathcal{E}(t_f)$ at t_f . Based on this definition, the energy distribution function is given by

$$f(\mathcal{E}(t_f)) = \int_0^\infty dt_a P(\mathcal{E}(t_f) | t_a) f(t_a), \quad (4.6)$$

where $f(t_a)$ is the overall probability distribution function for the anti-atoms to annihilate at time t_a . For each experimental annihilation event, I deduce the simulated annihilation events within a temporal bandwidth, t_{band} , of the experimental annihilation time. For those within t_{band} , I randomly select a fixed number, N_{samples} , of annihilations; which effectively samples $P(\mathcal{E}(t_f) | t_a)$. Taking the energy of each of the randomly selected trajectories for each observed annihilation event effectively integrates over t as properly weighed by $f(t_a)$. In Figure 4.11 I plot a histogram of these energies for $t_{\text{band}} = 0.01$ s and $N_{\text{samples}} = 100$ for each experimental trial. In Table 4.6 I indicate the ensemble-averaged energy, $\langle \mathcal{E}(t_f) \rangle$ as

predicted by the inverse method, which predicts the energy of the control populations (Trial B and C) are 1.81 ± 0.10 and 1.78 ± 0.12 times that of the adiabatically cooled (Trial A) population respectively. I have confirmed that using simulated annihilation data in place of the experimental data reproduces the simulated energy distributions as expected.

4.6 Final Energy Method Comparison

In Table 4.7, I summarise the final energy of Trials A, B and C as predicted by the different methods discussed in this thesis. The first method is the naive estimation of the final total energy which I presented in Chapter 2. The method simply involves calculating the trap depth (Figure 2.7) at each experimental annihilation time, $\Gamma(t_a)$. Assuming antihydrogen will escape the trap when the trap depth becomes equal to its total energy (naive assumption), the ensemble average of this quantity, $\langle\Gamma(t_a)\rangle$, provides a simulation-free estimate of the mean total energy of the distribution of trapped antihydrogen at the time of release from the magnetic trap. The naive method predicts that the mean total energy at the end of adiabatic expansion Trial A was 2.75 ± 0.72 times lower than that of Trial B and 2.14 ± 0.49 times lower than that of Trial C.

Trial	$\langle\Gamma(t_a)\rangle \pm 1\sigma_{\text{SE}}(\text{K})$	$\langle E(t_f)\rangle \pm 1\sigma_{\text{SE}}(\text{K})$	$\langle\mathcal{E}(t_f)\rangle \pm 1\sigma_{\text{SE}}(\text{K})$
A	0.081 ± 0.007	0.2226 ± 0.0002	0.183 ± 0.009
B	0.223 ± 0.009	0.3598 ± 0.0003	0.331 ± 0.011
C	0.173 ± 0.011	0.3467 ± 0.0003	0.326 ± 0.014

Table 4.7: Table comparing final energy predicted by naive method, $\langle\Gamma(t_a)\rangle$, final energy extracted from the simulation, $\langle E(t_f)\rangle$, and reconstructed energy, $\langle\mathcal{E}(t_f)\rangle$. The naive method involves calculating the trap depth at each experimental annihilation time, $\Gamma(t_a)$, and taking the ensemble average (denoted by an angled bracket). The energy extracted from the simulation involves determining the total energy (sum of Equation 3.33 and Equation 3.34) of an individual particle as it crosses the magnetic minimum. $\langle\mathcal{E}(t_f)\rangle$ is determined using the method described in Section 4.5. $\sigma_{\text{SE}} = \sigma/\sqrt{N}$ is the standard error on the mean, σ is the sample standard deviation and N is the total number of experimental antihydrogen counts during diagnostic ramp down (see Table 2.1).

Also in Table 4.7, I show the final mean total energy, $\langle E(t_f)\rangle$, as extracted from detailed Monte Carlo simulations of the experimental procedure, and presented in Section 4.3. The total energy is determined at t_f (~ 50 ms before diagnostic ramp down) and shows that the mean total energy at the end of adiabatic expansion Trial A was 1.616 ± 0.002 times less energetic than Trial B anti-atoms and 1.558 ± 0.002 times less energetic than Trial C anti-atoms.

The third energy prediction compared in Table 4.7 is the reconstructed energy, $\langle \mathcal{E}(t_f) \rangle$, as presented in Section 4.5. The method uses the simulated map of total energy prior to diagnostic ramp down to annihilation time, in combination with the experimental annihilation time data, to reconstruct the energy of the experimental distribution of trapped antihydrogen and is described in detail in [1]. The reconstructed Trial A final mean energy is 1.81 ± 0.10 times smaller than control Trial B and 1.78 ± 0.12 times smaller than control Trial C.

One of the reasons that $\langle \Gamma(t_a) \rangle$ is smaller than $\langle E(t_f) \rangle$ and $\langle \mathcal{E}(t_f) \rangle$ is that it relies on the naive assumption that antihydrogen atoms will escape the trap as soon as the trap depth becomes equal to their total energy during Octupole FRD. As I will show in Section 4.8, the annihilation time during Octupole FRD depends on an anti-atom's specific components of axial and transverse energy, rather than the sum of the components. As a result, antihydrogen atoms can remain trapped during Octupole FRD with total energy greater than the trap depth. In addition, the method estimates the energy of the anti-atom at the time of its release from the trap, rather than just prior to diagnostic ramp down. Therefore, the method estimates the energy of the distribution after it has been radially adiabatically cooled during the dump. I should note that, since the transverse energy distribution of the Trial A, B and C anti-atoms is similar (see Figure 4.8), the extent of radial cooling is expected to be similar between experimental trials.

In addition, the observed temporal shift between experiment and simulation contributes to $\langle \Gamma(t_a) \rangle$ being smaller than $\langle E(t_f) \rangle$ since, as I showed in Section 4.1, the experimental annihilation times tend to occur 10s of milliseconds later (when the trap depth is lower) than the annihilation times in the primary (linear FRD) simulation. Similarly, since the reconstruction method used to generate $\langle \mathcal{E}(t_f) \rangle$ uses the experimental annihilation times, the reconstructed energy is also lower than $\langle E(t_f) \rangle$ on-average. Since this shift has more of an effect in Trial A than in control Trials B or C, the reconstruction method predicts a greater (than the simulation) difference in mean energy of the adiabatically cooled anti-atoms compared to the control samples.

4.7 Convergence Studies

The fixed timestep, dt , used in the simulation (see Section 3.4 for a description of the symplectic stepper) must be small enough that the simulation results do not differ significantly if the timestep is reduced further, but it must be large enough that I can run simulations on a feasible timescale. Simulations presented in this thesis use a timestep, $dt = 3.5 \times 10^{-6}$ s (see Table 4.1). In this section I will describe the convergence studies that led to this choice.

For convergence studies, I simulate Trial A and use two quantities to test convergence: the ensemble-averaged annihilation time, $\langle t_a \rangle$, and the ensemble-averaged total energy at the end of adiabatic expansion, $\langle E(t_f) \rangle$. I vary the simulation timestep from 3.5×10^{-7} s

to 1.0×10^{-4} s, simulating 200,000 trajectories with identical initial velocities and positions for each timestep, and take the ensemble-average over the trajectories that survive until t_f (which may differ in each case). In this study, the convergence parameters, $\langle t_a \rangle$ and $\langle E(t_f) \rangle$, are both functions of the simulation timestep and hence I denote them as $f(dt)$. For a given convergence parameter, I define the measure of convergence as

$$\Delta f(dt) = \left| \frac{f(dt) - f(3.5 \times 10^{-7})}{f(3.5 \times 10^{-7})} \right|. \quad (4.7)$$

In other words, I calculate the difference in $\langle t_a \rangle$ and $\langle E(t_f) \rangle$ at a given timestep relative to its value at the smallest timestep in the range (3.5×10^{-7} s).

In Figure 4.12 I show $\Delta f(dt)$ for convergence parameters, $\langle t_a \rangle$ and $\langle E(t_f) \rangle$, as a function of simulation timestep. At $dt = 3.5 \times 10^{-5}$ s, results converge at the $\sim 1\%$ level. I choose a value of $dt = 3.5 \times 10^{-6}$ s, since both $\langle t_a \rangle$ and $\langle E(t_f) \rangle$ converge better than 0.1%. At this level, it takes ~ 1 day to simulate 200,000 trajectories running 100 simulations in parallel (2000 particles in series).




Figure 4.12: Absolute differences (%) relative to $dt = 3.5 \times 10^{-7}$ s, $\Delta \langle t_a \rangle$ (solid magenta), and $\Delta \langle E(t_f) \rangle$ (solid cyan), as a function of simulation timestep, dt . Each black cross represents a simulation of 200,000 Trial A trajectories with fixed timestep, dt .

4.8 Octupole FRD as an Adiabatic Expansion Diagnostic

As explained in Section 2, during experimental Trials A, B and C, after the axial expansion of the volume of the magnetic trap or corresponding control hold, the Octupole magnet was ramped to 0 A in ~ 1.5 s to release the antihydrogen atoms from magnetic confinement. In typical antihydrogen trapping experiments, magnetic release involves ramping both the




Figure 4.13: Energy components for mix (red) and no-mix (black) simulated antihydrogen trajectories, where trajectories are categorised based on their normalised axial energy 50 ms prior to the onset of Octupole FRD (mix have $\epsilon_{\parallel} \geq 0.1$, and no-mix have $\epsilon_{\parallel} < 0.1$). Axial, E_{\parallel} (a), (c) and (e), and transverse E_{\perp} (b), (d) and (f), energy components for Trials A, B and C respectively.

Octupole and Mirrors A and E to 0 A in the same period of time. If we employed the typical FRD protocol here, there would be significant axial adiabatic cooling during FRD itself on account of the Mirror Coils ramping down. This adiabatic cooling would be more significant in Trial B than in Trials A and C since Trial B ramps down from the tight trap between Mirrors A and C, whereas Trials A and C ramp down from confinement between Mirrors A and E. Therefore, to avoid this obvious energy-asymmetry between trials during the dump, the chosen technique was to release antihydrogen atoms from the trap by ramping down only the Octupole. However, since antihydrogen is lost radially during Octupole FRD and since mixing is not expected to equilibrate energies on the timescale of the dump, Octupole FRD raises the question: do the annihilation time distributions really reflect the axial energy distribution of the trapped antihydrogen atoms? Or alternatively, is the difference in annihilation times between adiabatic expansion Trial A and the control trials caused only by radial energy reduction as a result of energy mixing during the axial adiabatic cooling?

I will investigate this question by searching for a link between axial energy and annihilation time. This will be achieved by splitting the simulated particles into mix and no-mix categories (as defined in Section 3.3), analysing their components of axial and transverse energy, as well as the time they annihilate and the axial location of annihilation. I will determine parameters of the specific magnetic field configuration that I will use to hypothesise an explanation for the observed annihilation locations and times of both categories of particle. I then test this hypothesis by using the axial and transverse energy components of the simulated trajectories to predict axial annihilation locations and annihilation times based upon the hypothesis. Finally, I will interpret the link between axial energy and annihilation time during Octupole FRD from the hypothesis, backed by its ability to reproduce features in the simulated data.

For this discussion, I define mix particles as having normalised axial energy, $\epsilon_{\parallel} > 0.1$ (Section 3.3), and the rest of the trajectories are considered to be in the no-mix category. I sample ϵ_{\parallel} and categorise the simulated particles at t_f (50 ms before the dump). In Figure 4.13, I show independent distributions of E_{\parallel} and E_{\perp} for those in the mix (red) and no-mix (black) categories, where mix and no-mix trajectories have significantly different distributions of E_{\parallel} and E_{\perp} . Figure 4.14 shows the simulated axial annihilation locations, z , ((a) Trial A, (e) Trial B and (i) Trial C) and simulated annihilation times, t , ((b) Trial A, (f) Trial B and (j) Trial C) for the two categories; they annihilate at significantly different locations and times. In particular, mix particles tend not to annihilate at axial locations close to the trap center, whereas no-mix particles do. Also, no-mix particles tend to annihilate earlier in time than mix particles.

The hypothesis for the mechanism of antihydrogen loss during Octupole FRD relies on three principal assumptions:

A1. The FRD timescale (~ 1.5 s) is shorter than the timescale of energy mixing between

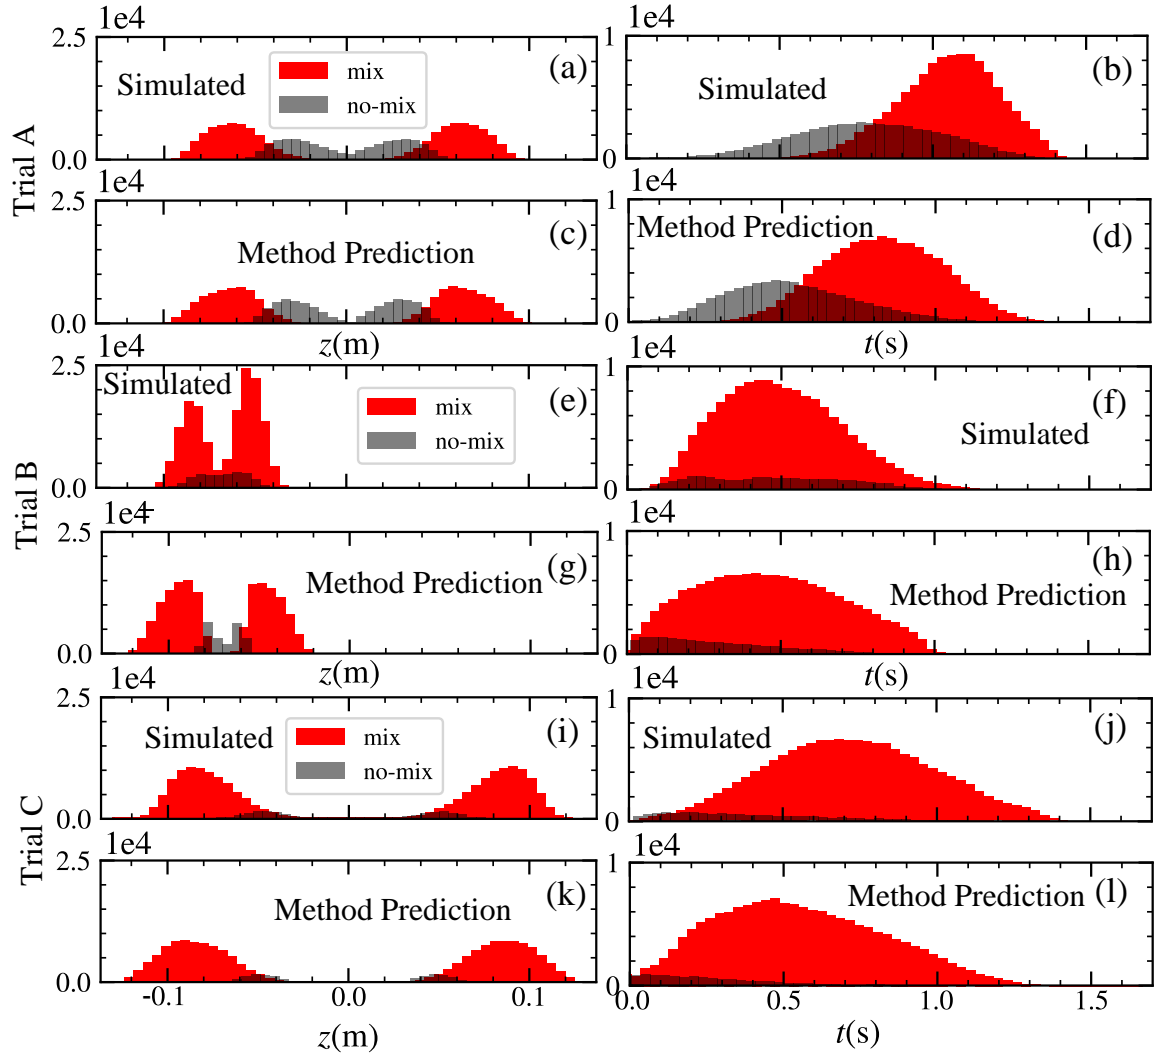


Figure 4.14: Axial annihilation location, z , (left) and annihilation time relative to the start of Octupole FRD, t , (right) for simulated Trial A [(a) and (b)], simulated Trial B [(e) and (f)] and simulated Trial C [(i) and (j)] respectively. Shown below each simulated z histogram is the predicted (labelled “Method Prediction”) z annihilation distribution based on E_{\parallel} of the simulated trajectories 50 ms before the start of Octupole FRD and $U_z(z)$ (Figure 4.15) for Trial A (c), Trial B (g) and Trial C (k). Shown below each t histogram is the predicted (labelled “Method Prediction”) t annihilation distribution based on E_{\perp} of the simulated trajectories 50 ms before the start of Octupole FRD and $\Gamma_r(\phi, z)$ (Figure 4.16) for Trial A (d), Trial B (h) and Trial C (l). ‘Mix’ and ‘no-mix’ categories are shown as red and black histograms respectively.




Figure 4.15: On-axis axial potential, $U_z(z)$, as a function of axial position, z , during Octupole FRD for Trial A (a), Trial B (b) and Trial C (c).

axial and transverse dimensions but much longer than the axial (~ 10 ms) and radial (~ 1 ms) period of oscillation of trapped antihydrogen. Note that Figure 3.10 estimates the mixing timescale in traps comparable to those during Octupole FRD for Trial A and C (standard trap, red curve) and for Trial B (tight trap, brown curve) but that this may differ on account of the antihydrogen energy distribution and that the Octupole field-strength varies during FRD.

- A2.** An antihydrogen atom quickly (compared to the FRD timescale) samples regions of the trap that can be reached, given its specific components of axial and transverse energy. As a result, an antihydrogen atom will escape the magnetic trap when a path to the trap wall becomes available, that can be reached given an antihydrogen atom's specific axial and transverse energy components.
- A3.** I neglect radial adiabatic cooling during Octupole FRD. Quantifying the severity of this assumption would require developing an adiabatic model similar to those described in Section 5, which will depend on the radial energy of the particle. Since I currently have no means of extracting the radial energy of the simulated anti-atom directly, I do not quantify this effect here.

Our model for antihydrogen loss during Octupole FRD requires us to define three quantities: the on-axis potential, the radial potential and the radial trap depth. For the following discussion I will use cylindrical coordinates (r, ϕ, z) . The on-axis potential, U_z , is defined as the potential difference (at $r = 0$) between the magnetic field at an axial location z and the axial location of the magnetic field minimum, z_{\min} i.e.,

$$U_z(z) = U(r = 0, z) - U(r = 0, z_{\min}). \quad (4.8)$$

Figure 4.15 shows U_z as a function of z for Trials A, B and C. An important feature of $U_z(z)$ is that it is exactly constant on-axis (and largely constant off-axis) during Octupole FRD. The reason for this is that, at $r = 0$, the vector field of the Octupole exactly cancels and so its current has no effect on the axial potential. I define the radial potential, $U_r(r, \phi, z, t)$ as

the potential difference between the magnetic field at radius r and at $r = 0$ i.e.

$$U_r(r, \phi, z, t) = U(r, \phi, z, t) - U(r = 0, \phi, z, t), \quad (4.9)$$

and radial trap depth, $\Gamma_r(\phi, z, t)$, is defined as the potential difference, at z , between the magnetic field at the radius of the electrode wall, r_{trap} , and the magnetic field at $r = 0$ i.e.,

$$\Gamma_r(\phi, z, t) = U_r(r = r_{\text{trap}}, \phi, z, t), \quad (4.10)$$

$\Gamma_r(\phi, z, t)$ is shown in Figure 4.16. $\Gamma_r(\phi, z, t)$ can be interpreted as the radial energy required for a particle to escape the trap at a given z . In Figure 4.16, $\Gamma_r(\phi, z, t)$ is shown for Trial A (and equivalently Trial C) at different times during Octupole FRD. The most important feature of $\Gamma_r(\phi, z, t)$ for this discussion is the appearance of minima (blue regions) at relatively high z at certain azimuths, where U_z is relatively high. The minima are caused by the vector sum of the radial magnetic field of the confining Mirror Coils and the Octupole field. This particular field quality is also responsible for mixing axial and transverse energies as discussed in Section 3.3.

We will use the energy components, E_{\parallel} and E_{\perp} , sampled just before the dump (Figure 4.13) to determine the region of the trap accessible to a given particle. When this region intersects with the trap wall, I will assume the particle annihilates. Since I am using the energy components before the dump, I am invoking assumption **A1** (see above list).

At a given time, an individual trajectory can climb to a maximum value of axial potential when $U_z(z) = E_{\parallel}$, which sets the axial region accessible to the particle. In addition, the trajectory can climb to a maximum value of radial potential when $U_r(r, \phi, z) = E_r$, where E_r is the maximum radial kinetic energy of the particle. Therefore, a particle can access regions of the trap where both $U_z(z) < E_{\parallel}$ and $U_r(r, \phi, z, t) < E_r$ are true.

Unfortunately, as E_{\parallel} and E_{\perp} are determined at crossings of the axial minimum, z_{min} , particle crossings occurring at $r \neq 0$ cross with some transverse potential energy, and hence the technique does not allow us to determine E_r and E_{ϕ} independently. Therefore, I will make a further assumption that the transverse energy is entirely radial. As a result, I hypothesise that the particle can access regions of the trap where both $U_z(z) < E_{\parallel}$ and $U_r(r, \phi, z, t) < E_{\perp}$ are true.

Given that this is the accessible region of the trap, I need to determine the time and axial location of the intersection of this region with the trap wall during Octupole FRD. I know that the radial potential that must be overcome to reach the trap wall is the radial trap depth, $\Gamma_r(\phi, z, t)$, and I have seen that minima in $\Gamma_r(\phi, z, t)$ exist at high z (Figure 4.16). As time increases during Octupole FRD, the magnitude of $\Gamma_r(\phi, z, t)$ decreases but the minima remain at high z . As a result, equipotentials of $\Gamma_r(\phi, z, t)$ gradually move inwards (to lower z) over the course of Octupole FRD. This means that escape routes to the trap wall gradually become available at lower z over the course of Octupole FRD.




Figure 4.16: Radial trap depth $\Gamma_r(\phi, z, t)$ (K), as function of axial position, z , and azimuthal angle, ϕ , at different times (relative to start of Octupole FRD) (a) 0.0 s, (b) 0.5 s and (c) 1.0 s during Adiabatic Expansion Trial A, and equivalently during Control Trial C.

