

Sympathetically Laser-Cooled Positron Plasmas for Antihydrogen Formation

Jack McCauley Jones

A thesis presented for the degree of
Doctor of Philosophy
March 18th 2022



Swansea University
Prifysgol Abertawe

Abstract

Answering the question of why we live in a matter-dominated universe is of great interest to contemporary physicists, as the Standard Model of Particle Physics predicts that matter and antimatter should only ever be produced in equal parts. Antihydrogen is a good candidate for searches for asymmetries between matter and antimatter as it is the simplest antimatter bound state, and it has an extremely well-understood matter counterpart: the hydrogen atom. The ALPHA collaboration at CERN can now routinely trap several hundred antihydrogen atoms in a magnetic trap, allowing precise measurements of the fundamental properties of antihydrogen.

ALPHA currently traps around 20 antihydrogen atoms every few minutes, so accumulating enough antiatoms to perform precision measurements can take many hours. Increasing this trapping rate would allow for faster or more precise measurements. Simulations and experimental data show that there is a strong correlation between the temperature of the positron plasma used in antihydrogen formation and the trapping rate that can be achieved. This thesis describes work towards using laser-cooled beryllium ions to obtain colder positron plasmas in the ALPHA trapping apparatus. Singly-charged beryllium ions are liberated from the surface of a solid beryllium target inside the apparatus via laser ablation, where the ions can be trapped, laser-cooled and mixed with positron plasmas. By mixing the ions and positrons together, they can exchange energy, and the ions can provide cooling to the positrons. Careful tuning of laser and trap parameters, as well as use of other techniques, has allowed for the successful sympathetic cooling of positron plasmas in ALPHA.

Declaration

This work has not previously been accepted in substance for any degree and is not being concurrently submitted in candidature for any degree.

[REDACTED]

Jack McCauley Jones, 18th March 2022

This thesis is the result of my own investigations, except where otherwise stated. Other sources are acknowledged by footnotes giving explicit references. A bibliography is appended.

[REDACTED]

Jack McCauley Jones, 18th March 2022

I hereby give consent for my thesis, if accepted, to be available for photocopying and for inter-library loan, and for the title and summary to be made available to outside organisations.

[REDACTED]

Jack McCauley Jones, 18th March 2022

The University's ethical procedures have been followed and, where appropriate, that ethical approval has been granted.

[REDACTED]

Jack McCauley Jones, 18th March 2022

Acknowledgements

Thanks to everyone at CERN who made my time there worth it, and special thanks to Adam Powell, Andrew Evans, Asia Peszka, Danielle Hodgkinson, Graham Stutter, James Brydges-Harrington, Joseph McKenna, Mark Johnson, Nathan Evetts, Patrick Mullan, Pete Knapp, Pierre Grandemange, Muhammed Sameed, and Steve Jones, for the Fun Times. Thanks to the people who ran the Beryllium Experiment with me: Asia Peszka, Dan Maxwell, and Kurt Thompson. Without you, this thesis would not have been possible (literally). Thanks to my supervisor, Niels Madsen, for keeping the Beryllium team on the straight and narrow, and for your invaluable guidance.

Thanks to my non-physics friends, family, and the bois for your support, and very special thanks to my partner, Samantha.

Contents

List of Figures	vii
1 Introduction	1
1.1 Antimatter and Symmetry	1
1.2 Studying Antimatter	2
1.3 Contemporary Experiments with Cold Antimatter	4
1.4 Laser Cooling	4
1.4.1 Sympathetic Cooling	5
1.5 Author Contributions	6
1.5.1 Publications	7
1.6 Thesis Outline	7
2 Apparatus	8
2.1 The Antiproton Decelerator	8
2.2 The Catching Trap	9
2.3 ALPHA-2	10
2.3.1 The Neutral Trap	10
2.4 The Positron Accumulator	11
2.5 ALPHA-g	12
2.6 The ALPHA-g Beamline	13
2.7 Diagnostic Tools	13
2.7.1 Faraday Cup	13
2.7.2 Microchannel Plate and Phosphor Screen Assembly	14
2.8 Vacuum Translators	16
2.9 Annihilation Detectors	16
3 Dynamics of Trapped, Charged Particles	19
3.1 Single-Particle Motion in a Penning-Malmberg Trap	20

3.2	Characteristics of a Plasma	22
3.3	Trapped Plasma Dynamics in a Penning-Malmberg trap	25
3.4	Perturbations and Heating in a Real Trap	26
3.5	Rotating Wall Compression	27
3.6	Cyclotron Cooling	28
3.7	Evaporative Cooling	29
3.7.1	Strong Drive Regime Evaporative Cooling	31
3.7.2	Adiabatic Cooling	32
3.8	Multispecies Plasmas	33
3.8.1	Centrifugal Separation	33
3.8.2	Sympathetic Cooling	34
3.8.3	Electron Kick	35
3.9	Plasma Characterisation	35
3.9.1	Plasma Dynamics During Potential Well Manipulations	35
3.9.2	Plasma Temperature Measurements	36
3.9.3	Systematic Uncertainties in Temperatures	37
3.10	Summary	38
4	Antihydrogen Production and Trapping	39
4.1	Antihydrogen Production	40
4.1.1	Quantifying Antihydrogen Production	42
4.2	Formation Mechanisms	43
4.3	Antihydrogen Trapping	44
4.4	Antihydrogen Accumulation	45
4.5	Simulations of Antihydrogen Formation Mechanisms	45
4.6	Experimental Plasma Parameter Tuning	48
4.6.1	Evaporative Cooling of Antiprotons	49
4.6.2	Antiproton Plasma Radius	49
4.6.3	Positron Temperature	52
4.6.4	Positron Density	53
4.6.5	Positron EVC and Antihydrogen Formation	54
4.7	Summary	55
5	Laser Cooling	56
5.1	Theory of Doppler Cooling	56
5.2	Cooling Requirements in an Atom	60
5.3	Laser Cooling Be^+	61

5.4	Laser Cooling of Ions in a Penning Trap	62
5.5	Laser Access in a Penning-Malmberg Trap	66
6	Experiments with Beryllium	68
6.1	Production of Beryllium Ions by Laser Ablation	68
6.1.1	Laser Ablation	69
6.1.2	Experimental Ablation Setup	69
6.1.3	Ablation Threshold	71
6.1.4	Ablated Ion Energy Distributions	72
6.2	Beryllium Trapping	74
6.2.1	Trapping Instability	75
6.2.2	Stacking	76
6.2.3	Be ⁺ Plasma Expansion	77
6.3	Experimental Cooling Setup	78
6.3.1	The Cooling Laser	78
6.3.2	Wavemeter Frequency Control	79
6.3.3	Acousto-Optic Modulator	80
6.3.4	Fluorescence Detection	82
6.3.5	Computer Control Hardware and Software	82
6.4	Typical Cooling Sequence	83
6.5	Experimental Results	84
6.5.1	Varying Initial Detuning	84
6.5.2	Fluorescence Temperature Measurements	84
6.5.3	Polarisation Dependence	87
6.5.4	Be ⁺ Heating Rate	87
6.5.5	MCP Distortion	88
7	Sympathetic Cooling of Positrons	91
7.1	Positron Preparation	92
7.2	Beryllium Preparation	93
7.2.1	Arrested Rotating Wall Ion Compression	93
7.2.2	Centre Extraction	94
7.3	Merging Ions and Positrons	96
7.3.1	Compression of the Mixture	97
7.4	Positron Temperature Measurements	97
7.5	Simulations of Sympathetic Cooling	98
7.6	Experimental Results	100

7.6.1	Varying Final Detuning	100
7.6.2	Varying number of ions	101
7.6.3	Centrifugal Separation Measurements	102
7.6.4	Extracting Temperatures from Radial Profiles	107
7.6.5	Final Remarks About Positron Temperatures	107
7.6.6	Improving Temperature Measurements	108
7.7	Mixture Heating Rate	109
7.8	Mixture Expansion Rate	110
7.9	Summary	110
8	Future Work and Conclusions	112
8.1	Towards Beryllium-Antihydrogen Compatibility	112
8.1.1	Number of Positrons	113
8.1.2	Number of Be ⁺ ions	113
8.1.3	Cutting Out the Arrested Rotating Wall	113
8.1.4	Further Reductions in Positron Temperature	115
8.2	The Neutral Trap and Large Plasmas	115
8.2.1	Sympathetic Cooling in the Neutral Trap	115
8.3	Improved Fluorescence Diagnostics	117
8.4	On-axis Cooling Laser	118
8.5	Fibre Delivery of the Cooling Beam	119
8.6	Beryllium for Magnetometry	120
8.7	Conclusions	120
8.7.1	Final Remarks	121
	Bibliography	122

List of Figures

1.1	The AD complex as of 2020 (approximately to scale). Protons from the Proton Synchrotron hit the Target, producing antiprotons that are captured and decelerated within the AD. Before the installation of ELENA (2019), antiprotons from the AD would be delivered directly to the experiments. Now, antiprotons will be delivered from the AD to ELENA for further deceleration before being sent to the experiments. The five experiments that are currently running and their positions within the AD are shown. The sixth experiment, ATRAP, ceased operation in 2020.	3
2.1	Reproduced from M.A. Johnson’s thesis [26]. A labelled schematic of the whole ALPHA experimental apparatus, after the installation of ALPHA-g in 2018.	9
2.2	Three-stage potentials in the positron accumulator’s Penning-Malmberg trap. Shaded/unshaded regions show the axial extents of the trap electrodes. Positrons enter from the Na-22 source on the right and enter the trap, where they lose energy through collisions with the nitrogen buffer gas. The pressure gradient is created by differential pumping and by using electrodes with increasing radii from stage 1 to stage 3. After the desired number of positrons has been accumulated in stage 3, the remaining buffer gas is quickly pumped out, and the particles can be ejected towards the rest of the experiment.	12
2.3	A typical FC signal. The large peak corresponds to a bunch of charged particles hitting the FC, and the long exponential decay is due to the RC circuit connected to the FC. The total charge of the particle bunch is equal to the maximum charge measured by the FC.	14

2.4	a) A single MCP channel. Here, a single incident electron causes a cascade of secondary electrons. A potential difference of up to 1000V is applied between the MCP's front and back, which provides acceleration for the cascade. b) An MCP and phosphor screen assembly currently used in ALPHA. The rectangular tabs are the connection points for the phosphor screen, MCP back, and MCP front (going clockwise).	15
2.5	a) A positron plasma imaged with an MCP and phosphor screen assembly. b) A large Be ⁺ plasma imaged on an MCP. The aperturing in the bottom-left and bottom-right sections of the image are caused by electrostatic potentials on the MCP's tabs and bolts – a hardware defect that is discussed in chapter 6.5.5.	16
2.6	Figure reproduced from ALPHA collaboration [32]. Reconstructed events from ALPHA's SVD. Left is an antiproton annihilation event, while right is a cosmic ray. The inner circle shows the location of the inner surface of the Penning-Malmberg trap electrodes. The outer three concentric layers are the layers of the SVD modules. Red points show where a high-energy particle has been detected passing through a layer, and red lines are fitted tracks. The blue diamond is the reconstructed position of the annihilation vertex.	17
3.1	A Penning trap, formed from two endcap electrodes and a single ring electrode, as well as a constant magnetic field along z . For a positive particle, the endcaps are held at a positive potential $+V_0$, while the ring is held at ground. This configuration produces a saddle electric potential where the particles are confined in the z -direction and anti-confined in the r -direction. A magnetic field $\mathbf{B} = B_0\hat{z}$ provides a trapping force in the r -direction to overcome the force from the electrodes, producing a 3-D electrostatic potential well.	20
3.2	Single-particle motion in a Penning trap. In the radial $x - y$ plane, the particle undergoes large, slow magnetron rotations around the centre of the trap, as well as small, fast cyclotron rotations. The particle also undergoes simple harmonic oscillation in the axial (z) direction (not pictured).	22
3.3	A sketch of a simple five-electrode Penning-Malmberg trap, with the middle three electrodes cut away for viewing clarity. A spheroidal plasma (blue) is trapped between the two biased end electrodes, while the central three electrodes are held at ground. A magnetic field pointing along the z -axis is provided by an external solenoid, not pictured due to its size. Spacing between electrodes has been exaggerated for clarity – in a real trap, the electrodes are built to overlap slightly, and are physically separated by ruby spacers.	23

3.4	<i>I</i> ($\bar{\kappa}$) as a function of magnetisation parameter. The asymptotic expressions for the strongly magnetised (blue) and weakly magnetised regime (red) are shown along with numerical calculations (green). Both expressions and the data are reproduced from Glinsky et al. [34].	24
3.5	An end-on sketch of a segmented RW electrode. Each segment is biased independently with a sinusoidally-varying voltage, which is phase-shifted for each segment. The phase-shift is labelled for a) a dipole drive and b) a quadrupole drive. The arrows represent the direction of the electric field.	28
3.6	On-axis potentials used during evaporative cooling of antiprotons. The much deeper initial well (blue) is gradually morphed into a much shallower well (dashed red) during the EVC. The inset shows a zoomed-in version of the final EVC well, with energetic antiprotons (shown as blue circles) being lost to the left. The transition between the initial well and the final well is smooth, and takes 10s.	30
3.7	On-axis potentials used during adiabatic cooling of positrons before being used for antihydrogen formation. Grey shaded/unshaded regions show axial extent of trap electrodes. Positrons are initially held in a deeper, narrower well (blue), which is changed into a shallower well (red dashed) over 10ms. The well becoming shallower effectively lengthens it axially, which causes adiabatic cooling.	32
3.8	Image reproduced from the ALPHA collaboration [45]. An MCP image of a mixed antiproton-electron plasma exhibiting centrifugal separation in a a) 1T and b) 3T magnetic field. While the two species are concentric in the trap, their different charge-to-mass ratios give them different trajectories as they are being ejected from the trap, causing them to image in different places. The faint halos visible near the antiproton plasmas are thought to be optical reflections occurring in the MCP.	33
3.9	A typical temperature diagnostic signal. In this example, the plasma being measured was an $\sim 800\text{K}$ positron plasma. The algorithm that chooses the linear region to fit essentially attempts to minimise the error in the fit, and is described in more detail by L. T. Evans [46].	37
4.1	Antihydrogen trapped per hour over the past several years in ALPHA. Since first trapping antihydrogen atoms in 2009, the trapping rate has increased drastically. ALPHA's first publication of trapped antihydrogen [49], which was based on data from 2010, had 0.11 trapped antihydrogen atoms per experiment. ALPHA's more recent publications report on experimental runs where 10-20 were captured per experiment [10, 22, 24].	40

4.2	Schematic of the Catching Trap’s Penning-Malmberg trap. a) On-axis potentials during particle catching. Antiprotons enter from the left, and are reflected by the potential applied to HVB (blue solid line). Before the reflected particles have enough time to leave the trap, a potential is raised on HVA to complete the trap. A plasma of electrons is prepared before capture and held in the trap to cool the antiprotons upon their capture. b) A schematic of the Penning trap electrode stack. Axial positions of a) and b) are matched. Electrodes 6 and 16 are segmented rotating wall electrodes.	41
4.3	Schematic of the ALPHA-2 Penning-Malmberg trap. Antiprotons are brought from the left and captured in the antiproton preparation region, and positrons are brought in from the right. The plasmas can be independently manipulated in either end of the trap before being brought together to the central antihydrogen synthesis region. The Neutral Trap magnets are shown in blue (mirror coils, solenoids) and green (octupole), which can be energised to form a 3-D magnetic field minimum region to trap antihydrogen atoms.	42
4.4	Potentials used during the merge. Shaded/unshaded regions show axial extent of trap electrodes. a) Potentials before the merge, with red and green shaded regions representing the space charge and physical extent of the plasmas. b) How the potentials change from the start of the merge (red dashed line) to the end (blue solid line). The merge step takes 1 second.	43
4.5	Figure reproduced from ALPHA collaboration [20]. The number of antihydrogen atoms detected after the Neutral Trap is ramped down after consecutive mixing cycles. Each mixing cycle is separated by ~ 4 minutes. The error bars are statistical, and the number of replicates is indicated in blue above each data point. The dashed linear fit gives a trapping rate of 10.5 ± 0.6 detected antihydrogen atoms per mixing cycle.	46
4.6	Antiproton production efficiency, $f_{\bar{H}}$, (upper) and trapping fraction, f_{trap} , (lower) as a function of antiproton EVC depth. Blue points are with EVC, and the red line is without EVC, with shaded regions representing error bars. Each data point is made up of several repeat measurements, and the error bars are standard error of the mean for each point. The calculation of $f_{\bar{H}}$ uses the number of antiprotons in the trap <i>before</i> EVC.	50
4.7	Production efficiency, $f_{\bar{H}}$, (upper) and trapping fraction, f_{trap} , (lower) as a function of antiproton plasma radius.	51

4.8	Production efficiency, $f_{\bar{H}}$, (upper) and trapping fraction, f_{trap} , (lower) as a function of positron temperature after mixing. Runs were binned by positron temperature (blue), and error bars are statistical standard errors for each bin. Raw data points are shown in red.	53
5.1	Atom behaviour as it absorbs and emits photons. For each event, momentum is conserved.	58
5.2	The scattering force, F_{scatt} , of a laser on a Be^+ ion as a function of the laser's detuning from resonance with the Be^+ cooling transition ($f_0 = 957$ THz). The detuning is shown in multiples of the natural linewidth of the cooling transition, $\Gamma = 2\pi \times 19.6$ MHz. The force is plotted for different values of the ion's velocity, v , and the laser intensity, I	59
5.3	Figure adapted from [67]. The periodic table, with elements that have been laser cooled highlighted (note: may not be exhaustive). Many elements are cooled on an $S \rightarrow P$ transition. Group 1 has a lone electron in their outer S states while in their ground state, allowing cooling on the $S \rightarrow P$ transition. Group 2 elements typically have an electron removed, leaving a single electron in their outer S state, and are cooled as ions. The noble gases are cooled by exciting an electron into a metastable S state, and then cooling with an $S \rightarrow P$ transition.	60
5.4	Energy levels relevant when laser cooling Be^+ ions. The levels are split in a high magnetic field, and are labelled by their m_J and m_I quantum numbers. The $2p$ m_I states have a splitting significantly smaller than the natural linewidth of the cooling transition, and are not shown. The laser cooling transition (purple) and repump transition (blue) are shown.	62
5.5	Figures created by Dan Maxwell [72]. Energy levels of the $2p^2P_{3/2}$ state (upper) and $2s^2S_{1/2}$ state (lower) in Be^+ as a function of magnetic field for small magnetic fields (left) and large magnetic fields (right). All energy levels are relative to the $2s^2S_{1/2}$ state at zero magnetic field in the absence of hyperfine interactions.	63
5.6	Figure created by Dan Maxwell [72]. Frequency difference between the cooling and repump transitions as a function of magnetic field.	64

5.7 Momentum-position phase space plot of a toy simulation of a particle undergoing simple harmonic oscillation in one axis with a resonant damping force, qualitatively similar to a particle in a Penning trap being laser cooled along the z -axis of the trap. The resonance with F_{scatt} occurs at $p = 0.5$, marked with a red line. In this plot, the particle starts at high radius and moves inwards as it cools. The cooling rate increases as the particle cools into resonance with F_{scatt} , and then drops significantly once the particle's momentum is too low to be Doppler-shifted into resonance at any point in its oscillation ($|kv_{max}| < |\omega - \omega_0|$). The magnitude of the cooling force has been greatly exaggerated in this toy simulation for clarity. 65

5.8 A particle's motion in the radial plane of a Penning trap. The laser is offset from the centre of the plasma to apply a positive torque to the plasma, reducing the velocity and radius of its magnetron oscillations. 66

5.9 A sketch of a Penning-Malmberg trap electrode stack with a segmented central electrode, allowing radial laser access. An axial laser can enter through the open ends of the stack. This geometry would be difficult to implement in ALPHA-2 due to the very limited space outside the trap. 67

6.1 a) the ablation source (without covering lid). The beryllium target and acceleration plate are electrically isolated from each other and from the rest of the source. The ablation source is mounted on a linear translator inside the trap vacuum chamber. b) optics setup for transporting laser light from the ablation laser to the source. Laser light first passes through optics in the laser lab: a half-waveplate and polarising beam splitter (PBS) for power control, and a telescope to radially expand and collimate the beam. The light is then guided along a roughly 5m path to the zone, where a final lens focuses the beam onto the beryllium target before it enters the trap's vacuum chamber. Components and distances not to scale. Some additional mirrors are present in the optical setup that are not pictured. 71

6.2 Schematic view of the ALPHA-2 Penning-Malmberg trap, including the location of the beryllium ablation source and cooling laser access. The cooling laser is tightly constrained, only attaining good overlap with trapped species close to the centre of the trap. 72

6.3	Number of ions ablated as a function of laser fluence. Multiple readings were taken for each laser fluence, with means of each point plotted. Error bars are standard errors of the mean. The fit to extract the ablation threshold is shown as a solid red line, which becomes dotted when extrapolating beyond points used in the fit.	73
6.4	Percentage of the ablated ion bunch that is able to pass all the way through the trap as a function of blocking voltage applied in the trap.	74
6.5	Magnetic field as a function of distance from the centre of the ALPHA-2 Penning Trap. The position of the ablation source and Penning-Malmberg trap electrode stack are shown above. Two booster solenoids (blue), which are sometimes energised during Be^+ catching, increase the magnetic field from 1T to 3T.	75
6.6	On-axis potentials used for catching beryllium ions. Shaded and numbered regions represent Penning trap electrodes. Ions enter from the left, then the gate electrode (E1) is raised from ground (blue) to 40V (red). Ions are cooled into the catching well, and then transferred over to the reservoir.	76
6.7	Number of ions caught as a function of number of ablation pulses used. Error bars are standard errors of the mean of multiple repeat measurements for each data point.	77
6.8	Number of ions loaded over several trial runs using the 3T stacking sequence. The shaded blue region indicates the standard deviation.	78
6.9	Be^+ plasma radius as a function of hold time, in 1T (blue) and 3T (red) magnetic fields. Radially small plasmas of Be^+ ions were prepared, and then held for a variable amount of time. After the hold time, the plasma was ejected from the trap to measure its size. Each data point consists of several repeat measurements of different Be^+ plasmas. Error bars are standard errors of the mean of each point.	79
6.10	Figure adapted from [25]. A schematic representation of the full laser setup used for Be^+ ablation and laser cooling. The ablation laser path (blue) utilises a PBS and half-waveplate for power attenuation, and a telescope for beam shaping. The ablation laser setup is described in more detail above (section 6.1.2). The cooling laser light passes through a telescope, for beam shaping, before being double-passed through an acousto-optic modulator. The cooling laser light is then sent to the zone, where it passes through a quarter-waveplate before entering ALPHA-2.	80

6.11	Change in laser frequency as a function of time as the setpoint is changed from an initial setpoint of -0.3GHz to a final setpoint of -0.1GHz . The setpoint was updated at $t = 0$, and is shown in green. The laser frequency is changed using active feedback from the wavemeter. Frequencies are plotted relative to $f_0 = 957.41\text{THz}$, the resonance frequency of the Be^+ cooling transition. . . .	81
6.12	An acousto-optic modulator. An RF signal is applied to the piezoelectric transducer, which converts the signal into acoustic vibrations in the crystal. These acoustic oscillations dynamically change the refractive index in the crystal and cause diffraction to occur.	82
6.13	Outline of a Be^+ -only cooling sequence. Much of the data taken with only Be^+ ions was taken before the arrested RW technique was developed, and so ion plasmas are left uncompressed.	84
6.14	Be^+ plasma temperature, as a function of initial detuning. The red line shows the temperature when the laser is off, while the blue points show the temperature when the laser is on. The cooling efficiency drops as the laser is brought closer to the resonance. Presumably, this is because the ions are quite hot to begin, and so require quite a large detuning to effectively cool them from their initial high temperature. Error bars are standard errors of the mean of repeat measurements at each detuning.	85
6.15	Fluorescence signal from a Be^+ plasma as the frequency of the laser is swept over the cooling transition. Signal data (blue) was taken with ions in the trap, as described in the text. Background data was taken after ions were removed from the trap. The signal data has been fit with a Gaussian on the left-hand side of the peak. The solid region of the fit indicates the region of the data that was used for the fit. On the right-hand side, the sharp drop-off of the signal is attributed to heating caused by the laser when the detuning is positive. This fit corresponds to a temperature of $\sim 160\text{mK}$, and assumes that the width of the peak is dominated by Doppler broadening.	86
6.16	Be^+ plasma temperature, measured with the plasma temperature diagnostic, as a function of laser polarisation. Each plasma contained around 2×10^5 ions. Error bars are standard errors of the mean of repeat measurements.	87
6.17	Be^+ temperature after laser cooling and followed by a variable hold with the laser off. Each plasma contained around 2×10^5 ions. Error bars are standard errors of the mean of repeat measurements.	88

6.18	a) Simulated contour plot of the maximum potential experienced by a particle approaching the MCP. Coloured lines show equipotentials (V). The MCP front, back, and phosphor were charged with -100V , 800V , and 1000V respectively. The circle represents the edge of the active area of the MCP. b) A picture of a real MCP (removed from the apparatus), with an image of a large plasma digitally overlaid on top. The bolts and phosphor screen align very well with the locations of the apertures.	90
7.1	A typical sympathetic cooling sequence that includes the arrested RW. Be^+ ions (pink) are prepared in tandem with positrons (green), before being merged. The Be^+ preparation involved several interleaved laser cooling (LC) and rotating wall (RW) compression holds, which form the arrested RW process.	94
7.2	a) A plasma of Be^+ ions after the arrested RW process, and after the centre extraction process, imaged on a microchannel plate (MCP). The axes show the calculated size in a 1T field. The calculation assumes that the particles follow the magnetic field lines as they are ejected from the high-field region of the trap to the lower-field region of the MCP. b) A plasma that has been compressed, but without centre extraction. A high-radius, uncompressed halo remains. Image distortion from electrostatic effects caused by the MCP housing are clearly visible. These electrostatic distortion effects are discussed in detail in chapter 6.5.5.	95
7.3	On-axis potentials used for merging Be^+ ions with positrons. Be^+ ions are initially in the left-hand well, with positrons in the right-hand well. The two species are merged over 30s by raising the bottom of the positron well, allowing positrons to flow into the Be^+ well. This process is generally performed with the cooling laser on, to actively cool the Be^+ during the merge.	96
7.4	Left: An MCP image of a mixed Be^+ -positron plasma, after RW compression. This image is dominated by positrons. Right: Same as left, but positrons have been removed with an ekick prior to imaging.	97
7.5	Reproduced from [82]. Calculated minimum temperature of 2×10^6 positrons sympathetically cooled with a variable number of Be^+ ions. The calculations were performed for three different plasma radii: 0.81mm, 1mm, and 1.22mm. The corresponding peak densities of these three plasmas were 9.1, 6, 2, and $4.3 \times 10^{13} \text{m}^{-3}$ respectively. The external heating rate used was 52K s^{-1}	99

- 7.6 Axial positron temperature in a mixed Be^+ -positron plasma as a function of final laser detuning. The mixed plasma contained 2.6×10^6 positrons, and $3.8 \pm 0.1 \times 10^5$ Be^+ ions. The laser was chirped from an initial detuning of -104Γ to the final detuning in approximately 40s. The error bars give the standard error of multiple measurements at each detuning. The green line shows the positron temperature in the absence of any Be^+ ions ($17.3 \pm 0.5\text{K}$). The red line shows the temperature when the laser is off, and the Be^+ ions are not laser cooled ($841 \pm 135\text{K}$). The shading around the lines shows standard error. 101
- 7.7 Figure reproduced from ALPHA collaboration [25]. Axial positron temperature in a mixed Be^+ -positron plasma as a function of the number of ions loaded. The data has been binned by ion number, and the vertical and horizontal error bars showing the standard errors of the values of positron temperature and Be^+ ion number in each bin. Loads containing 1.4×10^6 positrons (blue points) and 2.6×10^6 positrons (red points) were used. During sympathetic cooling, the cooling laser was chirped from an initial detuning of -129Γ to a final detuning of -7Γ in approximately 40s. The green line shows the positron temperature in the absence of any Be^+ ions for the 2.6×10^6 load, and the shading around the line represents the standard error. 103
- 7.8 Axially-integrated radial density profiles of Be^+ from a cold, mixed Be^+ -positron plasma. $5\mu\text{s}$ before imaging the Be^+ ions, the positrons are ejected from the plasma with an ekick in the opposite direction. The left-hand plots show the axially integrated radial density profiles of the Be^+ ions, while the right-hand plots show the raw MCP image. The red points show the experimental density data, while the green dot-dashed line shows results of the best-fit distribution extracted from the N2DEC code described in section 7.5. The laser detunings used were a) -20Γ , b) -36Γ , and c) -128Γ . The measured axial positron temperatures were a) $7.1 \pm 0.5\text{K}$, b) $10.1 \pm 0.2\text{K}$, and c) $370 \pm 100\text{K}$. The fit results were a) $6.2 \pm 0.6\text{K}$, b) $19.2 \pm 2\text{K}$, and c) $253 \pm 54\text{K}$. The corresponding positron densities extracted from the N2DEC code were a, b) $6.2 \pm 0.1 \times 10^{13} \text{m}^{-3}$, and c) $1.2 \pm 0.2 \times 10^{14} \text{m}^{-3}$. The slightly elliptical nature of the MCP images is assumed to be due to distortions from electrostatic fields from the MCP, and is corrected for in the analysis of the radial profiles (see section 7.6.4). 105

7.9	MCP images of centrifugally-separated Be^+ ions after the positrons have been removed. There were 2×10^5 Be^+ ions and 1.9×10^6 positrons. The delay between the positron removal and imaging the ions was: Left: $50\mu\text{s}$, Right: 1ms. After the $50\mu\text{s}$ delay, there are no clear signs that the hollow plasma is collapsing. After the 1ms delay, the hollow Be^+ structure appears to be part-way through collapsing.	106
7.10	Time evolution of the positron temperature after laser cooling the Be^+ -positron mixture. The mixture contained 1.4×10^6 positrons and $4.7 \pm 0.1 \times 10^5$ Be^+ ions, and was cooled by chirping the laser from -110Γ to a final detuning of -7Γ over approximately 40s. Red points show the evolution when the laser was turned off after cooling, which resulted in the positrons heating at a rate of around 25K s^{-1} initially, before stabilising at around 125K. The blue points show experiments where the cooling laser was left on after the laser cooling sweep, at a fixed detuning of -7Γ . The temperature of the positrons remained steady at a value of $6.7 \pm 0.3\text{K}$ for over 20s in this case.	109
7.11	Time evolution of the radius of three different plasmas without laser cooling in a 1T magnetic field. The red points show the evolution of a Be^+ -positron mixture, the blue points show the evolution of Be^+ ions only, and the green points show the evolution of positrons only.	110
8.1	A sympathetically compressed Be^+ -positron plasma is shown on the left (blue). The image is dominated by positrons. In the right-hand image, the positrons have been removed after compression, leaving only the sympathetically compressed ions (orange). Sizes given are the calculated sizes in a 3T field. The plasma contained 10^5 Be^+ and 10^6 positrons, and was driven with a 400kHz, 5V RW for 30s.	114
8.2	Figure created by Dan Maxwell [72]. Positron temperature as a function of octupole current. The mixture of 1.9×10^6 positrons and $\sim 5 \times 10^5$ Be^+ ions was prepared in a manner similar to the previous chapter. The mixture was then recompressed with the RW while the octupole was being ramped to the current setpoint. The mixed plasma was then shuffled back to the central region of the trap, which is influenced by the octupole magnet, to be laser cooled. . .	116
8.3	The layout of the ALPHA-2 Penning-Malmberg trap, and the current position of the SiPM used for detecting fluorescence photons.	117

1/ Introduction

Antimatter has inspired both scientists and the general public alike since its discovery, capturing the imagination with its exotic otherness, rarity, and generally sci-fi-esque properties. Antimatter was first hinted at by Dirac in 1928 [1], and was discovered experimentally by Anderson 4 years later when he discovered the positron, which is the antimatter counterpart of the electron. Extension of Dirac's theory to other particles and advances in experiment predicted and then discovered the antiproton and the antineutron some 20 years later [2] [3], solidifying the role of antimatter in modern physics.