Let us take an example particle that is trapped at some early time, t_1 , during Octupole FRD. Given that the particle is trapped, it samples only regions of the trap where both $U_z(z) < E_{\parallel}$ and $U_r(\phi, z, t) < E_{\perp}$ are true. As time goes on, the Octupole current decreases. At a later time t_2 , when the particle is still trapped, there may be regions of the trap where $\Gamma_r(\phi, z, t) < E_{\perp}$, but since these regions first appear at high z (where U_z is also high), the particle does not have sufficient E_{\parallel} to reach them. At a later time t_3 , the trap depth reduces such that there exists a region of the trap where both $U_z(z) < E_{\parallel}$ and $\Gamma_r(\phi, z, t) = E_{\perp}$ are true, and the particle annihilates.

We know that $U_z(z)$ is constant on-axis during Octupole FRD, so the axial region of the trap accessible to the particle is fixed as the Octupole current decreases. Since minima in $\Gamma_r(\phi, z)$ exist at high z and equipotentials in $\Gamma_r(\phi, z)$ move inwards as the Octupole current decreases, the first intersection of the accessible region with the trap wall will occur at the maximum z within the accessible axial region (we have invoked assumption **A2**). Therefore, I have found our predicted axial annihilation location, z_{predict} , where

$$U_z(z_{\text{predict}}) = E_{\parallel}. \quad (4.11)$$

Given that the particle annihilates at z_{predict} , I predict that annihilation will occur at a time when the trap depth at z_{predict} becomes equal to E_{\perp} . In other words,

$$\Gamma_r(\phi, z_{\text{predict}}, t_{\text{predict}}) = E_{\perp}. \quad (4.12)$$

The minimum $\Gamma_r(\phi, z_{\text{predict}})$ in ϕ (on account of assumption **A2**), gives us the predicted time of annihilation, t_{predict} .

Based on this hypothesis, I use the simulated energy components, E_{\parallel} and E_{\perp} , of the mix category at t_f to predict the axial annihilation location and the annihilation time of the antihydrogen atoms. The result is compared with the simulated annihilation times and axial annihilation locations in Figure 4.14. Because the mix and no-mix categories have significantly different distributions of E_{\parallel} and E_{\perp} , the method predicts the difference in axial annihilation

locations and annihilation times between the mix and no mix categories that I previously determined from the simulation.

There is a temporal offset between the model prediction for the annihilation times and the simulated annihilation times for all three trials. This is likely a result of the fact that I have neglected the radial adiabatic expansion during Octupole FRD (assumption **A3**) and the fact that I have assumed $\Gamma_r(\phi, z_{\text{predict}}, t_{\text{predict}}) = E_{\perp}$ rather than $\Gamma_r(\phi, z_{\text{predict}}, t_{\text{predict}}) = E_r$. Both effects, if included, would reduce the radial extent of the accessible region, making anti-atoms annihilate later in time. A goal of future study will be to determine the radial energy components of the trajectory and then to apply an adiabatic model of the radial cooling similar to those that will be described in Chapter 5. The goal of this study will be to remove both effects and to redetermine the accuracy of the method in predicting the simulated annihilation-time distribution.

By comparing t_{predict} in Trial A to those of Trial B and Trial C (t_{predict} is on average later in Trial A than Trial B or Trial C), I can deduce an important conclusion: based on the energy distributions predicted by the simulation and details of the magnetic fields, I can predict a difference in annihilation time distribution between adiabatically cooled populations and control holds. In other words, I have found an effect which couples the axial energy of the anti-atoms to the annihilation time during Octupole FRD and I have sufficiently shown that Octupole FRD is a good diagnostic of both the axial and transverse energy components of the trapped antihydrogen atoms.

The link between E_{\parallel} and axial annihilation location, and E_{\perp} and annihilation time could actually be used in reverse as a diagnostic of the axial and transverse energy components of an experimental distribution of antihydrogen. During Octupole FRD, the annihilation location would lead us directly to E_{\parallel} , and then in combination with the annihilation time it would be possible to predict E_r . This could provide a simulation-free route to accurately determining the energy components of an experimental distribution of antihydrogen from the observable annihilation times and locations. ALPHA has recently used techniques involving lineshape and time-of-flight measurements to determine (experimentally) the radial and axial energy components of trapped antihydrogen atoms [7]. A future experiment may involve comparison of the results of the two methods. One potential benefit of the Octupole FRD method is its simplicity, since it does not require laser interactions.

4.8.1 Known Magnetic Field Inaccuracy as a cause of Temporal Shifts between Experiment and Simulation

In Section 4.1, I used the Two-Sample Kolmogorov-Smirnov (KS) test to estimate a temporal offset somewhere between approximately -50 ms and -100 ms of the experimental annihilation time data relative to the primary (linear FRD) simulation in all three

experimental trials. I raised the possibility that this temporal offset is a diagnostic shift i.e. that it represents an inaccuracy of the modelling of the magnetic ramp down in the simulation. A possible source of inaccuracy is in our models of the magnetic fields that are described in Section 3.4.2.

The Truncated Vector Approximation Model (TAVP) that is used to determine the field from the two current loops that form our Mirror Coil model, as discussed in Section 3.4, is known to be inaccurate off-axis [95]. However, within the electrode radius ($r \approx 2$ cm), the model is expected to be inaccurate only at the $\sim 1\%$ level [95]. In the context of particle loss during Octupole FRD in this section, I found that, especially for the 2/3 of trajectories in the mix category, the annihilation time is highly sensitive to the radial field of the Mirror Coils at the electrode wall. This raises the possibility that the inaccuracy of the magnetic field model is leading to the offset between experimental and simulated annihilation times as described in Section 4.1. Although the McDonald model presented in [95] was shown to be faster and more accurate than TAVP, when implemented in our simulation it is much slower than TAVP for reasons unknown to us. Therefore, at the time of writing this thesis, I was unable to re-simulate the adiabatic expansion trials to directly extract the correction to the annihilation time distribution that a more accurate model may provide.

However, I can use techniques developed in Section 4.8 to predict the temporal offset between simulations of Trial A with TAVP and with McDonald. I re-predict the axial annihilation location and annihilation time of the trajectories using the simulated E_{\parallel} and E_{\perp} 50 ms before Octupole FRD (as discussed in Section 4.8), but this time using the 3rd order McDonald model for a circular current loop as described in [95]. I observed no difference on the required order (~ 10 ms) between the model prediction of the annihilation times with the TAVP model and the 3rd order McDonald model, and hence I conclude that this uncertainty in the magnetic field model is unlikely to be the source of the observed temporal shift between experiment and simulation. However, I cannot rule out other differences between the simulated magnetic fields and the physical magnetic fields as a source of the observed temporal offset between simulated and experimental annihilation times.

Adiabatic Cooling Models

Adiabatic cooling relies on the very good approximation that an adiabatic invariant is conserved during the slow expansion of the trap volume. In Section 3.2 I derived the adiabatic invariant for a particle oscillating between the walls of a box that expands axially, slowly compared to the axial period of the particle's oscillation. In the reality of antihydrogen motion in a magnetic trap like ALPHA's, the adiabatic invariant is not as simple to derive since the length of the magnetic trap is not a constant function of radius, the volume sampled by the particle is a function of the particle's energy, and the axial and transverse components of the particle's energy can randomly mix to some extent during the expansion. In this section I will show that the amount of mixing between axial and transverse energies changes the power law of an adiabatic model. Given that there exist two categories of antihydrogen trajectory that exhibit different degrees of energy mixing, a single adiabatic invariant cannot be accurately applied to the total ensemble of trajectories. I will derive two adiabatic models that do not rely on particle tracking simulations, by making limiting assumptions about the degree of energy mixing of the whole ensemble. The adiabatic models are expected to bound the energy decrease during adiabatic expansion, given that the cooling is consistent with an adiabatic process. The adiabatic models will depend on the volume and axial length of an antihydrogen trajectory, hence I will describe a method to approximate the volume and length of the trap sampled by a particle of a given energy. Finally, I will use a modified adiabatic model to estimate the extent of the radial compression during Trial A.

5.1 Derivation of Adiabatic Cooling Models

For a particle oscillating in a box that slowly expands in the z -direction, as I derived in Section 3.2,

$$v_z L_z = \text{constant}, \tag{5.1}$$

where v_z is the axial particle speed and L_z is the axial length of the box. Since an antihydrogen atom with higher energy can reach regions of higher magnetic potential in the trap, the volume sampled by an antihydrogen atom is a function of the particle energy, as is the axial trap length. For *Model 1* I assume the axial and transverse dimensions are uncoupled, such that there is no mixing between axial and transverse energy and that the magnetic trap expands only in the z -direction (along the trap axis, see Figure 2.1 for definition of trap coordinates). In this case the adiabatic invariant simply follows from Equation 5.1 such that,

$$E_{\parallel}(t)L_z^2(t, E(t)) = \text{constant}, \quad (5.2)$$

where E_{\parallel} is the axial particle energy (Equation 3.33).

Model 2 assumes energy equilibrates between all three spatial dimensions, which I will refer to as 3D energy mixing. For this derivation, I assume the trap expands only axially. According to Equation 5.1, the adiabatic trap expansion results in incremental changes in the maximum axial velocity (axial velocity at the trap minimum, since the integral in Equation 3.27 is performed over one period) given by

$$dv_{z,\text{pre-mix}} = -v_z \frac{dL_z}{L_z}, \quad (5.3)$$

where dL_z is the infinitesimally small change in trap length and I use the subscript “pre-mix” to indicate that the expression does not include changes in v_z that arise from energy mixing between axial and transverse degrees of freedom. Using $d(v_z^2) = 2v_z dv_z$, I find an expression for the incremental change in axial energy (Equation 3.33) as the trap expands,

$$dE_{\parallel,\text{pre-mix}} = \frac{M}{2} d(v_{z,\text{pre-mix}}^2) = -Mv_z^2 \frac{dV}{V}, \quad (5.4)$$

and assuming the trap expands only axially,

$$dE_{\perp,\text{pre-mix}} = 0, \quad (5.5)$$

where I have replaced $\frac{dL_z}{L_z}$ by $\frac{dV}{V}$, where V is the volume of the box, since I have assumed the box expands only axially and $dE_{\perp,\text{pre-mix}}$ is the incremental change in transverse energy due to the expansion (not including energy mixing). As energy mixes between axial and transverse degrees of freedom, the axial and transverse energy components are exchanged, but the total energy does not change on account of the energy mixing. For this reason, the incremental change in total energy (including mixing) can be equated to the total energy change that does not include mixing,

$$dE = dE_{\parallel,\text{pre-mix}} + dE_{\perp,\text{pre-mix}} = dE_{\parallel,\text{pre-mix}} = -Mv_z^2 \frac{dV}{V}, \quad (5.6)$$

where M is the antihydrogen mass. Since energy is mixing between axial and transverse degrees of freedom, the axial kinetic energy changes stochastically as the anti-atom oscillates

in the trap, thus at an instant in time the axial energy is some fraction, f , of the total energy,

$$\frac{M}{2}v_z^2 = fE, \quad (5.7)$$

where $0 \leq f \leq 1$. Equation 5.6 becomes

$$dE = -2fE \frac{dV}{V}, \quad (5.8)$$

which can be integrated,

$$\int \frac{1}{E} dE = -2 \int f \frac{1}{V} dV. \quad (5.9)$$

Since energy mixing results in stochastic changes in f during the axial trap expansion, the integral is modulated by the random variable f . Performing the integral gives

$$\ln E = -2\bar{f} \ln V + C, \quad (5.10)$$

where \bar{f} is the average f over the trap expansion and C is a constant. Since Model 2 assumes energy equilibrates between all three spatial dimensions, $\bar{f} = \frac{1}{3}$ and hence I arrive at Model 2

$$E(t)(V(t, E(t)))^{\frac{2}{3}} = \text{constant}, \quad (5.11)$$

where I have indicated the dependence of the trap volume on the particle energy. This derivation assumed axial trap expansion only, but I am now free to morph the trap volume into any shape I like (conserving volume) without changing the particle energy because the energy components are assumed to equilibrate between all three dimensions [90]. Therefore, Equation 5.11 is valid for any change in trap shape. For example, it can include the small radial compression that takes place in Trial A.

We can then equate the energy components at the end, t_f , of the adiabatic expansion experiment to those at the start, t_i , based on each adiabatic model. For Model 1 I have,

$$E_{\parallel}(t_f) = E_{\parallel}(t_i) \left(\frac{L_z(t_i, E(t_i))}{L_z(t_f, E(t_f))} \right)^2, \quad (5.12)$$

and similarly for Model 2,

$$E(t_f) = E(t_i) \left(\frac{V(t_i, E(t_i))}{V(t_f, E(t_f))} \right)^{\frac{2}{3}}. \quad (5.13)$$

The dependence of trap volume and axial trap length on the final energy of the adiabatically cooled anti-atom makes the calculation non-trivial. Assuming $L_z(t, E(t))$ and $V(t, E(t))$ are monotonic (I will show that this is the case in Section 5.2), I could step over possible values of $E_{\parallel}(t_f)$ and $E(t_f)$ until the energy solutions are found. However, this simple method is limited in precision by the number of energy samples taken over the range of all possible energies. If, instead, I take an initial guess that $L_z(t_f, E(t_f)) = L_z(t_f, E(t_i))$ and similarly $V(t_f, E(t_f)) = V(t_f, E(t_i))$, I can then apply each adiabatic model, obtain a new energy,

E_m , and then repeat the process until convergence is reached. This equates to solving the fixed-point iterative equation,

$$E_{\parallel,m+1}(t_f) = E_{\parallel}(t_i) \left(\frac{L_z(t_i, E(t_i))}{L_z(t_f, E_m(t_f))} \right)^2, \quad (5.14)$$

where the zeroth iteration is given by the overestimate,

$$E_{\parallel,m=0}(t_f) = E_{\parallel}(t_i) \left(\frac{L_z(t_i, E(t_i))}{L_z(t_f, E(t_i))} \right)^2, \quad (5.15)$$

for Model 1. Equivalently for Model 2,

$$E_{m+1}(t_f) = E(t_i) \left(\frac{V(t_i, E(t_i))}{V(t_f, E_m(t_f))} \right)^{\frac{2}{3}}, \quad (5.16)$$

and,

$$E_{m=0}(t_f) = E(t_i) \left(\frac{V(t_i, E(t_i))}{V(t_f, E(t_i))} \right)^{\frac{2}{3}}, \quad (5.17)$$

where m denotes iteration order, $t_i = -23.05$ s is the start time of the expansion which is chosen to be 500 ms into the simulation to ensure every particle has time to cross the magnetic minimum before its energy is sampled and $t_f = -50$ ms is the end time of the expansion, a short time before the magnetic ramp down.

For each simulated antihydrogen atom, I solve the fixed-point iterative equations for Model 1 and Model 2 (Equations 5.14 and 5.16), with the initial energies, $E(t_i)$ and $E_{\parallel}(t_i)$, extracted from the simulation. Simulated energy components are calculated as described in Section 3.3 at the most recent crossing of the trap minimum to the sample time (t_i or t_f). The ratio of the ensemble averaged energy component at t_f as predicted by the model to the ensemble averaged energy component at t_i as extracted from the simulation, e.g. $\langle E_{\parallel}(t_f) \rangle / \langle E_{\parallel}(t_i) \rangle$, gives us the predictions of the model. Trap volume and length are determined as described in the following section.

5.2 Trap Volume and Axial Trap Length

Determining the energy predicted by our adiabatic models requires solving Equations 5.14 and 5.16, which each involve a characteristic of the trap that varies as a function of energy. For Model 1 this is the axial length of the trap and for Model 2 this is the trap volume. I use an algorithm to determine the volume of the trap that is accessible by a particle of energy E at time t , and average the axial length of this volume over radius to determine the axial trap length.

To find the volume of the trap at a sampling time, t , I sum the magnetic field of the individual magnets using the methods described in Section 3.4, scaled by the current in each magnet (Figure 2.4) at time t , then find the location of the minimum, \mathbf{x}_{\min} , of the magnetic potential (Equation 3.2) using ROOT's Minuit2 Minimization Library as mentioned in Section 3.4. I define a grid of 500 points equally spaced in each spatial dimension over a cuboid volume




Figure 5.1: Potential, $U(r_{\text{trap}}, \phi, z)$ (K), on electrode wall ($r_{\text{trap}} = \sim 2.2$ cm) as a function of azimuthal angle, ϕ (rad), and axial position, z (mm). (a) Trial A $t = -50$ ms (before start of Octupole FRD) and (b) Trial A $t = 450$ ms (during Octupole FRD). Both subfigures are scaled to the same colour axis.

slightly larger than the volume of the trap. Since the trap has larger axial extent than radial extent, the distance between grid points in the x and y directions is $dx = dy = 80 \mu\text{m}$ and in the z direction is $dz = 800 \mu\text{m}$. The distance between grid points remains fixed.

The algorithm is initialised at the closest grid point to \mathbf{x}_{min} and is assigned a sample energy, E . Once initialised, the algorithm then steps recursively in each of the six total directions (positive and negative x , y and z). When it moves to a grid point, the algorithm determines whether the magnetic potential at the grid point is greater than E . If so, the algorithm does not recurse at this location and returns back to the previous grid point where it steps in a different direction. The algorithm starts to sample a boundary at which the magnetic potential is too high for E to overcome and the volume enclosed by this boundary is the volume of the trap that is accessible by a particle of energy E at time t i.e. it determines $V(t, E(t))$.

Starting with small E , the algorithm runs with gradually increasing E , to find the functional form of $V(t, E)$. As E increases, $V(t, E)$ also increases much like a container filling up with water. When E becomes equal to the trap depth ($E = \Gamma$), the volume sampled by the algorithm will intersect the trap wall, r_{trap} , at which point $V(t, E)$ becomes nonsensical since the antihydrogen atom would annihilate on the trap wall. It is worth noting that quasi-trapped antihydrogen trajectories are known to be able to survive on long timescales with energies greater than the trap depth, since they do not sample the entire volume accessible to a particle with their given energy, which represents a limitation of this model.

If the sample energy is high enough, the volume sampled by the algorithm will hit the trap wall. The lowest energy path to r_{trap} depends on the particular magnetic field configuration

i.e. the ratio of the Octupole and Mirror Coil currents. To visualise this, I take the example of Octupole FRD in which the ratio of Octupole and Mirror Coil current decreases linearly as a function of time. In Figure 5.1 I compare the magnetic potential at r_{trap} for Trial A (and equally Trial C) at $t = -50$ ms (just before the start of FRD, Figure 5.1(a)) and at $t = 450$ ms (during Octupole FRD, Figure 5.1(b)). At $t = -50$ ms the lowest path to r_{trap} lies out at high axial positions past the peak of the Mirror Coil field (± 0.137 m), through what are known as the ‘holes’. Much like the algorithm is able to sample this route to the trap wall, antihydrogen atoms with sufficient energy can escape the trap by effectively tunneling through the Mirror Coil field but it can take some time for the particles to find the path. This is the method of slow escape of the majority of quasi-trapped antihydrogen atoms. The ‘holes’ are formed by the sum of the radial Mirror Coil field and the radial Octupole field, which point in opposite directions at certain azimuths, creating minima. At $t = 450$ ms, the Octupole current has begun to decrease relative to the Mirror Coil current, at which point the holes are no longer the minimum potential on the trap wall. Instead the minimum path to the trap wall is at smaller z , between confining Mirrors A and E.

When applying the adiabatic models, I ignore quasi-trapped particles, which I define as those that annihilate before the onset of magnetic ramp-down (at $t = 0$ s). However, some trajectories that have $E(t_f) > \Gamma(t_f)$ still remain. These trajectories have not had sufficient time to sample the trap ‘holes’ during the adiabatic expansion experiment. For these trajectories, applying the adiabatic models (Equation 5.2 and Equation 5.11) requires calculating $V(t, E(t) > \Gamma(t))$. As mentioned above, it does not make sense to calculate $V(t, E(t) > \Gamma(t))$ since the volume intersects the trap wall (and hence antihydrogen sampling this volume would annihilate). Since, by definition, particles surviving on long timescales with energy greater than the trap depth do not sample the entire volume accessible to a particle of their energy, I make the assumption $V(t, E(t) > \Gamma(t)) = V(t, E(t) = \Gamma(t))$.

The ratio of the trap volume at t_f to t_i for the three experimental trials, $V(t_f)/V(t_i)$, is shown in Figure 5.2 up to $E = \Gamma$. Also shown is its axial length, averaged over radius, $L_z(t_f)/L_z(t_i)$. Notice small changes in the length and volume ratios in Trials B and C. In both trials, they result from the ramp down of Solenoid A in the first 5 s and in Trial C they result from the gradual ramp of Mirror B, as shown in Figure 2.4.

An interesting characteristic of the volume ratio for Trial A is the sharp increase in the ratio as energy decreases for $E < \sim 10$ mK. $V(t_i)$ and $V(t_f)$, as sampled by the algorithm, are shown for $E = 0.4$ K (Figures 5.3 (a) and (b) respectively) and $E = 10$ mK (Figures 5.3 (c) and (d) respectively). The sharp change in the gradient of the volume ratio coincides with the appearance of hollowing (hole in center of trap volume in Figure 5.3a) in $V(t_i)$. To understand the features of the magnetic field that lead to hollowing in the trap volume, in Figure 5.4 I plot the modulus of the magnetic field strength, $|B|$, in an xy -plane at $z = 0$,




Figure 5.2: Ratios of trap volume and length at the end, t_f , to the start, t_i , of the three trials of adiabatic expansion as a function of anti-atom total energy, $E(K)$. (a) Trap volume ratio, $\frac{V(t_f)}{V(t_i)}$, and (b) axial trap length ratio, $\frac{L_z(t_f)}{L_z(t_i)}$, are shown for Trial A (red solid), Trial B (blue dashed) and Trial C (green dotted). The trap volume is calculated using the algorithm described in the text and the axial trap length ratio is obtained by averaging the length of this trap volume over radius. Cyan and magenta dotted lines in (a) are used to indicate the energies referenced in Figure 5.3.

which shows that $|B|$ has a local maximum at $r = 0$, a minimum at $r \sim 3$ mm and then increases out to the trap wall. These features are caused by the quadrature sum of the axial Mirror Coil field, which is peaked at $r = 0$, and the radial Octupole field, which is peaked at high radius. The hollowing is a result of the sampling energy being lower than the crown of what is often referred to as the *Mexican-hat potential* shown in Figure 5.4. Although the volume sampled by a particle of energy 10 mK at t_i in Trial A is hollow, the volume sampled by a particle of energy 10 mK at t_f in Trial A is not; as a result, particles trapped within the hollow volume at t_i could undergo large changes in volume during adiabatic expansion.

If one wished to exploit this, one could consider attempting to laser cool trajectories into this volume prior to adiabatic expansion. Antihydrogen laser cooling has recently been demonstrated (see Section 1.6.1 for a description), achieving mean energy parallel to the cooling laser (approximately parallel to the trap axis) of ~ 20 mK and mean energy transverse to the cooling laser of ~ 56 mK [7]. The absolute lower limit to which antihydrogen can be cooled via Doppler laser cooling is set by the heating on account of photon recoil (Doppler limit) which is ~ 4 mK for ground state antihydrogen. In the reality of trapped antihydrogen, however, lack of three-dimensional trap access and the extent of energy mixing between axial and transverse dimensions limits achievable mean energies in the ALPHA trap to ~ 20 mK in practice [102], but that may still be low enough that a reasonable number of trajectories

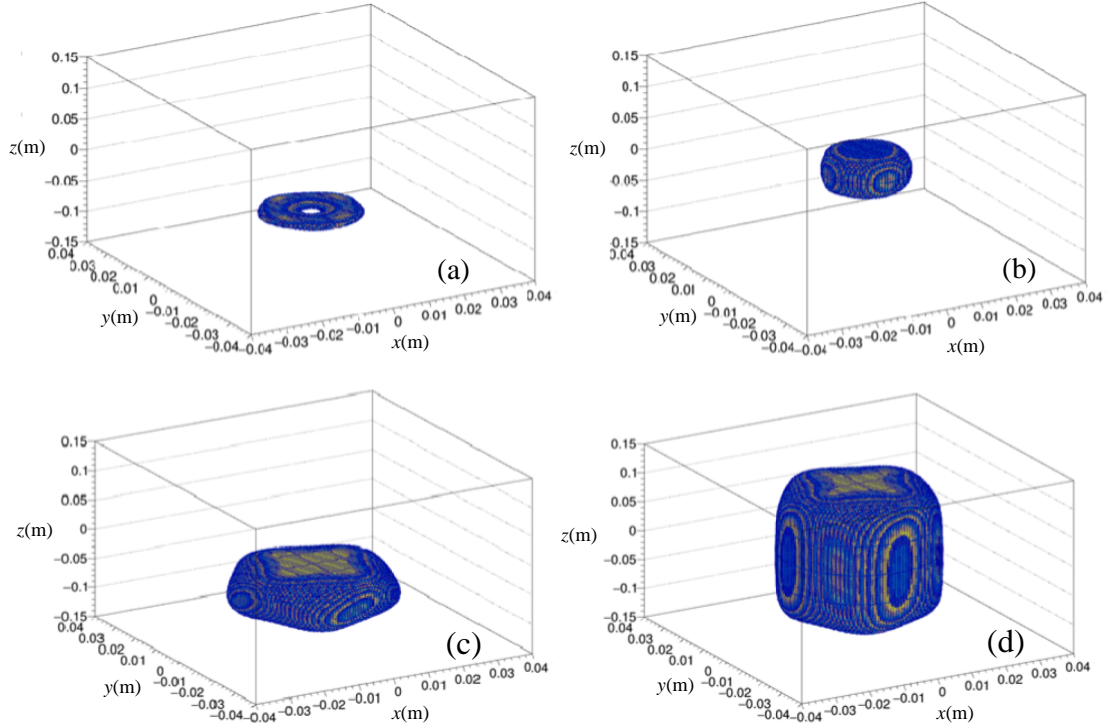


Figure 5.3: Trappable volume in Cartesian coordinates, (x, y, z) , accessible by a particle of energy E at time t during adiabatic expansion Trial A, found by recursively searching a grid of 500 equally spaced points. Since the trap has larger axial extent than radial extent, the distance between grid points in the x and y directions is $dx = dy = 80 \mu\text{m}$ and in the z direction is $dz = 800 \mu\text{m}$. (a) $E = 0.01 \text{ K}$, $t = t_i$ (b) $E = 0.01 \text{ K}$, $t = t_f$, (c) $E = 0.4 \text{ K}$, $t = t_i$, (d) $E = 0.4 \text{ K}$, $t = t_f$.




Figure 5.4: The *Mexican Hat Potential* in the ‘standard’ trap configuration (Mirrors A and E energised to 600 A, internal Mirror Coils at 0 A and the Octupole energised to 900 A) at $z = 0$. $|B(T)|$ is shown on the vertical axis, with blue and red colours indicating smaller and larger relative values of $|B(T)|$ respectively. The shape of the potential is dominated by the quadrature sum of the axial Mirror Coil field which is peaked at $r = 0$ and the radial Octupole field which has a minimum at $r = 0$.

could be confined within the hollow volume. In addition to these constraints, laser cooling antihydrogen into this volume would require shifting the cooling laser ~ 3 mm off-axis.

Even if it were possible to laser cool a large number of antihydrogen trajectories into this hollow volume, it is not obvious how much adiabatic cooling could be achieved. One of the reasons for this is that adiabatic cooling reduces energy during expansion, which, in turn, reduces the volume sampled by the antihydrogen atom (since volume is a function of energy). Further, at energies ~ 10 mK, the larger the change in volume over the course of the expansion, the larger the change in energy; but the larger the change in energy, the smaller the change in volume over the course of the expansion. In other words, the increase in cooling starts to cancel itself out. In addition, the hollow volume undergoes a relatively rapid expansion, thus it is not clear whether the anti-atoms would experience an adiabatic change in energy. With all of these considerations in mind, combined simulations of laser cooling and adiabatic expansion cooling are required to assess the feasibility of enhanced cooling by combination of laser cooling and adiabatic expansion cooling.

5.3 Adiabatic Modelling Results

	$\frac{\langle E(t_f) \rangle}{\langle E(t_i) \rangle}$ (%)	$\frac{\langle E_{\parallel}(t_f) \rangle}{\langle E_{\parallel}(t_i) \rangle}$ (%)	$\frac{\langle E_{\perp}(t_f) \rangle}{\langle E_{\perp}(t_i) \rangle}$ (%)
Simulation	62.08 ± 0.08	31.95 ± 0.09	85.78 ± 0.20
Model 1	63.12 ± 0.10	16.22 ± 0.05	100.00
Model 2	59.47 ± 0.08	45.04 ± 0.09	70.82 ± 0.10

Table 5.1: Percentage change of ensemble averaged energies between the referenced simulation (*Simulation*) and the adiabatic model, Model 1 and Model 2) just before the dump ($t_f = -50$ ms). Quasi-trapped particles that are lost between t_i and t_f are excluded from the ensemble-average. Stated errors are propagated to the percentage ratios from the standard error on the sample means. No error is stated for the transverse ratio for Model 1 since the model assumes no energy mixing and axial expansion only and therefore the transverse energy is exactly constant.

In the following section I use adiabatic models, Model 1 and Model 2, to predict the energy loss over the course of Adiabatic Expansion Trial A according to the assumptions of the models. I solve fixed-point iterative Equations 5.14 and 5.16 with the energy components at t_i extracted from a simulation of Adiabatic Expansion Trial A with simulation parameters shown in Table 4.1. The trap volume and axial trap length are determined using the method described in the previous section.

Table 5.1 compares the ensemble averaged energy components extracted from the simulation at t_f to the predictions of Model 1 and Model 2, where the predictions are given as a percentage of the energy components at t_i . The absolute values of the ensemble-averaged energy components at t_i are given in Table 4.3. Recall that Model 1 assumes no mixing between axial and transverse energy components and that the trap expands only axially. Therefore, as expected it overpredicts the amount of axial adiabatic cooling and hence it predicts the ensemble to have a lower final parallel energy where $\frac{\langle E_{\parallel}(t_f) \rangle}{\langle E_{\parallel}(t_i) \rangle} = 16.22 \pm 0.05$ % compared to 31.95 ± 0.09 % for the simulation. In addition, the assumptions of Model 1 lead immediately to their being no change in the transverse energy of the population and hence $\frac{\langle E_{\perp}(t_f) \rangle}{\langle E_{\perp}(t_i) \rangle} = 100.00$ %.

Model 2 provides the alternative assumption of complete 3D energy mixing. As expected, the model underpredicts the fractional decrease in axial energy with $\frac{\langle E_{\parallel}(t_f) \rangle}{\langle E_{\parallel}(t_i) \rangle} = 45.04 \pm 0.09$ %, compared to a simulation value of 31.95 ± 0.09 %. For this model, however, mixing reduces the transverse energy such that $\frac{\langle E_{\perp}(t_f) \rangle}{\langle E_{\perp}(t_i) \rangle} = 70.82 \pm 0.10$ %, beyond the corresponding simulation result of 85.78 ± 0.20 %. In addition, the total energy loss by the two models, $\frac{\langle E(t_f) \rangle}{\langle E(t_i) \rangle} = 63.12 \pm 0.10$ % and $\frac{\langle E(t_f) \rangle}{\langle E(t_i) \rangle} = 59.47 \pm 0.08$ % for Model 1 and Model 2 respectively,

bound the corresponding simulation result of $62.08 \pm 0.08\%$. Therefore, I can conclude the simulated decrease in energy during adiabatic expansion Trial A is consistent with adiabatic cooling within the range of possible energy mixing dynamics. Note that the lower final energy predicted by Model 1 compared to Model 2 shows that the largest energy reduction in an axial adiabatic expansion is achieved by maximising the degree of mixing between axial and transverse energies during the expansion, as can be easily deduced from the power law of the adiabatic model. However, increases in cooling are minimal at the level of a few percent.

5.3.1 Radial Compression during Adiabatic Expansion

As stated previously, in addition to the axial adiabatic expansion that takes place in Trial A, there is also a small amount of radial compression. To understand the extent of this compression, in Figure 5.5 I plot $\frac{V(t_f)L_z(t_i)}{V(t_i)L_z(t_f)}$ which represents the change in area (averaged over z) in the xy -plane as a function of energy during the three experimental trials. Note that a value less than 1 indicates a radial compression.




Figure 5.5: Ratio of area of trappable volume (averaged over axial position), $\frac{V(t_f)L_z(t_i)}{V(t_i)L_z(t_f)}$, at the end, t_f , to the start, t_i , of Trials A, B and C as a function of total anti-atom energy, E (K). Note that $\frac{V(t_f)L_z(t_i)}{V(t_i)L_z(t_f)} < 1$ indicates radial compression over the course of the Trial.

The reason for the compression can be understood by imagining the Mexican hat potential shown in Figure 5.4. Since the central peak of the Mexican hat is set by the axial Mirror Coil field, the radial position of the field minimum is set by the strength of the Mirror Coil. As the Mirror Coil field effectively weakens during expansion, because the length of the magnetic trap increases, the radius of the magnetic trap (averaged over axial positions) contracts adiabatically. As a result, the radial energy of the anti-atoms increases. The maximum amount of radial heating is experienced by particles with total energy of approximately 10 mK, which corresponds to the anti-atoms which sample the hollow volume before adiabatic

expansion as shown in Figure 5.3.

This effect is included in Model 2 since its expression is volume dependent. To investigate the size of the effect of the radial heating, I define Model 2* as an adiabatic model, which assumes complete 3D energy mixing during adiabatic expansion but it assumes the trap expands only axially i.e. that $\frac{V(t_i, E(t_i))}{V(t_f, E(t_f))} = \frac{L_z(t_i, E(t_i))}{L_z(t_f, E(t_f))}$. Model 2* is then,

$$E(t_f) = E(t_i) \left(\frac{L_z(t_i, E(t_i))}{L_z(t_f, E(t_f))} \right)^{\frac{2}{3}}, \quad (5.18)$$

and I can again solve for $E(t_f)$ using the same method as above. The results of this are shown in Table 5.2, where I have duplicated Table 5.1 but added a new line for the results of Model 2*. There is a difference in $\frac{\langle E(t_f) \rangle}{\langle E(t_i) \rangle}$ of $\approx 7\%$ between Model 2 (includes radial compression) and Model 2* (excludes radial compression). Note that this is the predicted difference in the ratio of total energies as a result of the radial compression given that there is 3D energy mixing. It's likely to be possible to counteract this radial compression by gradually decreasing the Octupole current to induce an equal and opposite radial adiabatic expansion.

	$\frac{\langle E(t_f) \rangle}{\langle E(t_i) \rangle} (\%)$	$\frac{\langle E_{\parallel}(t_f) \rangle}{\langle E_{\parallel}(t_i) \rangle} (\%)$	$\frac{\langle E_{\perp}(t_f) \rangle}{\langle E_{\perp}(t_i) \rangle} (\%)$
Simulation	62.08 ± 0.08	31.95 ± 0.09	85.78 ± 0.20
Model 1	63.12 ± 0.10	16.22 ± 0.05	100.00
Model 2	59.47 ± 0.08	45.04 ± 0.09	70.82 ± 0.10
Model 2*	52.18 ± 0.06	39.52 ± 0.08	62.13 ± 0.10

Table 5.2: Percentage change in ensemble averaged energy components in a simulation of adiabatic expansion Trial A and as predicted by adiabatic models, Model 1 and Model 2. Note that the first three rows are identical to that of Table 5.1; the third row gives the prediction of the modified adiabatic model that does not include radial compression, Model 2*. Stated errors are propagated to the percentage ratios from the standard error on the sample means. No error is stated for the transverse ratio for Model 1 since the model assumes no energy mixing and axial expansion only and therefore the transverse energy is exactly constant.

Optimisation and Applications of Adiabatic Cooling

The ultimate goal of experimental demonstration and simulations of adiabatic cooling is to improve precision antihydrogen physics. For use in such applications, it will be necessary to optimise the technique. In this chapter I will start to explore optimisation techniques by an analysis of the effect of changing the expansion duration on the energy of the adiabatically cooled population. I will then explore two possible applications of adiabatic expansion that have not yet been experimentally implemented, through detailed simulation. I will first discuss the application of adiabatic cooling to experiments that place upper bounds on the electric charge of antihydrogen, where slower anti-atoms allow us to reduce the depth of the magnetic confining potential whilst retaining large numbers of trapped antihydrogen. Anti-atoms confined in the shallower well can be more readily ejected from the trap if they have charge and therefore experience a force in an electrostatic field. Secondly, I will explore the compatibility of adiabatic cooling and experiments that measure the 1S-2S transition in antihydrogen without detailed atomic physics simulation, by investigating the spatial nature of orbits occupied by adiabatically cooled antihydrogen.

6.1 Modifying the Expansion Rate

As explained in Chapter 2, the kinetic energy of trapped antihydrogen atoms, subject to a longitudinally expanding magnetic trap, decreases adiabatically if the timescale of the expansion is much longer than the axial bounce period of the anti-atoms, which is ~ 10 ms. The upper limit on the expansion duration is set by the finite lifetime of the antihydrogen atoms as a result of collisions on background gas which was recently shown to be >66 h [69]. For the 2016 adiabatic expansion experiment, a somewhat arbitrary duration of ~ 23 s was chosen. In this section I will describe the effect of changing the duration of the adiabatic




Figure 6.1: (a) Superconducting magnet currents as a function of time during Adiabatic Expansion Trial A. (b) Superconducting magnet currents as a function of time during a 10.6 s-long adiabatic expansion, where the timescale of the Solenoid A, Mirror B, Mirror C and Mirror D current ramps are scaled down by a factor of 2. Time, t (s), is given relative to the onset of Octupole FRD.




Figure 6.2: Ensemble-averaged total energy, $\langle E(t) \rangle$ as a function of time, t , relative to the start of magnetic ramp-down for a 230.2s duration adiabatic expansion (red solid curve) and a 10.6 duration expansion followed by a hold in the standard trap configuration (defined in Figure 2.3(a)) (blue solid curve). Particles that annihilate before $t = 0$ s (quasi-trapped particles) have been excluded. 5000 trajectories were simulated in total, 4272 of which make it to $t = 0$ s for the 230.2s expansion, compared to 4576 for the 10.6s expansion.

expansion on the simulated energy distribution of populations of trapped antihydrogen.

The currents in the superconducting magnets as a function of time during Adiabatic Expansion Trial A are shown in Figure 6.1(a). To change the adiabatic expansion duration, I scale each of the Solenoid A, Mirror B, Mirror C and Mirror D current ramps by the same factor, whilst retaining the relative timing of the ramps. For example, in Figure 6.1(b) I show an adiabatic expansion with approximately half the duration of Adiabatic Expansion Trial A, resulting in a ≈ 10.6 second duration adiabatic expansion. In the following section, I compare different adiabatic expansion durations via simulation. The initial energy distribution of the total population of (5000) simulated anti-atoms is the same in each case, as the simulated trajectories are initialised in the same magnetic configuration.

For each expansion duration I will produce the equivalent of Figure 4.9: the ensemble-averaged energy as a function of time during the expansion. I simulate antihydrogen with total energy up to 0.75 K, allowing simulation of quasi-trapped anti-atoms. Since quasi-trapped anti-atoms escape the trap gradually over time, and they have the highest energies in the ensemble, the energy of the trapped distribution of antihydrogen gradually decreases over time even without an adiabatic expansion. For this reason, I exclude anti-atoms that annihilate before the adiabatic expansion has finished from the ensemble average since they result in a skew to higher energies at earlier times; making the magnitude of adiabatic cooling appear larger than it is.




Figure 6.3: Ensemble-averaged total energy, $\langle E(t) \rangle$ (K) as a function of scaled time, τ , in which the total expansion duration is scaled between 0 and 1. The initial energy of the 5000 simulated trajectories is truncated at 0.47 K (just below the magnetic trap depth) to prevent particle losses (quasi-trapped) before $\tau = 0$.

To make an effective comparison between expansion durations, in Figure 6.2 I compare a 230.2 s expansion to a 10.6 s expansion followed by a hold for a further 219.6 s in the standard trap configuration (see Figure 2.3(a) for definition of standard trap). In both cases the ensemble average only includes particles that survive until magnetic ramp-down at 0 s, which equates to 4576 and 4272 particles for the 10.6 s and 230.2 s expansions respectively. This means more quasi-trapped particles are lost in a longer expansion duration (recall that both 10.6 s duration followed by a hold and 230.2 s duration are simulated for the same total length of time). The total energy difference between the 10.6 s and 230.2 s duration at $t = 0$ s is 0.024 ± 0.002 K.

There are two possible reasons for the greater loss of quasi-trapped particles in the longer expansion duration. The first is that the shorter trap volume (present for longer in the 230.2 s expansion) has a magnetic configuration that quasi-trapped particles are more likely to escape from, possibly because the radial magnetic field of the Mirror Coils is stronger in the shorter magnetic trap, which reduces the radial confining potential at certain azimuths (see Figure 4.16). The second is that particles that begin with energy higher than the trap depth (which would be quasi-trapped) are quickly (before they have chance to escape) cooled below the trap depth. Whatever the reason, if one is trying to achieve the lowest final energy possible from adiabatic expansion, they may consider increasing the expansion duration to allow more quasi-trapped particles to escape. However, further simulations are required to determine whether a slower population is achieved by a slow adiabatic expansion or by first holding in the short magnetic volume (to allow quasi-trapped particles to escape) before adiabatically expanding.

Of course, one must also consider the reduction in the number of trapped antihydrogen atoms that results from quasi-trapped particle loss. If one is trying to use adiabatic expansion to retain larger numbers of trapped antihydrogen atoms (rather than trying to achieve the lowest mean energy), a shorter expansion duration is likely to be beneficial.

To isolate effects other than quasi-trapped particle losses that may cause a difference in final energy with different expansion durations, I compare simulated populations with initial energies sampled from a Maxwell-Boltzmann distribution truncated at 0.47 K, rather than 0.75 K (see Section 3.4 for a description of the typical initial energy distribution). Since every simulated trajectory will have energy below the magnetic trap depth, there will be no quasi-trapped particle losses. In Figure 6.3, I compare adiabatic expansion durations of 10.6 s, 22.1 s, 45.2 s and 230.2 s with mixing disabled. In this case, the difference in total final energy between the 10.6 s and 230.2 s expansion is 0.017 ± 0.002 K. Note that I have normalised the timescale of the simulated expansion in the figure so that the final energy differences are clear.

The reason for the difference in final energy as a function of expansion duration with the initial energy truncated at 0.47 K can be understood with a simple example of a box expanding adiabatically only in the z -direction by a factor of 2. For simplicity I will drop the dependence of trap volume and axial trap length on particle energy. Let us estimate the difference in total final energy in the two extreme limiting cases of energy mixing between axial and transverse dimensions. Assuming no mixing, Model 1 (Equation 5.2) gives the axial energy in the expanded box,

$$E_{\parallel}(t_f) = E_{\parallel}(t_i) \left(\frac{L_z(t_i)}{L_z(t_f)} \right)^2 = \frac{1}{4} E_{\parallel}(t_i), \quad (6.1)$$

where t_i and t_f are times before and after the adiabatic expansion respectively and the trap length $L_z(t_f) = 2L_z(t_i)$ since the box expands axially by a factor of 2. Since I have assumed L_z is not a function of energy, it is simple to take the ensemble average,

$$\langle E_{\parallel}(t_f) \rangle = \langle E_{\parallel}(t_i) \rangle \left(\frac{L_z(t_i)}{L_z(t_f)} \right)^2 = \frac{1}{4} \langle E_{\parallel}(t_i) \rangle. \quad (6.2)$$

The fact that the velocities in each dimension are sampled from an independent normal distribution sets the initial conditions, $\langle E_{\parallel}(t_i) \rangle = \frac{1}{3} \langle E(t_i) \rangle$ and $\langle E_{\perp}(t_i) \rangle = \frac{2}{3} \langle E(t_i) \rangle$. Therefore, the total energy loss assuming no mixing is given by,

$$\langle E(t_f) \rangle = \langle E_{\parallel}(t_f) \rangle + \langle E_{\perp}(t_f) \rangle = \frac{1}{4} \langle E_{\parallel}(t_i) \rangle + \langle E_{\perp}(t_i) \rangle = \frac{1}{4} \left(\frac{1}{3} \langle E(t_i) \rangle \right) + \frac{2}{3} \langle E(t_i) \rangle = 0.75 \langle E(t_i) \rangle, \quad (6.3)$$

where I have used $\langle E_{\perp}(t_i) \rangle = \langle E_{\perp}(t_f) \rangle$, on account of our assumption of no mixing. On the other hand, Model 2 (Equation 5.11) gives the change in energy assuming energy equilibrates between all three dimensions,

$$\langle E(t_f) \rangle = \langle E(t_i) \rangle \left(\frac{V(t_i)}{V(t_f)} \right)^{\frac{2}{3}} = \left(\frac{1}{2} \right)^{\frac{2}{3}} \langle E(t_i) \rangle \approx 0.63 \langle E(t_i) \rangle, \quad (6.4)$$




Figure 6.4: Time-averaged normalised axial energy, $\langle\langle\varepsilon_{\parallel}\rangle\rangle$ as a function of time, $t(s)$, during a 100s hold in the standard trap with mixing artificially turned off i.e. with an infinite length Octupole and the radial component of the Mirror magnetic field set to zero at all positions. Each line is an individual simulated trajectory. A non-constant $\langle\langle\varepsilon_{\parallel}\rangle\rangle$ over time indicates energy mixing between axial and transverse dimensions.

where I have used $V(t_f) = 2V(t_i)$ since I have assumed the box expands only axially.