In more recent years, several experiments have begun to trap and study antimatter with unprecedented precision. These were largely made possible by the Antiproton Decelerator at CERN, which is a facility that can produce slow, cold, and easily trappable antiprotons. One of the experiments that is fed by the Antiproton Decelerator is called ALPHA¹, and is the experiment where all of the work discussed in this thesis took place.

1.1 Antimatter and Symmetry

Antimatter is often described as a “twin” of regular matter, with each type of matter particle having a corresponding antiparticle with the same mass, but opposite charge and spin. These antiparticles are generally named after their matter counterpart, and their symbol is denoted with a bar: the proton (p) has an antimatter counterpart called the antiproton (\bar{p}), the neutron (n) has the antineutron (\bar{n}), and so on. The electron is the only exception to this naming convention: its antiparticle is called the positron (e^+) for historical reasons. When a matter particle and its antimatter counterpart come into contact they annihilate, converting the mass of both particles into energy. Annihilation products can be photons (γ) or massive particles such as pions (π^\pm, π^0), depending on the particles that are annihilating.

This symmetry between matter and antimatter is formalised in the Standard Model (SM) of particle physics, which includes three fundamental symmetries: charge conjugation, parity

¹Antihydrogen Laser PHysics Apparatus

inversion and time reversal. The transformation of charge conjugation (C) inverts the sign of all fundamental charges, such as electromagnetic charge and colour charge. Parity inversion (P) inverts the spatial coordinates used to describe a particle (like a mirror-image), and time reversal (T) reverses time. When antimatter was first discovered, it was thought that all natural phenomena were invariant under any one of the three symmetries, which would imply that matter and antimatter should be exact mirrors of each other, and that the amount of matter and antimatter in the universe should remain constant and exist in equal quantities. Despite this, the universe that we live in appears to be made almost entirely out of regular matter. The only explanations that can explain this asymmetry between matter and antimatter are some breaking of C, P, or T symmetries, or that the universe was just created with more matter than antimatter. This matter-antimatter asymmetry problem is still unsolved today, and is a core motivation for contemporary studies with antimatter.

Although it was later found that both P [4] and combined C-P symmetries [5] (and, by implication, T symmetry) are broken in nature, this does not provide enough symmetry-breaking to account for the asymmetry that we observe in nature. The CPT transformation, i.e. the simultaneous application of C, P, and T transformations, is the last remaining combination of these three symmetries that remains an unbroken symmetry in nature. If breaking of CPT symmetry was found, it could provide clues about the solution of the matter-antimatter asymmetry problem, as well as having implications for other fundamental principles, such as Lorentz Invariance [6].

1.2 Studying Antimatter

Studies that directly measure properties of antimatter are of particular interest, as these allow comparisons to these same properties in matter. This allows for direct tests of fundamental principles such as CPT symmetry. The antihydrogen atom, the bound state of an antiproton and a positron, is an attractive candidate for these kind of studies, both because it is the simplest and therefore easiest to synthesise antimatter atom, and because its matter counterpart, hydrogen, is a very well-understood atom that plays a fundamental role in physics.

In order to create antihydrogen atoms, both constituent parts – antiprotons and positrons – must be synthesised and trapped before being brought together to form antihydrogen atoms. Formation occurs by an antiproton capturing a positron into a stable orbit to form an (anti)atom.

In order to increase the amount of antihydrogen that can be trapped, it is very important to keep the temperature of both the antiprotons and positrons low, as this produces colder antihydrogen with lower kinetic energy, which is significantly easier to trap. Currently, a large part of the cooling performed on plasmas in the trap comes from keeping the trap at cryogenic

temperatures ($\lesssim 10$ kelvin) using liquid helium.

Other techniques, such as evaporative cooling [7] and adiabatic cooling [8] are also utilised to keep temperatures of particles within the trap as low as possible. This thesis discusses recent efforts towards implementing sympathetic cooling of positrons into the experiment, a technique which uses laser-cooled beryllium ions to cool down the positrons used for antihydrogen synthesis. Laser cooling is a powerful cooling technique that is widely used in physics to obtain ultra-low-energy particles. Being able to harness its power for this application has the potential to provide large reductions in the kinetic energy of positrons used in ALPHA, which in turn could lead to a significant increase in the amount of antihydrogen that can be trapped.




Figure 1.1: The AD complex as of 2020 (approximately to scale). Protons from the Proton Synchrotron hit the Target, producing antiprotons that are captured and decelerated within the AD. Before the installation of ELENA (2019), antiprotons from the AD would be delivered directly to the experiments. Now, antiprotons will be delivered from the AD to ELENA for further deceleration before being sent to the experiments. The five experiments that are currently running and their positions within the AD are shown. The sixth experiment, ATRAP, ceased operation in 2020.

1.3 Contemporary Experiments with Cold Antimatter

Antiprotons are extremely difficult to create in useful quantities, which is an unfortunate reality for an experiment that wishes to study antihydrogen. Antiprotons are produced at the Antiproton Decelerator (AD) Hall at CERN by colliding a high-energy proton beam with a solid metal target, which produces a shower of high-energy particles, some of which are antiprotons. These antiprotons are far too energetic to be trapped and studied in detail, and so must first be decelerated. This is achieved using a dedicated synchrotron called the Antiproton Decelerator, which decelerates antiprotons from an initial energy of 2.7 GeV down to 5.3 MeV. The Antiproton Decelerator operates in a manner similar to other synchrotrons, such as CERN’s Large Hadron Collider, by guiding particles around its 188-metre circumference using magnets. The crucial difference between the Antiproton Decelerator and other synchrotrons, however, is that it is designed to decelerate particles to low energies, rather than accelerate them to higher energies. These decelerated antiprotons are then delivered to several experiments within the AD Hall, where they can be trapped and studied [9].

There are 6 experiments within the AD Hall²: AEGIS, ALPHA, ASACUSA, ATRAP, BASE, and GBAR. Measurements on the spectrum of antihydrogen is the primary goal of ALPHA, ASACUSA and ATRAP. ALPHA [10] and ATRAP’s [11] designs employ traps to hold antihydrogen atoms for these measurements, while ASACUSA intends to study antihydrogen atoms in the form of a beam [12]. Recently, measurements of the effect of gravity on antimatter are being undertaken by ALPHA [13], AEGIS [14] and GBAR [15]. ALPHA and GBAR’s measurements will involve trapped antihydrogen atoms in a freefall-style experiment, while AEGIS will measure the deflection of a beam of antihydrogen atoms due to gravity. The BASE collaboration doesn’t produce antihydrogen at all, and instead makes extremely precise measurements of the properties of antiprotons [16]. The layout of the AD Hall is shown in figure 1.1.

1.4 Laser Cooling

In this context, energy and temperature are often used interchangeably, and so “cooling” a particle is synonymous with reducing its kinetic energy. For reference, our trap can trap charged particles with energies up to around 100 electronvolts³ (eV), which corresponds to around 1.2 million kelvin in temperature units. Room temperature, 300 kelvin, is roughly 0.026 eV. Laser cooling has been a widely-used method of obtaining cold, low-energy particles since it was first demonstrated in 1978 by Wineland, Drullinger and Walls [17], and has the potential to cool

²Or, more accurately, there were 6 until ATRAP stopped operation in 2020.

³1 electronvolt is roughly 1.6×10^{-19} joules

particles below a thousandth of a kelvin.

Laser cooling exploits two principles in order to work. First, photons carry momentum, despite the fact that they have no mass, and so an atom that scatters a photon will change its momentum based on the momenta of the incoming and outgoing photons due to conservation of momentum. The second is that an atom in a trap can be made to scatter more photons as it is moving towards the laser than when it is moving away from the laser by carefully tuning the frequency of the laser, causing a net reduction in the kinetic energy of the atom. These principles are discussed in more detail in chapter 5.

With these two principles in place, a trapped atom's kinetic energy can be lowered to sub-kelvin levels by illuminating it with laser light of a specific frequency. The frequency used depends on the atom that is being laser cooled, as the laser's frequency is chosen such that it matches the frequency of an atomic transition within the atom.

Direct laser cooling is limited to atoms and ions with simple level structures that can be interrogated by readily available lasers, which means that most of the periodic table is not directly coolable in any practical sense. Species with a more complicated level structure such as molecules are possible to cool directly [18], but the increased complexity of the level structure can mean that this is impractical. In order to cool a species with no level structure such as the electron, positron, or proton, other techniques must be used.

1.4.1 Sympathetic Cooling

It is still possible to harness the power of laser cooling with a non-coolable species by allowing a laser-cooled species to exchange energy with it. This method of indirect laser cooling is known as sympathetic cooling. In a Penning-Malmberg trap such as the one used at ALPHA, two species with the same sign of charge can easily be co-trapped in the same region of space, allowing Coulomb interactions to act as the energy exchange mechanism. If one of the species is actively laser cooled, the other species will be cooled via Coulomb interactions. Sympathetic cooling was first demonstrated in 1986 by sympathetically cooling Hg^+ ions with laser-cooled Be^+ ions in a Penning trap [19]. Since then, sympathetic cooling has been demonstrated with a plethora of species.

This thesis discusses a sympathetic cooling scheme where a positron plasma is sympathetically cooled by mixing it with laser-cooled beryllium ions. Since positrons are light, and sympathetic cooling is most efficient when the mass difference between the two species is minimised, beryllium is the natural choice as the lightest readily laser-coolable ion. This mixed, cooled plasma of positrons and beryllium ions could then be merged with a plasma of antiprotons to form antihydrogen atoms, in a similar manner to other antihydrogen work in ALPHA [20]. It

is beneficial to cool the positrons used in antihydrogen formation as much as possible, as there is a strong negative correlation between the temperature of the positrons and the trapping rate of antihydrogen atoms when synthesising antihydrogen atoms. Sympathetic cooling has the potential to produce significantly colder positron plasmas than other techniques that are currently employed in ALPHA.

Sympathetic cooling of positrons using beryllium ions has been demonstrated previously by B. M. Jelenković et al. in 2003 [21], where a few thousand positrons were cooled to $< 5\text{K}$ using around $10^5 - 10^6$ beryllium ions in a Penning-Malmberg trap. In order to implement this in ALPHA, the number of positrons being sympathetically cooled would need to be increased by a factor of ~ 1000 when compared to this previous work, as up to 3 million positrons are used when synthesising antihydrogen in ALPHA. The number of beryllium ions should be kept to a minimum, as antiprotons that are merged with the positron plasma can annihilate with beryllium ions. Practically, this means that any implementation of this scheme in ALPHA would not only need to use far more positrons, but also use a very different ratio of positrons to beryllium when compared to this previous work.

1.5 Author Contributions

As a fairly large group, work at ALPHA is often a collaborative effort. A lot of the work that I undertook during my time there was aided, directly or indirectly, by other members of the group.

The majority of the experimental work carried out for this thesis was undertaken by a subgroup within ALPHA (the “Beryllium Team”). This consists of myself and my supervisor, Niels Madsen, as well as other members:

When I joined in 2017, the group also included a Postdoc, Daniel Maxwell, and a Master’s student, Pete Knapp. Pete left in mid-2018, and we were joined by a new PhD student, Joanna Peszka, in 2019. Dan left the collaboration in 2020, and was replaced by another Postdoc a few months later - Kurt Thompson. I was directly involved in all of the experimental data-taking beryllium runs that took place during my PhD.

I was involved in setting up, commissioning and maintaining the majority of the hardware used in the beryllium experiment. Writing the software that acts as a control system, as well as integrating it with the main ALPHA control system, was a collaborative effort between myself, Dan, and later Joanna and Kurt.

During antihydrogen experimental data-taking runs, most members of the ALPHA collaboration are expected to attend some antihydrogen run shifts each year. As a student who was based at CERN full-time during my PhD, I took part in shifts for the majority of the 2017 and 2018 an-

tihydrogen runs. Shifts typically involved data-taking, and basic troubleshooting/maintenance of the experiment's hardware and software.

1.5.1 Publications

During my time with ALPHA, I was a co-author on the following publications:

- Characterization of the 1S–2S transition in antihydrogen (2018) [10]
- Observation of the 1S–2P Lyman- α transition in antihydrogen (2018) [22]
- Investigation of the fine structure of antihydrogen (2020) [23]
- Laser cooling of antihydrogen atoms (2021) [24]
- Sympathetic cooling of positrons to cryogenic temperatures for antihydrogen production (2021) [25]

1.6 Thesis Outline

In the rest of this thesis, I will discuss the work I undertook towards the goal of sympathetically cooling positrons within the ALPHA apparatus. Chapter 2 gives an overview of the experimental apparatus used. Chapter 3 discusses the dynamics of trapped, charged particles, particularly within Penning-Malmberg traps. Chapter 4 gives an overview of the antihydrogen production process. Chapter 5 covers the theory of laser cooling, and how this applies to particles trapped within a Penning-Malmberg trap. Chapters 6 and 7 discuss the bulk of the experimental work carried out in the thesis, first focusing on laser-cooling beryllium ions within the trap (6), and then the results of sympathetic cooling experiments with positrons and beryllium ions together (7). The final chapter (8) provides some discussion of potential future work related to the sympathetic cooling experiment, and provides some concluding remarks.

2/ Apparatus

This chapter gives an overview of the apparatus used during this PhD project, as well as the ALPHA apparatus in general. The ALPHA apparatus' primary purpose is to synthesise, trap and study antihydrogen atoms. In order to do this, plasmas of antiprotons and positrons are separately prepared, before being brought together and merged to produce low-temperature antihydrogen atoms which can be trapped. There are three main parameters of the antiproton and positron plasmas that affect how much antihydrogen is produced and subsequently trapped when the two plasmas are merged: the radial size, the temperature, and the number of particles in each plasma. The apparatus is designed to be able to control these plasma parameters in order to be able to optimise antihydrogen production and trapping.

Charged particles are trapped with Penning-Malmberg traps, an open-ended cylindrical variant of a Penning trap, and ALPHA uses several of these traps to perform the required plasma manipulations for antihydrogen synthesis and other experimental work. The open-ended nature of the Penning-Malmberg trap allows for easy particle loading and transport between traps. The section of the apparatus in which antihydrogen atoms are produced is an atom trap called ALPHA-2, which consists of two traps overlaid in the same region of space: a charged particle trap and a neutral particle trap. The charged particle trap is a 27-electrode Penning-Malmberg trap, and the neutral particle trap is a magnetic minimum trap. Charged species such as electrons, positrons, and antiprotons are contained by the Penning-Malmberg trap, while the Neutral Trap is used to trap antihydrogen atoms. A labelled schematic of the entire experiment is shown in figure 2.1

2.1 The Antiproton Decelerator

Antiprotons are created at CERN by colliding high-energy proton beams from the Proton Synchrotron with a solid Iridium target. These high-energy collisions produce a shower of particles, a small fraction of which are antiprotons. These antiprotons are then isolated and injected into the Antiproton Decelerator (AD), a synchrotron designed to decelerate antiprotons from 2.7 GeV down to 5.3 MeV, and deliver them in bunches to the various experiments within the AD




Figure 2.1: Reproduced from M.A. Johnson’s thesis [26]. A labelled schematic of the whole ALPHA experimental apparatus, after the installation of ALPHA-g in 2018.

hall. A single bunch contains $\sim 3 \times 10^7$ antiprotons. In ALPHA, antiprotons are caught and further decelerated with a dedicated Catching Trap, described in section 2.2.

A secondary decelerator, the Extra Low Energy Antiproton ring (ELENA), has been completed and is currently being commissioned to be ready for 2021 [27]. ELENA will take bunches of antiprotons from the AD and decelerate them from 5.3 MeV down to 100 keV, allowing for much better catching efficiencies by the individual experiments.

2.2 The Catching Trap

The Catching Trap (CT) is a Penning-Malmberg trap designed to catch antiprotons coming from the AD, and is labelled in figure 2.1. It can then further cool the antiprotons before transferring them to the rest of the experiment. Antiprotons entering the CT have an energy of 5.3 MeV, and first pass through a thin degrader foil which acts as a moderator. The small fraction of the incoming antiprotons that successfully pass through the degrader and have their energy reduced to less than ~ 4 keV are able to be trapped in the CT. The CT’s Penning trap is a 9-electrode trap with a special pair of high-voltage electrodes dedicated to catching high-energy antiproton bunches, which can be raised to a potential of ~ 4 kV. The rest of the electrodes (including those in other parts of the apparatus, such as the ALPHA-2 Penning-Malmberg trap) are typically able to go to ± 140 V, which means that in order to transfer the antiprotons into other regions of the apparatus, they first need to be cooled further from ~ 4 keV to below 140 eV. This

cooling is achieved by pre-loading an electron plasma into the CT, which self-cools by emitting cyclotron radiation, and can sympathetically cool the antiprotons. The physics behind this is discussed in more detail in chapter 3.

After cooling the antiprotons, the electrons are removed to leave a pure antiproton plasma, which can be ejected from the CT to be used in other parts of the apparatus. In normal operation, the antiprotons are re-caught in the ALPHA-2 Penning-Malmberg trap, where they can be used for antihydrogen formation. The role of the CT in antihydrogen production is discussed in more detail in chapter 4.

2.3 ALPHA-2

ALPHA-2 (sometimes called the “Atom Trap”) is the part of the apparatus where antihydrogen formation, trapping and study takes place. It consists of a cryogenic Penning-Malmberg trap, for storing and manipulating charged particles (primarily antiprotons and positrons), as well as a magnetic minimum trap for storing neutral antihydrogen atoms. ALPHA-2 also has laser access ports, which allow laser light to be injected into the central region of the trap to illuminate trapped antihydrogen atoms for spectroscopy measurements, or to laser cool particles within the trap. ALPHA-2 is surrounded by an annihilation detector (section 2.9), which can detect antiproton or antihydrogen annihilation events inside the trap.

2.3.1 The Neutral Trap

The Neutral Trap is an octupolar magnetic-minimum trap, and is used to trap antihydrogen atoms. It is formed by a set of superconducting magnets that produce a three-dimensional magnetic field minimum to trap neutral atoms: an octupole which provides radial confinement, and “mirror coils” (short solenoids) that provide axial confinement. These magnets are wound directly onto the vacuum chamber that contains the Penning-Malmberg trap, which superimposes the Neutral Trap’s trapping region in the same space as the Penning-Malmberg trap, so that antihydrogen atoms produced inside the Penning-Malmberg trap are already inside the Neutral Trap when they are born, and are trapped if their energy is low enough.

The trapping force is produced by the coupling of the magnetic moment of an atom to an inhomogeneous magnetic field, and is given by $\mathbf{F} = -\nabla(-\boldsymbol{\mu} \cdot \mathbf{B})$, for an atom with a magnetic moment $\boldsymbol{\mu}$ in a magnetic field \mathbf{B} .

An antihydrogen atom’s magnetic moment is dominated by the angular momentum of the positron, due to its mass being ~ 2000 times smaller than the mass of the antiproton. This means that the atom’s magnetic moment is well-approximated as $\boldsymbol{\mu} = -g_J \mu_B \mathbf{J} / \hbar$, with Landé g -

factor g_J , the Bohr magneton μ_B , and total angular momentum \mathbf{J} . In the ground state, $g_J \approx 2$, and the orbital angular momentum is 0, and so the only contribution to \mathbf{J} is the positron spin, which can take values of $\pm \frac{\hbar}{2}$. This gives the magnetic moment of an antihydrogen atom as $\mu = \pm \mu_B$. When μ is aligned with \mathbf{B} , the atom is an untrappable “high-field seeker”, and is attracted to regions of high magnetic field. When μ and \mathbf{B} are anti-aligned, the atom is a “low-field seeker”. As the Neutral Trap forms a 3-D magnetic minimum, these low-field seeking atoms are pulled towards the centre of the trap, where the magnetic field is lowest.

2.4 The Positron Accumulator

ALPHA’s positron accumulator is a buffer gas Penning trap based on a design by C. Surko and co-workers [28], and its primary purpose is to produce positrons that can be used for antihydrogen formation. Positrons are produced by decay of ^{22}Na into ^{22}Ne , which proceeds around 90% of the time via emission of a positron:



Positrons are continuously produced by the source at energies of up to a few hundred keV, which is too high to be effectively trapped and cooled in the accumulator. The accumulator’s role is therefore to slow the positrons down to the $< 100\text{eV}$ range such that they can be trapped and used for antihydrogen formation, and to transform the continuous flux from the radioactive source into a pulsed source, which can produce a tight bunch of positrons on-demand. A schematic of the positron accumulator can be seen in figure 2.1.

The accumulator consists of the ^{22}Na positron source, and a three-stage Penning-Malmberg trap which acts as an accumulation region. The source is cooled to cryogenic temperatures, and a layer of solid neon is frozen directly onto the surface of the source, which acts as a moderator to reduce the kinetic energy and energy spread of positrons passing through it. Most positrons annihilate in the neon moderator, but around 1% pass through and exit the moderator with a kinetic energy of around 80eV, and a much smaller energy spread than the incoming unmoderated positrons. These moderated positrons are then magnetically guided into the staged Penning trap. Nitrogen buffer gas is introduced into the trap through a gas inlet in stage 1, where it diffuses down through stages 2 and 3 before being pumped out. Radial size changes along the length of the trap, as well as differential pumping, cause a decreasing pressure gradient from stage 1 to stage 3. As positrons bounce along the length of the trap, they lose energy through collisions with nitrogen molecules and, due to the stepped potentials applied to the electrodes across the three sections, are eventually confined to stage 3. This process is shown in figure 2.2.




Figure 2.2: Three-stage potentials in the positron accumulator’s Penning-Malmberg trap. Shaded/unshaded regions show the axial extents of the trap electrodes. Positrons enter from the Na-22 source on the right and enter the trap, where they lose energy through collisions with the nitrogen buffer gas. The pressure gradient is created by differential pumping and by using electrodes with increasing radii from stage 1 to stage 3. After the desired number of positrons has been accumulated in stage 3, the remaining buffer gas is quickly pumped out, and the particles can be ejected towards the rest of the experiment.

In stage 3, the positrons are continuously compressed by a rotating wall electrode during an accumulation period, which typically lasts from a few tens of seconds to a few minutes.

Once the desired number of positrons has been accumulated in stage 3, the nitrogen is quickly pumped out of the system by a pair of cryopumps and the positrons are ejected from the positron accumulator by dropping the confining potential at the end of stage 3, and are then guided into the main experiment with solenoids.

2.5 ALPHA-g

ALPHA-g is an ongoing extension to the experiment, with the goal of measuring the force of gravity on antihydrogen atoms. Similarly to ALPHA-2, ALPHA-g consists of a Penning-Malmberg trap with an overlaid magnetic minimum trap such that it can confine both neutral antihydrogen atoms and non-neutral particles. The measurement scheme in ALPHA-g involves accumulating antihydrogen atoms in a similar manner to ALPHA-2, and then ramping down the magnetic trap and watching how they escape. Although ALPHA-g is largely irrelevant to the experimental work described in this thesis, its installation necessitated the installation of the ALPHA-g Beamline, a magnetic beamline designed to transport particles between the positron accumulator, ALPHA-2 and ALPHA-g. As such, ALPHA-g won’t be discussed in much detail.

2.6 The ALPHA-g Beamline

With the installation of ALPHA-g, the positron accumulator was moved several metres away from ALPHA-2 to make room for ALPHA-g, and the space was filled by a magnetic beamline which allows transport of charged particles between the different parts of the experiment (refer to figure 2.1). The beamline is responsible for transporting particles between the positron accumulator, ALPHA-2, and ALPHA-g. Since transfers through the beamline are between Penning traps, the energy scale of the particles is the same as for trapped particles, which is $\lesssim 100\text{eV}$. The beamline uses solenoids of varying length to produce axial magnetic fields that guide charged particles along its length. In order to steer particles around the 90-degree bend into the vertical ALPHA-g trap, a more complex array of magnets is used. A more thorough description of the beamline is given by M.A. Johnson [26].

Prior to the installation of ALPHA-g and the Beamline in 2018, positrons had only a short distance to travel from the accumulator into ALPHA-2. The new configuration, however, has positrons travelling through several metres of beamline as they travel from the accumulator to ALPHA-2, giving the positron bunch a much longer time to spread out axially as they travel. This results in a much lower catching efficiency for positrons when compared to the old configuration as the increased bunch length is longer than the trap it is being caught in. Current transfer efficiency of positrons from the accumulator to ALPHA-2 is around 16%: up to 30M positrons are accumulated in ~ 100 seconds, but only around 5M are caught in ALPHA-2.

2.7 Diagnostic Tools

The ALPHA apparatus has several different tools for measuring various properties of the different particles that are trapped in ALPHA. This section will initially discuss the Faraday Cup and Microchannel Plate, which are the primary tools used for plasma diagnostics and were used for the bulk of the experimental measurements reported in this thesis, before touching on other diagnostics used within ALPHA for other measurements.

2.7.1 Faraday Cup

A Faraday Cup (FC) is a device that can be used to count charged particles. The operation principle of the FC is simple: a conductive plate is connected to an RC circuit, and when particles strike the plate they induce a voltage in the plate. As long as the bunch length of the incident particles is much shorter than the RC decay time of the circuit, the total charge, Q , of the particle bunch can be determined. A typical signal from a FC is shown in figure 2.3, with a sharp




Figure 2.3: A typical FC signal. The large peak corresponds to a bunch of charged particles hitting the FC, and the long exponential decay is due to the RC circuit connected to the FC. The total charge of the particle bunch is equal to the maximum charge measured by the FC.

peak produced by an incident bunch of charged particles, followed by a slower RC decay of the circuit. The circuit response is of the form

$$V(t) = V_{max}e^{-t/RC}, \quad (2.2)$$

and the total charge, and therefore total number of particles, can be extracted by looking at the peak voltage,

$$qN = CV_{max}, \quad (2.3)$$

where q is the charge of a single particle and N is the number of particles.

A typical measurement with a FC involves ejecting all of the particles out of the trap and allowing them to strike the FC, where they are counted.

2.7.2 Microchannel Plate and Phosphor Screen Assembly

A microchannel plate (MCP) is a device that converts a single incident charged particle (or high-energy photon) into a cascade of secondary electrons. This provides effective gain on the signal, as the cascade of electrons is much easier to detect than the original incident particle. The MCP is a disk a few millimetres thick of an insulating material, typically glass, which contains an array of tiny holes (channels). The inner wall of each channel is then coated with a semiconductor layer that allows for production of secondary electrons within each channel when a charged particle strikes the wall. The front and back of the MCP are coated with a thin metal layer, allowing the front and back to be electrically biased separately. When a particle enters one of the MCP

channels, it strikes one of the inner walls and produces secondary electrons, which are accelerated through the MCP by the potential difference between the front and back and can strike the walls again, causing a cascade of secondary electrons from the single incident particle. Since there are many independent channels, spatial information about where the incident particle reached the MCP is retained. Figure 2.4 shows a sketch of a single MCP channel and an electron cascade, as well as a picture of an MCP assembly used in ALPHA.

The MCPs used at ALPHA are mounted in an assembly with a phosphor screen, which is used to detect the electrons exiting the MCP back. The phosphor screen can be imaged with a CCD camera, allowing for an image of the radial profile of an incident plasma to be produced, as in figure 2.5. Alternatively, the total charge arriving at the plate can be analysed in a Faraday Cup-style measurement, which loses the spatial information from the MCP but can measure particle arrival times at the MCP with quite a high resolution. A Silicon PM can also be used to measure the amount of light emitted by the phosphor screen to perform a Faraday Cup-style measurement. Looking at the light coming off the phosphor screen can produce a much cleaner signal than measuring the charge reaching the screen [29].




Figure 2.4: a) A single MCP channel. Here, a single incident electron causes a cascade of secondary electrons. A potential difference of up to 1000V is applied between the MCP's front and back, which provides acceleration for the cascade. b) An MCP and phosphor screen assembly currently used in ALPHA. The rectangular tabs are the connection points for the phosphor screen, MCP back, and MCP front (going clockwise).




Figure 2.5: a) A positron plasma imaged with an MCP and phosphor screen assembly. b) A large Be^+ plasma imaged on an MCP. The aperturing in the bottom-left and bottom-right sections of the image are caused by electrostatic potentials on the MCP’s tabs and bolts – a hardware defect that is discussed in chapter 6.5.5.

2.8 Vacuum Translators

In order to make a FC or MCP measurement, particles need to be ejected along the trap’s axis, where they can strike the FC or MCP. This causes a conflict, as when initially loading the trap with particles, they also need to enter on-axis. The solution implemented in ALPHA is a linear translator inside the vacuum chamber (referred to as a “stick”) that can move the diagnostic tools away from the trap axis for particle loading and back in place for a measurement.

There are many sticks within ALPHA, and a stick is typically equipped with at least one FC and MCP for measurements, and a hollow, cylindrical passthrough electrode to allow particle loading, which is typically grounded. Other tools may be found on a stick depending on its location: an electron gun for loading electrons into the trap, a beryllium source for loading beryllium ions, or mirrors to allow laser access along the axis of the trap.

2.9 Annihilation Detectors

One benefit of working with antimatter particles is that they are much easier to detect than regular matter, since they produce very energetic annihilation events when they come into contact with regular matter (usually by coming into contact with the trap walls). The Silicon Vertex Detector (SVD) is the main tool used for measurements with antihydrogen atoms. It is able to




Figure 2.6: Figure reproduced from ALPHA collaboration [32]. Reconstructed events from ALPHA’s SVD. Left is an antiproton annihilation event, while right is a cosmic ray. The inner circle shows the location of the inner surface of the Penning-Malmberg trap electrodes. The outer three concentric layers are the layers of the SVD modules. Red points show where a high-energy particle has been detected passing through a layer, and red lines are fitted tracks. The blue diamond is the reconstructed position of the annihilation vertex.

resolve single antiproton or antihydrogen annihilation events, and provides precise spatial and temporal information about the annihilation event.

The SVD consists of three concentric layers of silicon microstrip modules which surround the Neutral Trap. When an antihydrogen atom inside the trap annihilates, it typically produces between 2 and 4 charged pions [30], which pass through the trap walls and the SVD. Each charged pion is detected as it passes through the three layers of the detector, and the track it followed can be reconstructed. By analysing all the tracks produced by an annihilation event, the position of the annihilation event can be determined. Other charged annihilation products can also produce tracks in the SVD, but the majority of reconstructed tracks are produced by charged pions [31]. Figure 2.6 shows a sketch of the SVD, as viewed while looking down the axis of the trap.

The main source of background is from cosmic rays passing through the detector. Cosmic tracks qualitatively differ in appearance when compared to annihilation product tracks, and are filtered out using multivariate analysis [10]. Figure 2.6 shows two example events: one antiproton annihilation and one cosmic ray.

Scintillators are also used in ALPHA for detecting annihilation products. They can be placed outside of the trap, and produce light when radiation from positron or antiproton annihilation products pass through them. They are typically used for measuring particle losses during transport through the experiment, bunch structure measurements and particle counting. Many small caesium iodide (CsI) detectors are placed around the experiment, which con-

tain a CsI scintillator and a photodiode for readout. These can be used to make sure that too many particles are not annihilating during transport around the experiment and, since there are many of them, the position of losses can be roughly triangulated. Larger plastic scintillators are used for measuring bunch structure and particle counting. These scintillators are connected to a PMT and placed near to one of the sticks, where a bunch can be intentionally annihilated on one of the stick components (an MCP for example), and the time structure of the annihilation signal can be read off.

3/ Dynamics of Trapped, Charged Particles

Practically all particles used for physics within ALPHA are charged, with the only notable exception being the antihydrogen atom. Penning-Malmberg traps [33] (figure 3.3) are the workhorses of ALPHA, and several of them are used throughout the apparatus to trap and manipulate many different charged species. As such, understanding how these charged particles behave within these traps is critical to understanding any experiment performed at ALPHA.

In a charged particle trap, the motion of a particle is determined by the electromagnetic fields imposed by the trap. When more than one particle is present, the particles interact with each other through the Coulomb force, which can change their dynamics within the trap. When the interactions between particles begin to play a significant role in particle dynamics, the particles form a plasma. Unlike individual particles, a plasma is a highly-coupled system that behaves more like a fluid, and exhibits collective phenomena such as waves. Most collections of trapped, charged particles used in experiments at ALPHA are strongly-coupled enough to be considered a plasma.

Most plasmas found in nature are neutral plasmas, consisting of ionised atoms and free electrons, and have a net charge of zero. Plasmas used in ALPHA, however, are non-neutral plasmas that consist entirely of particles with the same sign of charge (such as a pure electron plasma), giving the plasma as a whole a non-zero net charge.

In this chapter, I will discuss the physics behind single-particle and non-neutral plasma dynamics in a Penning-Malmberg trap, and go into detail about some important manipulations that are performed on trapped charged particles in ALPHA's trap.

3.1 Single-Particle Motion in a Penning-Malmberg Trap

An ideal Penning trap is formed by a solenoidal magnetic field \mathbf{B} and a quadratic electrostatic potential ϕ [33]:

$$\mathbf{B} = B_0 \hat{z}, \quad (3.1)$$

$$\phi(r, z) = \frac{V_0}{R^2}(2z^2 - r^2), \quad (3.2)$$

where V_0 is the applied trap potential, and R is a measure of the size of the trap. This can be produced by a set of three electrodes: two endcaps and a ring, shown in figure 3.1, in which case $R^2 = r_0^2 + 2z_0^2$, where r_0 and z_0 are the distances from the trap centre to the ring and endcap electrodes respectively. In an ideal Penning trap, the shapes of these electrodes are hyperboloids of revolution in order to produce a harmonic electrostatic potential.

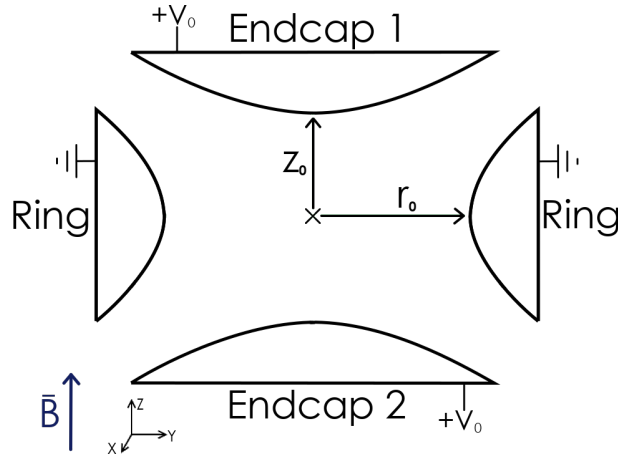


Figure 3.1: A Penning trap, formed from two endcap electrodes and a single ring electrode, as well as a constant magnetic field along z . For a positive particle, the endcaps are held at a positive potential $+V_0$, while the ring is held at ground. This configuration produces a saddle electric potential where the particles are confined in the z -direction and anti-confined in the r -direction. A magnetic field $\mathbf{B} = B_0 \hat{z}$ provides a trapping force in the r -direction to overcome the force from the electrodes, producing a 3-D electrostatic potential well.

The equation of motion for a particle with charge q and mass m is given by the Lorentz force:

$$m\ddot{\mathbf{r}} = q(-\nabla\phi + \dot{\mathbf{r}} \times \mathbf{B}), \quad (3.3)$$

with $\mathbf{r}=(x, y, z)$. Split into the x , y , and z components, and with a purely axial magnetic field $\mathbf{B} = B_0 \hat{z}$, this gives

$$\ddot{x} = \frac{\omega_z^2}{2}x - \omega_c \dot{y}, \quad (3.4)$$

$$\ddot{y} = \frac{\omega_z^2}{2}y + \omega_c \dot{x}, \quad (3.5)$$

$$\ddot{z} = -\omega_z^2 z. \quad (3.6)$$

Here, $\omega_c = qB_0/m$ is the cyclotron frequency. The axial motion is just simple harmonic motion with a frequency $\omega_z = \sqrt{4qV_0/mR^2}$, but the radial motion is slightly more complex. By making the substitution $u = x + iy$, the \hat{x} and \hat{y} motions can be combined into a single equation,

$$\ddot{u} - i\omega_c \dot{u} - \frac{\omega_z^2}{2}u = 0, \quad (3.7)$$

which has a solution of the form $u = e^{-i\omega_{\pm}t}$, where

$$\omega_{\pm} = \frac{1}{2}(\omega_c \pm \sqrt{\omega_c^2 - 2\omega_z^2}). \quad (3.8)$$

This splits into two superimposed motions in the $x - y$ plane. The lower-frequency solution, ω_- (or ω_m), is known as the magnetron motion, which comprises a slower, larger rotation around the trap axis, with a drift velocity given by $\mathbf{v} = \mathbf{E} \times \mathbf{B}/|\mathbf{B}|^2$. For this reason, it is also referred to as $\mathbf{E} \times \mathbf{B}$ drift. The higher-frequency solution is the modified cyclotron motion, and is more commonly referred to by $\omega'_c = \omega_c - \omega_m$. These cyclotron orbits are much smaller circular orbits around magnetic field lines.

The result is that there are three distinct oscillatory motions that a particle undergoes in a Penning trap, which are shown in figure 3.2. The particles undergo an axial bounce motion along the trap at a frequency ω_z , and undergo two superimposed motions in the radial direction: small cyclotron oscillations at a frequency ω'_c , and the slower magnetron oscillations at a frequency ω_m . Typically, $\omega_m < \omega_z < \omega'_c$.

The Penning-Malmberg trap is a variant of the regular Penning trap that uses hollow cylindrical electrodes instead of hyperbolic electrodes. The potentials formed by these electrodes are no longer the ideal harmonic potentials produced in a regular Penning trap, but this configuration is still easily able to produce a stable trapping potential. The cylindrical electrode stack is open at both ends, allowing for easy particle transport in and out of the trap along its axis. This configuration, shown in figure 3.3, allows for many more electrodes to be used, allowing particles to be moved dynamically along the length of the trap by altering the electrostatic potentials applied. With many electrodes, several trapping regions can also be formed within the trap, allowing, for example, two different species to be trapped in neighbouring potential wells and then merged together into a single potential well.




Figure 3.2: Single-particle motion in a Penning trap. In the radial $x - y$ plane, the particle undergoes large, slow magnetron rotations around the centre of the trap, as well as small, fast cyclotron rotations. The particle also undergoes simple harmonic oscillation in the axial (z) direction (not pictured).

3.2 Characteristics of a Plasma

Looking at single-particle dynamics gives a good idea of what is happening in the trap, but does not take into account interparticle interactions. When there are many particles in the same trap, and these interparticle interactions play an important role in the particle dynamics, it is sensible to instead treat them as a plasma.

An important feature of a plasma is that it produces an electric field that shields externally-applied electric fields. If a test charge is introduced into a plasma, the other particles will rearrange themselves so that they counteract the change in potential caused by the test charge, which will shield the rest of the plasma from the test charge. This is called Debye shielding, and is characterised by the Debye length,

$$\lambda_D = \sqrt{\frac{\epsilon_0 k_B T}{n q^2}}, \quad (3.9)$$

where ϵ_0 is vacuum permittivity, k_B is the Boltzmann constant, T is the plasma temperature,




Figure 3.3: A sketch of a simple five-electrode Penning-Malmberg trap, with the middle three electrodes cut away for viewing clarity. A spheroidal plasma (blue) is trapped between the two biased end electrodes, while the central three electrodes are held at ground. A magnetic field pointing along the z -axis is provided by an external solenoid, not pictured due to its size. Spacing between electrodes has been exaggerated for clarity – in a real trap, the electrodes are built to overlap slightly, and are physically separated by ruby spacers.

and n is the plasma density. The Debye length describes the distance over which the plasma shields an applied potential by a factor of $1/e$ (where e is Euler's number), and is one measure of whether a collection of charges is a plasma or not: a plasma must be significantly larger than λ_D in all dimensions. A related parameter is the plasma frequency, given by

$$\omega_p = \sqrt{\frac{nq^2}{\epsilon_0 m}}, \quad (3.10)$$

which sets the time scale over which plasma behaviour is observed. A typical particle in the plasma will travel one Debye length in a time $1/\omega_p$. A typical positron plasma used for antihydrogen production might have 2.5M positrons, $n = 2 \times 10^{14} \text{ m}^{-3}$, and $T = 20 \text{ K}$. This plasma has a Debye length of $\lambda_D = 20 \mu\text{m}$, and a plasma frequency of $\omega_p = 0.13 \text{ GHz}$.

In a trapped plasma, energy can be exchanged between the axial and radial motions through collisions. The rate at which energy is transferred between the two motions is parameterised by the magnetisation parameter [34],

$$\bar{\kappa} = \frac{\bar{b}}{r_c} \frac{1}{\sqrt{2}}, \quad (3.11)$$

where $r_c = \sqrt{\frac{k_B T}{m}} / \omega_c$ is the cyclotron radius and $\bar{b} = 2q^2 / 4\pi\epsilon_0 k_B T$ is twice the classical distance of closest approach, and the Glinsky collision rate [34]

$$\nu_{\perp\parallel} = n\bar{v}\bar{b}^2 I(\bar{\kappa}), \quad (3.12)$$

where $\bar{v} = \sqrt{2k_B T/m}$ is the thermal velocity. Asymptotic expressions for $I(\bar{\kappa})$ are provided by theory for the limits $\bar{\kappa} \gg 1$ and $\bar{\kappa} \ll 1$. The strongly magnetised regime occurs when $\bar{\kappa} \gg 1$, and the collision rate is exponentially small ($I(\bar{\kappa}) \propto \exp[-5(3\pi\bar{\kappa}^{2/5}/6)]$). The weakly magnetised regime occurs when $\bar{\kappa} \ll 1$, and in this regime $I(\bar{\kappa}) \propto \ln(\bar{\kappa})$. Both expressions, as well as data points from numerical calculations that bridge the $\bar{\kappa} \sim 1$ region, are discussed in detail by Glinsky et al. [34]. Results from the strongly and weakly-magnetised expressions, as well as numerical calculations performed by Glinsky et al., are shown in figure 3.4.

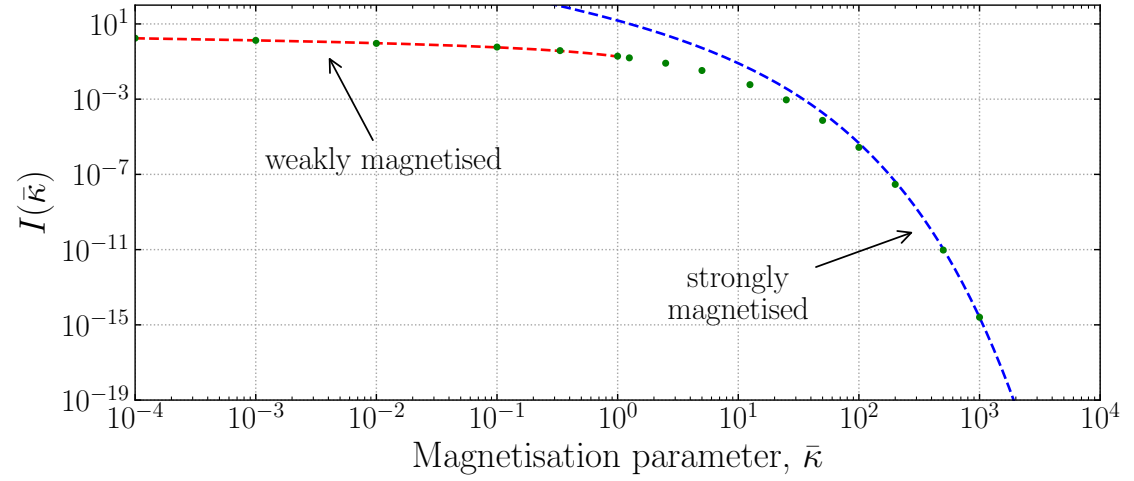


Figure 3.4: $I(\bar{\kappa})$ as a function of magnetisation parameter. The asymptotic expressions for the strongly magnetised (blue) and weakly magnetised regime (red) are shown along with numerical calculations (green). Both expressions and the data are reproduced from Glinsky et al. [34].

The magnetisation parameter $\bar{\kappa}$ is proportional to the magnetic field strength, and inversely correlated with the plasma temperature ($\bar{\kappa} \propto T^{-3/2}$). As $\bar{\kappa}$ gets very large, for example with a very cold plasma, the cyclotron radius r_c gets very small, and the plasma can form a crystal lattice structure, with each particle within the crystallised plasma performing small cyclotron orbits around its lattice site. In this very strongly magnetised regime, $\nu_{\perp\parallel}$ is very small, and the plasma is slow to reach thermal equilibrium. This can give rise to plasmas with non-thermal velocity distributions, and plasmas where the rate of transfer of energy between the axial motion and the radial motion is very low. A plasma in this regime can have different temperatures in the axial and radial directions, i.e. $T_{\parallel} \neq T_{\perp}$.

The typical positron plasma mentioned above ($N = 2.5 \times 10^6$, $n = 2 \times 10^{14} \text{ m}^{-3}$, $T = 20\text{K}$) would have a magnetisation parameter of $\bar{\kappa} = 12$ when immersed in a 1T magnetic field, and a collision rate $\nu_{\perp\parallel} = 83\text{kHz}$.

3.3 Trapped Plasma Dynamics in a Penning-Malmberg trap

In the plasma regime, the charged particles that make up the plasma move according to the imposed fields of the trap, but the presence of the other charges in the trap modifies the electric field around the plasma. For a plasma in thermal equilibrium, the electric potential and charge distribution must satisfy Poisson's equation

$$\nabla^2\Phi(r, z) = -\frac{qn(r, z)}{\epsilon_0}, \quad (3.13)$$

and the Maxwell-Boltzmann distribution

$$n(r, z) = n_0 \exp\left[-\frac{q\Phi(r, z)}{k_B T}\right]. \quad (3.14)$$

By considering only the direction parallel to the magnetic field (the axial direction), it is easy to see that an electric field along the plasma will cause a flow of charge. Charge will flow until there is no electric field along the plasma, which flattens the potential well holding the plasma. The level at which the potential well is flattened is referred to as the “space charge” (or “self-charge”). The solutions for the three-dimensional case are qualitatively similar.

A plasma in a quadratic potential has an ellipsoidal “cigar” shape, but is often approximated as a long cylinder with a uniform density. A cylindrical non-neutral plasma has a radial self-field of

$$E_r = \left(\frac{qr}{2\epsilon_0}\right)n, \quad (3.15)$$

with radius r , length L , and density n . Inside a Penning-Malmberg trap, the trap's magnetic field in combination with the self-field from the plasma cause the plasma to rotate at the $\mathbf{E}_r \times \mathbf{B}$ frequency, which is given by [33],

$$f_{E \times B} = \frac{E_r}{2\pi r B} \quad (3.16)$$

$$= \frac{qn}{4\pi\epsilon_0 B}. \quad (3.17)$$

This is very similar to the single-particle $\mathbf{E} \times \mathbf{B}$ motion, except the electric field of the plasma dominates over the field of the trap. Note that with a uniform density, the rotation is independent of radius: the plasma rotates as a rigid rotor, with a rotation velocity proportional to the density.

3.4 Perturbations and Heating in a Real Trap

In a real trap, there are many effects that disturb a trapped plasma, and these tend to manifest as plasma heating and expansion. Small positional imperfections in the trapping fields, known as patch potentials, can interfere with plasmas in the trap, and have recently been quantified in the ALPHA trap by C. Baker et al. [35]. These patch potentials occur when the electrode surfaces are not perfect equipotentials, which can occur due to surface contamination or other effects. When the inside of the ALPHA-2 Penning-Malmberg trap is continuously illuminated with UV light, which occurs when performing measurements of the 1S-2S transition in antihydrogen, it has been observed that the magnitude of the patch potentials increases, leading to increased plasma heating and a consequently lower antihydrogen trapping rate. The patch potentials can be reset back to their initial magnitude by warming the trap up to room temperature, and then cooling it back down to cryogenic temperatures. The effect that this has on experimental procedures within ALPHA can be seen in recent publications [10]: after a few hours of cumulative illumination with the 243nm laser, the antihydrogen trapping rate drops enough to warrant the rather slow process of heating and re-cooling the apparatus. Although this effect is not well-understood, its disappearance when the trap is heated and re-cooled hints that it could be due to surface contaminants, that evaporate when the trap is warmed from its cryogenic state.

Other effects, such as thermal radiation from room-temperature surfaces entering the trap from the open cylindrical ends and electronic noise on the trap electrodes are also likely to cause plasma heating and expansion, although these are less well-quantified. The effects of these imperfections are typically much more noticeable when the trapping fields are small (i.e. when they are in a shallow potential well), as the magnitude of the disturbance becomes closer to the magnitude of the trapping fields. Plasmas are generally kept in short, deep potential wells for as long as possible to minimise heating and radial expansion caused by these kinds of effects.

Another effect that can lead to heating and expansion in the trap is the addition of the Neutral Trap magnetic fields. The Neutral Trap uses an octupole to create radial magnetic fields that are used to radially confine antihydrogen atoms, but this also causes the Penning-Malmberg trap's solenoidal field lines to bend outwards towards the electrodes. Above a critical radius these field lines will intersect the trap electrodes, and charged particles following these field lines will collide with the electrodes and be lost from the trap. As well as this sharp cutoff, plasma heating and radial expansion is commonly observed in plasmas significantly smaller than the critical radius, which may be due to resonant effects between the plasma oscillations and the octupole field [36]. Interference from the octupole field on trapped plasmas is minimised in practice by keeping plasma radii significantly smaller than the critical radius, and generally avoiding holding plasmas for prolonged periods of time inside the octupole's field.

3.5 Rotating Wall Compression

Rotating Wall (RW) compression is a technique used to manipulate the radial size of a trapped plasma, first demonstrated by Huang et al. in 1997 [37]. A rotating electric field is introduced that can apply torque to a plasma, which changes its rotation frequency and therefore, by equation 3.16, its radial size. It was initially developed as a tool to increase plasma confinement times: effects such as collisions with background gas and trap imperfections tend to apply a drag to a rotating plasma, slowing and expanding the plasma until particles are eventually lost radially. Counteracting this drag with a positive torque allows for close control of a plasma's size and density, and an essentially unlimited lifetime in the trap.

In principle, any electric field rotating at a higher frequency than a plasma will apply a torque, but Debye shielding can mean that the actual torque applied is fairly negligible. The amount of torque that is produced can be greatly enhanced by coupling to rotating modes within the plasma (such as Trivelpiece-Gould modes [38]), which can be resonantly driven by the rotating wall, allowing a much higher torque to be transferred from the rotating wall to the bulk of the plasma. In practice, it is often easier to couple to these modes from one end of the plasma, and so plasmas are normally axially offset from the RW electrode such that only one end is immersed in the RW field. There are two distinct regimes that the system can be in, which depends on the amplitude of the rotating electric field. In the “weak drive” regime, compression is observed only at discrete RW frequencies that correspond to rotating plasma modes. In the “strong drive” regime [39], which uses a significantly higher-amplitude RW field, plasma compression is observed over a wide frequency range, and is unaffected by resonances with plasma modes. In the strong drive regime, the plasma rotation frequency is locked to that of the RW drive.

These rotating electric fields are achieved using a modified electrode (figure 3.5) which is azimuthally segmented, allowing a voltage to be applied to each segment independently. A rotating electric field is produced by applying a sinusoidally-varying voltage to each segment, and phase-shifting the sine wave by $\theta_j = 2\pi j/N$, where N is the number of segments and j is an integer that labels each segment. RW electrodes typically have 4, 6 or 8 segments, and ALPHA's RW electrodes have 6 segments.

To produce a rotating electric field, the voltage applied to the j^{th} electrode is given by

$$\phi_j = A \cos[m(2\pi f_{RW}t \pm \theta_j)], \quad (3.18)$$

with voltage amplitude A , frequency F_{RW} , and phase shift θ_j . m corresponds to the number of the mode of the plasma: $m=1$ is the dipole mode, and $m=2$ is the quadrupole mode.




Figure 3.5: An end-on sketch of a segmented RW electrode. Each segment is biased independently with a sinusoidally-varying voltage, which is phase-shifted for each segment. The phase-shift is labelled for a) a dipole drive and b) a quadrupole drive. The arrows represent the direction of the electric field.

The rotating electric field does work on and heats the plasma. As plasma temperatures increase, plasma modes are damped and plasma collisions are reduced, which both work to reduce coupling of a RW to the plasma [33]. In order to provide consistent and effective compression, a cooling mechanism is required to counteract the heating caused by the RW. For light species, such as electrons and positrons, cyclotron cooling (see section 3.6 below) is usually sufficient in the high magnetic fields used in ALPHA. Heavier species, such as antiprotons or beryllium ions, require an additional cooling mechanism - such as sympathetic cooling from a lighter species' cyclotron cooling, or direct laser cooling.

In a two-species plasma, sympathetic compression can take place: as one of the species is directly compressed, the other species is compressed via Coulomb interactions with the directly-compressed species. This has been observed with antiproton-electron mixed plasmas in ALPHA [40].

3.6 Cyclotron Cooling

Maxwell's equations tell us that charged particles that are accelerating emit energy in the form of radiation, and so trapped particles must also emit radiation as they oscillate. This radiation is dominated by emission from the cyclotron motion in a Penning-Malmberg trap, as the cyclotron

frequency is typically much larger than the magnetron and axial bounce frequencies.

The power emitted by an accelerating charge can be given by the Larmor formula, which can be used to produce an equation for characteristic cooling time due to cyclotron cooling,

$$\tau = \frac{3\pi\epsilon_0 m^3 c^3}{q^4 B^2}. \quad (3.19)$$

For reasonable magnetic fields ($B = 1T$), an electron can cool in a matter of seconds. Unfortunately, the cooling time is proportional to the cube of the particle's mass, so an (anti)proton in the same trap would take a few hundred years to cool via cyclotron emission alone. Practically, this means that only electrons and positrons efficiently self-cool via cyclotron radiation – all other particles that are trapped in ALPHA have a significantly higher mass, and their cyclotron cooling is negligible and can be ignored. Given the strong $1/B^2$ scaling, the magnetic field is maximised when high levels of cyclotron cooling are required. ALPHA uses a 1T solenoid for the ALPHA-2 trap, but has several regions where the magnetic field can be increased to 3T, to increase the effect of cyclotron cooling.

Electrons or positrons will cyclotron cool until they reach thermal equilibrium with black-body radiation coming from the trap walls, ignoring other heating sources. The ALPHA Penning-Malmberg trap electrodes are in thermal contact with liquid helium, and so slightly above 4K. In practice, it is observed that electron and positron plasmas will reach thermal equilibrium at around 20-30K, due to other heating mechanisms outlined in section 3.4.

3.7 Evaporative Cooling

Evaporative cooling (EVC) is a method of cooling a plasma that works by selectively removing high-energy particles from the trap. For a plasma in thermal equilibrium, the energy of the particles will follow a Maxwell-Boltzmann distribution. By removing particles above a certain cutoff energy, the distribution is truncated at this energy. Particles that remain in the trap re-thermalise through collisions to a lower temperature, which happens continuously as the evaporation is taking place. EVC is performed in a Penning-Malmberg trap by lowering one side of the potential well holding a plasma until the most energetic particles are able to escape the well, as in figure 3.6.

Lowering the potential barrier further causes more particle loss, and a colder plasma. EVC can be a very effective cooling method, and was instrumental in the first realisations of Bose-Einstein condensates [41, 42], but it does have some major drawbacks. The most obvious is that in order to achieve more cooling, more particles must be removed from the trap, which can be especially unwanted when dealing with very limited amounts of antimatter. EVC in a Penning-

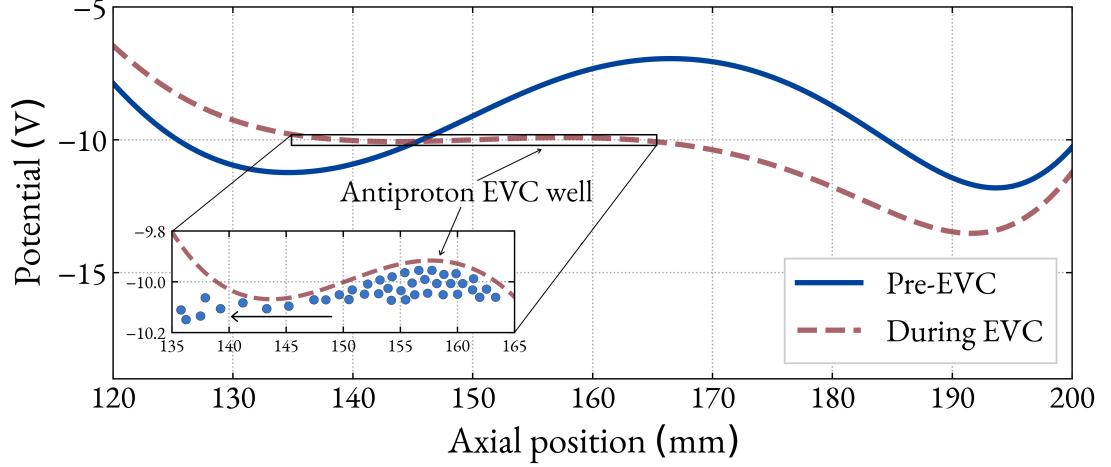


Figure 3.6: On-axis potentials used during evaporative cooling of antiprotons. The much deeper initial well (blue) is gradually morphed into a much shallower well (dashed red) during the EVC. The inset shows a zoomed-in version of the final EVC well, with energetic antiprotons (shown as blue circles) being lost to the left. The transition between the initial well and the final well is smooth, and takes 10s.

Malmberg trap also causes radial expansion of the plasma due to particles being preferentially lost from the centre of the plasma, where the confining potential is lowest. Particles lost from close to the the centre of the plasma carry away very little angular momentum, and so in order to conserve angular momentum, the plasma must increase in radius. In the ideal case, the expansion of an evaporatively-cooled plasma can be given by

$$\frac{N_0}{N} = \frac{\langle r^2 \rangle}{\langle r_0^2 \rangle}, \quad (3.20)$$

where N and N_0 describe the final and initial number of particles, and r and r_0 describe the final and initial plasma radius.

The evaporation timescale characterises the timescale over which evaporation takes place, and is derived in Ref. [7] as

$$\tau_{ev} \nu_{\perp\parallel} = \frac{\sqrt{2}}{3} \eta e^\eta, \quad (3.21)$$

where τ_{ev} is the evaporation timescale, $V = \eta k_B T$ is the final well depth, and $\nu_{\perp\parallel}$ is the plasma collision rate (equation 3.12).

3.7.1 Strong Drive Regime Evaporative Cooling

Strong Drive Regime Evaporative Cooling (SDREVC) is a technique that combines a strong drive RW (3.5) with EVC (3.7) to enable precise control of the number of particles in a plasma. It was recently pioneered in ALPHA by C. Carruth et al. [43,44]. RW compression is considered to be in the strong drive regime when the plasma's rotation frequency is locked to the frequency of the RW. This occurs with high-amplitude RWs, and is typically able to control a plasma's rotation frequency (and therefore plasma density) across a large frequency range [39].