Since Equation 6.4 predicts a lower final total energy than Equation 6.3, I can conclude that increased mixing during adiabatic expansion can lead to a lower final total energy. I hypothesise that this is the reason for the difference in final total energy for the different expansion durations shown in Figure 6.3. The longer the duration of the expansion, the more likely the trajectories are to sample regions of the trap that mix axial and transverse energies. The difference in final energy of each of the 22.1s, 45.2s and 230.22s adiabatic expansion durations compared to the final energy of the 10.6s duration are approximately 0.007 ± 0.002 K, 0.010 ± 0.002 K and 0.017 ± 0.002 K. Note that the difference in final energy shown in Figure 6.3 cannot be directly interpreted as the proportion of the energy difference in Figure 6.2 that comes from energy mixing since the initial energy distributions are truncated at different values.

6.1.1 Artificially Disabling Energy Mixing

In this section, I will describe a technique to artificially modify the simulated magnetic fields to ‘disable’ energy mixing between axial and transverse dimensions. Using this technique, I will compare the energy decrease during different durations of adiabatic expansion to test the hypothesis that the observed difference in final energies for simulated populations truncated at 0.47 K result from energy mixing.

In Section 3.3 I introduced the idea that energy mixing between axial and transverse




Figure 6.5: (a) Ensemble-averaged total energy, $\langle E(t) \rangle$ (K), for simulated trajectories where mixing has been artificially turned off, as a function of scaled time τ , in which the total expansion duration is scaled between 0 and 1. The initial energy of the 5000 simulated trajectories is truncated at 0.47 K (just below the magnetic trap depth) to prevent particle losses (quasi-trapped) before $t = 0$ s. Mixing is disabled by replacing the Octupole model with an infinitely long octupole and setting the radial field of the mirror coils to zero at all positions. (b) An energy-zoomed version of (a).

dimensions is caused by breaking of the azimuthal symmetry of $|\mathbf{B}(\mathbf{x})|$ by the quadrature sum of the magnetic fields of the Mirror Coils and Octupole in the radial and axial dimensions. Energy mixing can be artificially disabled by replacing the model for the Octupole field by an infinitely long octupole (solution for magnetic field given in Appendix A) and artificially setting the radial field of the Mirror Coils to zero at all positions.

Recall that in Section 3.3, I presented the time-averaged normalised axial energy, $\langle\langle \varepsilon_{\parallel} \rangle\rangle$, of a population of simulated antihydrogen atoms as a function of time (Figure 3.9). For antihydrogen atoms that mix a significant portion of their axial and transverse energy components, $\langle\langle \varepsilon_{\parallel} \rangle\rangle$ varies significantly as a function of time. In Figure 3.9, I show the equivalent for the same (standard) magnetic trap configuration, but with energy mixing between axial and transverse degrees of freedom artificially disabled. The majority of simulated anti-atoms have an $\langle\langle \varepsilon_{\parallel} \rangle\rangle$ that is approximately constant as a function of time, meaning the majority of the population do not mix energies in 100s, but a subset of the population with $\varepsilon_{\parallel} > \sim 0.8$ exchange some of their axial energy, though the reason for this is not known.

In Figure 6.5 I compare different durations of adiabatic expansion with mixing artificially disabled. The difference in the energy between the different duration adiabatic expansions at $t = 0$ s is around 0.001 ± 0.002 K and therefore artificially removing energy mixing between axial and transverse dimensions has removed the difference in the final energy of the different expansion durations within error. I conclude that the difference in final energy of the different




Figure 6.6: Diagram of a magnetic field sequence, in which antihydrogen trajectories are initialised in the short (tight) trap configuration at (1), before a slow adiabatic expansion takes place to arrive at the standard trap configuration at (2). After the expansion, the trap contracts back to the same tight configuration in (1) but the adiabatic contraction is faster than the slow expansion. More energy mixing is expected to occur during the slow expansion than the fast contraction. The length of the trap at each stage is indicated as a guide, although the true length of any one trajectory is energy and radius dependent.

expansion durations with an initial energy distribution truncated at 0.47 K is dominated by effects of energy mixing between axial and transverse dimensions.

6.1.2 Adiabatic Expansion and Contraction Cycling

The observation that an increased adiabatic expansion duration can lead to a lower final energy as a result of energy mixing (Figure 6.3) might cause one to imagine what happens in the following situation: what if I expand the trap with a very long duration adiabatic expansion, and then contract the trap back to the original configuration, with the contraction being fast compared to the expansion but sufficiently slow that the process is still adiabatic? Figure 6.6 depicts this particular sequence of magnetic trap manipulations. During the slow expansion, I expect more energy mixing between axial and transverse dimensions than during the fast contraction. Therefore, one might naively wonder whether it is possible to end up with $\langle \frac{E_3}{E_1} \rangle < 1$, where E_3 and E_1 are the antihydrogen energies in stages 3 and 1 of Figure 6.6 respectively and the angled bracket denotes the ensemble average. If this were the case, the process could be repeated, driving the energy of the trapped population lower and lower each time.

In this section I will prove that $\langle \frac{E_3}{E_1} \rangle \geq 1$ for an arbitrary trap expansion and an arbitrary degree of energy mixing and therefore that it is not possible to achieve cooling via this method. The proof will require us to derive general adiabatic models which reduce to Model 1 and Model 2 under the assumptions outlined in Chapter 5, and which also may be useful in future for approximating energy loss under certain trap manipulations and energy mixing conditions.

Further, the result is relevant for certain proposals of adiabatic expansion of antihydrogen atoms which involve an adiabatic expansion followed by an adiabatic contraction. One such

proposal [103] has been considered for use in ALPHA-g (see Section 1.4 for a description of ALPHA-g). Currently (at the time of writing this thesis), the ALPHA-g superconducting octupole magnet often quenches at high currents after being energised for ~ 20 minutes (the reason for this is currently unknown). If this behaviour continues, it will prevent future attempts to accumulate ('stack' - see Section 1.2.3) antihydrogen atoms in the same magnetic trap over long periods of time, since it is likely that the ALPHA-g octupole will quench at some point and cause the antihydrogen atoms to escape and annihilate. To avoid this issue, it may be possible to mix antiproton and positron plasmas for ~ 1 s with the Octupole magnet at high current, before slowly ramping down the Octupole magnet to some lower current (at which quenches are improbable). Then, the proposal involves ramping the ALPHA-g octupole magnet back up to high current in time to mix antiproton and positron plasmas on the next antiproton shot from the AD [103]. The proposal is intended to minimise the amount of time spent with the octupole magnet at high current and therefore to reduce the risk of an octupole quench. In theory, the antihydrogen atoms undergo a radial adiabatic expansion, leading to an energy reduction which may enable trapping of larger numbers of antihydrogen than it is possible to accumulate with the ALPHA-g octupole at a constant low current. The important question is whether the continuous cycling of an adiabatic expansion followed by an adiabatic contraction causes the antihydrogen atoms to gain energy over time, and hence to escape the magnetic trap and annihilate.

I will start with an example of a magnetic trap that has axial length L_1 , then expands slowly to length L_2 , before contracting quickly (compared to expansion, but still slow compared to axial bounce time such that it is still adiabatic) to return to a trap which also has axial length L_1 , as shown in Figure 6.6. I will assume the other dimensions of the trap remain constant (no radial expansion) and that the trap length is not a function of the particle energy. If I assume 3D energy mixing during the slow expansion and no mixing during the fast contraction, the changes in energy of the particle are governed by Model 2 (Equation 5.11) and Model 1 (Equation 5.2) respectively. In terms of the energy at stages (1), (2) and (3) (Figure 6.6), (E_1 , E_2 and E_3 respectively), Models 2 and 1 state that,

$$\begin{aligned} \text{3D mixing expansion: } E_1 L_1^{2/3} &= E_2 L_2^{2/3}, \\ \text{no mixing contraction: } E_{||2} L_2^2 &= E_{||3} L_1^2, \end{aligned} \tag{6.5}$$

where Model 2 involves L_1 and L_2 (rather than volume, V , as usual) on account of the box expanding only axially and therefore $V \propto L$. In the following paragraph, I will find an expression for E_3 in terms of E_1 using Equation 6.5. At (2), a single particle in the box can have an axial energy between 0 and E_2 . Therefore,

$$E_{||,2} = f_2 E_2 \tag{6.6}$$

where $0 \leq f_2 \leq 1$. Assuming no mixing during the fast contraction, I find $E_{\perp,3} = E_{\perp,2}$. Using

this and Equation 6.6, I can find an expression for $E_{\parallel,3}$ in terms of E_2 , E_3 and f_2 ,

$$\begin{aligned} E_{3\parallel} &= E_3 - E_{\perp,3} \rightarrow E_{\parallel,3} = E_3 - E_{\perp,2}, \\ E_{\perp,2} &= E_2 - E_{\parallel,2}, \\ E_{\perp,2} &= E_2 - (f_2 E_2) = (1 - f_2) E_2, \\ \therefore E_{\parallel,3} &= E_3 - (1 - f_2) E_2. \end{aligned} \tag{6.7}$$

Therefore, Equation 6.5 becomes,

$$\begin{aligned} \text{3D mixing expansion: } E_1 L_1^{2/3} &= E_2 L_2^{2/3}, \\ \text{no mixing contraction: } f_2 E_2 L_2^2 &= [E_3 - (1 - f_2) E_2] L_1^2, \end{aligned} \tag{6.8}$$

where I have substituted Equations 6.6 and 6.7 to eliminate parallel energy components. I can then eliminate E_2 from the problem entirely and I find,

$$\frac{E_3}{E_1} = \left(\frac{L_1}{L_2} \right)^{2/3} \left[f_2 \left(\frac{L_2}{L_1} \right)^2 + (1 - f_2) \right], \tag{6.9}$$

for a single particle.

Since an individual particle can have any f_2 ranging from 0 to 1 at stage (2), Equation 6.9 tells us that it is possible for an individual particle to lose energy during the slow expansion followed by the fast contraction. For example, take the extreme case where $f_2 = 0$, which means just as the trap volume began to contract, the particle exchanged all of its axial energy for transverse energy. Equation 6.9 reduces to

$$\frac{E_3}{E_1} = \left(\frac{L_1}{L_2} \right)^{2/3}, \tag{6.10}$$

meaning this single particle loses energy given that $L_2 > L_1$. This expression makes sense because, if a particle has no axial energy just before the fast contraction, it cannot gain any energy during the contraction and hence the expression reduces to the adiabatic invariant for the slow expansion. However, this is the case for a single particle only so I must now take the ensemble average over N particles,

$$\begin{aligned} \frac{1}{N} \sum_i \frac{E_{3,i}}{E_{1,i}} &= \frac{1}{N} \sum_i \left(\frac{L_1}{L_2} \right)^{2/3} \left[f_{2,i} \left(\frac{L_2}{L_1} \right)^2 + (1 - f_{2,i}) \right] \\ \left\langle \frac{E_3}{E_1} \right\rangle &= \left(\frac{L_1}{L_2} \right)^{2/3} \left[\langle f_2 \rangle \left(\frac{L_2}{L_1} \right)^2 + (1 - \langle f_2 \rangle) \right]. \end{aligned} \tag{6.11}$$

where

$$\langle f_2 \rangle = \frac{1}{N} \sum_i f_{2,i}. \tag{6.12}$$

Energy mixing during the slow expansion results in stochastic changes in f_2 over time during the expansion, and hence the ensemble distribution of f_2 at (2) ranges from 0 to 1 but, since I have assumed 3D energy mixing the mean of f_2 is,

$$\langle f_2 \rangle = \frac{1}{3}, \tag{6.13}$$




Figure 6.7: Ratio between ensemble-averaged energies (green solid curve) at (3) and (1) in Figure 6.6, $\langle E_3/E_1 \rangle$, as a function of the ratio between trap lengths at (2) and (1), L_2/L_1 . I assume the trap expands only longitudinally and that energy equilibrates between all three dimensions between (2) and (3). Therefore, Equation 6.14 gives the functional form of the green curve. The blue region indicates an overall cooling achieved by the magnet manipulations shown in Figure 6.6, whereas the red region indicates overall heating. Note that $L_2/L_1 < 1$ is actually a slow contraction between (1) and (2) and a fast expansion between (2) and (3).

which leads us to our expression for the ensemble-averaged energy ratio at stage (3) to at stage (1),

$$\left\langle \frac{E_3}{E_1} \right\rangle = \left(\frac{L_1}{L_2} \right)^{2/3} \left(1 + \frac{1}{3} \left[\left(\frac{L_2}{L_1} \right)^2 - 1 \right] \right). \quad (6.14)$$

In Figure 6.7, I show the average energy ratio, $\langle E_3/E_1 \rangle$, as a function of the trap length ratio, L_2/L_1 . Note that a value of $\langle E_3/E_1 \rangle < 1$ indicates that a given length ratio will lead to overall cooling during the process of slow expansion followed by a fast contraction, whereas $\langle E_3/E_1 \rangle > 1$ indicates heating. I can interpret a value of $L_2/L_1 < 1$ as a slow contraction followed by a fast expansion. The minimum lies at $\langle E_3/E_1 \rangle = 1$ when $L_2/L_1 = 1$, and therefore I can conclude that it is not possible to achieve cooling using this method. In fact, a manipulation of the trap in this way which achieves complete 3D energy mixing during the axial expansion and no mixing during the contraction, always heats the ensemble, except if the change in trap length tends to zero, in which case there is obviously no change in energy.

It is possible to generalise this proof to an arbitrary degree of mixing between axial and transverse dimensions during the slow expansion, which equates to a generic $\langle f_2 \rangle$. For the mix category of antihydrogen trajectories introduced in Section 3.3, energy shares equally between axial and transverse dimensions and therefore $\langle f_2 \rangle \sim 0.5$. I will still assume no energy mixing during the fast contraction. A further generalisation I will apply is that I have radial expansion and contraction as well as longitudinal. I will assume expansion (and contraction) in the x and y dimensions is self-similar and that the radial expansion is some

constant factor, R , of the axial expansion such that,

$$\frac{dL_x}{L_x} = \frac{dL_y}{L_y} = R \frac{dL_z}{L_z}, \quad (6.15)$$

where L_x and L_y are the lengths of the trap in the x and y dimensions respectively. I must then derive the adiabatic invariants for the expansion (with an arbitrary degree of energy mixing) and contraction (with no mixing) in order to arrive at the overall change in energy of the ensemble. Here I will state the expression but a full derivation is given in Appendix B. The change in energy, $\langle E_3/E_1 \rangle$, as a function of R and V_2/V_1 is,

$$\left\langle \frac{E_3}{E_1} \right\rangle = \frac{\left((1 - \langle f_2 \rangle) \left(\frac{V_2}{V_1} \right)^{\frac{2R}{2R+1}} + \langle f_2 \rangle \left(\frac{V_2}{V_1} \right)^{\frac{2}{2R+1}} \right)}{\left(\frac{V_2}{V_1} \right)^{\frac{2[R(1-\langle f_2 \rangle) + \langle f_2 \rangle]}{2R+1}}}. \quad (6.16)$$

where the two terms in the numerator arise from the adiabatic invariants of the radial and axial contraction where there is no mixing (hence two independent terms), and the denominator arises from the adiabatic invariant of the expansion with mixing to a degree f_2 . In Figure 6.8 I show a contour plot of the energy ratio, $\langle E_3/E_1 \rangle$, as a function of R and V_2/V_1 for three different values of $\langle f_2 \rangle$: (a) 0.1, (b) 1/3 and (c) 0.9. For any combination of f , R or V_2/V_1 , it is not possible to achieve $\langle E_3/E_1 \rangle < 1$. Therefore, I can conclude that in a generic radial and axial trap expansion, with an arbitrary degree of mixing between axial and transverse dimensions, it is not possible to effect cooling by first expanding the trap slowly and then contracting the trap quickly (but adiabatically).

The result shows that asymmetries in the degree of energy mixing between the adiabatic expansion and the adiabatic contraction cause the ensemble-averaged energy of the antihydrogen atoms to increase. Therefore, for the aforementioned ALPHA-g proposal which involves cycling of the octupole current [103] it is important that the rate of adiabatic expansion is equal to the rate of adiabatic contraction. Even if the ensemble-averaged energy of the population remains constant, energy mixing between axial and transverse degrees of freedom is likely to increase the energy of some anti-atoms in the population (and decrease the energy of others). Those with increased energy may be lost from the magnetic trap. Therefore, the proposal is likely to benefit from minimising the degree of axial to transverse energy mixing by, for example, minimising the timescale of the adiabatic expansion (and equally the contraction). Of course, this has experimental limitations since increasing the octupole ramp rate may, itself, increase the likelihood of an octupole magnet quench.

6.2 Charge Neutrality and Adiabatic Expansion

As introduced in Section 1.5, the antihydrogen charge can be measured experimentally, and compared with the hydrogen (neutral) charge as a precise test of CPT symmetry. The most recent experimental measurement of the antihydrogen charge was made in ALPHA-2 in 2016




Figure 6.8: Ratio between ensemble-averaged energies at (3) and (1) in Figure 6.6, $\langle E_3/E_1 \rangle$, for a generic radial and axial trap expansion and an arbitrary degree of energy mixing between axial and transverse dimensions. Equation 6.16 governs the functional form of the contour plot, where I plot R and V_2/V_1 on the x and y axes respectively. The subplots represent different degrees of energy mixing where (a) $\langle f_2 \rangle = 0.1$, (b) $\langle f_2 \rangle = 1/3$ and (c) $\langle f_2 \rangle = 0.9$.

[11]. The experiment involves using stochastically varying electrostatic potentials to force any hypothetically charged antihydrogen atoms to undergo a random walk in energy and escape the magnetic confining potential. Antihydrogen atoms with lower energy can be confined by a shallower magnetic trap, which is expected to increase measurement sensitivity. This energy reduction could be achieved by first adiabatically cooling the trapped anti-atoms before lowering the trap depth and then perhaps making a more precise measurement of the antihydrogen charge. In the following sections I will describe the experimental protocol for the measurement of the antihydrogen charge that was made in 2016 [11], before demonstrating possible increases in the precision of this measurement by simulating adiabatic expansion of antihydrogen prior to the charge neutrality measurement.

6.2.1 Chaos in Magnetic Traps

In a chaotic system, small changes to initial conditions lead to large differences over time. Simulations of trajectories of antihydrogen atoms trapped in ALPHA-2 have been thought to be spatially chaotic [12] [104], but it has recently been formally established [9]. Knowledge of the chaotic nature of trajectories informs experiments that use stochastic acceleration to place bounds on the charge of antihydrogen, since these experiments require an element of randomness in the timing of the electrostatic kicks to force the hypothetically charged anti-atoms to undergo a random walk in energy. It turns out that spatial chaos in the trajectories is almost enough to provide this but, because some trajectories are periodic for short timescales (tens of milliseconds), additional stochasticity in the drive frequency is required [12]. In this section, I will give a brief overview of the chaotic nature of antihydrogen trajectories, but a full description is given in [9].

The Largest Lyapunov Exponent (LLE) is a measure of the size of the divergence of a trajectory subject to a small perturbation to its initial position. The LLE of a simulated




Figure 6.9: Largest Lyapunov Exponent (LLE (s^{-1})) as a function of anti-atom total energy, $E(K)$, calculated for ~ 3000 trajectories simulated for 100 s in the standard trap configuration (Figure 2.3(a)). The points above the red line are chaotic since they have an LLE greater than 0. Figure from [9].

trajectory is determined by perturbing its initial position by $\delta\mathbf{x}_0 = 1 \times 10^{-8}$ m in the \hat{x} -direction. Each timestep, the separation vector between the spatial position of the perturbed and unperturbed trajectory, $\delta\mathbf{x}(t)$, is determined. The perturbed trajectory deviates in magnitude and direction from the unperturbed trajectory, but at each timestep we scale down the magnitude of the perturbation to match the original size (1×10^{-8} m), but leave the direction of the perturbation to deviate. The LLE is then given by,

$$\lambda = \lim_{t \rightarrow \infty} \lim_{|\delta\mathbf{x}_0| \rightarrow 0} \frac{1}{t} \ln \frac{|\delta\mathbf{x}(t)|}{|\delta\mathbf{x}_0|}, \quad (6.17)$$

where $\lambda > 0$ is an indicator that a trajectory is chaotic. Figure 6.9 shows λ for an ensemble of trajectories with total energies ranging from 0 K to 0.5 K. The average LLE is $\lambda = 178 \pm 97 \text{s}^{-1}$ and only 4% of trajectories have $\lambda < 10 \text{s}^{-1}$, which may be compatible with zero. Recall from Chapter 3.3 that $\frac{1}{3}$ of all trajectories do not mix axial and transverse energies and therefore even the majority of no-mix trajectories are spatially chaotic. It has been confirmed that, in the non-chaotic harmonic oscillator potential, the average LLE is consistent with zero [9].