The theory behind SDREVC is fairly simple: for an infinitely long, zero-temperature cylindrical plasma, the on-axis space charge ϕ_c of the plasma can be described by [43]

$$\phi_c = \frac{nqr_p^2}{4\epsilon_0} \left[1 + 2 \ln \left(\frac{R_w}{r_p} \right) \right], \quad (3.22)$$

with R_w and r_p as the electrode and plasma radii respectively. There are only three free parameters in this equation: the space charge ϕ_c , the plasma radius r_p , and the plasma density n , and so controlling two of these allows control of the third. EVC can be used to control ϕ_c – it will simply be equal to the depth of the potential well (the cooling from EVC is not important in this application). Strong drive RW can control n by changing the drive frequency, and so there must be only one solution for r_p . Controlling both r_p and n allows control of the number of particles in the plasma, N , which is given by [43]

$$N = n\pi r_p^2 L, \quad (3.23)$$

where L is the plasma length, which is controlled by the axial length of the trapping region. This technique has so far been applied to positrons and electrons in ALPHA, but has proved much harder to apply to heavier species due to the lack of an efficient cooling mechanism to counteract heating from the RW. SDREVC could in theory be applied directly to beryllium ions, with laser cooling providing the requisite cooling, but this has yet to be achieved. The improvement on the stability on the number of particles achieved with SDREVC is quite drastic: electron plasmas prepared without SDREVC had a standard deviation on the number of particles of 14%, which was reduced to 3% with SDREVC [43], with similar (or even larger) improvements with positron plasmas. SDREVC is now always used when preparing positrons for antihydrogen formation in ALPHA.

3.7.2 Adiabatic Cooling

Adiabatic cooling occurs when the restoring force from the trap does negative work on the plasma to increase its volume and decrease its temperature, without a change in entropy [8]. In a slowly-varying 1-D square well the action, $L_m v_m$, is a constant, with L_m being the length of the well and v_m is the axial bounce velocity of a trapped particle. Given this, increasing L_m by stretching the length of the well will give a corresponding reduction in v_m , reducing the particle's temperature. Adiabatic cooling retains and cools all the particles in the trap, but is limited by the length of the trap. There are some practical considerations that further limit the use of adiabatic cooling: electron and positron plasmas that are adiabatically cooled will soon re-thermalise with the trap walls, and so have a fairly short time window in which they remain cool (this is also true for evaporative cooling). Very long wells, which would be necessary in more extreme adiabatic cooling, are also observed to produce high rates of particle heating, even for heavier particles.

Adiabatic cooling remains a useful tool, and positron plasmas used for antihydrogen formation are kept in a deep, narrow well until the final moments before being merged with antiprotons so that they can be adiabatically cooled into a longer, shallower well that is used during the merge, shown in figure 3.7.

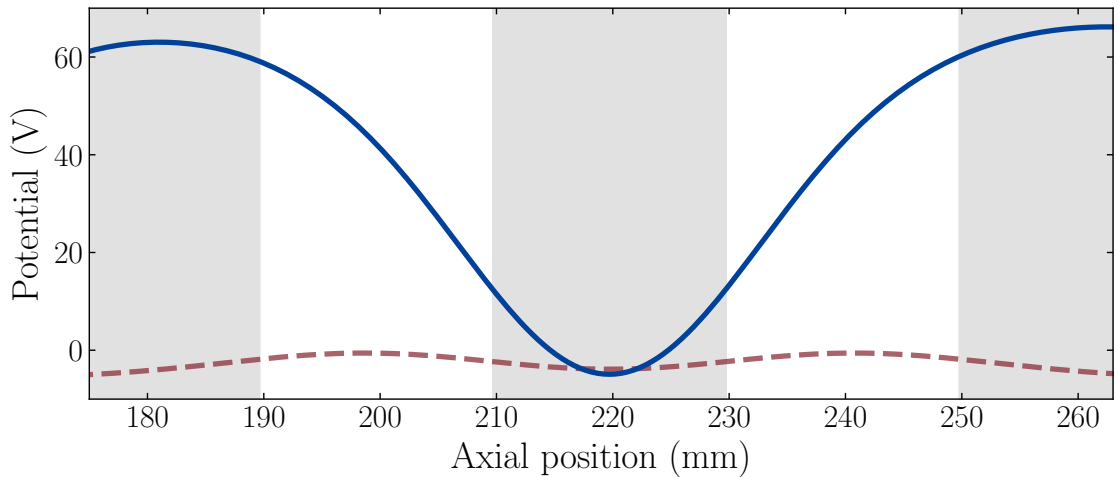


Figure 3.7: On-axis potentials used during adiabatic cooling of positrons before being used for antihydrogen formation. Grey shaded/unshaded regions show axial extent of trap electrodes. Positrons are initially held in a deeper, narrower well (blue), which is changed into a shallower well (red dashed) over 10ms. The well becoming shallower effectively lengthens it axially, which causes adiabatic cooling.

3.8 Multispecies Plasmas

Multispecies (“mixed”) plasmas are plasmas consisting of more than one species with the same sign of charge, which are co-trapped in the same potential well. The species are typically in constant contact with each other through Coulomb collisions, and so changing the parameters of one of the plasmas will affect the other. Two-species mixed plasmas are routinely used in ALPHA, and so their properties will be discussed here.

3.8.1 Centrifugal Separation

Centrifugal separation is when, in a mixed plasma, the heavier species will tend to migrate to a higher radius, with the lighter species migrating to a lower radius. In a completely separated plasma, the heavier species will form a halo around the lighter species. Figure 3.8 shows a centrifugally separated antiproton-electron plasma. The heavier antiprotons have formed a ring around the lighter electrons – note that the two species were concentric in the trap, but image in different places due to their different charge-to-mass ratios giving them different trajectories as they are ejected from the trap to be imaged [45].




Figure 3.8: Image reproduced from the ALPHA collaboration [45]. An MCP image of a mixed antiproton-electron plasma exhibiting centrifugal separation in a a) 1T and b) 3T magnetic field. While the two species are concentric in the trap, their different charge-to-mass ratios give them different trajectories as they are being ejected from the trap, causing them to image in different places. The faint halos visible near the antiproton plasmas are thought to be optical reflections occurring in the MCP.

The mechanism that drives centrifugal separation is momentum transfer in the ϕ -direction, in the cylindrical coordinate system (r, ϕ, z) . This momentum transfer is driven by the different charge-to-mass ratio of each species, causing them to tend to rotate at different rates. This causes the heavier species to migrate outwards, and the lighter species inwards, until the profiles are such that the mixed plasma as a whole rotates as a rigid body. When the two species are rotating at the same speed, there is no longer any momentum transfer in the ϕ -direction, and the two species are in thermal equilibrium.

The amount of separation between the two species can be parameterised by the “separation length”, l_{sep} , and is given by M. Knoop et al. [33] as

$$l_{sep} = \frac{k_B T}{|m_1 - m_2| \omega_r^2 R_p}, \quad (3.24)$$

where ω_r is the plasma rotation frequency, R_p is the plasma radius, and m_1, m_2 are the masses of the two species. l_{sep} can be understood as a comparison of the thermal energy to the rotational energy of the plasma, with units of length. As T decreases, l_{sep} gets small, and the two species become more separated. Separation occurs when $l_{sep} < R_p$, and complete separation occurs when $l_{sep} \ll \lambda_D$.

3.8.2 Sympathetic Cooling

When one species in a mixed plasma is cooled, the other species in the plasma can be cooled indirectly through Coulomb collisions between the two species – the two species in the mixed plasma are in thermal contact. Some very powerful cooling techniques are limited to specific species, with cyclotron cooling and laser cooling being the prime examples. An antiproton plasma does not significantly cyclotron cool, and cannot be laser cooled, as the antiproton lacks the required level structure, but it can easily be mixed with a plasma of electrons, which will self-cool via cyclotron cooling and bring the antiproton plasma’s temperature down with them. Sympathetic cooling is not a plasma-specific phenomenon, and can be utilised in traps with only a few particles: the main requirement is that the species are in thermal contact.

Sympathetic cooling of plasmas can be limited by centrifugal separation. As a mixed plasma gets colder, the degree of separation between the two plasmas increases (by equation 3.24), therefore decreasing the overlap and energy exchange between the two species. This can place an effective lower limit on the temperature of the sympathetically-cooled species, which depends on the mass difference and the plasma parameters.

Sympathetic cooling can also be used to drive a positron plasma’s temperature down below the 20-30K limit reached by cyclotron cooling alone by mixing in laser-coolable ions, and laser cooling them. As discussed in detail in later chapters, beryllium ions are the lightest readily laser-coolable ion, and so make a natural choice for this application. Minimising the mass difference between the two species helps in two ways: the collision rate (and therefore the amount of energy that can be exchanged between the species) increases as the mass difference decreases, and the limit of this sympathetic cooling is due to centrifugal separation, which is less pronounced when mass difference is minimised.

3.8.3 Electron Kick

The “electron kick” (ekick) is a technique used to separate multispecies plasmas where the two species have vastly different masses. This is the case for the mixed plasmas that are used in ALPHA: antiproton-electron and beryllium-positron mixed plasmas. In a mixed plasma where the two species are at similar temperatures, the lighter species will have a higher velocity than the heavier species.

By lowering one side of the confining potential well, waiting for the majority of the lighter species to escape, and then quickly re-raising the potential well, the lighter species can be removed from the trap while causing minimal loss of the heavier species. For the trap parameters in ALPHA, this corresponds to dropping the potential barrier for around 100 ns to remove most electrons or positrons from a mixed plasma. Particle loss of the heavier species is low, but an ekick is typically accompanied by heating of the heavier species.

A classic use of the ekick in ALPHA is for separating antiproton-electron plasmas before using the antiprotons for antihydrogen formation. The antiprotons are initially mixed with a large amount of electrons to provide much-needed cooling, and to allow RW compression of the antiprotons. The electrons must be removed before introducing positrons to avoid positron-electron annihilation, so the electrons are removed with ekicks.

3.9 Plasma Characterisation

Most of the measurements performed on plasmas in ALPHA are destructive measurements – the plasma is ejected from the trap by lowering one side of the potential well that it is confined in, and the particles are directed along the axis of the trap by the magnetic field lines. The plasma will then typically impact on a microchannel plate (MCP) or Faraday cup, where a measurement can be made. Understanding the dynamics of the plasma during all stages of this process are key to extracting useful information from these measurements.

3.9.1 Plasma Dynamics During Potential Well Manipulations

When ejecting (dumping) a plasma from a trap, there are two general methods: fast and slow. In a fast dump, the confining potential is lowered as fast as the electronics will allow, which doesn’t significantly affect the radial profile of the plasma. This type of dump is often used when imaging the plasma with a microchannel plate, or when measuring its charge with a Faraday cup.

The second type of dump, the slow dump, is a little more interesting because the potential manipulations are slow compared to the plasma frequency. This allows time for the plasma to “react” as particles are being lost from the trap. A consequence of the cylindrical geometry of the

Penning trap electrodes is that the confining potential is lowest along the trap axis, which means that during a slow dump, particles will be preferentially lost from the trap from the middle of the plasma, causing the plasma to hollow out in the centre. A hollow plasma is an unstable state, and higher-radius particles will start to collapse inwards to the centre. As the dump is nearing its end, the potential well will be quite shallow, which can cause expansion and heating, as discussed earlier in this chapter. This type of dump can be used to extract temperature information from a plasma, but is ill-suited to radial profile measurements with an MCP due to the hollowing-collapsing process changing the radial profile during the dump.

3.9.2 Plasma Temperature Measurements

A plasma's temperature can be measured by slowly lowering one side of its confining potential well, allowing particles to begin to escape from the well and leave the trap, with the highest energy particles escaping first. The temperature of the plasma can be reconstructed by looking at the rate at which particles escape the trap as a function of the confining potential barrier's height.

The number of particles that escape the plasma, $N(E_B)$, is measured as a function of the barrier height, $E_B(t, r)$, which is given by E. D. Hunter et al. [29] as

$$E_B(t, r) = -q [V(t, r) - \Phi(t, r)], \quad (3.25)$$

where r is the radius from the trap axis, $V(t, r)$ is the depth of the trap's electrostatic well created by the applied voltages on the trap electrodes, and $\Phi(t, r)$ is the plasma's self-consistent potential. The temperature can then be extracted from the rising edge of the signal, $N(E_B)$, by fitting to it an exponential of the form

$$N(E_B) \propto \exp \left[\frac{-E_B(t, r)}{k_B T} \right]. \quad (3.26)$$

This equation comes from an assumption that the particles in the plasma follow a Maxwell-Boltzmann distribution, and is valid only for the high-energy exponential tail of the distribution: the hot particles, which escape the trap first.

Particles that are closer to the trap axis see a lower potential barrier, and so escape the trap first. As these particles escape from the centre, the self-consistent plasma potential, $\Phi(t, r)$ begins to change, which in turn affects the barrier height, E_B . As more particles begin to leave the trap, the dependence of E_B on $N(E_B)$ increases, and $N(E_B)$ becomes less dependent on the temperature. Thus, only a small subset of particles that escape first strictly follow equation 3.26, and it is this subset that is used to extract the temperature. This subset forms a cylindrical core that has a radius on the order of one Debye length, which is extracted from the trap axis.

Because only a small fraction of particles that leave the trap carry the temperature information, an MCP (see section 2.7.2) is typically used to provide gain on the otherwise small signal from these particles. The output electrons from the MCP hit a phosphor screen, where they can be counted. This is either done by measuring the charge incident on the screen (i.e. using the screen as a Faraday cup), or by using a silicon photomultiplier (SiPM) to measure the amount of light coming from the phosphor screen. Use of a SiPM is a fairly recent development explored by E. D. Hunter et al. [29], which can give significantly better signal-to-noise ratios than the Faraday cup method.

Figure 3.9 shows an example signal taken from a positron plasma. As the on-axis depth of the applied potential well is known as a function of time (i.e. $V(t, r = 0)$ is known), the x-axis has been converted from time units into energy units. Particles begin to escape the trap when the self-consistent potential $\Phi(t, r = 0)$ exceeds $V(t, r = 0)$, which in figure 3.9 occurs at around $V(t, r = 0) = 4.5$ V. In order to be able to extract the temperature, the particles that form the central core, which carry the temperature information, must be identified. In practice, the particles that carry temperature information follow equation 3.26, and will appear as a linear region on the log-scale of figure 3.9, and so extracting the temperature is achieved by identifying and then fitting this region with equation 3.26.

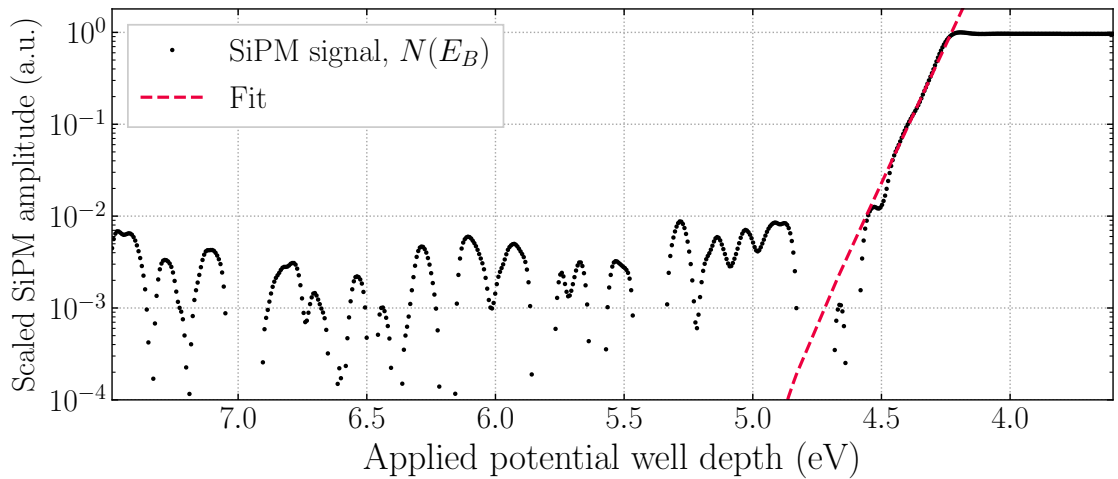


Figure 3.9: A typical temperature diagnostic signal. In this example, the plasma being measured was an ~ 800 K positron plasma. The algorithm that chooses the linear region to fit essentially attempts to minimise the error in the fit, and is described in more detail by L. T. Evans [46].

3.9.3 Systematic Uncertainties in Temperatures

This plasma temperature measurement technique can suffer from systematic uncertainties, from effects such as cooling from the expansion of the well that the particles are held in during the

ejection, radial variations in the length and temperature of the plasma, and systematic errors in fitting the history of the released positrons to the theoretically expected curve. Rather than carefully quantifying these systematic effects, the temperature measurements presented that use this measurement technique should instead be taken relative to other measurements taken using the same technique, and relative consistency can be ensured by comparing measurements that use the same potential well manipulations. This means that these temperatures are not absolute temperature measurements. These systematic effects may be as large as 50% on the absolute temperature determination [7, 29, 47], although particle-in-a-cell simulations for the relevant parameter regime find that they are around 10 – 15% in practice.

3.10 Summary

The techniques that have been described in this chapter are used in ALPHA for close control and diagnostics of plasmas trapped within the apparatus, and they form the building blocks upon which most measurements rely. The following chapters will discuss how these techniques are used for antihydrogen formation and trapping in ALPHA, and how they were used to create sympathetically cooled positron plasmas.

4/ Antihydrogen Production and Trapping

Efficient production of antihydrogen and being able to subsequently trap it is the basis of ALPHA's experimental program, and increasing the amount of antihydrogen that we can trap is always desirable, as it allows for experiments that are faster and/or more statistically significant. This is particularly important for experiments within CERN's Antiproton Decelerator (AD) hall, such as ALPHA, due to the limited amount of antiprotons delivered from the AD: each shot from the AD results in only $\sim 10^5$ trapped antiprotons inside the ALPHA apparatus, and the schedule of the AD is reliant on other parts of CERN's accelerator chain, which typically shuts down for a "technical stop" for a few months each year, and for "long shutdowns" for roughly two years out of every six [48]. Because of these factors, making the most out of every shot of antiprotons is of utmost importance.

In brief, antihydrogen is produced by bringing an antiproton and a positron together to form a single bound state. This is achieved in ALPHA by merging trapped plasmas of antiprotons and positrons. How much antihydrogen can be trapped depends largely on the three main parameters that describe both the positron plasma and the antiproton plasma: size, temperature and number of particles. These plasma parameters affect both the formation rate of antihydrogen atoms, and the temperature of the antihydrogen that is produced. Together, these are the dominant factors that determine how much antihydrogen can be trapped. Figure 4.1 shows how controlling these parameters (along with other improvements) has drastically increased the amount of antihydrogen available in ALPHA over the past decade or so.

The rest of this chapter will go into detail about the mechanics of producing and trapping antihydrogen atoms, and will discuss experimental and simulation-based efforts that have been performed in order to try to understand how changes in the different plasma parameters affect the antihydrogen formation and trapping process.

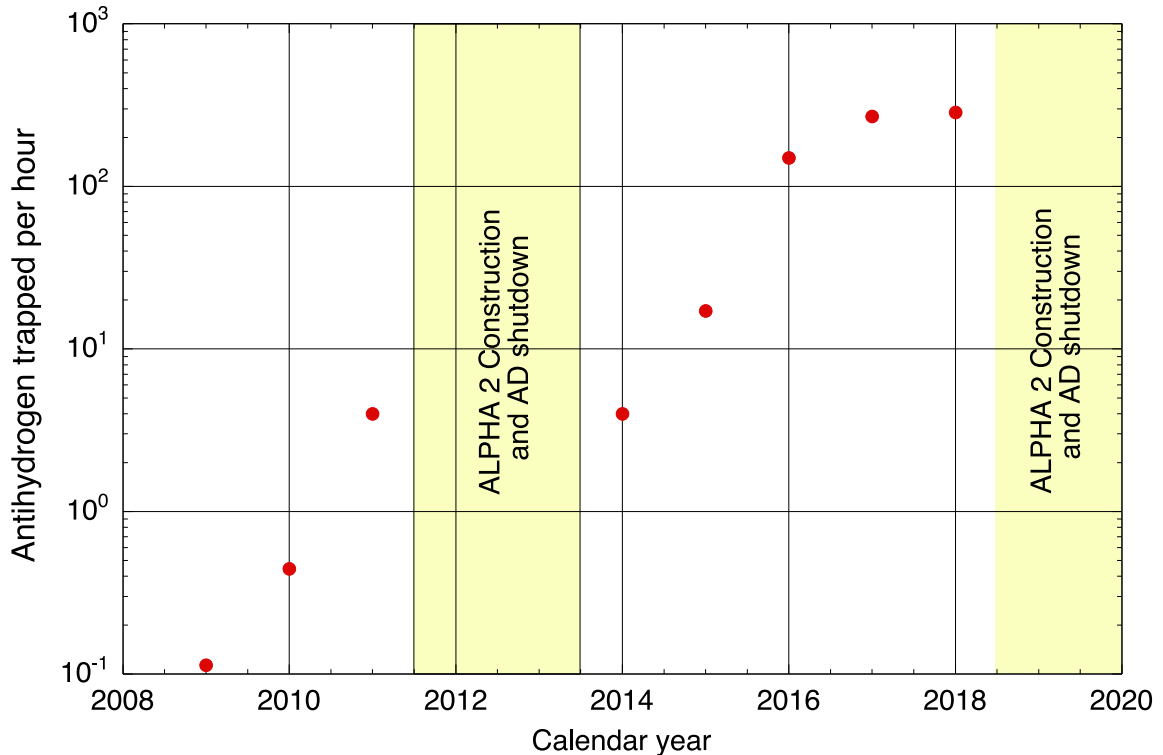


Figure 4.1: Antihydrogen trapped per hour over the past several years in ALPHA. Since first trapping antihydrogen atoms in 2009, the trapping rate has increased drastically. ALPHA’s first publication of trapped antihydrogen [49], which was based on data from 2010, had 0.11 trapped antihydrogen atoms per experiment. ALPHA’s more recent publications report on experimental runs where 10-20 were captured per experiment [10, 22, 24].

4.1 Antihydrogen Production

The first, and most challenging to obtain, ingredient for antihydrogen production is the antiproton. Antiprotons are delivered in bunches containing $\sim 3 \times 10^7$ antiprotons at an energy of 5.3 MeV from CERN’s AD. An overview of the AD was given in section 1.3. These antiprotons have far too much energy to be caught directly into the Catching Trap, and so must first pass through a thin degrader foil. A large fraction of the antiprotons annihilate on the foil, and a small fraction of the antiprotons that do not annihilate in the foil have their energy reduced to below ~ 4 keV, which allows them to be caught in the Catching Trap (CT). The catching process uses two dedicated high-voltage electrodes to dynamically trap the incoming antiproton bunch, which are shown in figure 4.2. Before bunch arrival a blocking potential is raised (electrode labelled HVB in figure 4.2) to around 4kV, which will reflect incoming antiprotons that have less energy than the blocking potential. To complete the catch, a ~ 4 kV potential is quickly raised on HVA after the bunch has entered the trap, confining the antiprotons between HVA and HVB. This dynamic catching method is used throughout ALPHA for transferring




Figure 4.2: Schematic of the Catching Trap’s Penning-Malmberg trap. a) On-axis potentials during particle catching. Antiprotons enter from the left, and are reflected by the potential applied to HVB (blue solid line). Before the reflected particles have enough time to leave the trap, a potential is raised on HVA to complete the trap. A plasma of electrons is prepared before capture and held in the trap to cool the antiprotons upon their capture. b) A schematic of the Penning trap electrode stack. Axial positions of a) and b) are matched. Electrodes 6 and 16 are segmented rotating wall electrodes.

particles between various parts of the apparatus.

These antiprotons are still far too energetic to be confined by any electrodes except the dedicated high-voltage catching electrodes (HVA and HVB), and so must first be further cooled before they are transferred to other parts of the apparatus. This is achieved through sympathetic cooling, by mixing the antiprotons with a plasma of $\sim 8.5 \times 10^7$ electrons, which cool via emission of cyclotron radiation. This brings the kinetic energy of a large fraction of the antiprotons below $\sim 100\text{eV}$. Any uncooled antiprotons are ejected, and a large fraction of the electrons are also ejected, leaving a mixture of $\sim 1 \times 10^5$ antiprotons and $\sim 1 \times 10^6$ electrons. These are compressed with a rotating wall and are allowed to re-cool before ejecting the remaining electrons and transferring the antiprotons to ALPHA-2. The antiprotons are dynamically caught directly into a second pre-prepared electron plasma in ALPHA-2, and a similar process of cooling, ejecting a large fraction of the electrons, compressing and re-cooling takes place. This produces an antiproton plasma containing $\sim 1 \times 10^5$ antiprotons with a temperature of $\sim 100\text{K}$.

At the same time, a plasma containing $\sim 30\text{ M}$ positrons is prepared in the positron accumulator before being ejected from the accumulator and dynamically caught in ALPHA-2. An overview of the accumulator is given in section 2.4. The number of positrons is controlled by using SDREVC (see section 3.7.1), and they are left to self-cool via cyclotron emission. This produces a plasma of up to 3×10^6 positrons, with a radial size of $\sim 0.6\text{mm}$ and a temperature of $\sim 30\text{K}$.

The two plasmas are then mixed in the centre of the trap in a process called the “slow merge”. A schematic view of the ALPHA-2 Penning-Malmberg trap is shown in figure 4.3, and the po-




Figure 4.3: Schematic of the ALPHA-2 Penning-Malmberg trap. Antiprotons are brought from the left and captured in the antiproton preparation region, and positrons are brought in from the right. The plasmas can be independently manipulated in either end of the trap before being brought together to the central antihydrogen synthesis region. The Neutral Trap magnets are shown in blue (mirror coils, solenoids) and green (octupole), which can be energised to form a 3-D magnetic field minimum region to trap antihydrogen atoms.

tentials used during the slow merge process are shown in figure 4.4. Just before the merge, the two plasmas are held in adjacent potential wells such that they almost overlap, and then over a period of 1 second, the plasmas are allowed to merge by lowering the potential barrier between them. Antiprotons are free to enter the positron plasma and positrons can drift to the left (and escape the trap). This allows the antiprotons to mix with the positrons without accelerating the antiprotons, and allows the positrons to continuously evaporatively cool during the process [20].

4.1.1 Quantifying Antihydrogen Production

Increasing the amount of antihydrogen that is trapped can be achieved in two ways. The first is to simply produce more antihydrogen. Since the number of antiprotons delivered by the AD is fixed, this can be achieved by increasing the fraction of antiprotons that successfully form antihydrogen atoms. I will introduce the “production efficiency” parameter to quantify this,

$$f_{\bar{H}} = \frac{\text{Number of produced antihydrogen atoms}}{\text{Number of trapped antiprotons}}. \quad (4.1)$$

A production efficiency of around 0.5 is typical in recent experiments: around half of the antiprotons in the trap are converted into antihydrogen atoms.

The second way to increase the amount of trapped antihydrogen is to increase the fraction of produced antihydrogen atoms that can be trapped. As we will see later, the vast majority of antihydrogen atoms that are produced are not trapped, as their kinetic energy is too high to

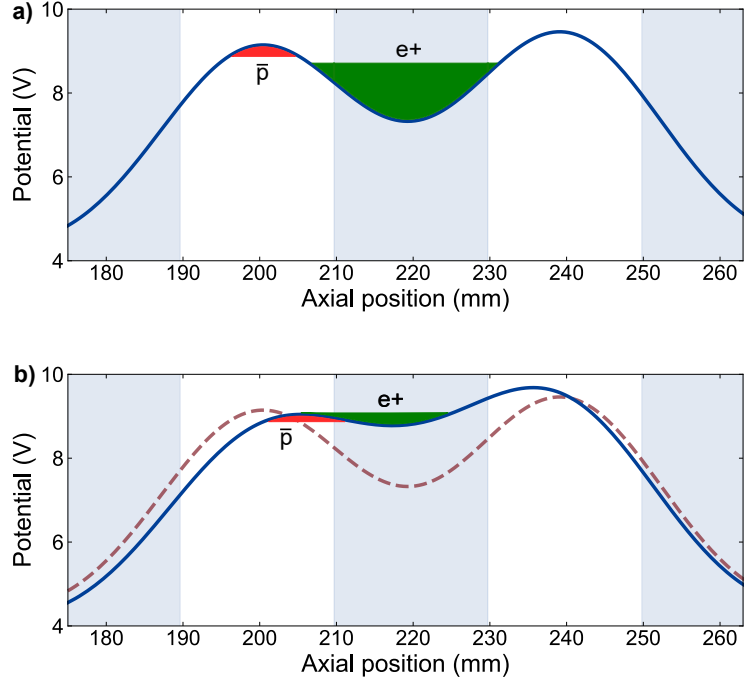


Figure 4.4: Potentials used during the merge. Shaded/unshaded regions show axial extent of trap electrodes. a) Potentials before the merge, with red and green shaded regions representing the space charge and physical extent of the plasmas. b) How the potentials change from the start of the merge (red dashed line) to the end (blue solid line). The merge step takes 1 second.

allow them to be confined by the shallow Neutral Trap. I will quantify this with the “trapping fraction” parameter,

$$f_{trap} = \frac{\text{Number of trapped antihydrogen atoms}}{\text{Number of produced antihydrogen atoms}}. \quad (4.2)$$

Trapping fractions of around 3×10^{-4} are typical: for every 3300 antihydrogen atoms produced, about 1 is successfully trapped.

4.2 Formation Mechanisms

The main antihydrogen formation mechanism in this environment is three-body recombination¹ [50],



¹The “recombination” part of the name is a bit of a misnomer as it implies that the antiproton and positron were previously combined - it should really be called “three-body capture”. I will continue to use “recombination” for consistency with literature.

This can be understood as an antiproton capturing a positron to form an antihydrogen atom, with the second “spectator” positron carrying away the additional energy. This typically results in a highly excited antihydrogen atom (\bar{H}^{**}) which can easily be ionised by further collisions with positrons. In order to form a stable antihydrogen atom, the weakly bound antiproton and positron must be “lucky” enough to avoid collisional ionisation (essentially the opposite of equation 4.3), and instead undergo de-excitation collisions with positrons,



These de-excitation collisions cause the antiproton and positron to become more strongly bound, and it can eventually exit the positron plasma as a bound state. The final hurdle that the bound state has to overcome is field ionisation once it has left the plasma and entered the non-zero electric field environment in the rest of the trap (recall that the electric field inside the plasma is zero). If it is strongly-bound enough to avoid becoming field-ionised, it can then slowly de-excite via photon emission down to a ground state antihydrogen atom. If the excited antihydrogen atom is instead ionised at some point during this process, either due to a collision with a second positron inside the plasma, or due to field ionisation after it has left the plasma, the antiproton will most likely remain trapped. It then has the opportunity to capture another positron, and has another chance to form a stable antihydrogen atom. Simulations show [51,52] that a single antiproton will typically undergo many formation and break-up cycles before finally forming a stable antihydrogen atom.