Using the method described in [9], local linearisations of the magnetic force, $\mu_B \nabla |\mathbf{B}(\mathbf{x})|$, can be used to find the *local exponential divergence*, Ψ , which involves double derivatives of the magnetic field ($|\mathbf{B}|_{yx} = \partial^2 |B(\mathbf{x})| / \partial y \partial x$, etc.) and is a measure of the growth rate of the perturbation at a particular magnetic field region. Non-zero values of Ψ indicate divergence leading to chaos. These regions occur at high radii and axially close to the confining Mirrors where the radial magnetic field is the vector sum of the radial field of the Mirror Coils and the Octupole field, generating azimuthal asymmetry. More surprisingly, there is a region of non-zero Ψ at smaller axial locations and near to zero radius. This divergent region is thought

to be formed by the quadrature sum of the axial Mirror Coil field which is peaked at $r = 0$ and the radial Octupole field, which causes the magnetic minimum to be shifted approximately 3 mm off-axis, forming a potential which is sometimes referred to as the *Mexican Hat Potential* which I showed in Figure 5.4 [9].

6.2.2 Principles of Charge Neutrality Experiments

As introduced in Section 1.5, in 2016 the ALPHA collaboration made the measurement that currently sets the most precise bound on the antihydrogen charge, of $|Q| < 0.71\text{ppb}$, where Q is the antihydrogen charge in units of the unit charge, e [11]. In this section I will describe the experimental procedure which led to this bound, before considering how the precision of the measurement may be improved. To measure the antihydrogen charge, voltages are applied to the electrodes of the ALPHA-2 Penning-Malmberg trap, which produces an electrostatic potential, ϕ , within the antihydrogen trapping region. Hypothetically charged anti-atoms in the ALPHA-2 trap experience the Lorentz force,

$$\mathbf{F} = Qe(\mathbf{E} + \mathbf{v} \times \mathbf{B}), \quad (6.18)$$

where \mathbf{v} is the antihydrogen velocity, \mathbf{B} is the trap magnetic field and $\mathbf{E} = -\nabla\phi$ is the electric field. By stochastically varying the electrostatic potential in the trap, the anti-atoms undergo a random walk in energy and will tend to escape the magnetic trap if they have sufficient charge.

During the charge neutrality experiment, the voltage on each electrode of the Penning-Malmberg trap is chosen to produce the combined vacuum potential shown in Figure 6.10. As the particle undergoes its motion in the trap, which is presumably dominated by the force on its magnetic moment from the magnetic field, we repeatedly switch the sign of the vacuum potential (Figure 6.10(a)). Figure 6.10(b) shows an example of the variation of the electrode potential on a single electrode as a function of time. The time between sign switches is varied stochastically, with mean time between sign switches, \bar{t} , and variation in switch times given by a uniform distribution with standard deviation, σ . Hypothetically charged anti-atoms undergo a random walk in energy, which will cause those with a charge Qe to escape the trap if

$$|Q| \gtrsim \frac{\Gamma}{e\Delta\Phi\sqrt{\nu}}, \quad (6.19)$$

where $\nu = 84,900$ is the total number of sign switches during the experiment (which lasts 114.9 s), $\Delta\phi \approx 100\text{ V}$ is the average change in electrostatic potential experienced by a particle between sign switches and Γ is the trap depth [11].

The parameters of the variation in the electrostatic potential as a function of time, $\bar{t} = 0.3\text{ ms}$ and normalised drive randomisation, $\sigma/\bar{t} = 0.2$, are optimised via simulation, as described in [12]. If \bar{t} is too small, the anti-atom is effectively stationary as the polarity




Figure 6.10: (a) A snapshot of the stochastic potential, ϕ (V), as a function of Cartesian position coordinate, z (mm), generated by the ALPHA-2 electrode stack on the inner surface of the electrode wall (green solid curve), on the trap axis (brown dashed curve) and at $r = 15.7$ mm (orange dashed curve). The sign of the potential shown in (a) is flipped in stochastic intervals as a function of time as shown in (b), where the sign of the electrostatic potential indicates the voltage on a single electrode. Note that (b) is just an example of the time dependence on the potential applied to a single electrode; for some types of electrode it is not possible to switch the sign of the potential as a square wave with a stochastic period (see [11] for details). Figure from [11].

switches, meaning it does not move to a significantly different region of the vacuum potential before the polarity switches back. As a consequence, the particle energy tends to a constant. At the other extreme, where \bar{t} is too large, the number of kicks is reduced and the efficiency of the experiment decreases [56]. One may wonder why it is necessary to stochastically vary the switch time i.e. why it is necessary to have $\sigma/\bar{t} > 0$, since almost all of the trajectories are chaotic as discussed in Section 6.2.1, suggesting that stochasticity in the motion of the anti-atom is naturally occurring. However, on small timescales (~ 50 ms) the trajectories are periodic [9], and hence periodic sign-switches of the potential can make the energy of the hypothetically charged anti-atoms oscillate (and not increase) on these timescales [12]. The optimal magnitude of σ/\bar{t} depends on the timescale of this periodicity and experimental sensitivity plateaus beyond $\sigma/\bar{t} = 0.2$ [12].

Once the 114.9 s of stochastic acceleration is over, the magnetic confining potential is gradually removed and the anti-atoms that have remained trapped throughout are counted by the SVD. The experimental procedure involved interleaving the experiment described above with null trials, where the electrodes were grounded. 10 stochastic and 10 null trials were performed, with each observing 12 remaining anti-atoms [11].

Simulations of the experimental procedure are used to produce the survival probability, s , as a function of antihydrogen charge, as is shown in Figure 6.11. s is defined as the fraction of antihydrogen atoms of charge $|Qe|$ surviving the entire 114.9 s stochastic polarity switching phase, relative to the number of surviving particles with $Q = 0$. Note that this experimental technique is not sensitive to the sign of the antihydrogen charge. A Bayesian analysis based on the number of detected annihilations following the stochastic and null trials allows us to deduce the survival fraction for the experiment, which leads directly to a bound of 0.59 ppb. Systematic errors and statistical errors in the Monte Carlo simulation then increase the bound on the antihydrogen charge to 0.71 ppb [11].

To understand how to improve this bound, we can refer to Equation 6.19. A simple way to increase charge sensitivity is therefore to increase ν , which simply requires running the experiment for longer. A second parameter for optimisation might be to increase the size of each kick. Unfortunately, we are already limited by the maximum voltage that can be applied to the electrodes. A third parameter for optimisation is Γ . It is possible to reduce the trap depth, but often at the expense of particle loss. However, if we adiabatically cool the population before reducing the trap depth, we can keep a high number of particles confined. In addition, compared to 2016, many more antihydrogen atoms can be trapped in ALPHA-2. Ideally the antihydrogen atoms would have energies just below the trap depth, meaning a small electrostatic kick arising from a minute antihydrogen charge will lead to ejection from the trap.

To reproduce the survival probability curve (Figure 6.11) for the experiment described




Figure 6.11: Simulated survival probability, s , as a function of hypothetical antihydrogen charge, $|Q|$, for the experimental protocol described in [11]. The orange band of varying thickness is the 1σ confidence region. Figure from [11].

above, precise details of the time dependence of the electrode potentials as a function of time are required. For example, for certain electrodes it is not possible to switch the sign of the potential in a stochastic square wave [11]. Therefore, to determine possible increases in precision by first adiabatically cooling trapped antihydrogen atoms before performing the charge neutrality experiment, we will first re-simulate the method described in the theory paper [12] which preceded the 2016 experiment [12]. The general experimental procedure is the same as that described above, but the electrostatic potential is generated by applying ± 350 V to neighboring electrodes to generate the potential shown by the black curves in Figure 6.12. The time dependence of the sign of the potential is a square wave similar to that shown in 6.10(b) (but with a 350 V amplitude). Once we establish agreement between our simulated results and those described in the theory paper [12], we will determine the achievable improvement in precision relative to the theory result, and assume a similar relative improvement in precision could be achieved when all of the experimental details in [11] are included (although the realistic experimental proposal will need to eventually be simulated).

We will simulate the charge neutrality experiment using the method described in Section 3.4, but we will modify the method to include an antihydrogen charge. We are required to add the Lorentz force term to the equation of motion of the anti-atom and to calculate the electric field that enters into the Lorentz force law.




Figure 6.12: Electrostatic potential at $r = 0$ (black solid line), at $r = 0.6R_{\text{wall}}$, where $R_{\text{wall}} \approx 2$ cm is the inner radius of the electrode wall. See [12] for a description of the orange dashed and orange dashed-dotted line as they are not relevant for the work discussed in this chapter. Figure adapted from [12].

6.2.3 Modified Symplectic Leapfrog Integrator

We must modify the symplectic Leapfrog integrator described in Section 3.4 to include antihydrogen charge. An antihydrogen atom with a charge $q = Qe$ in our magnetic minimum trap, combined with some electric field, \mathbf{E} , experiences the Lorentz force and subsequently the coupled differential equation that governs the motion of the anti-atom (Equation 3.38) becomes

$$M \frac{d\mathbf{v}}{dt} = q [\mathbf{E}(\mathbf{x}) + \mathbf{v} \times \mathbf{B}(\mathbf{x})] - \mu_B \nabla |\mathbf{B}(\mathbf{x})|, \quad (6.20)$$

where M is the antihydrogen mass and μ_B is the Bohr magneton. Finding the solution to the anti-atom motion is complicated by the fact that the velocity is present on the right hand side of Equation 6.20, a complexity that was not present in the charge neutral antihydrogen simulations presented in Section 3.4. The Leapfrog algorithm evaluates the force on the particle half a timestep before the velocity to step the velocity forward (Equation 3.39). The algorithm transforms into,

$$\begin{aligned} \mathbf{x}_{m+1/2} &= \mathbf{x}_{m-1/2} + \mathbf{v}_m dt \\ \mathbf{v}_{m+1} &= \mathbf{v}_m + \left\{ \frac{q}{M} [\mathbf{E}(\mathbf{x}_{m+1/2}) + \mathbf{v}_{m+1/2} \times \mathbf{B}(\mathbf{x}_{m+1/2})] - \frac{\mu_B}{M} \nabla |\mathbf{B}(\mathbf{x}_{m+1/2})| \right\} dt, \end{aligned} \quad (6.21)$$

but $\mathbf{v}_{m+1/2}$ cannot be directly evaluated. Therefore, we replace it by the average of the velocities half a timestep before and after it i.e. we set,

$$\mathbf{v}_{m+1/2} = \frac{\mathbf{v}_m + \mathbf{v}_{m+1}}{2}, \quad (6.22)$$

which can be simplified by defining,

$$\boldsymbol{\Sigma} = \frac{q\mathbf{E}(\mathbf{x}_{m+1/2})}{M} - \frac{\mu_B}{M} \nabla |\mathbf{B}(\mathbf{x}_{m+1/2})|, \quad (6.23)$$

and

$$\boldsymbol{\Omega} = \frac{q\mathbf{B}(\mathbf{x}_{m+1/2})}{M}. \quad (6.24)$$




Figure 6.13: Sketch of the configuration used to calculate the electrostatic potential generated by the n^{th} electrode raised to 1 V; with all other electrodes grounded. The inner radius of the electrode is r_w and the length of the n^{th} electrode is L_n . Figure from [13].

Equation 6.21 becomes,

$$\begin{aligned} \mathbf{x}_{m+1/2} &= \mathbf{x}_{m-1/2} + \mathbf{v}_m dt \\ \mathbf{v}_{m+1} &= \mathbf{v}_m + \left\{ \boldsymbol{\Sigma} + \left(\frac{\mathbf{v}_m + \mathbf{v}_{m+1}}{2} \right) \times \boldsymbol{\Omega} \right\} dt, \end{aligned} \quad (6.25)$$

and it can be shown that,

$$\mathbf{v}_{m+1} = \frac{\mathbf{C} + \mathbf{A}\mathbf{A} \cdot \mathbf{C} - \mathbf{A} \times \mathbf{C}}{1 + A^2}, \quad (6.26)$$

where $\mathbf{A} = \boldsymbol{\Omega} \frac{dt}{2}$ and $\mathbf{C} = \mathbf{v}_m + (\boldsymbol{\Sigma} + \mathbf{v}_m \times \boldsymbol{\Omega}/2) dt$ [105]. Hence, we have arrived at the modified second-order symplectic Leapfrog integrator.

6.2.4 Solving for the Electrostatic Field

We are required to calculate $\mathbf{E}(\mathbf{x})$ at all \mathbf{x} within the trap, given the ± 350 V voltage on each electrode in the Penning-Malmberg trap. For the following calculation, we can ignore space charge effects on the potential because the antihydrogen charge is known to be small and the density of trapped antihydrogen is low.

We must solve Poisson's equation for the electrostatic potential, with the voltage on the inner radius of each electrode creating a boundary condition. To avoid recalculating the electrostatic potential each time the voltages are altered, we will calculate the electrostatic potential from each electrode, raised to 1 V; with all other electrodes grounded, as shown in Figure 6.13. Poisson's equation for the potential, ϕ_n , from the n^{th} electrode becomes,

$$\begin{cases} \nabla^2 \phi_n = 0 \\ \phi_n(r = r_n, z) = \begin{cases} 1 & \text{for } z \in [-L_n/2, L_n/2] \\ 0 & \text{otherwise} \end{cases} \end{cases} \quad (6.27)$$

where r_n and L_n are the radius and length of the n^{th} electrode respectively [13]. The total potential, ϕ , for a set of N electrode voltages $\{V_n\}$ is then,

$$\phi(r, z) = \sum_{n=1}^N V_n \phi_n(r, z - z_n), \quad (6.28)$$

where z_n is the axial centre of the n^{th} electrode. The solution can be found using separation of variables [97],

$$\phi_n(r, z) = \frac{2}{\pi} \int_0^\infty dk \frac{I_0(2kr/L_n)}{I_0(2kr_n/L_n)} \text{sinc}(k) \cos\left(\frac{2kz}{L_n}\right) \quad (6.29)$$

where $I_0(x)$ is the zeroth-order modified Bessel function of the first kind. Equation 6.29 can be integrated numerically using the simple trapezium rule [13]. This method allows the potential, and hence the electric field using $\mathbf{E} = -\nabla\phi$, to be found at the coordinates of the particle position at runtime during the simulation.

Unfortunately, we find this Bessel function solution is too slow to compute at runtime in simulations of the charge neutrality of antihydrogen [106]. As an alternative, the electrostatic potential, or even the electric field can be calculated on a grid in advance of running the simulation. At runtime, the corresponding quantity can be found at the particle location via interpolation. This opens up the possibility to use codes like COMSOL Multiphysics[®] (COMSOL): a software package that can solve for the electrostatic potential or electric field of the Penning-Malmberg electrodes on an arbitrary spatial grid, which makes it simpler to model details of the electrodes.

We use COMSOL to solve for the electric field within the trap when one electrode is at 1 V and the rest of the electrodes are grounded (as described above, Figure 6.13). Since the central electrodes are the same length, the solution for one electrode is approximately a translation of the other (inexact because of differing boundary conditions), which reduces the amount of data that needs to be stored and calculated. The two outer electrodes are longer (Figure 2.1), and therefore require a separate solution. By summing the electric field generated by each electrode, scaled by the voltage it is raised to (± 350 V, Equation 6.28), we arrive at the solution for the electric field within the trap volume. We find that the most accurate and fastest method is to calculate the electric field (rather than the electrostatic potential) in COMSOL, since it avoids calculating $\nabla\phi$ at simulation runtime, although this means all three (x, y, z) electric field components must be saved and imported into the simulation [106]. Due to the azimuthal symmetry of the solution, we only record the electric field solution at radial and axial coordinate locations (r, z) . We use a grid of 20 radial points \times 20 axial points per electrode, though simulation results converge with a grid spacing $\lesssim 5$ (radial and axial) divisions per electrode. During the simulation, linear interpolation is used to obtain the electric field at the particle position (we find it is faster and more accurate than quadratic interpolation) [106].

6.2.5 Charge Neutrality of Adiabatically Cooled Antihydrogen

In this section we compare results of simulations of the (theoretical) charge neutrality experiment with and without adiabatic expansion. We refer to the simulation in which no

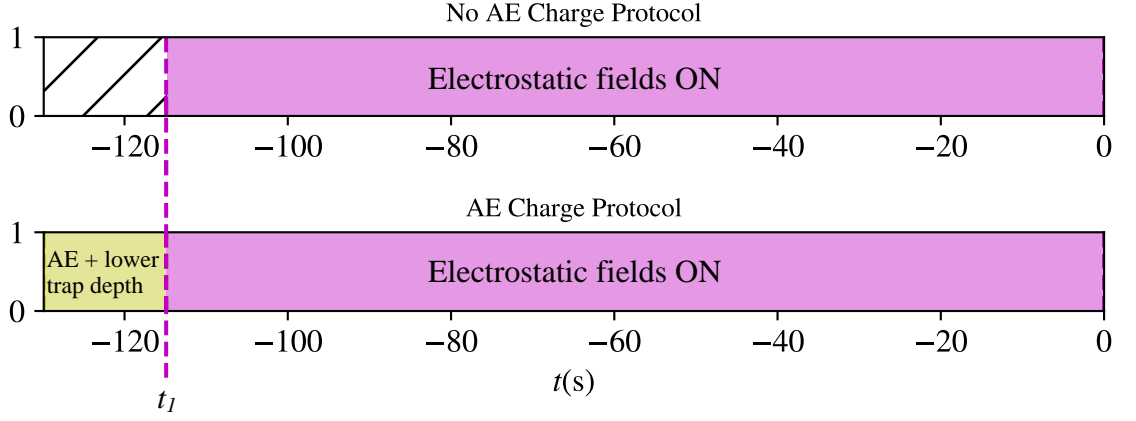


Figure 6.14: Definition of “No AE Charge Protocol” and “AE Charge Protocol”. Simulation time, $t(s)$, is along the x -axis relative to the start of magnetic ramp down (at $t = 0$). A hashed region indicates that the simulation has not started. A pink region indicates a period of stochastic electrode switching with $\sigma/\bar{t} = 0.2$ and \bar{t} given by that in Table 6.1. The yellow region indicates a period of axial adiabatic expansion and trap depth reduction in which the Mirror Coil and Octupole currents are reduced to those given in Table 6.1. The on-axis magnetic potential as a function of time during the ‘AE + lower trap depth’ phase is shown in Figure 6.15. The dotted magenta line at t_1 indicates the start of the electrostatic field switching phase.

adiabatic expansion takes place as “No AE Charge Protocol”, in which we simulate the same procedure described in the theory paper which preceded the most recent experimental bound on the antihydrogen charge [12], where neighboring Penning-Malmberg trap electrodes are at ± 350 V (Figure 6.12), with its sign switched as a stochastic square wave with $\sigma/\bar{t} = 0.2$ and $\bar{t} = 0.3$ ms (see [12] for details of the optimisation of these parameters). The electrostatic fields are switched for a total duration of 114.9s, as indicated by the pink region in Figure 6.14.

We compare the No AE Charge Protocol to the “AE Charge Protocol”, in which we simulate a phase of adiabatic expansion and lowering of the trap depth which takes 15 s (yellow region, Figure 6.14) prior to a 114.9 s electrostatic field switching phase. In Figure 6.15 we show the on-axis potential as a function of time during the AE Charge Protocol. Also shown is the Octupole current, I_{oct} , which decreases linearly from 900.1 A to 200.0 A from $t = -119.3$ s to -116.9 s. Note that in the AE Charge Protocol, the electrostatic fields are off until the expansion and trap depth lowering has finished. Since the adiabatically cooled trajectories will lose energy (and therefore traverse less potential in the same time period) we re-optimize \bar{t} using the same method as in [12], where \bar{t} is varied for a single charge and the minimum s indicates the optimal \bar{t} . Optimal values of \bar{t} are compared in Table 6.1.




Figure 6.15: Snapshots of on-axis potential, U_z , as a function of axial location, z , at different times, t , during the AE Charge Protocol (see Figure 6.14). Time increments in 3 s intervals from (a)-(f). I_{oct} is the Octupole current at time t .

	$I_{\text{mirr}}(t_1)$ (A)	$I_{\text{oct}}(t_1)$ (A)	Γ (K)	Γ_z (K)	$\langle E(t_1) \rangle \pm 1\sigma_{\text{SE}}$ (K)	\bar{t} (ms)	$f[Q = 0, t = 0]$
No AE	606.0	885.7	0.52	0.61	0.342 ± 0.005	0.30	0.95
AE	80.0	200.0	0.038	0.081	0.038 ± 0.001	0.95	0.076

Table 6.1: Comparison of simulation parameters for “No AE Charge Protocol” and “AE Charge Protocol” as defined in Figure 6.14. $I_{\text{mirr}}(t_1)$ and $I_{\text{oct}}(t_1)$ are the currents in Mirrors A and E at t_1 (as indicated by the dashed magenta line in Figure 6.14). Γ is the magnetic trap depth (defined as the lowest potential on the electrode surface). Γ_z is the minimum-maximum axial confining potential of the magnetic trap (the potential that must be overcome for an anti-atom to escape axially over the confining Mirror Coil). $\langle E(t_1) \rangle$ is the ensemble-averaged energy at t_1 , excluding particles that annihilate at $t < t_1$. $\sigma_{\text{SE}} = \frac{\sigma}{\sqrt{N}}$ is the standard error on the mean, where σ is the sample standard deviation and N is the number of simulated particles surviving until $t = t_1$ for $Q = 0$.

Once the adiabatic expansion is complete, the magnetic trap depth can be lowered by either reducing the current in Mirror Coils A and E or by reducing the Octupole current. When the Octupole and Mirror Coil currents are reduced, the trap volume expands and the trap depth reduces. The trap expands radially and axially on account of the reduction in Octupole and Mirror Coil currents respectively, which means the particles undergo both radial and axial adiabatic cooling to some extent. The trap depth reduction results in loss of relatively high energy particles, and hence the mean energy of the distribution also decreases due to this effect.

Recall that in Section 4.8, in studies of Octupole FRD we found that antihydrogen trajectories escape the trap through minima in radial potentials, $\Gamma_r(\phi, z, t)$ out at high axial locations (Figure 4.16). Our observation tells us that a particle’s ease of escape is dependent on its specific axial and radial energy components, as well as the magnetic potential in the axial and radial dimensions. Therefore, the optimal relative Mirror Coil and Octupole currents to reduce to after adiabatic expansion is likely to depend on the distribution of axial and transverse energies. In addition, particle loss during evaporative cooling sets a lower limit to which the trap depth can be reduced, since enough particles need to be trapped during the electrostatic field switching phase to have high experimental statistics. For all of the reasons raised above, the trap depth reduction will need to be optimised in detail. For this study we have not yet performed this optimisation and therefore the results presented here are simply a first demonstration of the increase in the bound that could be achieved by first adiabatically cooling the population.