4.3 Antihydrogen Trapping

Antihydrogen is trapped by the Neutral Trap, an octupolar magnetic-minimum trap that is discussed in some more detail in section 2.3.1. The magnets that form the Neutral Trap can be seen in figure 4.3. The antihydrogen trapping process itself is fairly simple: the Neutral Trap is energised at some point before the final merge step that brings the antiproton and positron plasmas together, and the overlapping geometry of the two traps ensures that antihydrogen atoms are produced directly inside the Neutral Trap. Low temperatures are critical for success here because of the shallow depth of the Neutral Trap, which can only trap antihydrogen atoms with a temperature below 0.54K. Since an antihydrogen atom will inherit its kinetic energy from the antiproton and the positron that formed it, and the antiproton and positron plasmas have typical temperatures of $\sim 100\text{K}$ and $\sim 20\text{K}$ respectively before merging them, it should be clear that most antihydrogen atoms will have an energy larger than the trap depth, and will escape the trap. Of the $\sim 50,000$ antihydrogen atoms that are produced in a typical mixing cycle,

around 10-20 will be trapped, giving a trapping fraction, f_{trap} , of $\sim 10^{-4}$ (equation 4.2). If we assume that the antihydrogen atoms follow a Maxwell-Boltzmann distribution, and that they have a temperature significantly higher than the trap depth, such that the trapped antihydrogen atoms belong to the exponential tail of the distribution, this gives the trapping fraction a scaling $\propto T^{-3/2}$ with the temperature of the antihydrogen atoms.

4.4 Antihydrogen Accumulation

A recent development that has increased the antihydrogen sample size that ALPHA is able to work with is antihydrogen accumulation (or “stacking”) [20]. Antihydrogen is synthesised and trapped as described above, but the Neutral Trap is left on, with any trapped antihydrogen atoms remaining in the trap. The cycle is then repeated many times, and each time more antihydrogen is added to the trap. Figure 4.5 shows accumulation of up to 5 stacks, totalling around 50 simultaneously trapped antihydrogen atoms. Antihydrogen has a lifetime of > 60 hours in the trap [53], which is easily long enough to allow stacking for several hours. In more recent experiments, stacking over several hours has yielded more than 1000 simultaneously trapped antihydrogen atoms [24]. Stacking can indirectly increase the number of antihydrogen atoms available for study. Without stacking, each batch of antihydrogen atoms is trapped and studied separately. This period of study can take several minutes in the case of spectroscopy experiments, and during this time no more antiprotons can be caught. Since the AD is on a fixed ~ 2 minute cycle, many shots from the AD are wasted. When stacking, however, many stacks of antihydrogen can be studied together, meaning that more shots from the AD are caught in total. Being able to study more simultaneously trapped antihydrogen atoms also increases signal-to-noise ratio on antihydrogen annihilation events.

In the Penning-Malmberg trap, charged particles follow the axial magnetic field lines that point along the axis of the trap, produced by the trap’s external solenoid. When stacking, the Neutral Trap magnets must remain on the whole time and transverse magnetic fields from the Neutral Trap magnets interfere with and distort these field lines. Above a critical radius, these field lines will guide charged particles into the walls of the trap. The critical radius for the ALPHA-2 Penning-Malmberg trap is around 4.5mm, and so in order to minimise plasma distortion and avoid particle loss, plasma radii must be kept well below this limit.

4.5 Simulations of Antihydrogen Formation Mechanisms

The rate at which three-body recombination occurs in a strongly magnetised, cryogenic plasma in a strong magnetic field was investigated by Glinsky and O’Neil [54], and they found that the




Figure 4.5: Figure reproduced from ALPHA collaboration [20]. The number of antihydrogen atoms detected after the Neutral Trap is ramped down after consecutive mixing cycles. Each mixing cycle is separated by ~ 4 minutes. The error bars are statistical, and the number of replicates is indicated in blue above each data point. The dashed linear fit gives a trapping rate of 10.5 ± 0.6 detected antihydrogen atoms per mixing cycle.

steady-state scaling of the formation rate, Γ_{tbr} , with the positron density n_e and temperature T_e is given by $\Gamma_{tbr} \propto n_e^2 T_e^{-9/2}$. In order to obtain a better understanding of the antihydrogen formation process, many simulations have been carried out, and it quickly becomes clear that the assumption of steady-state formation is not a good description of how antihydrogen formation occurs in ALPHA. As discussed by Jonsell et al. [52], the steady-state assumption requires the antiproton to be immersed in the positron plasma for longer than the time recombination takes to occur. This is not the case in the current experimental setup where, as discussed in section 4.2, antiprotons repeatedly enter and exit the positron plasma. This means that antihydrogen atoms that form within the plasma have a limited amount of time to undergo de-excitation collisions before they exit the plasma, where they are likely to be field-ionised by the large electric fields that are present outside of the plasma. This is sometimes referred to as arrested recombination [55], as the recombination process is effectively arrested part-way through when the antiproton leaves the positron plasma. This arrested recombination will have a different scaling rate with the positron temperature and positron density when compared to steady-state recombination.

The effect of the changing T_e on antihydrogen production rate, $\Gamma_{\bar{H}}$, has been measured experimentally by the ATHENA experiment [56](a predecessor of ALPHA) by monitoring antihydrogen formation as T_e was increased by heating the positrons with a radiofrequency source. This found the scaling of antihydrogen production rate as being $\Gamma_{\bar{H}} \propto T_e^{-1.1 \pm 0.5}$ – a long way

from the steady-state scaling.

This scaling is fairly close to the temperature scaling expected for radiative recombination of $R \propto T_e^{-0.63}$, where the extra energy is carried away by a photon instead of a second positron, but radiative recombination can be ruled out as a dominant formation mechanism as the experimental production rate was around an order of magnitude higher than the calculated radiative recombination rate [52]. It's also worth noting that in this data, they were unable to measure the base positron temperature, T_e (only ΔT_e was measured), which could have had a large effect on the scaling.

A series of simulations by S. Jonsell et al. [50, 52, 57] have been performed in order to shed light on the antihydrogen formation process. Here, I will give a short overview of this series of simulations and their underlying physics. The positrons are initially treated as a plasma, with the plasma either being treated as an infinitely long cylinder with a fixed radius and density, or of a finite size and shape which is calculated based on solving Poisson's equation self-consistently for both the electric field and the spatial profile, based on real trap conditions, a procedure that is outlined by Dubin et al. [58]. The positron plasma temperature is set at a fixed value. Antiproton trajectories are based on classic equations of motion, with the justification that quantum mechanical effects are only expected to play a role when interparticle distances are significantly shorter than in these simulations. A box of a predefined size is defined around each antiproton, and for each time step, the probability of a positron entering into this box is calculated. Once a positron is inside this box, its motion is explicitly solved, while the rest of the positron plasma is treated as a continuous medium. Trajectories for both antiprotons and positrons are obtained by integrating Newton's equations of motion for particles subject to the Lorentz force,

$$\mathbf{F} = q(\mathbf{E} + \mathbf{v} \times \mathbf{B}), \quad (4.5)$$

with electric and magnetic fields \mathbf{E} and \mathbf{B} , charge q and particle velocity \mathbf{v} . Each simulation only considers a single antiproton, and so antiproton-antiproton or antiproton-antihydrogen interactions are not included. These simulations have been able to produce results that show good agreement with experimental data across a wide range of plasma parameters and mixing schemes, for example the paper by Jonsell et al. in 2018 [50], which primarily investigates the effects of changing positron plasma parameters on the trapping fraction, shows remarkably good agreement with experiment for both the trapping fraction, f_{trap} , and the production efficiency, $f_{\bar{H}}$.

In Jonsell et al's 2009 paper [52], the velocity of the produced antihydrogen atoms is presented for various values of n_e . As n_e is increased, the average antihydrogen velocity also in-

creases, which leads to a lower trapping fraction. This is attributed to “epithermal” antihydrogen formation, which is where antiprotons that have not had sufficient time to thermalise with the positron plasma form antihydrogen atoms, leading to higher-temperature antihydrogen atoms that are harder to trap.

These results are not directly applicable to our current slow merge mixing scheme, as these simulations are based on the “ATHENA mixing” scheme, where antiprotons are injected with an energy of a few eV ($2 \text{ eV} \approx 1.5 \times 10^4 \text{ K}$ in this simulation), compared to the $\sim 100 \text{ K}$ antiprotons injected in the slow merge. Nonetheless, this highlights the important point that the antiproton cooling rate is not necessarily larger than the antihydrogen formation rate. As the primary goal of this thesis is to obtain significantly colder positron plasmas, it should be kept in mind that colder positrons will affect both the antiproton cooling rate and the antihydrogen production rate such that, with cold enough positrons, this epithermal regime could be re-entered.

Simulations by Robicheaux et al. [59] instead use Classical Trajectory Monte Carlo simulations along with level population evolution to extract useful information about the scaling of antihydrogen production with positron density and temperature. The most important things to highlight about the results of this simulation are that the positron density and temperature both play very important roles in antihydrogen formation, and that there don’t appear to be simple power laws that describe how these parameters affect the final antihydrogen production rate. These simulations find, for example, that the scaling of the antihydrogen formation rate with positron temperature is strongly dependent on the positron plasma’s density, and vice versa.

Taken as a whole, all of these simulations paint a picture of the antihydrogen formation process where the different parameters intertwine in a complex manner. Armed with knowledge from these simulations, we can try to balance these parameters in the real experiment.

4.6 Experimental Plasma Parameter Tuning

During 2017, several experimental studies were undertaken in order to optimise the antiproton and positron plasmas for producing antihydrogen with the slow merge technique. Although I was present for a large part of the 2017 run, the majority of the data presented in this section was taken by other people, and analysed by me. The focus of these experimental studies was to change the main three plasma parameters, number of particles N , density n , and temperature T . Ideally, each of these parameters would be controlled completely independently, but this is rarely possible in a real experiment: evaporative cooling (see section 3.7), for example, changes all three simultaneously. This leads to a large parameter space that can be difficult to move around in, and so the data collected is not as comprehensive or clean as their simulated counterparts.

4.6.1 Evaporative Cooling of Antiprotons

Evaporative cooling, discussed in detail in chapter 3, selectively removes the highest-energy particles from the plasma, causing the remaining plasma to be at a lower average kinetic energy, and therefore a lower temperature. The more particles removed, the lower the temperature of the remaining plasma. This is achieved in a Penning-Malmberg trap by lowering the potential on one side of the well confining the particles. The minimum depth that the potential well reaches, i.e. the EVC depth, is the parameter that has the most effect on how many particles are removed from the trap. A feature of evaporative cooling in a Penning-Malmberg trap is radial expansion of the plasma caused by particles being preferentially lost from the centre of the plasma, where the confining potential is lowest. Particles lost from the centre of the plasma carry away no angular momentum, and so in order to conserve angular momentum, the plasma must increase in radius. In the ideal case, the expansion of an evaporatively-cooled plasma can be given by $N_0/N = \langle r^2 \rangle / \langle r_0^2 \rangle$, where N and N_0 describe the final and initial number of particles, and r and r_0 describe the final and initial plasma radii. Evaporative cooling changes all three plasma parameters simultaneously, and so its effects on antihydrogen formation can be hard to predict.

In this first data set presented in figure 4.6, a plasma of $\sim 90,000$ antiprotons with an initial radius of ~ 0.5 mm was evaporatively cooled before being merged with a plasma of $\sim 3 \times 10^6$ positrons with a radius of ~ 0.6 mm. Figure 4.6 shows the effect of evaporative cooling (EVC) on antiprotons on both the $f_{\bar{H}}$ and f_{trap} . A shallower EVC causes a much lower $f_{\bar{H}}$, which is likely dominated by the lower number of antiprotons that make it into the merge, as more of them are lost in the EVC step. This large loss in $f_{\bar{H}}$ is not compensated by an increase in f_{trap} , which remains almost constant across the different EVC depths used. This implies that the lower-temperature, larger antiproton plasma does not produce colder antihydrogen atoms. This is consistent with the assumption that the antiprotons are reaching thermal equilibrium with the positrons before antihydrogen formation, i.e. that the system is in the thermal regime, rather than the epithermal regime, as changing the initial temperature of the antiprotons will have little effect if they have ample time to thermalise with the positron plasma.

4.6.2 Antiproton Plasma Radius

In this next data set presented in figure 4.7, the antiproton plasma's size was varied by changing the frequency of the rotating wall (chapter 3.5) used to compress it. The number of antiprotons and the temperature remained constant. These antiprotons were then used in a slow merge to form antihydrogen, similar to the previous data set. A naive assumption might be that matching the radial sizes of the positron and antiproton plasmas would be close to optimal for anti-

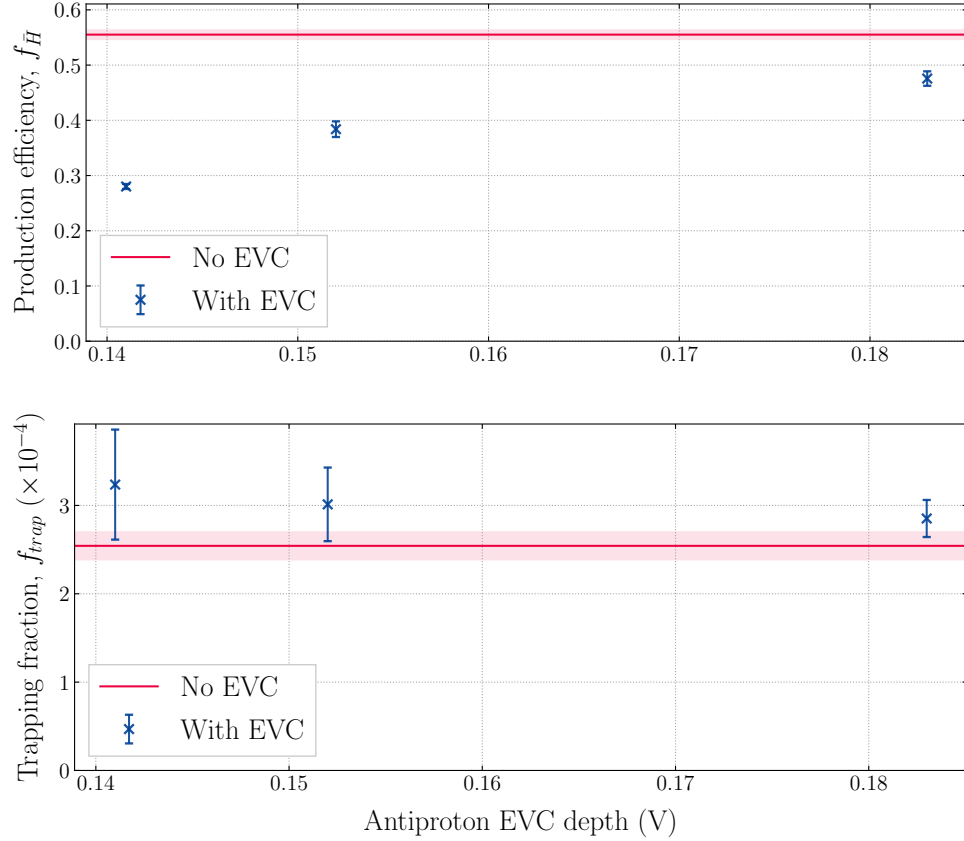


Figure 4.6: Antiproton production efficiency, $f_{\bar{H}}$, (upper) and trapping fraction, f_{trap} , (lower) as a function of antiproton EVC depth. Blue points are with EVC, and the red line is without EVC, with shaded regions representing error bars. Each data point is made up of several repeat measurements, and the error bars are standard error of the mean for each point. The calculation of $f_{\bar{H}}$ uses the number of antiprotons in the trap *before* EVC.

hydrogen formation, as this maximises overlap between the two plasmas. The nominal size of a positron plasma during antihydrogen formation is around 0.6mm, so this is a natural size to aim for with the antiproton plasma given the previous assumption.

Figure 4.7 shows the effect of the radial size of the antiproton plasma on the production efficiency, $f_{\bar{H}}$, and the trapping fraction, f_{trap} . It's quickly clear that the previous naive assumption that matching plasma sizes is the best is incorrect. $f_{\bar{H}}$ is essentially independent of plasma size, which is perhaps surprising given that in the most extreme case, the initial radius of the antiproton plasma ($\sim 1.7\text{mm}$) is almost 3 times larger than the initial radius of the positron plasma ($\sim 0.6\text{mm}$). This is accompanied by a marked improvement in f_{trap} as the antiproton plasma radius increases, implying that the larger antiproton plasma is able to produce colder antihydrogen atoms. This phenomenon has not been replicated by any simulations, and in fact Jonsell et al. [50] explicitly recommend the opposite – very small radius ($\sim 0.1\text{mm}$) antiproton plasmas –

in order to maximise the f_{trap} . The argument put forward is that antiprotons at non-zero radius will have an extra contribution to their velocity from their $\mathbf{E} \times \mathbf{B}$ drift, which increases with radius, meaning that higher-radius antiprotons will have a higher velocity, leading to a lower f_{trap} . The reason for this unexpected increase in f_{trap} with larger antiproton plasmas currently remains unexplained.




Figure 4.7: Production efficiency, $f_{\bar{H}}$, (upper) and trapping fraction, f_{trap} , (lower) as a function of antiproton plasma radius.

4.6.3 Positron Temperature

The positron temperature, T_e , is expected to play a major role in both the formation rate of antihydrogen, and on its final temperature, affecting both f_{trap} and $f_{\bar{H}}$. Experimental and simulated results have shown that decreasing T_e tends to give a monotonic increase in the amount of antihydrogen that we can trap, i.e. colder positrons are always better, at least for currently attainable parameter ranges, which makes reductions in positron temperature particularly valuable. Despite this, there have been relatively few clean data sets exploring the role of T_e in recent years. This can mainly be attributed to the scarcity of antiprotons: clean systematic studies take a long time and a lot of antiprotons that could otherwise be used for more interesting physics.

The following data set is taken from a subset of the runs undertaken in 2017, in which samples of antihydrogen atoms were illuminated with 243nm laser light. In these runs, it was found that the longer the laser had been shining inside the trap, the hotter the positron plasmas were, and the lower the antihydrogen trapping rate became. This effect occurred over several consecutive days of running the experiment. It is thought that exposure of the trap electrodes to 243nm light caused patch potentials to slowly appear on the electrodes (discussed in chapter 3.4), which lead to asymmetric electric fields and increased plasma heating. The effects of these patch potentials have only recently been studied in depth and measured in the ALPHA trap by C. Baker et al. [35], although the exact mechanism that causes 243nm light to create these patch potentials remains unknown. This manifested primarily as an increasing T_e from day-to-day as the experiment was being run. It was possible to reset the trap and get rid of this anomalous heating effect by warming the trap up to room temperature and cooling it back down again. This was an inconvenience for the measurement that was actually taking place (measurement of the 1S-2S transition frequency [60]), but allowed for this parasitic measurement of the role of T_e on antihydrogen formation.

In the slow merge sequence, some positrons are left in the trap after the merge (see fig. 4.4b), the temperature of which were measured after each merge cycle. Figure 4.8 shows the temperature of these positrons (“after mixing”) against $f_{\bar{H}}$ and f_{trap} . As the merge includes some evaporative cooling of the positrons and mixes them with a hotter antiproton plasma, this is not a measurement of the initial temperature, T_e , directly. T_e should be higher than this measured temperature, but strongly correlated with it. There is a clear trend towards lower positron temperatures increasing both $f_{\bar{H}}$ and f_{trap} , although the effect on f_{trap} is much more drastic than the effect on $f_{\bar{H}}$. Reducing the measured positron temperature from 22.5 K down to 15 K gave a factor of 2 increase in f_{trap} , and a $\sim 33\%$ increase in $f_{\bar{H}}$. The individual red data points show the large amount of scatter, which is produced by shot-to-shot variations in plasma temperature..




Figure 4.8: Production efficiency, $f_{\bar{H}}$, (upper) and trapping fraction, f_{trap} , (lower) as a function of positron temperature after mixing. Runs were binned by positron temperature (blue), and error bars are statistical standard errors for each bin. Raw data points are shown in red.

4.6.4 Positron Density

The density of the positron plasma, n_e , plays an important role in antihydrogen formation. It directly affects the antihydrogen production rate, $\Gamma_{\bar{H}}$, and the temperature of the produced antihydrogen.

Its effect on the produced antihydrogen's temperature is related to the rotation speed of the plasma around the axis. The rotation speed of the plasma, v_D , given by Jonsell et al. [50], is

$$v_D = \frac{E}{B} = \frac{n_e e r}{2\epsilon_0 B}, \quad (4.6)$$

with radial distance from trap centre r , fundamental charge e and vacuum permittivity ϵ_0 . This effectively means that positron plasmas with a higher density will have a higher linear velocity,

which will translate into antiprotons, and therefore antihydrogen atoms, with a higher velocity. This suggests that f_{trap} should decrease with increases in either n_e or r . This isn't exactly what simulations show – and actually that f_{trap} can increase with increasing positron density. The effects of changes in n_e are highly dependent on other plasma parameters, which means that n_e is an interesting candidate for experimental tuning. Despite this, there have been very few independent studies of the effects of n_e on antihydrogen production, again largely due to the scarcity of antiprotons.

4.6.5 Positron EVC and Antihydrogen Formation

EVC can also be performed on the positron plasma before mixing. As with the antiproton EVC, this gives a larger, less dense and colder positron plasma consisting of fewer positrons. There is some (limited) data on the effects of positron EVC on antihydrogen production, which is summarised in table 4.1. This EVC roughly halved the number of positrons and the density of the plasma, and reduced the temperature by around a third. The EVC caused $\sim 20\%$ decrease in $f_{\bar{H}}$, but a $\sim 7\%$ increase in f_{trap} . The net effect was a reduction in the amount of trapped antihydrogen. Although the EVC hindered antihydrogen production for these particular initial conditions, that does not necessarily mean that positron EVC will hinder antihydrogen production for all initial conditions.

	Number of e^+	n_e (m^{-3})	T_e (K)	$f_{trap} (\times 10^{-4})$	$f_{\bar{H}}$
Without EVC	3.2×10^6	2×10^{14}	21.3	3.0	0.38
With EVC	1.7×10^6	9.5×10^{13}	14.3	3.2	0.31

Table 4.1: Positron plasma density n_e , temperature T_e and number, with and without positron EVC, as well as trapping fraction, f_{trap} , and production efficiency, $f_{\bar{H}}$, when these positrons are used for antihydrogen formation. The experimental sequence used was very similar to the antiproton radius and EVC studies discussed earlier in the chapter.

By comparing with the data from figure 4.8, we can see that the reduction in the T_e from 21.3K to 14.3K, when ignoring changes to the other plasma parameters, should have increased f_{trap} by closer to 50%. It is therefore likely that the changes to the other parameters, particularly n_e and the number of e^+ , caused a decrease in f_{trap} , which was overcompensated for by the increase in f_{trap} caused by the lower T_e . A similar argument can be made for $f_{\bar{H}}$, with the difference being that the increase in $f_{\bar{H}}$ as T_e is reduced is much less pronounced (going from the data in figure 4.8 again), and so the net effect from the EVC was a reduction in $f_{\bar{H}}$.

4.7 Summary

The three main plasma parameters – temperature, density, and number of particles – for both the antiproton and positron plasmas are the main tuning knobs available to optimise antihydrogen production. Increasing the radius of the antiproton plasma (i.e. reducing the density) gave significant increases in the antihydrogen trapping rate (fig 4.7), but this was limited by the radial size of the trap. Decreasing the antiproton temperature likely has little or no effect, as the antiprotons thermalise with the positron plasma before forming antihydrogen, meaning that the temperature of the produced antihydrogen is effectively set by the temperature of the positron plasma.

Simulations suggest that the density of the positron plasma plays an important role in antihydrogen formation, but its effect has not been experimentally explored in detail. The temperature of the positron plasma stands out as the best parameter that can be improved for large, monotonic increases in antihydrogen trapping rates. This, coupled with the realistic possibility of using sympathetic cooling to produce significantly colder positron plasmas, means that pursuing reductions in positron temperatures is a promising method of increasing antihydrogen trapping rates.

The chapters that follow will go into more detail about how positron plasmas can be sympathetically cooled to temperatures far below what is currently attainable in ALPHA, and will present the experimental methods that have been developed to produce positron plasmas with comparable density and number to those currently used in ALPHA, except at much lower temperatures.

5/ Laser Cooling

The idea that electromagnetic radiation exerts a force has been around for more than a century, and was first investigated in the lab with macroscopic objects such as the Crookes radiometer (which famously ended up not being able to demonstrate radiation pressure), and the Nichols radiometer [61]. This force arises from the fact that photons carry both energy and momentum, and therefore impart a force on objects that absorb or reflect them. Around 1970, this effect was investigated on atoms by using a laser tuned to be resonant with the absorption line of a particular atom, and it was shown that each photon absorption or emission event causes a change in the atom's momentum. This force was soon being used to cool atoms, with Hänsch and Schawlow demonstrating cooling of a low-density gas in 1974 [62], and Wineland, Drullinger and Walls demonstrating laser cooling of trapped magnesium ions to below 40K in 1978 [17]. Modern laser cooling experiments can cool trapped atoms to close to absolute zero, allowing very precise spectroscopic measurements of trapped atoms, production of Bose-Einstein condensates [41], optical clocks [63], producing ultracold trapped ions that can be used for quantum computing [64], and more.

The rest of this chapter will cover the basics of laser cooling and the additional considerations that arise when laser cooling in a Penning-Malmberg trap. Chapters 6 and 7 will discuss how laser cooling is applied in ALPHA to create cold plasmas of Be^+ ions, and how these cold Be^+ ions can be used to sympathetically cool positron plasmas.

5.1 Theory of Doppler Cooling

An atom has quantised energy levels that its electrons can occupy. The lowest energy level is the ground state, g , and higher energy levels are excited states, e_i . These excited states are unstable, and have finite lifetimes, τ_{e_i} , after which they will decay back down to the ground state. The energy difference between an excited state and the ground state can be expressed in terms of photon energy,

$$E_e - E_g = \frac{hc}{\lambda_0}, \quad (5.1)$$

with Planck's constant h , speed of light c , and photon wavelength λ_0 . This gives the wavelength of light that is required to excite an atom from a particular ground state into a particular excited state. If a ground state atom is illuminated with light which has a wavelength $\lambda \approx \lambda_0$, the atom can absorb a photon and become excited. After an average time τ_e , the atom de-excites back to the ground state and emits a photon with wavelength $\lambda = \lambda_0$, which carries away the excess energy. A photon carries a momentum $p = \hbar k$, where $k = 2\pi/\lambda$ is the wavevector and $\hbar = h/2\pi$ is the reduced Planck's constant, and so each absorption and emission event must cause the atom to recoil to conserve momentum. Singly-charged beryllium ions are typically laser cooled by excitation with light of $\lambda = 313$ nm, which causes the ion to recoil with a velocity $v_{rec} = 0.15$ m/s each time a photon is absorbed or emitted.

Lasers produce monochromatic, collimated light. If a trapped atom is illuminated with a laser tuned to be resonant with a transition in the atom, the atom will continually scatter photons. Figure 5.1 shows a sketch of this process. Each absorption will cause the atom to recoil in the direction of the incident laser light. The direction of the photon released during emission depends on whether the emission was stimulated, in which case the photon is emitted in the same direction as the incident photon, or spontaneous, in which case the photon is emitted in a random direction. For stimulated emission, the total velocity change of the atom during one absorption/emission cycle is zero, as the atom will gain $+v_{rec}$ from the excitation and $-v_{rec}$ from the emission. For spontaneous emission, the contribution from the emission events will average to zero over many cycles as the emission is isotropic, leaving only the contribution from the absorption event. The total velocity change over N scattering events is therefore $N \times v_{rec}$.

So far, this isn't enough to actually cool the atom. Whether the atom gains or loses speed depends on whether the atom happens to be moving towards the laser or away from the laser when it absorbs a photon. An atom oscillating in a trap will be moving towards the laser half the time and away from it half the time, leading to no net change in velocity over many cycles. The final piece of the puzzle is the Doppler effect, which causes atoms that are moving towards the laser to "see" the light as having a slightly shorter wavelength, while atoms moving away from the laser see light of a slightly longer wavelength. The laser can then be detuned to a slightly longer wavelength (known as "red detuned"), and so atoms moving towards the laser will see light that has been Doppler shifted into resonance, while atoms moving away from the laser will see light that has been Doppler shifted away from the resonance. The atom is now more likely to scatter photons when it is moving towards the laser, and the effect is a net reduction in the atom's velocity over many photon scattering events. This is where the name "Doppler cooling" comes from, and this is the type of laser cooling that will be discussed for the rest of this thesis.

The scattering of many photons will, on average, produce a scattering force, which is given

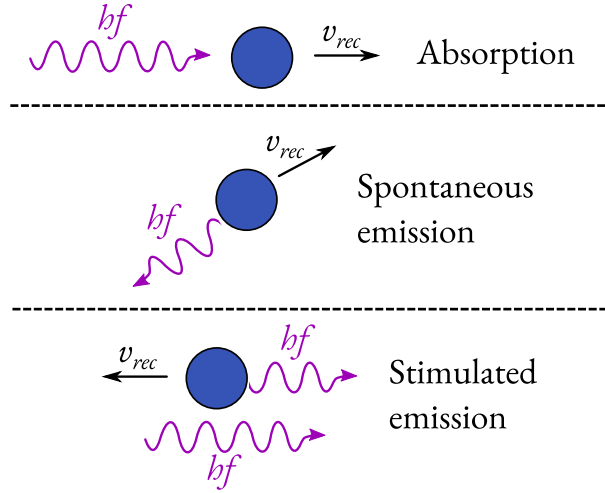


Figure 5.1: Atom behaviour as it absorbs and emits photons. For each event, momentum is conserved.

by the photon momentum multiplied by the scattering rate. C. Foot [65] gives the rate at which an atom will scatter photons from a laser as

$$R_{scatt} = \frac{\Gamma}{2} \frac{\Omega^2/2}{\delta^2 + \Omega^2/2 + \Gamma^2/4}, \quad (5.2)$$

where Γ is the natural linewidth of the transition, Ω is the Rabi frequency, and δ is the detuning. This detuning is given by

$$\delta = \omega - \omega_0 + kv \quad (5.3)$$

where $\omega - \omega_0$ is the difference between the laser's frequency and the transition's resonant frequency (i.e. the laser detuning), and kv is the Doppler shift due to the particle's velocity v . The Rabi frequency is related to the saturation intensity by $I/I_{sat} = 2\Omega^2/\Gamma^2$, where I is the laser intensity, and I_{sat} is the saturation intensity on resonance, which is given by [66]

$$I_{sat} = \frac{\pi \Gamma \hbar c}{3 \lambda_0^3}. \quad (5.4)$$

The saturation intensity is the intensity at which the atom spends $1/4^{th}$ of the time in the excited state.