The difference in Mirror Coil and Octupole currents just prior to the electrostatic field switching phase (at t_1) is compared for the No AE and AE Charge Protocols in Table 6.1. We also show the trap depth, Γ (lowest potential on electrode wall), as well as the minimax of the




Figure 6.16: (Total) Energy distributions at the start of the electrostatic field switching phase, $E(t_1)(K)$, for the “No AE Charge Protocol” (black histogram) and “AE Charge Protocol” (blue histogram). The total histogram area is normalised to 1 in both cases.

Mirror Coil potential, Γ_z , which gives an idea of the magnetic potential depth corresponding to the Mirror Coil current.

In both protocols, we sample the energy of the simulated trajectories at t_1 . For the AE Charge Protocol, this equates to sampling the energy after the adiabatic expansion and trap depth reduction phase. Figure 6.16 shows a histogram of the energy of each trajectory t_1 (ignoring trajectories that annihilate at $t < t_1$). The corresponding ensemble-averaged energies, $\langle E(t_1) \rangle$, are compared in Table 6.1 and show that the non-adiabatically-cooled distribution is 9.0 ± 0.3 times more energetic than the adiabatically cooled distribution.

In Table 6.1 we show the fraction of trajectories, f , that survive until $t = 0$ relative to the initial number of confined trajectories for $Q = 0$. For example, if we experimentally trap 1000 trajectories in the short initial adiabatic expansion well (and antihydrogen is charge neutral), and perform the AE Charge Protocol experiment, we expect to detect 76 annihilations (excluding background) during FRD. The number of trajectories that can be trapped in the short initial adiabatic expansion well is a parameter that will require experimental analysis, but the fact that we observed a similar number of annihilation counts during Adiabatic Expansion Trial A, Control Trial B and Control Trial C (see Table 2.1), and even more counts in Trials A and B than in Trial C, is a good sign for our ability to trap large numbers of antihydrogen atoms in the short magnetic volume. For the No AE Charge Protocol we simulate 1000 trajectories per charge (and hence 950 survive until $t = 0$ for $Q = 0$), and for the AE Charge Protocol we simulate 30,000 trajectories (and 2280 survive until $t = 0$ for $Q = 0$).

In Figure 6.17 we show the survival probability, s , (relative to the number surviving at




Figure 6.17: Survival probability, s , (at $t = 0$) as a function of normalised anti-atom charge, $|Q|$, (antihydrogen charge = Qe), relative to $Q = 0$. We compare the “No AE Charge Protocol” (black curve) to the “AE Charge Protocol” (blue curve). Note that a *lower* s at a given Q indicates an *increased* sensitivity to the hypothetical antihydrogen charge. The dotted red line indicates the observed survival fraction given that $Q = 0$.

$t = 0$ for $Q = 0$) as a function of Q . The No AE Charge Protocol s (black curve) is the same as that presented in [12] within statistical uncertainty. The AE Charge Protocol s (blue curve) suggests an achievable factor of ~ 10 increase in precision by first adiabatically cooling trajectories before decreasing the depth of the magnetic confining potential to those shown in Table 6.1 and by simulating the same duration of stochastic electrostatic field switching as in the most recent experimental bound [11].

Further increases in precision are expected to come from optimisation and from increasing the duration of the electrostatic potential switching phase. Increasing the duration is made possible by long antihydrogen confinement times [69], and results in a linear increase in the number of electrostatic kicks given to the hypothetically charged anti-atoms. Equation 6.19 then suggests a square-root increase on the bound on Q . We have begun to investigate the optimal trap depth reduction via simulation, which suggests, interestingly, that the most precise bound on Q may be achieved by reducing only the Octupole current (and leaving the Mirror Coil current high, see [106] for details).

6.3 Compatibility of Adiabatic Cooling and 1S-2S Spectroscopy

In Section 1.3.3 I introduced ALPHA’s high precision 1S-2S spectroscopy experiment, in which the 1S-2S transition in antihydrogen can be excited by two counter-propagating photons in




Figure 6.18: Diagram of the experimental setup for 1S-2S spectroscopy and laser cooling of antihydrogen. 121.6 nm laser light for antihydrogen laser cooling enters the trap at 2.3° to the trap axis, indicated by a purple beam. The 243.1 nm laser cavity for 1S-2S spectroscopy also at 2.3° to the trap axis is shown as a pink beam. Figure from [7], see [7] for a detailed description.

measurements of the transition for precise comparison with hydrogen as a fundamental test of CPT symmetry. As described in Section 1.3.3, this two-photon transition is free of first-order Doppler broadening and instead the dominating effect that broadens the observed linewidth is transit-time broadening which contributes 50 kHz full-width at half maximum and is sensitive to the velocity of the anti-atom perpendicular to the laser beam [56].

Narrowing of the 1S-2S spectral linewidth has already been observed by laser cooling the trapped antihydrogen atoms in the flat/spectroscopy magnetic trap prior to performing the 1S-2S experiment [7], suggesting the prospect of an improved measurement of the 1S-2S transition frequency in the near future. Since the lack of 3D trap access and details of the degree of energy mixing between axial and transverse dimensions limit the mean achievable energy via laser cooling to ~ 20 mK [102], one could consider first laser cooling in the shorter magnetic trap (tight configuration, Figure 2.3) before expanding the trap volume adiabatically into the flat/spectroscopy magnetic trap, making an extremely high precision measurement of the 1S-2S transition in trapped antihydrogen possible. Ultimately, it will be necessary to optimise the combination of laser cooling and adiabatic cooling for 1S-2S spectroscopy. In this section, however, I focus on simulations of (only) adiabatically cooled antihydrogen in the flat/spectroscopy trap.

The spatial details of orbits occupied by adiabatically cooled trajectories are of interest because the 1S-2S spectroscopy laser occupies a relatively small volume of the trap, and it is critical for the measurement that the antihydrogen atoms traverse this volume. In this section, I will simulate adiabatic expansion of antihydrogen, followed by a hold in the magnetic trap used for spectroscopy experiments (flat trap configuration, Figure 2.3(c)). During the hold I will determine the time spent by the trajectories within the region of the 1S-2S spectroscopy laser beam.

In Figure 6.18 I show the experimental setup for 1S-2S spectroscopy experiments at

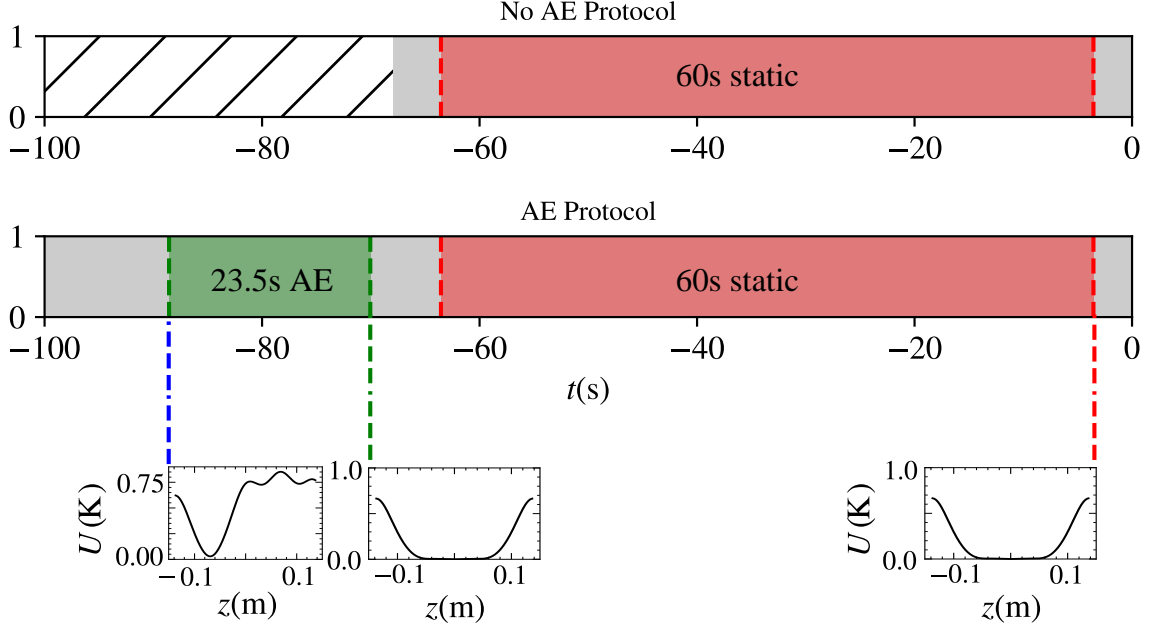


Figure 6.19: Definition of ‘No AE Protocol’ and ‘AE Protocol’. Each protocol refers to a specific set of magnetic field manipulations as indicated by the plots of magnetic potential, $U(K)$, at the bottom of the figure. The vertical dashed lines trace down to graphs of $U(K)$ which represent the status of the potential at that simulation time. Time, $t(s)$, is along the x -axis, relative to the start of magnetic ramp-down. A hatched region indicates that the simulation has not started at that time. The red region indicates a 60 s hold in the flat trap configuration, and the green region indicates a 23.5 s adiabatic expansion (labelled “AE”) of the same duration as Trial A, as defined in Chapter 2, but ending in the flat trap (Figure 2.3(c)) rather than the standard trap (Figure 2.3(a)).

ALPHA; including the 243.1 nm 1S-2S spectroscopy beam, obtained by twice frequency doubling light from a 972 nm diode laser, which enters the trap off-axis at 2.3° to the trap axis to enable on-axis plasma preparation. The laser beam has a waist, $\omega_0 = 200 \mu\text{m}$ and a linewidth of 100 MHz [56].

We model the laser beam region as a Gaussian beam, as shown in Figure 6.20, with radial extent as a function of z defined as,

$$\omega(z_{\text{laser}}) = \omega_0 \sqrt{1 + \left(\frac{z_{\text{laser}}}{z_R} \right)^2}, \quad (6.30)$$

where $z_R = \frac{\pi\omega_0^2}{\lambda}$ is the Rayleigh range and $\lambda = 243.1 \text{ nm}$ is the wavelength of the laser. Since the laser beam enters the trap at an angle of 2.3° , the central axis of the laser beam, z_{laser} can be expressed in terms of the axial trap coordinate, z , as

$$z_{\text{laser}} = \frac{z}{\cos 2.3^\circ}. \quad (6.31)$$

Although the laser beam waist is $200 \mu\text{m}$, simulations have shown that ionisations via absorption of a further 243 nm photon occur at locations that are smeared around the laser




Figure 6.20: Gaussian beam profile. ω_0 is the waist radius of the Gaussian beam, $\omega(z_{\text{laser}})$ is the radius of the beam at a given axial coordinate, z_{laser} , and z_R is the Rayleigh range. $\omega(z_{\text{laser}})$ is defined as the radial location at which the intensity of the beam drops to $1/e$ of the intensity at zero radius. Figure modified from [14].

beam within a wider region of around 1 mm [107]. Therefore, for simulations discussed here I use $\omega_0 = 1$ mm. Unless stated otherwise, simulations involve 10,000 antihydrogen trajectories, initialised and simulated using the same procedure described in Section 3.4.

We simulate the same (~ 23.5 s duration) adiabatic expansion procedure (Trial A) described in Chapter 2, of which simulation results are presented in Chapter 4. But here, rather than Octupole FRD taking place 50 ms after the expansion has finished, I instead hold the adiabatically cooled antihydrogen atoms in the flat trap configuration for a further 70 s, which I will refer to as the *AE Protocol* and is defined in Figure 6.19. Note that this trap expansion is predominantly axial and therefore adiabatic cooling mainly reduces the axial particle energy (as I showed in Figure 4.9). Since the precision of antihydrogen 1S-2S spectroscopy experiments is limited by the antihydrogen velocity perpendicular to the laser beam (Equation 1.10), 1S-2S spectroscopy experiments following axial adiabatic expansion protocols rely on energy mixing between axial and transverse dimensions to reduce the energy transverse to the trap axis (which is almost perpendicular to the 1S-2S spectroscopy laser beam). Although Trial A simulations in this thesis have shown small decreases in transverse energy during a 23.5 s duration axial expansion, energy mixing is known to be more significant in the flat trap configuration (see Figure 3.10).

During the hold in the static flat trap, over a period of 60 s (labelled ‘60 s Static’) I track the motion of each simulated trajectory, determine the time spent by the trajectory within the Gaussian beam region of waist 1 mm. I compare this simulation to what I will refer to as the *No AE Protocol*: a static hold in the flat trap configuration for 70 s in which I determine the




Figure 6.21: Total time $\hat{t}(s)$ spent by $\sim 10,000$ simulated antihydrogen trajectories within the region of the 1S-2S laser beam during the 60 s hold of the (a) No AE Protocol and (b) AE Protocol. Both protocols are defined in Figure 6.19.

length of time spent by each particle in the period from -63.55 s to -3.55 s. Quasi-trapped particles are ignored in each simulation. For the AE Protocol, 14.5 % of trajectories are lost before ramp down, compared to 11.8 % of trajectories lost before ramp down for the No AE Protocol.

In Figure 6.21, I show a histogram of the time, $\hat{t}(s)$, spent by each trajectory in the beam region during the 60 s hold in the flat trap configuration for the No AE Protocol (a) and AE Protocol (b) populations. Note the bin at $\hat{t} = 0$ s in each case: this bin indicates the number of particles that interact very briefly (or not at all) with the 1S-2S spectroscopy laser during the 60 s static hold. In fact, for the AE Protocol 23.6 % of trajectories did not enter the spatial region of the beam at all in 60 s, which is compared to only 2.1 % of the No AE Protocol particles which did not enter the spatial region of the laser beam in 60 s. Excluding particles with $\hat{t} = 0$ s, the average time spent within the laser beam for No AE Protocol particles is 0.26 s, compared to 0.28 s for AE Protocol particles. Therefore, the reduction in kinetic energy due to adiabatic expansion does seem to increase the time spent within the laser beam slightly for this subset of trajectories.

To further diagnose the trajectories, I determine the radial distribution of the simulated antihydrogen orbits. To do this, I split the x and y axes, each ranging from -3.0 mm to 3.0 mm into 1000 bins, forming a grid of 10^6 square bins. Before simulating any trajectories, each bin is initialised to 0, then at each timestep I determine which bin in the xy -plane the particle is in. For a simulation of 10,000 trajectories for the 60 s static hold, I trace out the




Figure 6.22: Time-summed xy -position, $\Sigma(s)$, shown as the greyscale axis is scaled to the minimum and maximum time spent in any single bin for No AE Protocol (a) and AE Protocol (b). (c) and (d) show the same data as (a) and (b) respectively but the greyscale axis is modified to make the features more prominent. A total of 10^6 square bins lie within a square grid of 3.0×3.0 mm. Σ is the total time spent by the $\sim 10,000$ trajectories in each square bin during the 60s hold. No AE Protocol and AE Protocol are defined in Figure 6.19

total time spent by the 10,000 trajectories in each xy -bin (flattened in z), which I define as the time-summed xy -position, $\Sigma(x, y)$.

Figure 6.22 shows $\Sigma(x, y)$ for the 60s static hold of the No AE Protocol (a) and AE Protocol (b) simulated populations, scaled to the minimum and maximum time spent by any particle in any one bin. Figures 6.22 (c) and (d) are the same simulated data as in (a) and (b) respectively but with the greyscale axis scaled to highlight differences in $\Sigma(x, y)$ between the No AE Protocol (c) and AE Protocol (d) simulated populations. The key feature of the adiabatically cooled trajectories (AE Protocol) is that they spend significantly less time on-axis than the non-adiabatically cooled populations, whose $\Sigma(x, y)$ is peaked at zero radius. For the adiabatically cooled population, rings appear, indicating that the trajectories spend more time at distinct radii.

The trajectories which tend to avoid regions of the trap in which they could interact with the 1S-2S laser could be anti-atoms whose energy is mainly azimuthal, meaning they have both small axial and radial kinetic energy components and hence orbit the laser beam at some roughly fixed radius. Since this longitudinal adiabatic expansion mainly reduces the axial energy of the trajectories, one could imagine this effectively increases the number of these trajectories that are present. In addition, trajectories with low axial energy components relative to their total energy ($\langle\langle\varepsilon_{\parallel}\rangle\rangle < 0.1$, see Section 3.3 for a full discussion) tend not to mix axial and transverse energy components and could be the reason these orbits do not get kicked onto an alternative trajectory which crosses the laser beam over long periods of time.

One might wonder whether the trajectories that do not cross the laser beam within 60s would cross the beam if held in the flat trap for a longer period of time: a simulation that I have not yet performed. Recall from Figure 3.10 that $\sim 80\%$ of trajectories are expected to have mixed a significant portion of axial and transverse energy during a hold in the flat trap; which increases to $\sim 90\%$ after 100s, but it takes around 1000s for almost all of the trajectories to mix energies [9]. Therefore, simulating a 100s hold in the flat trap could gain significant insight into whether the trajectories that do not enter the beam region within 60s would mix energies in the remaining 40s and be kicked into the beam region. Beyond that, it would likely take a ~ 1000 s simulation to gain further insight. I should also remark that energy mixing simulations used to determine these mixing timescales were conducted with non-adiabatically cooled populations, which have not only a significantly different magnitude of energy, but also a significantly reduced ratio of axial and transverse energy. Equivalent mixing timescales have not yet been determined for adiabatically cooled antihydrogen.

Although I have mainly explored axial adiabatic expansion in this section, the 1S-2S experiment (being dominated by transit-time-broadening) lends itself to radial adiabatic expansion cooling, achieved by slow (compared to the radial bounce time of the trajectories, ~ 1 ms) reduction of the Octupole current. One might also imagine that a combined axial and

radial adiabatic expansion could reduce the number of particles that do not enter the region of the 1S-2S laser beam within 60 s, since the reduction in radial energy could increase the normalised axial energy ($\varepsilon_{\parallel} = E_{\parallel}/(E_{\perp} + E_{\parallel})$), and hence could cause more particles to mix (see Section 3.3 for more detail).

I have simulated a modified AE Protocol in which there is an additional radial adiabatic expansion in which the Octupole current is reduced from 900.1 A at $t = -78.55$ s to 300 A at $t = -70.05$ s (an 8.5 s radial expansion mid way through the 23.5 s axial adiabatic expansion). Although this simulation resulted in 84.4% of particles being lost (largely evaporated during the radial adiabatic expansion), the number of particles that did not enter the beam region within the 60 s static hold remained high at 23.1%. A possible explanation for this is that if the non-interacting trajectories have large azimuthal but small radial energy components their energy will remain largely unchanged during the radial adiabatic expansion. Although this is not promising for combined axial and radial adiabatic protocols, I have not simulated a radial adiabatic expansion without an axial adiabatic expansion. In this case the radial energy reduction would increase ε_{\parallel} and likely increase the degree of energy mixing, which may prove a suitable protocol for use in 1S-2S spectroscopic measurements.

Conclusion

In this thesis I have analysed the dynamics of antihydrogen trajectories confined in the ALPHA-2 magnetic minimum trap, when the trap volume expands slowly compared to the speed of trapped anti-atoms. Throughout, I have tested the hypothesis that trapped antihydrogen atoms lose energy as the volume of their magnetic confining potential slowly increases.

In ALPHA, the fundamental properties of trapped antihydrogen are measured and compared to those of hydrogen in the search for matter-antimatter asymmetries. Increases in experimental precision are required to set ever tighter bounds on the possible differences between hydrogen and antihydrogen. In Chapter 1, I introduced ALPHA's spectroscopy experiments which probe the energy levels of trapped antihydrogen as a stringent test of charge-parity-time (CPT) invariance. I stated that the current dominant uncertainty in measurements of the 1S-2P transition [57] and in measurements of the fine structure of antihydrogen [58] is Doppler broadening, which is approximately linearly proportional to the axial energy of trapped antihydrogen. The most precise comparison of hydrogen and antihydrogen is currently set by measurements of the 1S-2S transition [56], with uncertainty dominated by transit time broadening which is approximately linearly proportional to the transverse energy of trapped antihydrogen. I also described ALPHA's most recent experiment - ALPHA-g - which intends to measure the effect of the Earth's gravitational field on trapped antihydrogen [59] as a test of Einstein's Weak Equivalence Principle (WEP). The experiment relies on adiabatic cooling, since lower energy antihydrogen atoms can be confined by a shallower magnetic trap that is more sensitive to the gravitational potential. I also introduced experiments that measure the antihydrogen charge as a fundamental test of CPT invariance, by stochastic acceleration of hypothetically charged anti-atoms [11]. I suggested sensitivity to the antihydrogen charge may be increased by reducing the trap depth. By adiabatically cooling prior to trap depth reduction, larger numbers of antihydrogen atoms can be retained.

In Chapter 2 I described a proof-of-principle adiabatic cooling experiment that took place

in ALPHA-2 in 2016. The experiment involved formation and trapping of antihydrogen in a relatively small magnetic volume, before the volume of the magnetic trap was expanded predominantly in the axial direction in ~ 23 s. After the slow expansion, the magnetic trap was gradually removed in ~ 1.5 s by linearly ramping down the Octupole field, during which time antihydrogen annihilations with internal trap structures were detected. The experiment relied on the assumption that antihydrogen atoms with *lower* energy annihilate when the magnetic trap is *shallower* which happens *later* in time during linear magnetic ramp-down. In the experiment, annihilation times during magnetic ramp-down of antihydrogen atoms subject to a slowly expanding trap volume (Trial A) were compared to control samples held in static traps in the small (Trial B) and large (Trial C) magnetic volumes. I presented the experimental data, which showed that on-average the annihilation times of Trial A were 1.85 ± 0.08 and 1.35 ± 0.06 times *later* (relative to magnetic ramp-down) than Trials B and C respectively. I used a simple (and naive) method to estimate the energy corresponding to the annihilation times, by calculating the average trap depth at the annihilation time, $\langle \Gamma(t_a) \rangle$. I found that the naive method suggested the mean energy of the adiabatically cooled population was 2.75 ± 0.72 times lower than Trial B and 2.14 ± 0.49 times lower than Trial C.