Photons carry momentum $p = \hbar k$, and so the scattering force can be written as

$$\hbar k R_{scatt} = F_{scatt} = \hbar k \frac{\Gamma}{2} \frac{I/I_{sat}}{1 + I/I_{sat} + 4\delta^2/\Gamma^2}. \quad (5.5)$$

As $I \rightarrow \infty$, equation 5.5 becomes $F_{max} = \hbar k \Gamma/2$.




Figure 5.2: The scattering force, F_{scatt} , of a laser on a Be^+ ion as a function of the laser's detuning from resonance with the Be^+ cooling transition ($f_0 = 957$ THz). The detuning is shown in multiples of the natural linewidth of the cooling transition, $\Gamma = 2\pi \times 19.6$ MHz. The force is plotted for different values of the ion's velocity, v , and the laser intensity, I .

Increasing the laser power gives rise to so-called power broadening of the transition's lineshape. Power broadening has a Lorentzian lineshape, with a FWHM [65]

$$\Delta\omega_{FWHM} = \Gamma \sqrt{1 + \frac{I}{I_{sat}}} \quad (5.6)$$

Figure 5.2 shows the scattering force on a beryllium ion as a function of laser detuning for a few values of the ion's velocity and laser intensity. As the particle's velocity increases, the magnitude of the required laser detuning increases (and the peak shifts to the left). As the laser intensity is increased, the scattering force increases and the peak broadens due to power broadening.

The lowest temperature reachable by Doppler cooling is the Doppler cooling limit, T_D , and is given by

$$T_D = \frac{\hbar\Gamma}{2k_B}, \quad (5.7)$$

where Γ is the natural linewidth of the transition, and k_B is the Boltzmann constant. Typical values of T_D are on the order of 1 mK. This arises from the stochastic nature of the absorption

Group	→1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	
1	1 H																	2 He	
2	3 Li	4 Be											5 B	6 C	7 N	8 O	9 F	10 Ne	
3	11 Na	12 Mg											13 Al	14 Si	15 P	16 S	17 Cl	18 Ar	
4	19 K	20 Ca	21 Sc	22 Ti	23 V	24 Cr	25 Mn	26 Fe	27 Co	28 Ni	29 Cu	30 Zn	31 Ga	32 Ge	33 As	34 Se	35 Br	36 Kr	
5	37 Rb	38 Sr	39 Y	40 Zr	41 Nb	42 Mo	43 Tc	44 Ru	45 Rh	46 Pd	47 Ag	48 Cd	49 In	50 Sn	51 Sb	52 Te	53 I	54 Xe	
6	55 Cs	56 Ba	57-70 * Lu	72 Hf	73 Ta	74 W	75 Re	76 Os	77 Ir	78 Pt	79 Au	80 Hg	81 Tl	82 Pb	83 Bi	84 Po	85 At	86 Rn	
7	87 Fr	88 Ra	89-102 ** Ac	104 Rf	105 Db	106 Sg	107 Bh	108 Hs	109 Mt	110 Ds	111 Rg	112 Cn	113 Nh	114 Fl	115 Mc	116 Lv	117 Ts	118 Og	
													66 Dy	67 Ho	68 Er	69 Tm	70 Yb		

Figure 5.3: Figure adapted from [67]. The periodic table, with elements that have been laser cooled highlighted (note: may not be exhaustive). Many elements are cooled on an $S \rightarrow P$ transition. Group 1 has a lone electron in their outer S states while in their ground state, allowing cooling on the $S \rightarrow P$ transition. Group 2 elements typically have an electron removed, leaving a single electron in their outer S state, and are cooled as ions. The noble gases are cooled by exciting an electron into a metastable S state, and then cooling with an $S \rightarrow P$ transition.

and emission events, leading to a Brownian-like random walk in velocity space. The Doppler cooling limit is not a fundamental temperature limit of laser cooling, and other techniques such as Sisyphus cooling or evaporative cooling can be used to reduce the temperature of a sample to sub-Doppler temperatures.

5.2 Cooling Requirements in an Atom

An ideal atom for laser cooling would have two states: the ground state g and a single excited state e . After being excited from g to e , this ideal atom would have no option but to decay back to g , and the scattering process could be repeated many times to achieve laser cooling. In reality, however, no atoms have a perfect two-level system. Most atoms have additional states that they can decay into, which can require additional lasers to “repump” the atom back into the cooling transition. The more states that a particular atom can decay into that are outside of the two-level cooling transition, the less convenient that atom is to laser cool, as each additional state can require an additional laser frequency to repump the atom back into the cooling transition.

The availability of tuneable lasers at the required frequency can also be a limiting factor, and several atoms have cooling transitions in the deep UV range, where sufficiently powerful lasers are much more difficult to manufacture. Figure 5.3 shows the species that have been laser cooled so far.

5.3 Laser Cooling Be⁺

Be⁺ is the ion of choice for producing sympathetically-cooled positron plasmas because it is the lightest easily laser coolable ion. Before the mechanics of producing these sympathetically-cooled plasmas are discussed (in the following two chapters: 6, 7), it is important to understand the level structure of Be⁺, and how the laser cooling transition is driven.

Figure 5.4 shows energy levels that are relevant when cooling Be⁺. Be⁺ is cooled by exciting the $2s^2S_{1/2}(m_I, m_J) = (\frac{3}{2}, \frac{1}{2}) \rightarrow 2p^2P_{3/2}(\frac{3}{2}, \frac{3}{2})$ transition with σ_+ -polarised $\lambda = 313\text{nm}$ light. Repump light exciting the $2s^2S_{1/2} = (\frac{3}{2}, -\frac{1}{2}) \rightarrow 2p^2P_{3/2}(\frac{3}{2}, \frac{1}{2})$ can be used to optically pump atoms into the cooling transition, as atoms in the $2p^2P_{3/2}(\frac{3}{2}, \frac{1}{2})$ state will decay into the $2s^2S_{1/2}(\frac{3}{2}, \frac{1}{2})$ state (the ground state of the cooling transition) with a probability of 2/3 [68].

The repump transition could be driven by a dedicated repump laser at a different frequency to the cooling laser, but in the experiments described in later chapters, repumping is reliant on off-resonant excitation from the cooling laser. The cooling transition is also off-resonantly depleted by the cooling laser, due to imperfect σ_+ -transitions from the lower state of the cooling transition to other m_J states in the $2p$ manifold. Importantly, the frequency difference between the cooling and repump transitions is smaller than the difference between m_J states in the $2p$ manifold, so ions are repumped into the cooling transition faster than they are lost from imperfect σ_+ -transitions. For example, in a 1T magnetic field, the frequency difference between the repump and cooling transitions is around 8GHz, while the frequency difference between the $2p^2P_{3/2} m_J = \frac{3}{2}$ and $m_J = \frac{1}{2}$ states is around 18 GHz.

In order to find these transition frequencies, the energy levels of Be⁺ in a magnetic field must be calculated. These energy levels are given, to a good approximation, by solving the effective Hamiltonian [69]

$$H = hA\mathbf{J} \cdot \mathbf{I} + \mu_B g_J \mathbf{J} \cdot \mathbf{B} + \mu_B g_I' \mathbf{I} \cdot \mathbf{B}. \quad (5.8)$$

The first term is the magnetic dipole hyperfine interaction, and includes Planck's constant h , the magnetic hyperfine constant A , the total electronic angular momentum \mathbf{J} , and the nuclear spin $\mathbf{I} = 3/2$ [70]. The second term is the electronic Zeeman interaction, which includes the Bohr magneton μ_B , the magnetic field \mathbf{B} , and the g-factor $g_J = -\mu_J/(J\mu_B)$. The final term is the nuclear Zeeman interaction, which includes the g-factor $g_I' = -\mu_I/(I\mu_B)$. Higher-order terms, such as the electric quadrupole hyperfine interaction, are ignored. The eigenvalues of this equation give the relative energies of the hyperfine states within the $2s^2S_{1/2}$ and $2p^2P_{3/2}$ states. The energy difference between these states at zero magnetic field [71] can then be used to determine the transition frequencies between the hyperfine states.

Figure 5.5 shows the energy levels of the $2s^2S_{1/2}$ and $2p^2P_{3/2}$ states as a function of mag-




Figure 5.4: Energy levels relevant when laser cooling Be⁺ ions. The levels are split in a high magnetic field, and are labelled by their m_J and m_I quantum numbers. The $2p$ m_I states have a splitting significantly smaller than the natural linewidth of the cooling transition, and are not shown. The laser cooling transition (purple) and repump transition (blue) are shown.

netic field. The energies are given relative to the $2s^2S_{1/2}$ state at $\mathbf{B} = 0$ in the absence of hyperfine interactions. Magnetic fields in ALPHA are typically either 1T or 3T.

5.4 Laser Cooling of Ions in a Penning Trap

As discussed in chapter 3, a particle in a Penning trap undergoes three superimposed oscillatory motions: the axial bounce motion along the length (z -axis) of the trap at a frequency ω_z , small cyclotron orbits around magnetic field lines in the radial direction at the modified cyclotron frequency ω'_c , and slower magnetron orbits in the radial direction around the centre of the trap, at a frequency ω_m .

Let us first consider cooling in the axial direction. A particle in a Penning trap in a harmonic potential undergoes simple harmonic oscillation in the z -direction, and this motion can be cooled by using a laser beam with a component in the z -direction. This particle's behaviour can be visualised with a momentum-position plot, as in figure 5.7. For a given laser detuning, there is one value of the particle's momentum which will give a doppler shift that cancels out the laser detuning, where the particle will be resonant with the laser (i.e. $\omega - \omega_0 = -kv$, see equation 5.3). This particle passes through this resonant region twice per oscillation as long as its maximum velocity gives it a large enough Doppler shift to cancel out the laser detuning (as long as $|kv_{max}| > |\omega - \omega_0|$). Once the particle has cooled below this point, it is no longer travelling

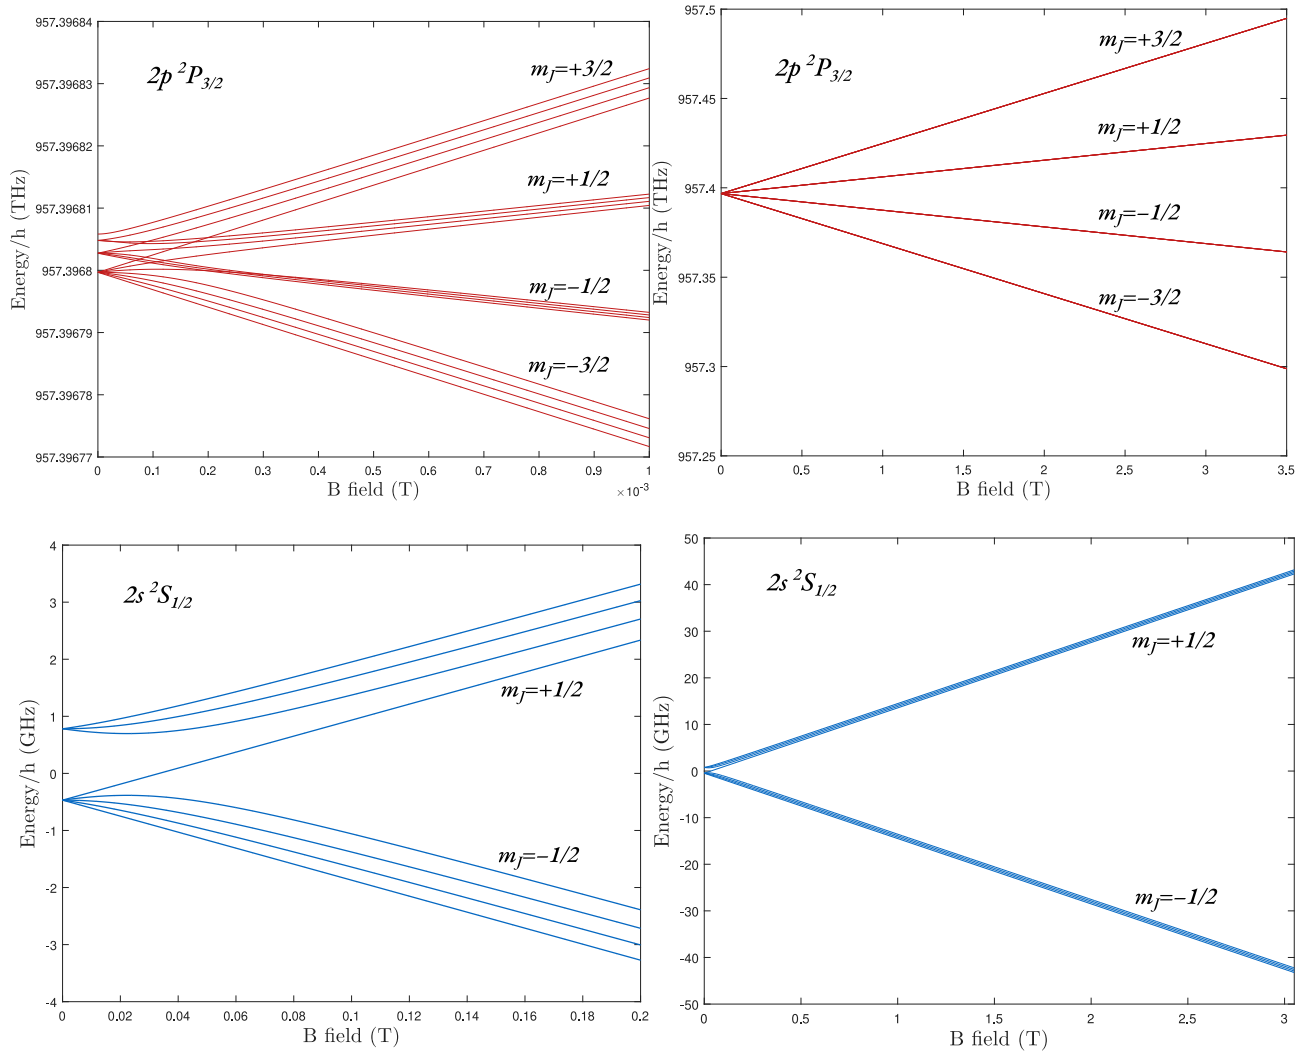


Figure 5.5: Figures created by Dan Maxwell [72]. Energy levels of the $2p^2P_{3/2}$ state (upper) and $2s^2S_{1/2}$ state (lower) in Be^+ as a function of magnetic field for small magnetic fields (left) and large magnetic fields (right). All energy levels are relative to the $2s^2S_{1/2}$ state at zero magnetic field in the absence of hyperfine interactions.




Figure 5.6: Figure created by Dan Maxwell [72]. Frequency difference between the cooling and repump transitions as a function of magnetic field.

fast enough at any point in its oscillation to Doppler shift the light into resonance, and cooling will slow drastically. In order to effectively cool the particle further, the laser detuning can be reduced to bring the particle back into resonance so that it can be cooled further, a process that is known as “chirping” the laser. By carefully choosing the speed of the chirp, the particle can remain close to resonance over a large temperature range. Care needs to be taken to not chirp the laser too quickly, as the chirp can “overtake” the particle and leave the particle stranded at high temperature while the laser is tuned to be resonant with low-temperature particles.

In the radial direction, the cyclotron mode can be cooled in a manner similar to the axial cooling, except using a laser beam with some component in the radial direction.

Cooling of the magnetron mode is more complex in a Penning trap, as the total energy of the magnetron mode is negative. This means that in order to “cool” the motion, e.g. to reduce its radius and velocity, energy needs to be supplied to the magnetron mode. Cooling of the magnetron mode can be achieved in a Penning trap by spatially offsetting the radial laser from the centre of the trap. The offset allows the laser to supply energy to the magnetron mode, thus lowering its temperature. When cooling a plasma, this has the effect of applying a positive torque to the plasma, compressing it. This principle is shown in figure 5.8.

When cooling a plasma, there can be significant mixing between the axial and radial modes, allowing cooling of one degree of freedom to affect the other degrees of freedom. As discussed in chapter 3, the equipartition rate between the degrees of freedom depends on the collision rate within the plasma, which in turn depends on a few parameters including the plasma tempera-




Figure 5.7: Momentum-position phase space plot of a toy simulation of a particle undergoing simple harmonic oscillation in one axis with a resonant damping force, qualitatively similar to a particle in a Penning trap being laser cooled along the z -axis of the trap. The resonance with F_{scatt} occurs at $p = 0.5$, marked with a red line. In this plot, the particle starts at high radius and moves inwards as it cools. The cooling rate increases as the particle cools into resonance with F_{scatt} , and then drops significantly once the particle's momentum is too low to be Doppler-shifted into resonance at any point in its oscillation ($|kv_{max}| < |\omega - \omega_0|$). The magnitude of the cooling force has been greatly exaggerated in this toy simulation for clarity.

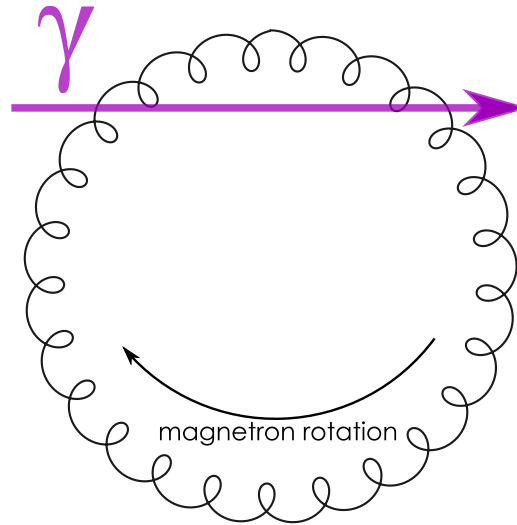


Figure 5.8: A particle’s motion in the radial plane of a Penning trap. The laser is offset from the centre of the plasma to apply a positive torque to the plasma, reducing the velocity and radius of its magnetron oscillations.

ture. As the plasma temperature gets lower, the equipartition rate slows until there is very little mixing between the degrees of freedom at low plasma temperatures.

5.5 Laser Access in a Penning-Malmberg Trap

To be able to directly cool all three Penning trap motions therefore requires laser light with components in both the axial and radial directions. This can be achieved with two separate lasers – one axial and one radial – or by using a single laser at an angle with respect to both the axial and radial motions. Laser access in a Penning-Malmberg trap can be quite restricted, due to the geometry of the trap. The open-ended cylindrical electrode stack allows easy access for an axial laser, but requires a more careful design to allow radial laser access. This could be done, for example, by using a segmented electrode with gaps between the segments that allow laser light to pass into the trapping region. Figure 5.9 shows a generic Penning-Malmberg trap with a segmented electrode to allow radial laser access into the trapping region (a similar trap is used at NIST [73]). A similar setup that allows the radial laser to be significantly angled in the axial direction could provide cooling of all degrees of freedom with a single laser.

The ALPHA-2 Penning-Malmberg trap was not designed to have radial laser access. Even if a segmented electrode was introduced into an updated trap design, the presence of the Neutral Trap magnets and the liquid helium volume surrounding the trap severely limits the amount of space that is available radially around the trap. The trap’s primary function is the production of antihydrogen atoms, and so any design that adds radial laser access must also not interfere with




Figure 5.9: A sketch of a Penning-Malmberg trap electrode stack with a segmented central electrode, allowing radial laser access. An axial laser can enter through the open ends of the stack. This geometry would be difficult to implement in ALPHA-2 due to the very limited space outside the trap.

this in a negative way. As such, laser access in ALPHA-2 is restricted to the dedicated laser access ports, which are offset by 2.4° from the axis of the trap. This means that only the axial mode receives a significant amount of direct cooling, and cooling of the radial modes can only occur through interparticle collisions that transfer energy between the modes.

6/ Experiments with Beryllium

The primary experimental work described in this thesis involves introducing a laser-cooled species into a positron plasma, providing sympathetic cooling to the positrons. These colder positron plasmas have the potential to greatly improve the trapping rate of antihydrogen in ALPHA, as discussed in chapter 4. This species must be positively charged so it can be co-trapped with a positron plasma, it needs to have a convenient laser cooling transition to facilitate easy laser cooling, and its mass should be as low as possible to maximise its ability to sympathetically cool positrons. With these considerations in mind, Be^+ is the natural choice for sympathetically cooling positrons, as it is the lightest easily laser-coolable ion. The main experimental results are split into two chapters. This chapter deals with the production, manipulation and laser cooling of Be^+ plasmas within the ALPHA apparatus. The next chapter (7) discusses how Be^+ and positrons are merged, and subsequently cooled, to produce sympathetically-cooled positron plasmas that could be used in antihydrogen formation.

6.1 Production of Beryllium Ions by Laser Ablation

Ion production in a Penning trap is typically performed using a heated oven to produce neutral atoms, which are then ionised [74, 75]. Ovens reach high temperatures ($> 500\text{K}$) and produce a continuous flux of atoms while in operation. The excess heat and potential vacuum contamination from an oven make it incompatible with the cryogenic, ultra-high vacuum environment of the ALPHA-2 trap, and so another production method must be used. Laser ablation is able to produce a short pulse of ions with minimal heating and contamination of the environment, making it a good candidate for ion production in ALPHA.

The viability of using laser ablation for ion production in ALPHA was investigated previously using a Penning-Malmberg trap external to the primary ALPHA-2 setup (the “ablation test chamber”) by M. Sameed et al. [76, 77]. This ablation test chamber used a very similar ablation source to the one currently mounted inside the ALPHA apparatus. These previous studies demonstrated ablation and loading of Be^+ ions into a Penning-Malmberg trap, which led to a Be^+ source being installed inside the main apparatus.

6.1.1 Laser Ablation

Laser ablation is the process of removing material from a surface using short, high-intensity laser pulses. When a laser pulse is focused onto a target, a few different mechanisms can cause material to be removed from the target. Thermal heating can cause both vaporisation and induce stress in the material that can lead to material ablation from the surface. With high enough photon energy, the laser light can also cause bond-breaking in the material, causing single atoms, clusters, or fragments to desorb from the surface. During this process, atoms that have been removed from the surface can also be ionised. The interplay between these mechanisms, as well as other phenomena, such as plasma formation and surface electric fields generated from electron and ion ejection, cause the ablation process to be rather complex. A more complete description of the ablation process is given by Bäuerle [78].

Although the ablation process is complex, there are some important principles that are relevant to ablation as an ion production method (adapted from [76]):

- The ablation process is dependent on a parameter called the “fluence”, which is a measure of energy deposited per unit area.
- There is a minimum fluence required for the production of ions with laser ablation, which is called the ablation threshold. More highly-charged ions (Be^{2+} , Be^{3+} etc.) have separate, higher ablation thresholds.
- The ion yield increases as fluence increases.
- The average ion kinetic energy increases as fluence increases.
- The ion yield is higher for shorter wavelength ablation pulses.

Using these principles, we can determine that there will be a balance when choosing the correct ablation fluence: too low and the ion yield will be low (or even effectively zero if below the threshold), too high and the ions will have additional kinetic energy, making the ions harder to trap and cool. High ablation fluences will also lead to production of unwanted, more highly-charged ions.

6.1.2 Experimental Ablation Setup

The experimental setup used for the production of Be^+ ions inside ALPHA-2 is shown in figure 6.1. It consists of a custom-built ablation source which is mounted inside the ALPHA-2 vacuum chamber, and a pulsed laser that produces the ablation pulses.

The ablation source (see figure 6.1 a)) consists of a small ($\sim 10\text{mm} \times 15\text{mm} \times 0.5\text{mm}$) piece of beryllium mounted inside an aluminium housing. This housing has a window to allow laser light to enter the source, and a mirror to reflect the laser light onto the beryllium target. Emitted ions are accelerated away from the target by applying a bias voltage to the target. An acceleration plate is positioned opposite the target, which is also electrically biased to accelerate ions away from the target. The acceleration plate has a small aperture, allowing accelerated ions to be ejected from the source, and into the trap. The target and acceleration plate are typically biased with $+25\text{V}$ and -10V respectively. The target, acceleration plate, and source housing are all electrically isolated from one another. The source is mounted on a linear translation stage which travels perpendicularly to the axis of the trap, which allows the source to be moved on-axis for ion production, allowing ions produced to enter the open end of the Penning-Malmberg trap where they can be captured (shown in figure 6.2). Once ablation is complete, the source can be moved away from the trap axis to allow other instruments on the stick to be operated, and to allow other particles to be loaded into the trap.

A single pulse can produce upwards of 10^8 Be^+ ions, but as only a small fraction of the ablated material is ionised (around 0.25% [77]), a single ablation pulse will also liberate up to around $10^{10} - 10^{11}$ neutral atoms from the target. Minimising vacuum contamination is very important, and so the source is designed with a small exit aperture which allows only a small fraction of the ablation plume to escape the source and enter the trap. A higher fraction of ions than neutrals can escape the source, as the ions are guided out of the source by the electric fields from the source bias voltages, and by the magnetic field of the trap solenoid.

The ablation laser is a Quantel Ultra 20 [79], a pulsed Nd:YAG laser with a frequency-tripling stage to produce $\lambda = 355$ nm light. 355 nm light is used as, unlike higher wavelengths (i.e. 532 nm), it will photoionise a significant fraction of the Be during the ablation process, avoiding the need for a second photoionisation laser. The ablation laser can produce pulses with energies of a few millijoules, and a pulse width of 6 ns. Figure 6.1 b) shows a sketch of the optics that guide the light to the ablation source. Pulses from the ablation laser first pass through an adjustable half-waveplate and polarising beam splitter (PBS), which together act as a variable attenuator [80]. Pulses are typically attenuated to $< 100 \mu\text{J}$ for ablation. For a beam waist of $22.5 \mu\text{m}$, this corresponds to fluences of up to $5\text{J}/\text{cm}^2$. The pulse then passes through a pair of convex lenses that form a telescope, which radially expands and collimates the pulse. The pulse passes through a final $f = 25$ cm lens positioned just outside the ALPHA-2 vacuum chamber which focuses the pulse onto the beryllium target. Figure 6.2 shows a schematic view of the ALPHA-2 Penning-Malmberg trap, including the location of the ablation source and the path of the cooling laser. The cooling laser is discussed in more detail later in section 6.3.1.

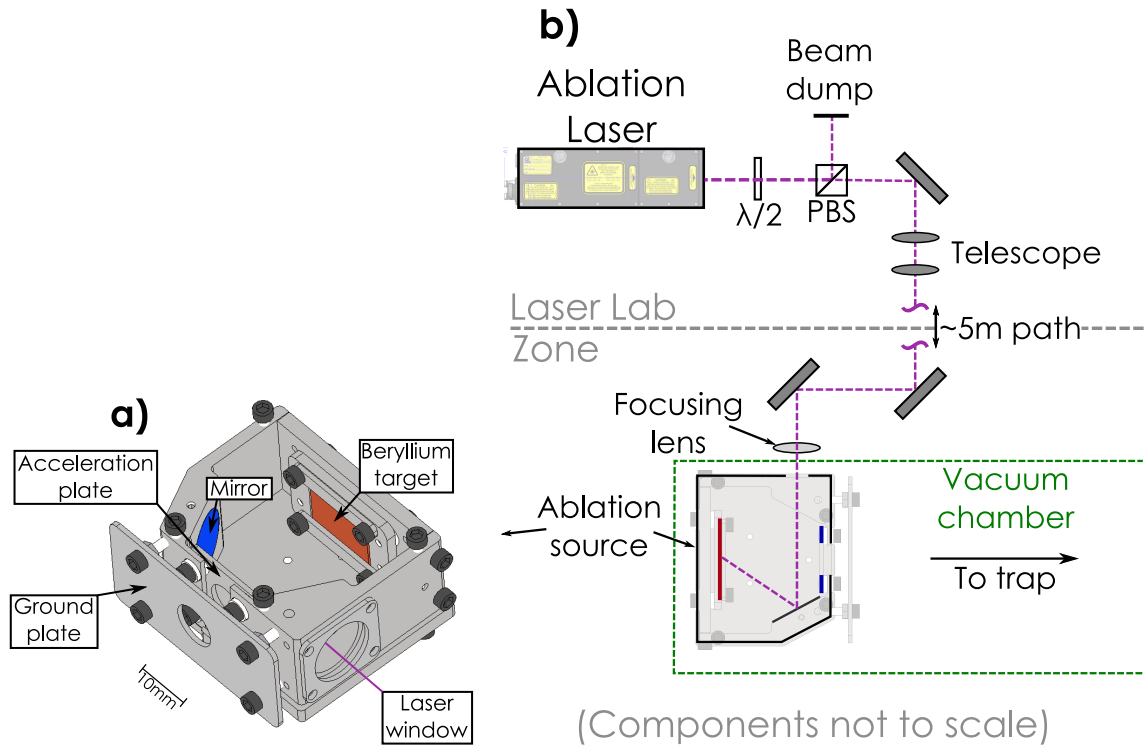


Figure 6.1: a) the ablation source (without covering lid). The beryllium target and acceleration plate are electrically isolated from each other and from the rest of the source. The ablation source is mounted on a linear translator inside the trap vacuum chamber. b) optics setup for transporting laser light from the ablation laser to the source. Laser light first passes through optics in the laser lab: a half-waveplate and polarising beam splitter (PBS) for power control, and a telescope to radially expand and collimate the beam. The light is then guided along a roughly 5m path to the zone, where a final lens focuses the beam onto the beryllium target before it enters the trap’s vacuum chamber. Components and distances not to scale. Some additional mirrors are present in the optical setup that are not pictured.

6.1.3 Ablation Threshold

The ablation threshold for singly-charged Be, ϕ_{th} , is the minimum laser fluence required to produce significant amounts of Be^+ during laser ablation. Ablation pulses with fluences below ϕ_{th} produce a negligible amount of ions. Above ϕ_{th} , the number of singly-charged ions produced scales linearly with fluence [77].

ϕ_{th} was measured for Be^+ by varying the laser pulse energy while measuring the amount of charge leaving the beryllium target. The amount of charge leaving the target was measured with a Faraday cup-style measurement (see chapter 2.7.1). To adjust the pulse energy, a power meter was placed in the beam path and the laser power was adjusted with the half-waveplate. The power meter was then removed from the beam path to allow ablation pulses to strike the Be target for




Figure 6.2: Schematic view of the ALPHA-2 Penning-Malmberg trap, including the location of the beryllium ablation source and cooling laser access. The cooling laser is tightly constrained, only attaining good overlap with trapped species close to the centre of the trap.

the measurements. Several ablation pulses were then sent to the experiment for each power setting, and the charge of the beryllium target was measured for each ablation event. Figure 6.3 shows two measurements of the ablation threshold on the same Be target at two points in time: data for figure 6.3 a) was taken in 2018 and figure 6.3 b) was taken in 2020. In order to extract an ablation threshold, the linear region above the ablation threshold was used to extrapolate back to $y = 0$. The linear regions that were fitted were chosen in an *ad hoc* manner, due to some irregularities in both data sets. Figure 6.3 a) shows a shoulder feature between $1.5 - 2 \text{ J/cm}^2$, which was not included in fluence calculations. This feature was present in other data sets taken at around the same time, and its origin is unknown. Figure 6.3 b) shows a leveling off of the ablated charge above around 2.5 J/cm^2 , which is possibly due to saturation in readout electronics. These points were omitted when calculating values for ϕ_{th} .

These data sets give values for ϕ_{th} of around 2 J/cm^2 . The calculation makes the assumption that the beryllium target is exactly at the focal point of the final focusing lens, where the beam waist has been measured to be $22.5 \mu\text{m}$. This assumption is difficult to verify as the target is enclosed inside the vacuum chamber, making an *in situ* measurement of the spot size on the target prohibitively difficult. As such, this should be considered more of an upper bound on ϕ_{th} . This is in rough agreement with previous measurements taken in the ablation test chamber [77].

6.1.4 Ablated Ion Energy Distributions

Ions produced by ablation are accelerated by the DC bias voltages on the target and acceleration plate, but this is not the only acceleration mechanism. Ablation forms a high-density plasma very close to the surface of the target, and interactions within this plasma cause the ablated ions to be accelerated, with higher pulse energies giving more ions, a higher charge density, and more acceleration [81].

The kinetic energy distribution of an ion bunch can be measured by raising a blocking voltage using one or more electrodes inside the Penning-Malmberg trap, firing an ablation pulse, and




Figure 6.3: Number of ions ablated as a function of laser fluence. Multiple readings were taken for each laser fluence, with means of each point plotted. Error bars are standard errors of the mean. The fit to extract the ablation threshold is shown as a solid red line, which becomes dotted when extrapolating beyond points used in the fit.

measuring the fraction of the ablated ion bunch that has enough kinetic energy to overcome the potential barrier and pass all the way through the trap. Ions with a kinetic energy lower than the potential barrier are reflected, while ions with a kinetic energy higher than the potential barrier are able to pass through it, and are subsequently detected at the far end of the trap with a Faraday Cup. By repeating this several times with different blocking voltages applied in the trap, the energy distribution of the bunch can be deduced. Figure 6.4 shows the results of this experiment for two different laser fluences, $\phi = 3.9 J/cm^2$, and $\phi = 5.0 J/cm^2$. The y-axis on this plot is percentage of the ablated ion bunch that is able to pass all the way through the trap. The relatively low overall efficiency ($< 1\%$) is mostly due to a large fraction of the ions not making it out of the ablation source's small exit aperture. Predictably, the lower fluence produces less energetic ions, and a lower fraction of the ions are able to pass through the trap. Ion bunches can be almost completely blocked using up to around 60V for the $\phi = 3.9 J/cm^2$ bunches, and around 120V for the $\phi = 5.0 J/cm^2$ bunches. Ion bunches with low kinetic energies are both easier to catch and easier to cool, and so laser fluences should be reduced where possible to keep ion kinetic energy to a minimum.




Figure 6.4: Percentage of the ablated ion bunch that is able to pass all the way through the trap as a function of blocking voltage applied in the trap.

6.2 Beryllium Trapping

After ablation, the ions need to be caught in the ALPHA-2 Penning-Malmberg trap. The ablation source is positioned inside the fringe region of ALPHA-2’s main solenoid, at a magnetic field of around 480 Gauss. Figure 6.5 shows the relative positions in the magnetic field of the ablation source and the Penning-Malmberg trap. As ions escape the source after ablation, they are guided by magnetic field lines into the electrode stack. The ions start in a magnetic field of $B = 480$ Gauss, and travel into the magnetic field of $B = 1 \times 10^4 - 3 \times 10^4$ Gauss. The ion bunch will tend to compress as it travels into higher magnetic field regions as the magnetic field lines compress.

Catching Be^+ ions uses the same dynamic catch method that is used for antiprotons, electrons and positrons in ALPHA, which was described in chapter 4.1. In short, ions enter the trap and are reflected by a blocking electrode. Before the ions have time to leave the trap, a second (gate) electrode is raised to complete the trap, and the ions are confined between the blocking and gate electrodes. In order to understand how to optimise the catching procedure, it is instructive to consider the role of a Be^+ plasma used for sympathetic laser cooling of positrons. Such a Be^+ plasma has a similar radius to a typical ALPHA positron plasma (below around 1mm), has a low temperature (below around 5K), and contains $10^5 - 10^6$ Be^+ ions [82]. Therefore, the ideal catching procedure would reproducibly capture $\sim 10^6$ ions, while keeping their temperature and radius as low as possible.




Figure 6.5: Magnetic field as a function of distance from the centre of the ALPHA-2 Penning Trap. The position of the ablation source and Penning-Malmberg trap electrode stack are shown above. Two booster solenoids (blue), which are sometimes energised during Be^+ catching, increase the magnetic field from 1T to 3T.

6.2.1 Trapping Instability

The dynamic catch trapping method tended to have very poor shot-to-shot stability in the number of ions trapped in these experiments. A lot of this instability is thought to stem from the ablation process. Recall from section 6.1.1 that changes in fluence (i.e. changes in the energy of the ablation laser pulse) cause not only changes in the amount of ions produced, but also their kinetic energy. Changes in the kinetic energy will change the time of flight of the ions travelling from the ablation source to the trap, which will change the timing for catching the ions in the trap. Although measured fluctuations in laser pulse energy from shot-to-shot were relatively small ($\pm 5\%$), this could lead to much larger variations in the amount of ions trapped – variations as large as a factor of 2 or 3 were not uncommon. The amount of instability was also quite sensitive to things such as the alignment of the ablation laser, which would drift from day-to-day and required daily adjustments.

With all of these factors, shot-to-shot instability was a real problem when running the experiment that was never completely solved, but some techniques discussed in the rest of this chapter were able to partially alleviate issues with instability in the number of ions trapped.




Figure 6.6: On-axis potentials used for catching beryllium ions. Shaded and numbered regions represent Penning trap electrodes. Ions enter from the left, then the gate electrode (E1) is raised from ground (blue) to 40V (red). Ions are cooled into the catching well, and then transferred over to the reservoir.

6.2.2 Stacking

The instability in the number of ions caught per shot can be partially alleviated by utilising stacking, where multiple shots are caught and instabilities are averaged out. This led to the development of a repeatable stacking sequence that has a few improvements over the basic catching procedure. The potentials used in this stacking procedure are shown in figure 6.6. The first addition is a “catching well”, which is a smaller potential well nested inside the main potential well formed by the barrier and gate electrodes. After ions are caught in the main well, they are allowed to thermalise through collisions (or the cooling laser) for a short period of time. The gate electrode is then returned to ground, and the only ions left in the trap are ones which have lost enough energy to remain trapped in the catching well. These ions are then transferred to the reservoir and the process can be repeated several times. As well as increasing the loading stability, the number of shots caught can be used as another method of varying the total number of ions caught.

One successful stacking procedure that was developed used the potentials shown in figure 6.6 as well as the cooling laser. The reservoir is located in the central electrodes of the trap, and so is continuously cooled during stacking.

Figure 6.7 shows the number of ions caught as a function of the number of ablation pulses used. The relation is roughly linear, but there is some dropoff at higher numbers of stacks, indicating that there is some particle loss during the process. Figure 6.8 shows the stability of a stacking sequence over 13 trials, each using 16 stacks of Be^+ ions. The rather large variation in




Figure 6.7: Number of ions caught as a function of number of ablation pulses used. Error bars are standard errors of the mean of multiple repeat measurements for each data point.

the number of ions shows that the loading process can still have fairly large fluctuations in the amount of ions caught, even when using stacking.

6.2.3 Be^+ Plasma Expansion

Trapped plasmas tend to radially expand, and ions tend to expand faster than lighter species, such as leptons, due to their higher mass and subsequently lower axial bounce frequencies. High expansion rates can cause particle loss as expansion drives the ions into the trap walls, and is accompanied by high amounts of plasma heating [83]. Large Be^+ plasmas are not desirable for sympathetic cooling: Be^+ plasmas should be approximately the same size as the positron plasmas that they are cooling, to maximise the interactions between the two species, and positron plasmas used for antihydrogen formation are typically quite small and quite dense.

This was investigated experimentally, and it was found that a Be^+ plasma will see significant expansion (radius increases of up to 50%) over ~ 10 s, while a positron plasma's radius is stable over > 60 s. Minimising the expansion rate of a plasma in the trap can be achieved in a couple of ways. The potential wells used to hold Be^+ ions should be made as short and deep as possible, as expansion is significantly slower in deeper wells. The magnetic field also plays a role, with higher magnetic fields producing plasmas with smaller radii and slower expansion rates. Both of these effects are linked to reducing the size of the plasma as much as possible, which helps to reduce the effects of imperfections in the trap and therefore reduce expansion rates, as discussed in chapter 3.4. Figure 6.9 shows expansion rates of be plasmas held in a 1T field (blue) and a 3T field (red). The sizes of both plasmas shown are the calculated sizes in a 1T field. This means that the 3T points have been scaled up to show how large the plasma would have been in a 1T




Figure 6.8: Number of ions loaded over several trial runs using the 3T stacking sequence. The shaded blue region indicates the standard deviation.

field, to allow an easier direct comparison between the expansion rates of both plasmas. The 1T plasma shows a markedly faster expansion rate than the 3T plasma, indicating that holding the Be^+ plasma in a higher magnetic field is an effective way of reducing its expansion rate. At longer hold times, the expansion rate appears to slow, but this is likely in part due to the limited size of the MCP, which means that particles above a certain radius are too far out to be imaged on the MCP.

While using deep potential wells is fairly easy to implement, using higher magnetic fields is not always possible. The maximum magnetic field at the centre of the trap is 1T, while the end regions of the trap (the antiproton and positron preparation regions in figure 6.2) are able to reach 3T with the booster solenoids energised (see figure 6.5). This means that ions cannot simultaneously be caught in 3T and be actively laser-cooled during the catching process, as the laser only overlaps the central 1T region of the trap.

6.3 Experimental Cooling Setup

6.3.1 The Cooling Laser

The cooling laser is a commercial Toptica TA-FHG pro. It consists of a master laser diode that produces 1252nm light, which then passes through two frequency doubling cavities to produce 626nm light, and subsequently (continuous wave) 313nm light. The 313nm light, which is used for laser cooling, is then sent through a series of optical components before being guided into the experiment. The laser was typically run with powers up to 10mW at 313nm going into the trap. The layout of the optics and other components used to control and deliver the 313nm




Figure 6.9: Be^+ plasma radius as a function of hold time, in 1T (blue) and 3T (red) magnetic fields. Radially small plasmas of Be^+ ions were prepared, and then held for a variable amount of time. After the hold time, the plasma was ejected from the trap to measure its size. Each data point consists of several repeat measurements of different Be^+ plasmas. Error bars are standard errors of the mean of each point.

light to the experiment are shown in figure 6.10.

313nm light from the cooling laser is first guided through a pair of lenses that form a telescope for beam shaping, and is then passed through an acousto-optic modulator in the double-pass configuration (see section 6.3.2 below), before being guided to the experimental zone. A mirror that is reflective at 313nm and transmissive at 355nm is used to overlap the cooling and ablation beams, allowing both to travel along the same path to the experimental zone¹. Once in the zone, a second identical mirror is used to separate the 313nm cooling beam and the 355nm ablation beam. The cooling laser finally passes through a quarter-waveplate ($\lambda/4$), which gives it its circular polarisation (σ^+) required to drive the laser cooling transition in Be^+ before being guided through the trap. The cooling beam's waist is around 1.2mm at the centre of the trap.

6.3.2 Wavemeter Frequency Control

A small fraction of the 626nm light is picked off and is sent via an optical fibre into a HighFinesse Wavelength Meter WS Ultimate 2 (wavemeter). The wavemeter measures the wavelength of the 626nm light, from which the 313nm light's wavelength is inferred. The wavemeter is regularly calibrated by referencing it to a stable HeNe laser. The wavemeter is also able to control the wavelength of the laser by changing the current going in to the 1252nm laser diode. This forms

¹This is due to space constraints – both lasers must pass through a small hole in the wall between the laser lab and the zone.

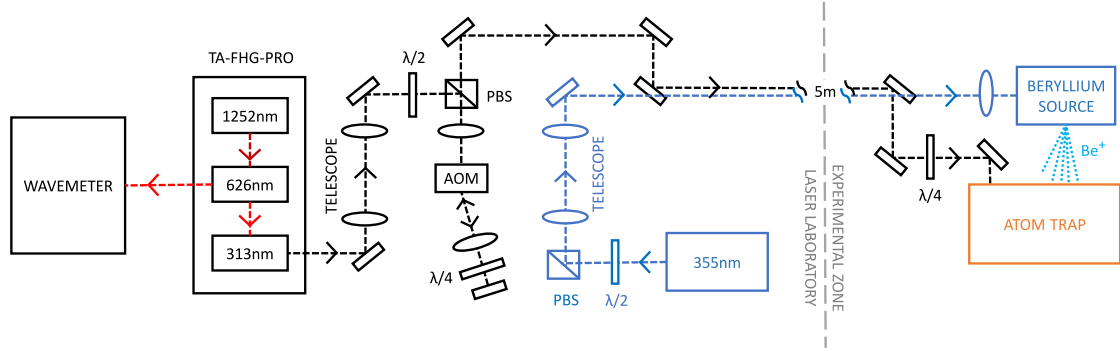


Figure 6.10: Figure adapted from [25]. A schematic representation of the full laser setup used for Be^+ ablation and laser cooling. The ablation laser path (blue) utilises a PBS and half-waveplate for power attenuation, and a telescope for beam shaping. The ablation laser setup is described in more detail above (section 6.1.2). The cooling laser light passes through a telescope, for beam shaping, before being double-passed through an acousto-optic modulator. The cooling laser light is then sent to the zone, where it passes through a quarter-waveplate before entering ALPHA-2.

a feedback loop, which can be used to regulate or stabilise the frequency of the laser. Tuning the laser frequency with the diode current is typically able to tune the frequency of the laser by around 2.5GHz at 313nm without mode-hopping in the current setup. In normal operation, a PID controller is used to lock the laser to a desired frequency setpoint. To change the laser frequency, the setpoint is changed discretely and the PID controller regulates the diode current to smoothly reach the new setpoint.

The wavemeter has a precision of around $\pm 1\text{MHz}$, and is able to measure the laser frequency around every 10ms (the exact readout rate depends on the input power of the laser). The slow readout rate means that frequency changes with this method are slow (typically $> 1\text{s}$), and the nature of a PID loop means that the rate of change in frequency slows as the frequency approaches the setpoint. This method is also subject to some hysteresis, which depends on the state of the laser meaning that the sweep will vary slightly from run-to-run. This method of changing the frequency is used when slow frequency changes over a large frequency range are desired. An example of a frequency sweep using this wavemeter control system is shown in figure 6.11.

6.3.3 Acousto-Optic Modulator

For faster, more deterministic sweeps, an acousto-optic modulator (AOM) is used. An AOM takes an input RF signal at a frequency f_{RF} and frequency shifts laser light passing through it by $m \times f_{RF}$, where $m = \dots, -2, -1, 0, 1, 2, \dots$ is the order of diffraction. The beam is also




Figure 6.11: Change in laser frequency as a function of time as the setpoint is changed from an initial setpoint of -0.3GHz to a final setpoint of -0.1GHz . The setpoint was updated at $t = 0$, and is shown in green. The laser frequency is changed using active feedback from the wavemeter. Frequencies are plotted relative to $f_0 = 957.41\text{THz}$, the resonance frequency of the Be^+ cooling transition.

deflected by an angle $m \times \theta(f)$. The AOM is used in the double-pass configuration [84], the layout of which is shown in figure 6.10. In the double-pass configuration, light passes through the AOM and the first-order diffracted beam is picked off. This is then reflected back through the AOM, where it is diffracted again by the same amount. In this geometry the incident beam and the double-passed outgoing beam are exactly overlapped, and so a polarising beam splitter and quarter-waveplate are used to separate the incoming and outgoing beam. Using the AOM in the double-pass configuration rather than single-pass doubles the frequency range that the light can be modulated over, and (in theory) eliminates the dependence of the beam position on the RF frequency applied to the AOM.

The AOM used in the setup is a model 3200-1210 from EQ Photonics, and is driven by a Gooch and Housego 3307 series RF AOM Driver. The AOM can nominally be driven with a frequency between 180MHz and 220MHz. This gives a total sweep range of 80MHz in the double-pass configuration. The driving frequency, f_{RF} , is controlled by an analog voltage, which is produced by an FPGA (section 6.3.5). This allows for fast and reproducible sweeps over the 80MHz range given by the AOM. The main loss of power in this setup is due to diffraction into other modes ($m \neq 1$). Efficiencies of 33% are typical in our setup.




Figure 6.12: An acousto-optic modulator. An RF signal is applied to the piezoelectric transducer, which converts the signal into acoustic vibrations in the crystal. These acoustic oscillations dynamically change the refractive index in the crystal and cause diffraction to occur.

6.3.4 Fluorescence Detection

When a Be^+ ion de-excites from the upper state of the cooling transition, it emits a photon (i.e. the atom fluoresces). Detecting these fluorescence photons can provide useful information about the Be^+ ions emitting them. In particular, the temperature of the Be^+ ions can be deduced by scanning the frequency of the cooling laser over the transition's resonance and measuring the linewidth of the transition. To pick up fluorescence photons with a high efficiency, the detector (or some light-collection optics) should be close to the ions to maximise the amount of light that can be picked up, and should be placed in a position that minimises background light.

Mounting some fluorescence collection hardware close to the centre of the ALPHA Penning-Malmberg trap would be very challenging due to the lack of space around the trap (discussed in section 5.5). As such, the fluorescence detection setup used in this thesis was mounted outside the trap, which severely limits the detection efficiency. In the Be^+ laser cooling setup in ALPHA, these fluorescence photons are detected with a silicon photomultiplier mounted outside the trap, on one of the laser access windows. This places the silicon photomultiplier around 1.6m from the centre of the trap, subtending a solid angle of around 10^{-6} steradians.

6.3.5 Computer Control Hardware and Software

The beryllium control system has a few main components that control the different parts of the beryllium experiment, and link them in with the rest of the ALPHA control system. At the heart of the beryllium control system is an FPGA, which communicates (mostly via TTL signals) with the main ALPHA sequencer, and with the beryllium setup's hardware. This includes both

the ablation laser and the cooling laser, as well as other pieces of hardware. The FPGA also handles interlocks for both lasers, preventing laser light from entering the main experimental zone when it may be unsafe to do so. The FPGA also exchanges information about the status of the hardware to a windows machine running a custom user interface (UI). This UI allows some manual control of the beryllium setup's hardware, as well as displaying status information about the hardware.

6.4 Typical Cooling Sequence

I will now give an overview of a typical ion-only laser cooling sequence, before discussing some experimental results. The goal of this type of sequence is to load some ions into the trap, laser cool them, and then perform measurements on them. The measurement can be a destructive method, such as dumping to an MCP to obtain the radial profile or the axial temperature, or a fluorescence measurement.

Ions were loaded into the trap using the method described in section 6.2. Up to 16 consecutive shots of ions were used, using catching potentials similar to those shown in figure 6.6. The ions were then transferred to the centre of the trap to be laser cooled.

For optimal laser cooling, the detuning of the laser should be close to the Doppler shift caused by the motion of the ions in the trap. Ions caught in the trap initially have an energy of $> 10\text{eV}$, which corresponds to a Doppler shift of $> 46\text{GHz}$. Ideally, the laser would initially be detuned close to -46GHz from resonance to more efficiently cool the hot ions, but this is not easy due to the limited modehop-free tuning range of the laser setup. The laser can only easily be tuned by around -2.5GHz during an experimental sequence, so this is the frequency that is initially chosen for cooling. The ions are held at this fixed initial detuning for up to 2 minutes, which reduces their temperature from their initial $\sim 10\text{eV}$ ($\sim 8 \times 10^4\text{K}$) down to around 300K . After this initial hold, the laser can be brought closer to resonance by chirping the frequency using the wavemeter (section 6.3.2). Chirping the laser down towards the resonance takes place over around 30s, and the laser is typically brought within -100MHz of the resonance. This can reduce the temperature of the ions to below 1K . There is an optional additional hold at this final laser frequency close to resonance before the ions are dumped out of the trap for a measurement. A typical example of the timeline of this kind of sequence is shown in figure 6.13.




Figure 6.13: Outline of a Be^+ -only cooling sequence. Much of the data taken with only Be^+ ions was taken before the arrested RW technique was developed, and so ion plasmas are left uncompressed.

6.5 Experimental Results

6.5.1 Varying Initial Detuning

As mentioned above, more closely matching the detuning to the thermal velocity of the ions should produce more efficient cooling of the ions. This was investigated by varying the initial detuning from the maximum -2.5GHz down to -1GHz in a series of experimental runs. In these runs the chirp step was omitted, so the sequence just involved loading 8 shots of ions, moving them to the centre of the trap and holding them for 150s in the laser at a fixed frequency. Their axial temperature was measured with the plasma measurement technique described in section 3.9.2. Figure 6.14 shows the results of these runs, and shows a strong dependence of the ions' temperature on the initial detuning. Data for "no-laser" runs is shown on the same plot – timing was identical to the runs with laser on, except the laser was blocked from the experiment with a shutter. The temperature without the laser was almost $15,000\text{K}$, while the temperature with the laser varied from over 6000K at -1GHz detuning to 300K at -2.5GHz detuning.

6.5.2 Fluorescence Temperature Measurements

Using the setup described in section 6.3.4, fluorescence light emitted by ions that are illuminated by the cooling laser can be picked up. Fluorescence signals from the ions were used to infer the temperature of the ions by preparing a sample of cold ions and sweeping the laser across the cooling transition's resonance. This allows a measurement of the width of the transition, from which the temperature can be inferred.

In figure 6.15, 8 shots of Be^+ were loaded, transferred to the centre of the trap without being compressed, and laser cooled for 60s at a detuning of -2GHz . The detuning was then reduced to 60MHz, before the AOM was used to sweep a further 80MHz across the resonance in 80ms. During the AOM sweep, the fluorescence signal was recorded with the fluorescence detection setup (section 6.3.4). After the sweep, the ions were ejected from the trap and a second, identical AOM sweep was performed to gather background data.




Figure 6.14: Be^+ plasma temperature, as a function of initial detuning. The red line shows the temperature when the laser is off, while the blue points show the temperature when the laser is on. The cooling efficiency drops as the laser is brought closer to the resonance. Presumably, this is because the ions are quite hot to begin, and so require quite a large detuning to effectively cool them from their initial high temperature. Error bars are standard errors of the mean of repeat measurements at each detuning.

Most of the background comes from laser light that is scattered inside the trap and picked up by the detector. The amount of scattered light is very dependent on the alignment of the laser within the trap, which can drift from day-to-day. Sweeping the AOM also causes slight changes in the angle of the laser beam coming out of the AOM, which can cause the background to change over the course of an AOM sweep. The magnitude of this effect is dependent on the alignment of the laser both through the AOM and through the trap. For these reasons, each fluorescence run is followed by a background measurement, which can be subtracted from the signal.

Extracting the temperature of the ions from the shape of the resonance is challenging, as the act of sweeping the laser across the cooling transition changes the cooling power, and thus the temperature of the ions, during the sweep. This is especially pronounced as the laser passes the resonance, moving from a negative detuning where the ions are cooled to a positive detuning where the ions are heated. In figure 6.15, this causes an asymmetric lineshape, with a much steeper right-hand side (positive detuning).

The left-hand side of the peak in figure 6.15 was therefore fit with a Gaussian function, and the full width at half maximum can be used to calculate the ion temperature. This assumes that the peak's width is Doppler dominated, and that the increased cooling as the laser is swept towards the resonance has a small effect on the temperature. Both of these effects work to increase the width of the Gaussian fit, so the calculated ion temperature is a conservative overestimate.

Figure 6.15 shows a typical fitted fluorescence sweep from a cold sample of ions, and the fit corresponds to a temperature of $\sim 160\text{mK}$.




Figure 6.15: Fluorescence signal from a Be^+ plasma as the frequency of the laser is swept over the cooling transition. Signal data (blue) was taken with ions in the trap, as described in the text. Background data was taken after ions were removed from the trap. The signal data has been fit with a Gaussian on the left-hand side of the peak. The solid region of the fit indicates the region of the data that was used for the fit. On the right-hand side, the sharp drop-off of the signal is attributed to heating caused by the laser when the detuning is positive. This fit corresponds to a temperature of $\sim 160\text{mK}$, and assumes that the width of the peak is dominated by Doppler broadening.

6.5.3 Polarisation Dependence

The cooling transition in Be^+ is driven by σ^+ -polarised light (chapter 5.3). A quarter-waveplate is used to rotate the linearly-polarised light coming from the laser into σ^+ -polarised light before it enters the trap. To verify that the light being sent into the trap was in fact σ^+ -polarised, the angle of the quarter-waveplate was adjusted during laser cooling, and the ion temperature was measured as a function of waveplate angle. The results of this scan are shown in figure 6.16, and show a large increase in ion temperature as the laser polarisation is moved far from the circular σ^+ -polarisation. There is a fairly large region, roughly -30° to $+10^\circ$ where the measured ion temperature is mostly independent of polarisation. The ion temperature is therefore not very sensitive to changes in polarisation, but will eventually increase once the light is far enough away from σ^+ -polarisation. The fact that the flat region is not centred on 0° implies that the direction of the linearly-polarised incoming laser light is several degrees off from the assumed direction of the polarisation.

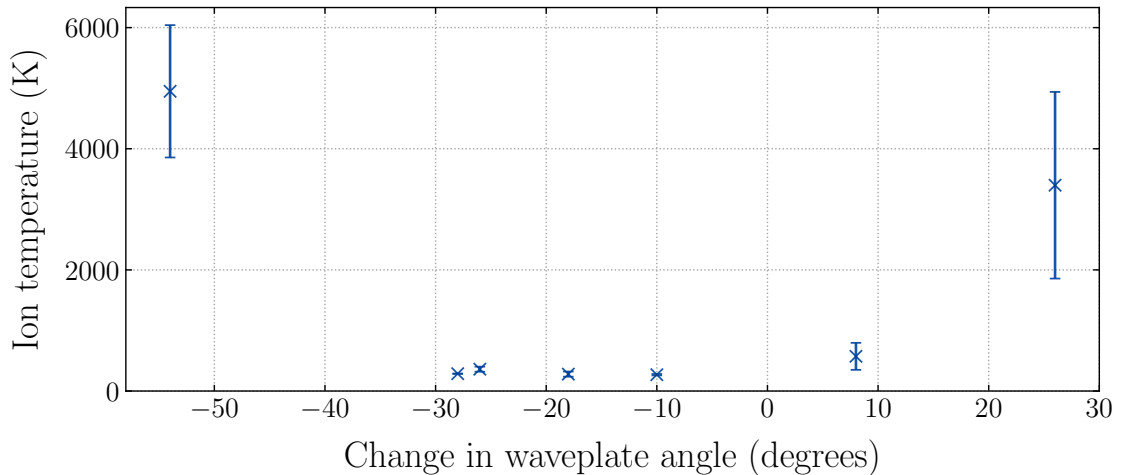


Figure 6.16: Be^+ plasma temperature, measured with the plasma temperature diagnostic, as a function of laser polarisation. Each plasma contained around 2×10^5 ions. Error bars are standard errors of the mean of repeat measurements.

6.5.4 Be^+ Heating Rate

Be^+ ions will tend to have a faster heating rate than lighter species in the trap, as discussed in section 6.2.3. Quantifying this heating rate is useful, as it allows comparisons to the heating rate of a mixed Be^+ -positron plasma in the following chapter 7. This can help determine how much influence the heating from the presence of the Be^+ ions affects the heating of the mixed plasma as a whole. It also allows us to determine how long Be^+ ions will remain cold after the cooling

laser is turned off.

Figure 6.17 shows an experiment where ions were cooled as described in section 6.2.3, and then the cooling laser was turned off and the ions were allowed to sit in the trap during a variable hold time. After the hold, the ions were ejected from the trap, and their temperature measured.

The results show that the Be^+ ions quickly heat, with a heating rate of around 9K/s in the first 40 seconds. This trend approximately continues to higher hold times, but there is an unexpected reduction in temperature after the 75s hold time.




Figure 6.17: Be^+ temperature after laser cooling and followed by a variable hold with the laser off. Each plasma contained around 2×10^5 ions. Error bars are standard errors of the mean of repeat measurements.

6.5.5 MCP Distortion

When imaging large positively-charged plasmas on an MCP, the edges of the plasma can become distorted. The cause of this was initially unknown, but after some investigation it was determined that these distortions are from electrostatic potentials on the MCP housing. During operation, the tabs that are used to connect the HV cables to the MCP back, MCP front, and phosphor screen (see figure 6.18 b)) are not shielded from the particles incident on the MCP front. This means that particles that are close to the edge of the MCP can see the high voltages that are applied to the tabs - particularly the high positive voltages that are applied to the MCP back and phosphor tabs. These high voltages can deflect incoming particles away from the edges of the MCP. Another electrostatic issue arises from the bolts that hold the different layers of the MCP together. These bolts are not grounded, and are electrically isolated from all other parts of the MCP by nonconducting washers. It was therefore hypothesised that these bolts could become charged to a high positive voltage during operation from their close proximity to the

MCP back and phosphor screen, and could deflect incoming particles in a similar way to the high voltage tabs.

Some basic simulations were performed using COMSOL (a piece of multiphysics simulation software [85]) to see if voltages on the MCP tabs could cause deflection of incoming particles, and it was found that electrostatic distortions of this type were a plausible explanation. The simulations used a basic model of the MCP, and did not include the bolts. Figure 6.18 a) shows a contour plot of the peak potential that a particle travelling towards the MCP would experience. In this example, the MCP front, back and phosphor were charged with -100V , 800V , and 5000V respectively. The main effect in this image is from the phosphor tab, which is located in the lower-right. The potentials that a particle close to the edge of the MCP and near the phosphor tab can experience are similar in magnitude to typical ejection energies from the trap.

Since the approximate magnitude of the potentials is enough to deflect particles, the next thing to check is whether the apertures seen on MCP images line up with the tabs on the MCP. Figure 6.18 b) shows an MCP image of a large plasma digitally overlaid on a photo of a real MCP used in ALPHA. The smaller apertures line up almost perfectly with the three bolts, and the large aperture in the lower right of the image lines up with the phosphor tab.

There are a few solutions to this problem: one is to just eject particles with higher energy, so they can get over the potential barriers from the voltages on the MCP housing. Another solution, which was implemented after most of the experimental work carried out in this thesis, is to include a metal shield that attaches to the MCP, which shields particles from the high-voltage tabs.