In Chapter 3 I described the theory required to analyse the dynamics of adiabatically cooled antihydrogen. I described the principles of antihydrogen trapping by the force on its magnetic moment in a magnetic field, before deriving a quantity that is conserved to very good approximation when the magnetic fields are slowly time-varying, known as an adiabatic invariant. I introduced the concept of energy mixing between axial and transverse dimensions, as its extent controls the transverse energy reduction during a predominantly axial trap expansion. I described a Monte Carlo simulation of trapped antihydrogen in ALPHA-2, in which the position and velocity of trajectories is pseudo-randomly initialised and the motion of the trapped anti-atoms is solved for, at fixed timestep intervals, using a symplectic Leapfrog integrator.

In Chapter 4, I presented results of simulations of the proof-of-principle experiment described in Chapter 2. I began by comparing experimental and simulated annihilation times in detail using the Kolmogorov Smirnov (KS) Two-Sample Test. I described the KS p-value: a measure of the confidence level at which I can conclude experimental and simulated annihilation times are in agreement, and found KS p values of 1.3×10^{-4} , 2.1×10^{-3} and 0.137 for Trials A, B and C respectively. The low agreement between simulation and data in Trials A and B led me to offset the experimental annihilation times with respect to the simulation, and recalculate the KS p value, to test whether the results were consistent with a constant temporal offset. In doing so, I found that the experimental data in all three trials was consistent with an offset ranging from approximately -50 ms (3%) to -100 ms (6%) with respect to the simulation.

I considered possible explanations for the temporal offset, such as software communication delays, uncertainties in the location of antihydrogen formation and magnetic field errors. One effect that was particularly significant came from nonlinearities in the Octupole current decay as controlled by its power supply. A few days after the experiment, each experimental trial was repeated without particles to obtain a high resolution Octupole current trace during ramp down. With this ramp down modelled in the simulation, the KS p -values became 0.709, 0.181 and 0.095 for Trials A, B and C respectively (recall that the Trial C high resolution DCCT trace was not recorded and therefore I used the Trial A trace for this analysis), indicating agreement between simulation and experiment. Since the high resolution trace was recorded a few days after the experiment, I cannot rule out changes in the current decay on account of laboratory temperature fluctuations and that the Trial C experimental Octupole current decay differed in some way from that of Trial A.

Although the KS analysis suggested agreement between simulation and data when the high resolution Octupole current trace was used, I found the simulation results for all three trials were consistent with an offset ranging from 0 ms to 60 ms (4 %, opposite direction). Given the timing resolution with which the currents in the superconducting magnets in the experiment itself were recorded, and the fact that a detailed study of the relative experimental timings was not carried out, I concluded reasonable agreement between simulated and experimental annihilation times.

Given that I established reasonable agreement between experimental and simulated annihilation times, I extracted the energy of the simulated distributions, which showed that adiabatic cooling resulted in a final mean total energy of 0.2226 ± 0.0002 K, which equates to a 37.9 ± 0.1 % mean total energy decrease during the trap expansion, and can be broken down into a 68.1 ± 0.2 % decrease in mean parallel energy and a 14.2 ± 0.2 % decrease in transverse energy. By comparison to the control trials, I showed that the mean final energy of the control populations, Trials B and C, were 1.616 ± 0.002 and 1.558 ± 0.002 times that of the adiabatically cooled (Trial A) simulated population respectively.

I then estimated an uncertainty on the simulated energy distributions based on the range of possible temporal offsets between experimental and simulated annihilation times. This uncertainty models the worst case scenario: that the observed temporal offset is indicative of a difference in energy between simulation and experiment. It is more likely that the temporal shift reflects inaccuracies in our models of the Octupole current during FRD, which is backed up by the agreement between simulation and data when the Octupole current decay follows that of the high resolution DCCT trace during FRD. To estimate the energy uncertainty, I calculated the trap depth at each simulated annihilation time for the range of offsets with KS p value > 0.05 , and determined the difference in energy compared to the result with zero offset. The method allowed me to place a lower bound on the ratio of the final energy in Trial

B (C) to that of Trial A of 1.47 (1.43).

I then used an established technique to determine the energy distribution of the (experimentally) trapped antihydrogen atoms, using the simulated annihilation times and energies. The technique gave a mean final energy of the Trial A experimental antihydrogen distribution of 0.183 ± 0.009 . The corresponding Trial B and Trial C mean final reconstructed energies were 1.81 ± 0.10 and 1.78 ± 0.12 times that of Trial A respectively.

The lowering of the radial Octupole magnetic field during diagnostic ramp down raised the question: does the annihilation time diagnose reductions in axial energy during adiabatic expansion? Or alternatively, is the observed shift in annihilation time of the adiabatically cooled population an indicator of reduction in radial energy resulting from energy mixing during the axial trap expansion? During the investigation, I split the simulated distribution into mix and no-mix categories, and found that the two categories annihilate at significantly different axial locations and times. I found an explanation, which described the annihilation location and time of a trajectory quite accurately based on its components of axial and transverse energy, and hence found an effect coupling the axial energy of the trapped antihydrogen atoms to the annihilation time.

In Chapter 5 I confirmed that the simulated energy decrease during Trial A was consistent with an adiabatic process, by comparison to the energy loss predicted by two adiabatic models. Model 1 assumed no mixing between axial and transverse energies, whereas Model 2 assumed energy equilibrates between all three dimensions. The two models therefore bound the energy decrease that is consistent with an adiabatic process, within the possible range of energy mixing dynamics.

In Chapter 6 I simulated the effect of changing the expansion duration on the mean energy of the adiabatically cooled distribution. I compared 10.6 s and 230.2 s duration adiabatic expansion simulations and found a difference in final mean total energy of 0.024 ± 0.002 K, which was caused by a combination of increased loss of quasi-trapped particles and increased mixing between axial and transverse energy components during the 230.2 s duration expansion.

I investigated the applications of adiabatic cooling, by simulating a proposed experiment to improve the current bound on the antihydrogen charge, by first adiabatically cooling the trapped antihydrogen atoms before reducing the trap depth and performing the charge neutrality measurement, which involves stochastic acceleration of hypothetically charged antihydrogen using an electric field. The simulation results suggested adiabatic cooling would allow a factor of ~ 10 increase in experimental precision. Finally, I used a simple analysis of the spatial region of the trap occupied by adiabatically cooled anti-atoms to assess the compatibility of adiabatic expansion cooling with 1S-2S spectroscopy experiments. I found that axially adiabatically cooled antihydrogen tended not to occupy low radius trajectories, and hence avoided the 1S-2S laser beam. During a 60 s hold in the spectroscopy magnetic

trap, 23.6% of adiabatically cooled trajectories did not enter the beam region, compared to 2.1% of non adiabatically cooled trajectories.

7.1 Future Work

The 2016 adiabatic expansion experiment has demonstrated that it is possible to use axial adiabatic cooling to reduce the mean energy of trapped antihydrogen atoms to around 62% (simulated result) of that of a typical distribution of trapped antihydrogen. The similar number of antihydrogen counts in the three experimental trials has shown that it is possible to confine similar numbers of trapped anti-atoms in a short magnetic well, as in a standard magnetic well, and therefore that adiabatic cooling does not lead to a significant reduction in the number of anti-atoms ALPHA is able to confine. This technique will need to be refined, and we will need to demonstrate that we are able to routinely adiabatically cool large numbers of trapped anti-atoms.

The fact that our detailed Monte Carlo simulations were able to reproduce the experimental annihilation distribution within reasonable experimental uncertainty is an encouraging result for the accuracy of the simulation. Since ALPHA is now capable of trapping larger numbers of antihydrogen, a future route may be to carry out a similar but more carefully characterised experiment which can be compared more precisely with simulation as a more stringent test of our ability to accurately simulate the dynamics of antihydrogen confined in the ALPHA magnetic trap.

Since I now understand that the axial energy of a trapped antihydrogen atom sets its axial annihilation location, and its radial energy sets its annihilation time, I have found a means of inferring the axial and radial energy components of a trapped anti-atom via a simple analysis of its patterns of annihilation during Octupole FRD. Future work may involve comparison of the energies predicted by the technique to methods which measure the energy of trapped anti-atoms via linewidth and time-of-flight measurements [7], as well as to simulation. One simple way I may achieve this is by using the annihilation data of the 2016 adiabatic expansion experiment to reconstruct the axial and radial energy components of the trapped anti-atoms for comparison to simulation. However, our uncertainty in the Octupole current as a function of time during Octupole FRD limits the accuracy of this comparison since the method relies on our ability to model the trap magnetic field as a function of time. In addition, the number of trapped antihydrogen atoms in each experimental trial is insufficient to characterise the double-peaked structure of the axial annihilation distribution. For these reasons, a future Octupole FRD experiment would involve a careful high resolution measurement of the Octupole current during ramp down in the experiment itself (rather than some days later), and would involve trapping of larger numbers of antihydrogen atoms. For accurate comparison with simulation, I would need to develop a means of determining

the radial (rather than transverse) energy of the simulated anti-atoms. Due to increases in ALPHA's capability to trap large numbers of antihydrogen atoms, it is likely that the Octupole ramp down rate during FRD could be decreased, whilst retaining a reasonable signal-to-noise detection ratio. As a result, the precision of our method to measure the energy components of antihydrogen would increase.

In addition to studying the ability of this method to accurately determine the radial and axial energy components of trapped antihydrogen, Octupole FRD experiments could enable the ratio of mix to no-mix antihydrogen atoms to be measured experimentally for comparison to simulation. To achieve this, I would want to optimise our FRD protocol to lead to the greatest separation possible between annihilation times and axial locations between mix and no-mix categories. Even without such experiments, the ability of the method to reproduce the double-peaked axial annihilation structure that is familiar of trapped antihydrogen atoms during FRD is an encouraging result for the accuracy of our knowledge of the energy mixing dynamics in the magnetic trap.

Future work will involve developing optimised adiabatic expansion protocols for use in experiment. This will involve optimising initial magnetic confining potentials to maximise the number of trapped antihydrogen atoms, as well as maximising the change in volume over the course of the expansion. Further increases in the precision of the charge neutrality experiment as a result of an improved adiabatic expansion protocol may be possible. The simulation results presented in this thesis used the same adiabatic expansion procedure as was demonstrated in the proof-of-principle 2016 experiment. Further studies have begun to investigate the application of negative currents on the internal Mirror Coils (B..D) during the expansion to increase the trap depth and the degree of energy mixing [108]. Increasing the trap depth during adiabatic expansion means more particles can be confined because the rate of quasi-trapped loss decreases. The quasi-trapped trajectories then adiabatically cool to energies below the regular trap depth, meaning more particles can be confined in the regular trap.

We are also interested in developing experimental protocols for combined laser cooling and adiabatic expansion cooling. Since laser cooling is limited to achieving mean antihydrogen energies ~ 20 mK, it may be beneficial to first laser cool in the small magnetic volume, before adiabatically expanding the magnetic trap. The author has recently developed a relatively simple laser cooling simulation, which largely follows the simulation method described in [102] but models the laser as a single frequency (although it has a 50 MHz-100 MHz bandwidth in reality) and models the laser beam to have constant intensity (although the beam is pulsed in reality). Due to the lack of modelling of the pulsed nature of the laser beam, the simple laser cooling simulation decreases the antihydrogen energy around 100 times faster than is expected experimentally. Using the simple simulation, I have simulated a period of laser cooling in the

small magnetic volume, followed by an axial adiabatic expansion and then a 60 s hold in the spectroscopy trap. Preliminary results show that the fraction of antihydrogen atoms which do not enter the spatial region of the 1S-2S spectroscopy laser beam during the 60 s hold increases further to $\sim 60\%$. Future work will involve increasing the accuracy of the relatively simple laser cooling simulation to rule out the fact that this result is due to simulation inaccuracies. The fast cooling rate of the simple simulation may model the dynamics inaccurately since the degree of energy mixing between axial and transverse dimensions is expected to be reduced relative to reality. If the effect persists, one may consider performing a radial (rather than axial) adiabatic expansion by slowly reducing the Octupole magnetic field, which would reduce the transverse energy and therefore the dominating uncertainty in the 1S-2S experiment: transit time broadening. A similar analysis of the spatial region of the trap occupied by such trajectories would be necessary.

A

Infinite Octupole Model

In Cylindrical coordinates, the magnetic vector potential, \mathbf{A} , for this infinite octupole is,

$$\mathbf{A}_\infty = Fr^4 \cos(4\phi)\hat{z} \quad (\text{A.1})$$

where F is a constant [1]. Since

$$\mathbf{B} = \nabla \times \mathbf{A}, \quad (\text{A.2})$$

we can find an expression for the Cartesian magnetic field components of the infinite octupole,

$$B_x = 8IKy(y^2 - 3x^2), \quad (\text{A.3})$$

$$B_y = -8IKx(y^2 - 3x^2), \quad (\text{A.4})$$

and

$$B_z = 0, \quad (\text{A.5})$$

where I is the current and K is a constant where $K = 19.6655 \text{ TA}^{-1}\text{m}^{-3}$ for an infinite octupole with a radius equal to that of the ALPHA Octupole.

B

Generic Adiabatic Expansion and Contraction

In Section 6.1.2 I derived the change in energy for an ensemble of particles undergoing a slow axial adiabatic trap expansion, followed by a faster (but still adiabatic) axial trap contraction, where the initial and final magnetic trap configurations are identical. I made the assumption that the slow expansion allowed enough time for energy to equilibrate between all three dimensions (3D mixing) during the expansion, and in that the fast (but adiabatic) contraction there was no mixing between energy components. The derivation showed that the final mean energy, $\langle E_3 \rangle$, is always greater than or equal to the initial mean energy, $\langle E_1 \rangle$.

In this appendix, I will derive the equivalent result for a generic (combination of radial and axial) trap expansion, and with a generic degree of energy mixing between axial and transverse degrees of freedom during the slow adiabatic expansion (Equation 6.16). I use the same labelling as in Figure 6.6, in which the particles have energy E_1 prior to the expansion, energy E_2 after the expansion but before the contraction and energy E_3 after the contraction. I will derive both adiabatic invariants for the slow expansion and the fast contraction, and combine the two to determine the ratio $\langle E_3/E_1 \rangle$, which is the measure of the ensemble-averaged change in total energy over the course of the slow expansion and fast contraction. For the following derivation I assume the particles bounce in a hard-walled box (not a magnetic potential) for simplicity.

I will start with the derivation of the adiabatic invariant for the slow trap expansion ((1)→(2) in Figure 6.6). Equation 5.3 gives the infinitesimal change in velocity in a given dimension on account of an adiabatic change in the trap length in a given dimension. By using $d(v_i^2) = 2v_i dv_i$, where v_i is the particle speed in a given direction ($i \in x, y, z$) and dv_i is the infinitesimal change in v_i , I find an expression for the infinitesimal change in v_i^2 caused

by the expansion

$$\begin{aligned} dv_{x,\text{pre-mix}}^2 &= -2v_x^2 \frac{dL_x}{L_x} \\ dv_{y,\text{pre-mix}}^2 &= -2v_y^2 \frac{dL_y}{L_y} \\ dv_{z,\text{pre-mix}}^2 &= -2v_z^2 \frac{dL_z}{L_z} \end{aligned} \quad (\text{B.1})$$

where L_i is the trap length in a given direction, dL_i is the infinitesimal change in length and I use the subscript ‘‘pre-mix’’ to indicate that the expressions do not include changes in speed components on account of energy mixing. I will assume expansion (and contraction) in the x and y dimensions is self-similar and that the radial expansion is some constant factor, R , of the axial expansion such that

$$\frac{dL_x}{L_x} = \frac{dL_y}{L_y} = R \frac{dL_z}{L_z}, \quad (\text{B.2})$$

which leads to

$$\begin{aligned} dv_x^2 &= -2v_x^2 R \frac{dL_z}{L_z}, \\ dv_y^2 &= -2v_y^2 R \frac{dL_z}{L_z}, \\ dv_z^2 &= -2v_z^2 \frac{dL_z}{L_z}. \end{aligned} \quad (\text{B.3})$$

As energy mixes between axial and transverse degrees of freedom, the x , y and z energy components are exchanged, but the total energy does not change on account of the energy mixing. Therefore, the incremental change in total energy (including mixing), dE , can be equated to the total energy change that does not include mixing, The infinitesimal change in total energy is given by

$$dE = \frac{1}{2}M(dv_{x,\text{pre-mix}}^2 + dv_{y,\text{pre-mix}}^2 + dv_{z,\text{pre-mix}}^2), \quad (\text{B.4})$$

where M is the particle mass. By replacing the infinitesimal changes in velocity components squared by the expressions above (Equation B.3) I find,

$$dE = -M \left(R(v_x^2 + v_y^2) + v_z^2 \right) \frac{dL_z}{L_z}. \quad (\text{B.5})$$

Since energy is mixing to some extent between degrees of freedom, at an arbitrary instant the axial kinetic energy is some fraction, f , of the total energy,

$$\begin{aligned} \frac{1}{2}Mv_z^2 &= fE \\ \frac{1}{2}M(v_x^2 + v_y^2) &= (1-f)E. \end{aligned} \quad (\text{B.6})$$

where $0 \leq f \leq 1$. Since energy mixing involves stochastic kicks which share energy randomly between the axial and transverse dimensions, for a single particle f constantly changes over the course of the expansion. For a particle which mixes energy equally between all three spatial dimensions, the time average of this fraction, $\bar{f} \rightarrow \frac{1}{3}$. Based on this definition I can eliminate the speed components from Equation B.5 such that

$$dE = -M \left(R(1-f)E + fE \right) \frac{dL_z}{L_z}. \quad (\text{B.7})$$

The trap volume, $V = L_x L_y L_z$, allows us to express the infinitesimal change in V in terms of infinitesimal changes in trap length,

$$\begin{aligned}\frac{dV}{V} &= \frac{dL_x}{L_x} + \frac{dL_y}{L_y} + \frac{dL_z}{L_z} \\ &= \frac{RdL_z}{L_z} + \frac{RdL_z}{L_z} + \frac{dL_z}{L_z} \\ &= (2R + 1) \frac{dL_z}{L_z},\end{aligned}\tag{B.8}$$

and to express the infinitesimal change in energy in terms of an infinitesimal change in trap volume,

$$dE = \frac{-2(R(1-f)E + fE)}{2R+1} \frac{dV}{V}.\tag{B.9}$$

This expression can be integrated,

$$\int \frac{1}{E} dE = \int \frac{-2(R(1-f)E + fE)}{2R+1} \frac{1}{V} dV.\tag{B.10}$$

To perform the integral, I must note that f is varying stochastically as the trap volume expands, meaning the integral is modulated by the variable f . As a result, performing the integral leads to

$$\ln E = \frac{-2(R(1-\bar{f})E + \bar{f}E)}{2R+1} \ln V + C,\tag{B.11}$$

where C is a constant and I have replaced f by its time average since I have integrated over the expansion. Rearranging this expression leads to the adiabatic invariant for the slow expansion

$$EV^{\frac{2[R(1-\bar{f})+\bar{f}]}{2R+1}} = \text{constant},\tag{B.12}$$

and hence E_2 can be expressed in terms of E_1 ,

$$E_2 V_2^{\frac{2[R(1-\bar{f})+\bar{f}]}{2R+1}} = E_1 V_1^{\frac{2[R(1-\bar{f})+\bar{f}]}{2R+1}},\tag{B.13}$$

where V_1 and V_2 are the trap volumes at stage (1) and stage (2) respectively (see Figure 6.6).

Then I am required to derive the adiabatic invariant for the fast (no energy mixing) contraction (stage (2) to stage (3) in Figure 6.6). The infinitesimal changes in velocity are equal to the expressions used above,

$$\begin{aligned}dv_x^2| &= -2v_x^2 R \frac{dL_z}{L_z} \\ dv_y^2| &= -2v_y^2 R \frac{dL_z}{L_z} \\ dv_z^2| &= -2v_z^2 \frac{dL_z}{L_z},\end{aligned}\tag{B.14}$$

where I have used Equation B.2. I can also use Equation B.8 to convert to volume,

$$\begin{aligned}dv_x^2 &= \left(\frac{-2R}{2R+1} \right) v_x^2 \frac{dV}{V}, \\ dv_y^2 &= \left(\frac{-2R}{2R+1} \right) v_y^2 \frac{dV}{V}, \\ dv_z^2 &= -\frac{2v_z^2}{2R+1} \frac{dV}{V}.\end{aligned}\tag{B.15}$$

Each expression can be integrated independently to give three adiabatic invariants

$$\begin{aligned} v_x^2 V^{\frac{2R}{2R+1}} &= \text{constant}, \\ v_y^2 V^{\frac{2R}{2R+1}} &= \text{constant}, \\ v_z^2 V^{\frac{2}{2R+1}} &= \text{constant}. \end{aligned} \tag{B.16}$$

To find an expression for E_3 in terms of E_2 , I use the fact that

$$E_3 = \frac{1}{2}M(v_{x,3}^2 + v_{y,3}^2 + v_{z,3}^2), \tag{B.17}$$

where $v_{i,3}$ is the velocity component at stage (3) (see Figure 6.6), and apply the adiabatic invariants to give an expression for E_3 in terms of E_2 ,

$$E_3 = \frac{M}{2} \left((v_{x,2}^2 + v_{y,2}^2) \left(\frac{V_2}{V_1} \right)^{\frac{2R}{2R+1}} + v_{z,2}^2 \left(\frac{V_2}{V_1} \right)^{\frac{2}{2R+1}} \right), \tag{B.18}$$

where $v_{i,2}$ is the velocity component at stage (2) (see Figure 6.6). Similar to the method for the slow expansion, I want to eliminate the speed components from Equation B.18 in favor of the total energy.

Recall that previously I used f to denote the instantaneous fraction of parallel energy of an individual particle during the mixing expansion. The stochastic nature of energy mixing means that, during the slow expansion, the distribution in time of the fractional axial energy of an individual particle is equal to the distribution of the fractional axial energy of the ensemble at an instant in time. Since the fast contraction follows the slow expansion, the energy mixing during the expansion sets the initial condition for the fast contraction, and hence the fractional axial energy at stage (2) is distributed according to

$$\begin{aligned} \frac{1}{2}Mv_{z,2}^2 &= fE \\ \frac{1}{2}M(v_{x,2}^2 + v_{y,2}^2) &= (1-f)E, \end{aligned} \tag{B.19}$$

where f has the same definition as above. This allows us to express the energy at stage (3), E_3 , in terms of the energy at stage (2), E_2 ,

$$E_3 = \left((1-f) \left(\frac{V_2}{V_1} \right)^{\frac{2R}{2R+1}} + f \left(\frac{V_2}{V_1} \right)^{\frac{2}{2R+1}} \right) E_2, \tag{B.20}$$

where I have used $V_3 = V_1$. Combining with Equation B.13 allows us to eliminate E_2 ,

$$\frac{E_3}{E_1} = \frac{\left((1-f) \left(\frac{V_2}{V_1} \right)^{\frac{2R}{2R+1}} + f \left(\frac{V_2}{V_1} \right)^{\frac{2}{2R+1}} \right)}{\left(\frac{V_2}{V_1} \right)^{\frac{2[R(1-f)+f]}{2R+1}}}, \tag{B.21}$$

and taking the ensemble-average (over particles) leads to,

$$\left\langle \frac{E_3}{E_1} \right\rangle = \frac{\left((1-\langle f \rangle) \left(\frac{V_2}{V_1} \right)^{\frac{2R}{2R+1}} + \langle f \rangle \left(\frac{V_2}{V_1} \right)^{\frac{2}{2R+1}} \right)}{\left(\frac{V_2}{V_1} \right)^{\frac{2[R(1-\bar{f})+\bar{f}]}{2R+1}}}, \tag{B.22}$$

where I denote the ensemble-average with an angled bracket. For reasons discussed above, $\bar{f} = \langle f \rangle$, and thus I arrive at our expression for the generic adiabatic invariant for a slow adiabatic expansion with an arbitrary degree of energy mixing and an arbitrary trap expansion,

$$\left\langle \frac{E_3}{E_1} \right\rangle = \frac{\left((1 - \langle f \rangle) \left(\frac{V_2}{V_1} \right)^{\frac{2R}{2R+1}} + \langle f \rangle \left(\frac{V_2}{V_1} \right)^{\frac{2}{2R+1}} \right)}{\left(\frac{V_2}{V_1} \right)^{\frac{2[R(1-\langle f \rangle) + \langle f \rangle]}{2R+1}}}. \quad (\text{B.23})$$

To put $\langle f \rangle$ into context, for the mix category of antihydrogen atoms, $\langle f \rangle \sim 0.5$.

REFERENCES

- [1] C Amole et al. Discriminating between antihydrogen and mirror-trapped antiprotons in a minimum-B trap. *New Journal of Physics*, 14(1):15010–, 2012.
- [2] Mark Johnson. *Design and commissioning of beamlines for the ALPHA antihydrogen experiment*. PhD thesis, 2019.
- [3] M Ahmadi et al. Antihydrogen accumulation for fundamental symmetry tests. *Nature Communications*, 8(1):681–681, 2017.
- [4] Steven Jones. Personal Communication, July 2022.
- [5] C. Amole et al. The ALPHA antihydrogen trapping apparatus. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 735:319–340, 2014.
- [6] G.B Andresen et al. Antihydrogen annihilation reconstruction with the ALPHA silicon detector. *Nuclear Instruments & Methods in Physics Research. Section A, Accelerators, Spectrometers, Detectors and Associated Equipment*, 684:73–81, 2012.
- [7] C.J Baker et al. Laser cooling of antihydrogen atoms. *Nature (London)*, 592(7852):35–42, 2021.
- [8] M Ahmadi et al. Observation of the hyperfine spectrum of antihydrogen. *Nature (London)*, 548(7665):66–69, 2017.
- [9] M Zhong, J Fajans, and A F Zukor. Axial to transverse energy mixing dynamics in octupole-based magnetostatic antihydrogen traps. *New Journal of Physics*, 20(5):053003–, 2018.
- [10] Peter Young. The Leapfrog method and other “symplectic” algorithms for integrating Newton’s laws of motion. <https://young.physics.ucsc.edu/115/leapfrog.pdf>, 2014. Physics 115/242 Lecture Notes UCSC.
- [11] M Ahmadi et al. An improved limit on the charge of antihydrogen from stochastic acceleration. *Nature (London)*, 529(7586):373–376, 2016.

- [12] M Baquero-Ruiz, A.E Charman, J Fajans, A Little, A Povilus, F Robicheaux, J.S Wurtele, and A.I Zhmoginov. Using stochastic acceleration to place experimental limits on the charge of antihydrogen. *New Journal of Physics*, 16(8):83013–, 2014.
- [13] Chukman So. *Antiproton and Positron Dynamics in Antihydrogen Production*. PhD thesis, University of California, Berkeley, 2014.
- [14] Vittorio. Degiorgio. *Photonics : A Short Course*. Undergraduate Lecture Notes in Physics. Springer International Publishing, Cham, 2014.
- [15] E SCHRÖDINGER. An undulatory theory of the mechanics of atoms and molecules. *Nippon Sugaku-Buturigakkwaishi*, 1(1):106–117, 1927.
- [16] David J. (David Jeffery) Griffiths. *Introduction to elementary particles*. Physics textbook. Wiley-VCH, Weinheim, 2nd, rev. ed. edition, 2008.
- [17] Carl D Anderson. The positive electron. *Physical Review*, 43(6):491–494, 1933.
- [18] NORWOOD RUSSELL HANSON. Discovering the positron (ii). *The British Journal for the Philosophy of Science*, 12(48):299–313, 1962.
- [19] P. M. S BLACKETT and G OCCHIALINI. Photography of penetrating corpuscular radiation. *Nature (London)*, 130(3279):363–363, 1932.
- [20] Owen Chamberlain, Emilio Segrè, Clyde Wiegand, and Thomas Ypsilantis. Observation of antiprotons [13]. *Physical Review*, 100(3):947–950, 1955.
- [21] Bruce Cork, Glen R Lambertson, Oreste Piccioni, and William A Wenzel. Antineutrons produced from antiprotons in charge-exchange collisions [5]. *Physical Review*, 104(4):1193–1197, 1956.
- [22] Simon Mitton. The expanding Universe of Georges Lemaître. *Astronomy Geophysics : the Journal of the Royal Astronomical Society*, 58(2):2.28–2.31, 2017.
- [23] E Hubble. A relation between distance and radial velocity among extra-galactic nebulae. *Proceedings of the National Academy of Sciences - PNAS*, 15(3):168–173, 1929.
- [24] G. Gamow. Expanding Universe and the Origin of Elements. *Phys. Rev.*, 70:572–573, Oct 1946.
- [25] R. A. Alpher, H. Bethe, and G. Gamow. The origin of chemical elements. *Phys. Rev.*, 73:803–804, Apr 1948.
- [26] R. H. Dicke, P. J. E. Peebles, P. G. Roll, and D. T. Wilkinson. Cosmic Black-Body Radiation. *Astrophysical Journal*, 142:414–419, July 1965.

- [27] Laurent Canetti, Marco Drewes, and Mikhail Shaposhnikov. Matter and antimatter in the Universe. *New Journal of Physics*, 14(9):95012–, 2012.
- [28] C.S Wu, E Ambler, R.W Hayward, D.D Hoppes, and R.P Hudson. Experimental test of parity conservation in beta decay [5]. *Physical Review*, 105(4):1413–1415, 1957.
- [29] J.H Christenson, J.W Cronin, V.L Fitch, and R Turlay. Evidence for the $\Delta S = 2$ decay of the K^0 meson. *Physical Review Letters*, 13(4):138–140, 1964.
- [30] M Adinolfi et al. First evidence of direct CP violation in charmless two-body decays of B^0 mesons. *Physical Review Letters*, 108(20), 2012.
- [31] J. Alcaraz et al. Search for antihelium in cosmic rays. *Physics Letters B*, 461(4):387–396, 1999.
- [32] Gary Steigman. Observational tests of antimatter cosmologies. *Annual Review of Astronomy and Astrophysics*, 14(1):339–372, 1976.
- [33] F.W. Stecker. On the nature of the baryon asymmetry. *Nuclear Physics B*, 252:25–36, 1985.
- [34] A. G. Cohen, A. De Rújula, and S. L. Glashow. A Matter-Antimatter Universe? *The Astrophysical Journal*, 495(2):539–549, mar 1998.
- [35] G Luders. Proof of the TCP theorem. *Annals of Physics*, 281(1-2):1004–1018, 2000.
- [36] V Alan Kostelecký and Robertus Potting. CPT and strings. *Nuclear Physics. B*, 359(2):545–570, 1991.
- [37] Ralf Lehnert. CPT symmetry and its violation. *Symmetry (Basel)*, 8(11):114–, 2016.
- [38] Michael Martin Nieto and T Goldman. The arguments against “antigravity” and the gravitational acceleration of antimatter. *Physics Reports*, 205(5):221–281, 1991.
- [39] Myron L Good. K^0 and the Equivalence Principle. *Physical Review*, 121(1):311–313, 1961.
- [40] Sandip Pakvasa, Walter A Simmons, and Thomas J Weiler. Test of equivalence principle for neutrinos and antineutrinos. *Physical Review. D, Particles and fields*, 39(6):1761–1763, 1989.
- [41] E.G Adelberger, B.R Heckel, C.W Stubbs, and Y Su. Does antimatter fall with the same acceleration as ordinary matter? *Physical Review Letters*, 66(7):850–853, 1991.
- [42] Michael Martin Nieto and T Goldman. The arguments against “antigravity” and the gravitational acceleration of antimatter. *Physics Reports*, 205(5):221–281, 1991.

- [43] A Benoit-Lévy and G Chardin. Introducing the Dirac-Milne Universe. *Astronomy and Astrophysics (Berlin)*, 537:A78–, 2012.
- [44] Dragan Slavkov Hajdukovic. Do we live in the Universe successively dominated by matter and antimatter? *Astrophysics and Space Science*, 334(2):219–223, 2011.
- [45] N. Bezginov, T. Valdez, M. Horbatsch, A. Marsman, A.C. Vutha, and E.A. Hessels. A measurement of the atomic hydrogen lamb shift and the proton charge radius. *Science (American Association for the Advancement of Science)*, 365(6457):1007–1012, 2019.
- [46] O Arnoult, F Nez, L Julien, and F Biraben. Optical frequency measurement of the 1s–3s two-photon transition in hydrogen. *The European Physical Journal. D, Atomic, Molecular, and Optical Physics*, 60(2):243–256, 2010.
- [47] Christian G Parthey, Arthur Matveev, Janis Alnis, Birgitta Bernhardt, Axel Beyer, Ronald Holzwarth, Aliaksei Maistrou, Randolph Pohl, Katharina Predehl, Thomas Udem, Tobias Wilken, Nikolai Kolachevsky, Michel Abgrall, Daniele Rovera, Christophe Salomon, Philippe Laurent, and Theodor W Hänsch. Improved measurement of the hydrogen 1s-2s transition frequency. *Physical Review Letters*, 107(20):203001–203001, 2011.
- [48] G. Baur et al. Production of antihydrogen. *Physics Letters B*, 368(3):251–258, February 1996.
- [49] M Hori and J Walz. Physics at CERN’s Antiproton Decelerator. *Progress in particle and Nuclear Physics*, 72:206–253, 2013.
- [50] J. S Hangst et al. Trapped antihydrogen. *Nature (London)*, 468(7324):673–676, 2010.
- [51] G B Andresen, MD Ashkezari, M Baquero-Ruiz, W Bertsche, P D Bowe, E Butler, CL Cesar, M Charlton, A Deller, S Eriksson, J Fajans, T Friesen, M C Fujiwara, DR Gill, A Gutierrez, J S Hangst, W N Hardy, R S Hayano, ME Hayden, A J Humphries, R Hydromako, S Jonsell, S L Kemp, L Kurchaninov, N Madsen, S Menary, P Nolan, K Olchanski, A Olin, P Pusa, COe Rasmussen, F Robicheaux, E Sarid, D M Silveira, C So, J W Storey, R I Thompson, DP van der Werf, J S Wurtele, and Y Yamazaki. Confinement of antihydrogen for 1,000 seconds. *Nature physics*, 7(7):558–564, 2011.
- [52] C AMOLE et al. Resonant quantum transitions in trapped antihydrogen atoms. *Nature (London)*, 483(7390):439–443, 2012.
- [53] C Amole et al. Description and first application of a new technique to measure the gravitational mass of antihydrogen. *Nature Communications*, 4(1):1785–1785, 2013.

- [54] C Amole et al. An experimental limit on the charge of antihydrogen. *Nature Communications*, 5(1):3955–3955, 2014.
- [55] M Ahmadi et al. Observation of the 1s-2s transition in trapped antihydrogen. 2017-01-26.
- [56] M Ahmadi et al. Characterization of the 1s-2s transition in antihydrogen. *Nature (London)*, 557(7703):71–75, 2018.
- [57] M Ahmadi et al. Observation of the 1s–2p Lyman- transition in antihydrogen. *Nature (London)*, 561(7722):211–215, 2018.
- [58] M Ahmadi et al. Investigation of the fine structure of antihydrogen. *Nature (London)*, 578(7795):375–380, 2020.
- [59] Jeffrey S Hangst et al. Addendum to the alpha proposal; the alpha-g apparatus. Report Number: CERN-SPSC-2016-031 ; SPSC-P-325-ADD-1.
- [60] S Jonsell, E A G Armour, M Plummer, Y Liu, and A C Todd. Helium–antihydrogen scattering at low energies. *New Journal of Physics*, 14(3):035013, March 2012.
- [61] G. B. Andresen, W. Bertsche, P. D. Bowe, C. C. Bray, E. Butler, C. L. Cesar, S. Chapman, M. Charlton, J. Fajans, M. C. Fujiwara, D. R. Gill, J. S. Hangst, W. N. Hardy, R. S. Hayano, M. E. Hayden, A. J. Humphries, R. Hydomako, L. V. Jorgensen, S. J. Kerrigan, L. Kurchaninov, R. Lambo, N. Madsen, P. Nolan, K. Olchanski, A. Olin, A. P. Povilus, P. Pusa, E. Sarid, S. Seif El Nasr, D. M. Silveira, J. W. Storey, R. I. Thompson, D. P. van der Werf, and Y. Yamazaki. Antiproton, positron, and electron imaging with a microchannel plate/phosphor detector. *Review of scientific instruments*, 80(12):123701–123701, 2009.
- [62] D. L. Eggleston, C. F. Driscoll, B. R. Beck, A. W. Hyatt, and J. H. Malmberg. Parallel energy analyzer for pure electron plasma devices. *Physics of fluids. B, Plasma physics*, 4(10):3432–3439, 1992.
- [63] CERN Bulletin. ELENA prepares a bright future for antimatter research. ELENA : de brillantes perspectives pour la recherche sur l’antimatière. page 4, Jul 2011.
- [64] Siara Fabbri. *Optimization of Antiproton Capture and Delivery for the ALPHA Antihydrogen Experiment*. PhD thesis, 2021-12-31.
- [65] C. M. Surko. The buffer-gas positron accumulator and resonances in positron-molecule interactions. Written in conjunction with the symposium honoring the retirement of Dick Drachman and Aaron Temkin (NASA Goddard Space Flight Center, Greenbelt,

MD; November 18, 2005); to be published by the NASA Conference Publication Service, A. Bhatia, ed. (2006).

- [66] J. Fajans and C. M. Surko. Plasma and trap-based techniques for science with antimatter. *Physics of plasmas*, 27(3):30601–, 2020.
- [67] Michael E Glinsky and Thomas M O’Neil. Guiding center atoms: Three-body recombination in a strongly magnetized plasma. *Physics of Fluids. B, Plasma Physics*, 3(5):1279–1293, 1991.
- [68] M.C Fujiwara et al. Temporally controlled modulation of antihydrogen production and the temperature scaling of antiproton-positron recombination. *Physical Review Letters*, 101(5):053401–053401, 2008.
- [69] A Capra. Lifetime of magnetically trapped antihydrogen in alpha. *Hyperfine interactions*, 240(1):1–6, 2019.
- [70] Joseph McKenna. Personal communication, 2021.
- [71] Stracka Simone. Real-time detection of antihydrogen annihilations and applications to spectroscopy. *EPJ Web of conferences*, 71:00126–, 2014.
- [72] A Capra. Machine learning for antihydrogen detection at alpha. *Journal of Physics. Conference series*, 1085(4):42007–, 2018.
- [73] *Trapped charged particles a graduate textbook with problems and solutions*. Advanced textbooks in physics. World Scientific Publishing Co. Pte. Ltd., Singapore ;, 2016 - 2016.
- [74] T. Pohl, H. R. Sadeghpour, Y. Nagata, and Y. Yamazaki. Cooling by spontaneous decay of highly excited antihydrogen atoms in magnetic traps. *Phys. Rev. Lett.*, 97:213001, Nov 2006.
- [75] C L Taylor, Jingjing Zhang, and F Robicheaux. Cooling of rydberg antihydrogen during radiative cascade. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 39(23):4945–4959, nov 2006.
- [76] Christopher J. Foot. *Atomic physics*. Oxford master series in physics. Atomic, optical and laser physics ; 7. Oxford University Press, Oxford, 2005.
- [77] C. Ø. Rasmussen, N. Madsen, and F. Robicheaux. Aspects of 1S-2S spectroscopy of trapped antihydrogen atoms. *J. Phys. B*, 50(18):184002, 2017.
- [78] *Quantum electrodynamics*. Advanced series on directions in high energy physics ; vol. 7. World Scientific, Singapore ;, 1990.

- [79] Savely G Karshenboim. Precision physics of simple atoms: QED tests, nuclear structure and fundamental constants. *Physics Reports*, 422(1):1–63, 2005.
- [80] Stanley J Brodsky and Ronald G Parsons. Precise theory of the Zeeman spectrum for atomic hydrogen and deuterium and the Lamb shift. *Physical Review*, 163(1):134–146, 1967.
- [81] A E Charman et al. Description and first application of a new technique to measure the gravitational mass of antihydrogen. *Nature Communications*, 4(1):1785–1785, 2013.
- [82] M. Ahmadi et al. Enhanced control and reproducibility of non-neutral plasmas. *Phys. Rev. Lett.*, 120:025001, Jan 2018.
- [83] C. J. Baker et al. Sympathetic cooling of positrons to cryogenic temperatures for antihydrogen production. *Nature Communications*, 12(1):6139–6139, 2021.
- [84] F Robicheaux. Simulations of antihydrogen formation. *Physical Review. A, Atomic, molecular, and optical physics*, 70(2):1–22510, 2004.
- [85] Miguel Cerqueira Bastos, Gunnar Fernqvist, Gregory Hudson, John Pett, Andrea Cantone, Francis Power, Alfredo Saab, Bjorn Halvarsson, and John Pickering. High accuracy current measurement in the main power converters of the large hadron collider: Tutorial 53 part 53 in a series of tutorials on instrumentation and measurement. *IEEE instrumentation measurement magazine*, 17(1):66–73, 2014.
- [86] Wolfgang Ketterle. 8.421 Atomic and Optical Physics I. MIT OpenCourseWare, March 2015.
- [87] M A Alarcón, C J Riggert, and F Robicheaux. Simulations of Majorana spin flips in an antihydrogen trap. *Journal of Physics. B, Atomic, Molecular, and Optical Physics*, 52(15):155004–, 2019.
- [88] Francis Robicheaux. Personal Communication, July 2022.
- [89] L. D. (Lev Davidovich) Landau. *Mechanics*. Course of theoretical physics ; v. 1. Butterworth-Heinemann, Oxford, 3rd ed. edition, 1996.
- [90] P. M Bellan. A microscopic, mechanical derivation of the adiabatic gas relation. *American Journal of Physics*, 72(5):679–682, 2004.
- [91] Jonathan Goodman. Lecture notes on monte carlo methods fall semester. Unpublished, Courant Institute of Mathematical Sciences, goodman@cims.nyu.edu, 2005.
- [92] Root minuit2 minimization library. https://root.cern/doc/master/group__Minuit.html.

- [93] I.P. Omelyan, I.M. Mryglod, and R. Folk. Optimized Forest–Ruth- and Suzuki-like algorithms for integration of motion in many-body systems. *Computer Physics Communications*, 146(2):188–202, 2002.
- [94] William Bertsche. Personal Communication, March 2022.
- [95] Peter Granum, Magnus Linnet Madsen, Joseph Tiarnan Kerr McKenna, Danielle Louise Hodgkinson, and Joel Fajans. Efficient calculations of magnetic fields of solenoids for simulations. *Nuclear instruments and methods in physics research. Section A, Accelerators, spectrometers, detectors and associated equipment*, 1034(C):166706–, 2022.
- [96] Teukolsky Saul A. Vetterling William T. Press, William H. and Brian P. Flannery. *Numerical recipes example book (C)*. Cambridge University Press, Cambridge, 2nd ed / william t.vetterling .. [et al.] edition, 1992.
- [97] J. D Jackson. Classical electrodynamics, 3rd ed. *American Journal of Physics*, 67(9):841–842, 1999.
- [98] Cameron R. Taylor. Finite difference coefficients calculator. <https://web.media.mit.edu/~crtaylor/calculator.html>, 2016.
- [99] K. McDonald. Expansion of an axially symmetric, static magnetic field in terms of its axial field. <https://physics.princeton.edu/~mcdonald/examples/axial.pdf>.
- [100] C. William (Charles William) Gear. *Numerical initial value problems in ordinary differential equations*. Prentice-Hall series in automatic computation. Prentice-Hall, Englewood Cliffs, 1971.
- [101] Prof. Dmitry Panchenko. MIT Mathematics Lecture 14, Kolmogorov-Smirnov Test. <https://ocw.mit.edu/courses/18-443-statistics-for-applications-fall-2006/resources/lecture14/>, 2006.
- [102] P H Donnan, M C Fujiwara, and F Robicheaux. A proposal for laser cooling antihydrogen atoms. *Journal of Physics. B, Atomic, Molecular, and Optical Physics*, 46(2):25302–, 2013.
- [103] Joel Fajans. Personal Communication, April 2022.
- [104] A I Zhmoginov, A E Charman, R Shalloo, J Fajans, and J S Wurtele. Nonlinear dynamics of anti-hydrogen in magnetostatic traps: implications for gravitational measurements. *Classical and quantum gravity*, 30(20):205014–, 2013.
- [105] Paul Murray Bellan. *Fundamentals of Plasma Physics*. Cambridge University Press, Cambridge, 2006.

- [106] Evan Deddo. Simulating stochastic acceleration of antihydrogen atoms with a hypothetical small charge. Unpublished, University of California, Berkeley, Bachelor Thesis, 2021.
- [107] Chris Orum Rasmussen. *Tests of Fundamental Symmetries with Trapped Antihydrogen*. PhD thesis, Aarhus University, 2016.
- [108] Nicolas Kalem. Personal Communication, March 2022.