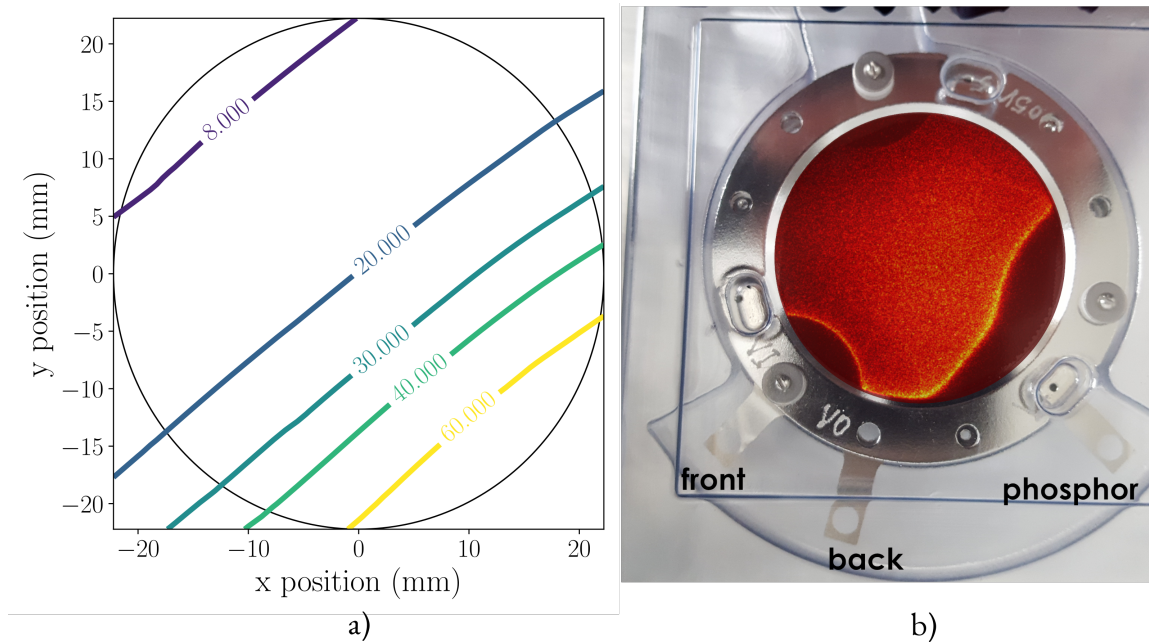


Figure 6.18: a) Simulated contour plot of the maximum potential experienced by a particle approaching the MCP. Coloured lines show equipotentials (V). The MCP front, back, and phosphor were charged with -100V , 800V , and 1000V respectively. The circle represents the edge of the active area of the MCP. b) A picture of a real MCP (removed from the apparatus), with an image of a large plasma digitally overlaid on top. The bolts and phosphor screen align very well with the locations of the apertures.

7/ Sympathetic Cooling of Positrons

Obtaining colder positron plasmas is important when making antihydrogen in ALPHA, as there is a strong correlation between the temperature of the positron plasma and the trapping rate of antihydrogen atoms (as discussed in chapter 4.6.3). Colder positron plasmas increase the production rate of antihydrogen, and reduce the temperature of any antihydrogen atoms produced, making them easier to trap.

This chapter will focus on the production of sympathetically-cooled positron plasmas, which is achieved by mixing a positron plasma with the laser-cooled beryllium ions discussed in chapter 6. In the mixed plasma, the Be^+ ions are cooled directly with laser cooling, while the positrons cool sympathetically through collisions with the Be^+ ions. Sympathetic cooling of positrons using laser-cooled Be^+ ions was investigated in 2003 by Jelenković et al [21]. In their study, a few thousand positrons were sympathetically cooled in a Penning-Malmberg trap to below 5K, and they inferred the positron temperature by measuring the radial distribution of the Be^+ ions (see section 7.6.4), rather than by performing a direct temperature measurement on the positrons themselves. In contrast, ALPHA uses between 1×10^6 and 3×10^6 positrons in each antihydrogen formation cycle, and so the experiments outlined in this chapter focus on sympathetically cooling around three orders of magnitude more positrons than this previous study. In addition, the plasma temperature diagnostic method discussed in chapter 3.9.2 allows for direct measurements of the positron temperature in ALPHA, which can be used in tandem with positron temperatures inferred from radial distributions of Be^+ .

Since the final goal here is to use these sympathetically-cooled positrons for antihydrogen formation, it is helpful to consider the constraints imposed by the antihydrogen production sequence when developing a sequence for production of a sympathetically-cooled positron plasma. The number of positrons used in current antihydrogen production sequences is between 1×10^6 and 3×10^6 , and so a sympathetic cooling sequence should use a similar amount. The number of Be^+ ions used should be minimised, as the presence of additional Be^+ ions in the mixed plasma can drive heating and radial expansion, and having Be^+ ions in the trap during antihydrogen formation could potentially lead to antiproton loss through Be^+ ions capturing antiprotons [82]. The next major constraint is on time, as an antihydrogen sequence takes around 4 minutes: the

sympathetic cooling sequence should be able to fit comfortably inside this time window so that it does not slow down antihydrogen production sequences. The final major constraint comes from the fact that the magnetic minimum trap must be energised for the duration of the sequence when stacking antihydrogen atoms. The magnetic minimum trap distorts the solenoidal field used for trapping charged species, which manifests as increased radial expansion and heating of radially large plasmas (see chapter 3.4). This means that Be^+ ions and positrons should be kept radially small wherever possible. Keeping radii small can be done through compression techniques, such as rotating wall (RW), and by keeping particle numbers low.

The constraints imposed during an antihydrogen production sequence should be considered guidelines for the rest of this chapter rather than strict rules or requirements, as this chapter is just discussing the production of sympathetically cooled positrons rather than their implementation into an actual antihydrogen production sequence. Further discussion of how sympathetically cooled positrons might be implemented into an actual antihydrogen production sequence can be found in the following chapter (8).

7.1 Positron Preparation

Positrons are accumulated in the positron accumulator before being transferred to ALPHA-2 through the ~ 6 metre long beamline (see chapter 2). The positrons then undergo a series of preparation steps, with the goal of reducing their temperature and radial size, before being mixed with Be^+ ions.

During the transfer from the accumulator, the positron plasma often picks up impurity ions from ionising background gas in the vacuum chamber. To remove these impurity ions, a broad-spectrum RF drive is applied to a nearby trap electrode, which autoresonantly heats any impurity ions in the positron plasma, driving them out of the trap. The positrons are largely unaffected, as the RF drive sweeps through frequencies that are outside of their axial motional frequencies. The positrons then undergo simultaneous strong drive regime RW compression and evaporative cooling (SDREVC - chapter 3.7.1). The procedure is able to set the number of positrons and their density to a set value, with a reproducibility of a few %. The positrons are then held for several more seconds to allow them to cyclotron cool to their equilibrium temperature with the trap walls. The entire positron accumulation and preparation takes around 155 seconds, as shown in figure 7.1, and produces a positron plasma containing $1 \times 10^6 - 3 \times 10^6$ positrons with a density of around $1 \times 10^{14} \text{m}^{-3}$ and a temperature of around 20K.

7.2 Beryllium Preparation

Sequences for beryllium catching and preparation for sympathetic cooling experiments are based on sequences from the previous chapter. The final version of the sympathetic cooling sequence is shown in figure 7.1, and is the sequence that was used for most of the experimental data collected for, and presented in, the rest of this chapter. This sequence consists of three main stages: the beryllium preparation, shown in purple, the positron preparation, shown in green, and the mixture manipulations, shown in blue.

Keeping the radial size of Be^+ plasmas small when mixing with positrons is important as the mixed plasma will inherit its angular momentum, and therefore radial size, from both species being mixed together. Rather than introducing a large Be^+ plasma into a positron plasma to produce a very large mixed plasma that is difficult to compress, it was instead found that keeping the Be^+ plasma as small as reasonably possible before mixing them with a positron plasma was a much more manageable route to a small mixed plasma.

In the beryllium preparation stage, ions are caught using the stacking procedure outlined in chapter 6.2. After each stack is caught, it is transferred to the reservoir well to be held while another stack is caught. The reservoir well in this sequence is in the centre of the atom trap, which importantly allows laser cooling during the stacking process. Each stack is held for ~ 5 seconds, which allows some time for the ions in the reservoir to cool before another shot is caught. During the stacking process, the laser is detuned far from resonance, by up to -2.5 GHz, for efficient cooling of the hot ions. This corresponds to a detuning of -129Γ , where $\Gamma = 2\pi \times 19.6\text{MHz}$ is the natural linewidth of the Be^+ cooling transition.

7.2.1 Arrested Rotating Wall Ion Compression

Compressing a plasma with a rotating wall (RW) requires a cooling mechanism, as the rotating wall will heat a plasma and cause it to decouple from the RW drive (see section 3.5). When compressing a pure Be^+ plasma, laser cooling is a great choice of cooling mechanism. Unfortunately, the geometry of the ALPHA-2 trap prevents simultaneous laser cooling and RW compression: there is one RW electrode close to each end of the trap, neither of which are close enough to the central laser cooling region to allow laser cooling while the RW is acting on the Be^+ plasma (shown in figure 6.2).

A solution to this problem was developed whereby the plasma is first laser cooled in the centre of the trap, before being moved to one of the RW electrodes where it can be compressed, and then moved back to the centre of the trap to be recooled. This ends up as a balancing act between cooling and compression, as the plasma will tend to expand during laser cooling holds

and heat up during RW holds. This process can be repeated multiple times to achieve higher levels of compression.

After a long period of trial and error, experimenting with different numbers and lengths of RW holds and laser cooling holds, the most consistent compression was produced from the sequence shown in figure 7.1. This sequence uses four laser cooling holds with durations of 30 – 120s interleaved with two RW holds with durations of 80s and 50s. The compression achieved was modest, but not nearly as strong as a directly compressed positron plasma. The time taken for arrested RW to work is several times longer than a standard RW compression, which severely limits its usefulness as a tool to be used in ALPHA. Its use should be reserved for times when other compression techniques are not viable. Despite this major drawback, arrested RW was used for most of the experimental data presented in this chapter.

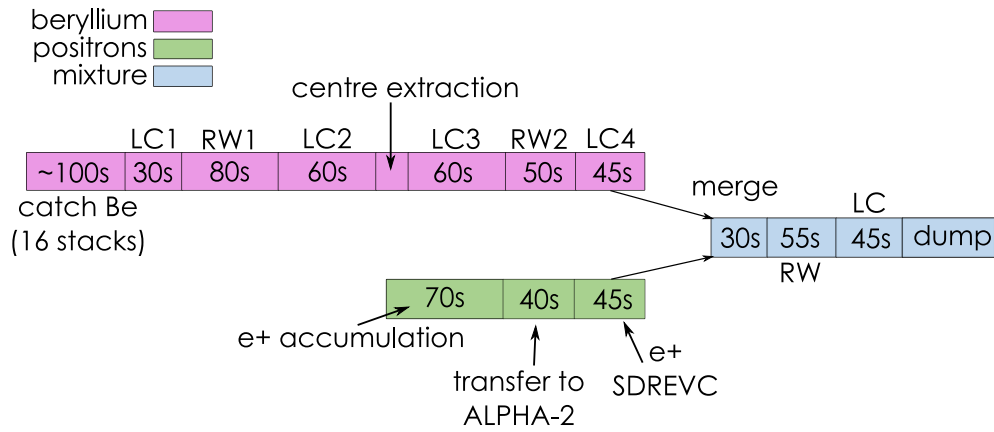


Figure 7.1: A typical sympathetic cooling sequence that includes the arrested RW. Be^+ ions (pink) are prepared in tandem with positrons (green), before being merged. The Be^+ preparation involved several interleaved laser cooling (LC) and rotating wall (RW) compression holds, which form the arrested RW process.

7.2.2 Centre Extraction

After the second laser cooling hold, the ions undergo a centre extraction process. The goal of centre extraction is to selectively remove high-radius ions from the trap, leaving only lower-radius ions. This effectively “compresses” the plasma. This technique was used in tandem with the RW, as it was found that the RW alone was able to produce a well-compressed core of ions, but often left a low-density halo of ions at higher radius, around the edge of the core. This halo is visible in figure 7.2 b), and is much more pronounced than in figure 7.2 a), where centre extraction took place.

The potential manipulations are very similar to the manipulations used in the EVC process. Recall from the discussion of EVC in chapter 3.7 that the first particles to leave the plasma during




Figure 7.2: a) A plasma of Be^+ ions after the arrested RW process, and after the centre extraction process, imaged on a microchannel plate (MCP). The axes show the calculated size in a 1 T field. The calculation assumes that the particles follow the magnetic field lines as they are ejected from the high-field region of the trap to the lower-field region of the MCP. b) A plasma that has been compressed, but without centre extraction. A high-radius, uncompressed halo remains. Image distortion from electrostatic effects caused by the MCP housing are clearly visible. These electrostatic distortion effects are discussed in detail in chapter 6.5.5.

the EVC process are both hot and low-radius, as the confining potential of the trap is lowest along the central trap axis (i.e. at low radius). The main difference with centre extraction is that instead of throwing away these hot, low-radius ions, they are instead siphoned into an adjacent potential well and kept. The ions that remain in the starting potential well, i.e. the colder and higher-radius ions, can then be ejected from the trap, thus leaving a plasma in the trap with a smaller radius than the initial plasma. The plasma is also effectively heated as a side effect of the centre extraction process, and needs to be re-cooled.

7.3 Merging Ions and Positrons

After both species have been prepared, they need to be merged so sympathetic cooling can take place. The Be^+ and positrons are held in adjacent potential wells, and are then merged by lifting the bottom of the positron well to “pour” the positrons into the Be^+ well. Lifting the positrons and then letting them fall into the Be^+ potential well gives them additional kinetic energy which must be cooled, but this approach was used because it avoids giving excess kinetic energy to the Be^+ ions. The cooling laser is left on for the duration of the merge, which provides cooling to counteract the additional energy gained by the positrons as they are merged with the Be^+ ions. The merge takes place over 30s, and uses the potentials shown in figure 7.3. It was found experimentally that merges that were too fast caused excessive plasma heating, while merges that were too slow caused excessive plasma expansion. The time taken for the merge was tuned until both the temperature and the radius of the resulting mixed plasma were reasonable.




Figure 7.3: On-axis potentials used for merging Be^+ ions with positrons. Be^+ ions are initially in the left-hand well, with positrons in the right-hand well. The two species are merged over 30s by raising the bottom of the positron well, allowing positrons to flow into the Be^+ well. This process is generally performed with the cooling laser on, to actively cool the Be^+ during the merge.

7.3.1 Compression of the Mixture

The mixture must be re-compressed after the merge step, as the slow merge process gives ample time for the mixture to expand (the expansion rate of the mixture is investigated below, in section 7.8). The re-compression of the mixture uses a 1MHz RW. Although not experimentally confirmed, the most likely compression mechanism is that the positrons in the mixture undergo strong drive regime compression, and the Be^+ ions are sympathetically compressed through collisions with the positrons. The compressed mixture can be seen in figure 7.4.




Figure 7.4: Left: An MCP image of a mixed Be^+ -positron plasma, after RW compression. This image is dominated by positrons. Right: Same as left, but positrons have been removed with an ekick prior to imaging.

7.4 Positron Temperature Measurements

Positron temperatures can be measured by using the plasma temperature measurement technique discussed in chapter 3.9.2. This technique involves slowly lowering a confining potential barrier, and measuring how many particles leave the trap as a function of the confining barrier's height. The temperature of the entire plasma can then be inferred from the hot, exponential tail of particles that leave the trap first. As mentioned in chapter 3.9.2, temperature measurements taken using this technique suffer from systematic effects, and so these temperature measurements should be taken relative to each other, rather than as an absolute measurement of positron temperatures. Particle-in-a-cell simulations indicate that these systematic effects are around 10 – 15%.

Since only the first particles that leave the trap contain temperature information, it is unnecessary to eject all of the positrons from the trap to measure their temperature. In a cold,

mixed plasma, the Be^+ ions will centrifugally separate from the positrons, and as this temperature measurement will initially extract a central core of particles that are used to determine the temperature, Be^+ ions are unlikely to be ejected in large quantities if the well depth is only lowered slightly below the point where particles start escaping. The majority of the particles that do escape the trap will be positrons, which have centrifugally separated to the centre of the trap. Additionally, positrons produce a significantly stronger signal than Be^+ (around $10\times$ higher) on an MCP. These effects work together to mean that these partial ejections produce signals that are almost entirely dominated by positrons.

This can be exploited to perform measurements of both the positron temperature and the number of Be^+ ions present in a mixed plasma in a single run. After the partial ejection to measure the positron temperature, the mixture with the beryllium and the remaining positrons is then re-compressed, before removing the rest of the positrons with an ekick 3.8.3. The remaining ions can then be counted by imaging them on an MCP. The re-compression before imaging the Be^+ ions is required as if the Be^+ plasma is too large, ions at a higher radius can be deflected by the electrostatic effects near the edge of the MCP, or miss the MCP entirely.

7.5 Simulations of Sympathetic Cooling

Simulations of sympathetic cooling of positrons using Be^+ ions in ALPHA were performed in 2014 by Madsen et al. [82]. Their simulation uses code based on the numeric, two-dimensional equilibrium code (N2DEC) which was previously used to simulate centrifugal separation between antiprotons and electrons in ALPHA [45]. The code solves Poisson's equations self-consistently for both the electric field and the spatial profile in 2 dimensions, (r, z) . The code was extended to include sympathetic cooling effects by fixing the temperature of the Be^+ ions and calculating the equilibrium temperature of the positrons. The collision rate between the two species was dependent on the overlap between them, which was given by the Poisson solver. It also included a cyclotron cooling term, and an additional term which was included to cover additional sources of heating in the trap that are observed experimentally.

In these simulations, 2×10^6 positrons were mixed with a varying number of Be^+ ions, and a plot of the positron equilibrium temperature as a function of the number of Be^+ ions is shown in figure 7.5. The minimum equilibrium temperature of the positrons decreases as the number of Be^+ ions is increased, but with strongly diminishing returns. These results can be used to inform how many Be^+ ions to use for the experiment. The number of Be^+ ions should be kept as low as possible, as discussed in section 7, but these simulations show that the effectiveness of the sympathetic cooling sharply drops off if there are too few ions.

There is therefore a sweet spot where the number of ions is high enough to effectively sym-

pathetically cool the positrons, but not too high, to avoid excessive heating and expansion from having too many ions. The simulations imply that this sweet spot is around $1 \times 10^5 - 5 \times 10^5$ for cooling 2×10^6 positrons. The results in figure 7.5 also show that positron plasmas with lower densities reach lower temperatures, although the dependence of temperature on positron density is not very pronounced. These simulations were used to inform the experiments discussed in section 7.6.2.

The N2DEC code also gives the spatial profile of the mixed plasma at a given temperature. By re-running the code with different input parameters, results from the simulation can be compared to spatial profiles obtained experimentally. By obtaining a best-fit result, the experimental plasma's parameters can be estimated. The input parameters of the code are the number and density of both species, and the equilibrium temperature. This is essentially identical to the method used by Jelenković et al [21] to extract positron temperatures from the radial profiles of their Be^+ ions. This technique is used to determine the temperature of centrifugally-separated mixed plasmas in section 7.6.3.

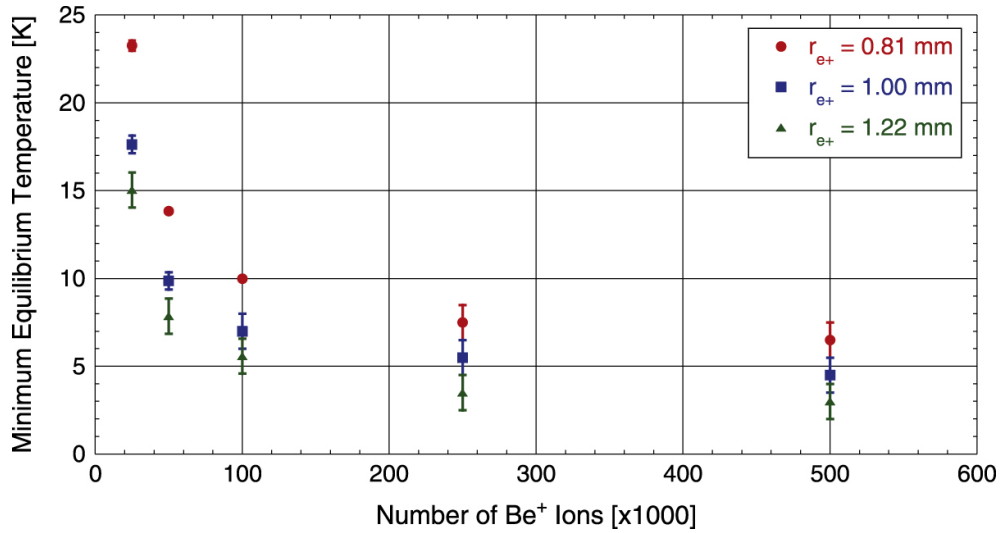


Figure 7.5: Reproduced from [82]. Calculated minimum temperature of 2×10^6 positrons sympathetically cooled with a variable number of Be^+ ions. The calculations were performed for three different plasma radii: 0.81mm, 1mm, and 1.22mm. The corresponding peak densities of these three plasmas were $9.1, 6, 2,$ and $4.3 \times 10^{13} \text{m}^{-3}$ respectively. The external heating rate used was 52K s^{-1} .

7.6 Experimental Results

7.6.1 Varying Final Detuning

During a sympathetic cooling sequence, two laser frequencies are used for cooling. The first is far-detuned from resonance, with a detuning of up to -129Γ . This is used to keep the temperature under control during the beryllium preparation, and is used to pre-cool the Be^+ ions before the merge and final cooling of the mixture occurs. After the Be^+ ions have been merged with the positrons and compressed, the laser is then chirped using the wavemeter control system (see section 6.3.2) to a final detuning, much closer to resonance. The value of this final detuning affects the temperature of the Be^+ ions, and should, if sympathetic cooling is occurring, affect the temperature of the positrons as well.

This was investigated in the following experiment: first, a plasma of $3.8 \pm 0.1 \times 10^5$ Be^+ ions was prepared as described in section 7.2, with the cooling laser at a fixed detuning of -129Γ . 2.6×10^6 positrons were prepared as described in section 7.1, and the Be^+ and positrons were merged with the cooling laser on, as outlined in section 7.3. After the merge, the mixture was re-compressed with a RW, which takes around 30s. Immediately after the mixture was re-compressed, it was transferred back to the middle of the trap to be laser cooled. The laser detuning was held fixed for 5s at a value of -100Γ , before the laser was chirped over around 40s from -100Γ to a final detuning, closer to resonance. When the laser's detuning reached the final detuning setpoint, a partial ejection was used to measure the positron temperature. This was repeated several times while varying the value of the final detuning from around -60Γ to $+5\Gamma$, the results of which are shown in figure 7.6. There is a clear reduction in the positron temperature as the final detuning is brought closer to resonance, with a minimum positron temperature of 6.6 ± 0.6 K. This is followed by a sharp increase in the temperature as the final detuning crosses the resonance, and passes onto the heating side of the transition. The increase in temperature occurs just before the laser crosses the resonance – this is likely due to a systematic uncertainty in the determination of the laser frequency. The two lines show useful reference values – the green line (17.3 ± 0.5 K) shows the temperature of a positron plasma prepared in the same manner, except without any beryllium ions loaded in the trap. The blue line (841 ± 135 K) shows the temperature of the positrons when the laser is blocked for the final mixture cooling stage (even in this “no-laser” case, the laser is used during the beryllium preparation stages of the sequence).

The minimum achieved positron temperature of around 7K is significantly lower than the 17K seen without sympathetic cooling, but is still much higher than the temperature of the Be^+ ions measured in the previous chapter of around 160mK. This suggests that there is a mechanism that prevents the positrons from completely thermalising with the laser-cooled Be^+ ions.

The dominant mechanism is likely to be centrifugal separation, which is discussed in detail later in this chapter (7.6.3). There is good agreement between the experimentally-measured minimum positron temperatures (figure 7.6) and the simulated positron temperatures (figure 7.5).




Figure 7.6: Axial positron temperature in a mixed Be^+ -positron plasma as a function of final laser detuning. The mixed plasma contained 2.6×10^6 positrons, and $3.8 \pm 0.1 \times 10^5$ Be^+ ions. The laser was chirped from an initial detuning of -104Γ to the final detuning in approximately 40s. The error bars give the standard error of multiple measurements at each detuning. The green line shows the positron temperature in the absence of any Be^+ ions ($17.3 \pm 0.5\text{K}$). The red line shows the temperature when the laser is off, and the Be^+ ions are not laser cooled ($841 \pm 135\text{K}$). The shading around the lines shows standard error.

7.6.2 Varying number of ions

It was discussed earlier (in section 7.5) that there is a balance when determining the amount of ions that should be used for sympathetic cooling: too few, and the cooling power will be too low, but too many and the Be^+ ions will drive excessive heating and radial expansion in the mixed plasma. Simulations predicted that the sweet spot for the amount of ions lies between 1×10^5 and 5×10^5 Be^+ ions to cool a plasma of 2×10^6 positrons. To investigate this experimentally, a sequence very similar to the one described in the previous section 7.6.1 was used, except with a fixed final detuning. The number of Be^+ ions loaded was varied by changing the number of ablation shots used in the Be^+ stacking procedure (see section 6.2.2), and by longitudinally splitting the plasma using potential well manipulations. The sequence was run multiple times with the same laser cooling parameters: a fixed detuning of -129Γ was used for Be^+ preparation, and the laser was chirped from -129Γ to a final detuning of -7Γ for the final cooling of the mixture.

In each run, the positron temperature and the number of ions used were both measured using the partial temperature diagnostic outlined in section 7.4. This data has been binned by

the number of ions, and is shown in figure 7.7. The experiment was repeated with two different positron loads. Loads containing 1.4×10^6 positrons and 2.6×10^6 positrons are shown in blue and red respectively. The number of positrons was varied by changing the well depth used during positron SDREVC.

The data shows a threshold value for the number of Be^+ ions, above which the temperature of the positrons is mostly independent of the number of Be^+ ions loaded. Below this threshold, the positron temperature decreases as the number of Be^+ ions is increased. This threshold value is different for the two positron loads, but the ratio of the number of positrons to the Be^+ threshold is similar for both positron loads – a ratio of roughly 10 : 1 positrons : Be^+ .

Below the threshold value for the number of ions loaded, the positron plasma reaches a significantly higher temperature than the case where no Be^+ ions were loaded at all. This higher than expected positron temperature can potentially be attributed to increased heating caused by the presence of Be^+ ions. Be^+ plasmas tend to expand and heat faster than lighter species, as discussed in chapter 6.2.3, and so interactions between the two species can cause an increased heating and expansion rate in the positrons. When ion numbers are low enough, this increased heating rate causes the positrons to be hotter than their baseline temperature (without ions loaded), but when ion numbers are higher this heating is beaten by sympathetic cooling, and positron temperatures reduce to below their baseline temperature.

This threshold behaviour is consistent with the simulations performed by Madsen et al. [82] (figure 7.5). Although the simulations don't show explicit threshold behaviour, they do show very strongly diminishing returns in positron temperature reduction when adding more ions, which is qualitatively similar to the threshold behaviour in figure 7.7. In their calculations, with ions and positrons that were cold enough to centrifugally separate, increasing the number of Be^+ ions decreased the minimum attainable positron temperature, but with diminishing returns as the number of ions was increased. This is consistent with the additional ions mostly being added on the outside of the centrifugally-separated mixture, where their interaction with the positrons in the centre is minimal.

7.6.3 Centrifugal Separation Measurements

Centrifugal separation is a process whereby a mixed non-neutral plasma radially separates into its components by mass, with the heavier species forming a “halo” around the lighter species. The physics underlying this process was described in section 3.8.1. Importantly, centrifugal separation is a temperature-dependent process: as the temperature of both species decreases, they become more and more separated (by equation 3.24). Measuring the amount of centrifugal separation can be used to infer the temperature of the positrons, which is independent of




Figure 7.7: Figure reproduced from ALPHA collaboration [25]. Axial positron temperature in a mixed Be^+ -positron plasma as a function of the number of ions loaded. The data has been binned by ion number, and the vertical and horizontal error bars showing the standard errors of the values of positron temperature and Be^+ ion number in each bin. Loads containing 1.4×10^6 positrons (blue points) and 2.6×10^6 positrons (red points) were used. During sympathetic cooling, the cooling laser was chirped from an initial detuning of -129Γ to a final detuning of -7Γ in approximately 40s. The green line shows the positron temperature in the absence of any Be^+ ions for the 2.6×10^6 load, and the shading around the line represents the standard error.

the measurements of positron temperatures using the plasma temperature diagnostic method in sections 7.6.1 and 7.6.2.

Centrifugal separation also provides an explanation of why the minimum positron temperature attained in the previous two sections ($\sim 7\text{K}$) is significantly higher than the measured minimum temperature of the Be^+ ions ($\sim 200\text{mK}$): as the two species cool down and begin to separate, the amount of Coulomb interaction decreases, and the rate of energy transfer between the Be^+ and the positrons decreases as a result. The positrons then reach an equilibrium when the cooling power from the Be^+ ions is equal to the ambient heating from the environment. The heating rate of the positrons was experimentally measured in section 7.7.

This was investigated by preparing a cold, mixed Be^+ -positron plasma as outlined above, and then using an ekick to remove the positrons from the trap. This just leaves the Be^+ ions, which can be imaged with an MCP. The delay between removing the positrons and imaging the Be^+ ions was $5\mu\text{s}$. This experiment used 2×10^5 Be^+ ions and 1.2×10^5 positrons. This is around 10 times fewer positrons than used in the previous two sections, but a comparable amount of Be^+ ions. The number of positrons was reduced for this experiment in order to reduce the overall

size of the mixed plasma. When a “normal” load of 1.9×10^6 positrons was used, the edge of the Be^+ plasma extended beyond the edge of the MCP, making accurate measurements of the radial distribution of the Be^+ ions difficult.

To investigate how the sharpness of separation is affected by the temperature, the final laser detuning was changed in a manner similar to section 7.6.1. For a given final laser detuning, the experiment was run several times to extract the positron temperature (using the plasma temperature diagnostic method 7.4), and then separately repeated to measure the radial profiles. Extracting both the positron temperature and the radial profile in the same run would perhaps be possible, but more challenging. The feasibility of performing both measurements on a single plasma in the same run is discussed in section 7.6.6.

The results of this experiment are shown in figure 7.8. For each setting of laser detuning, the left-hand plot shows the axially-integrated number density of Be^+ ions as a function of radius from the centre of the plasma, and the right-hand side shows the corresponding image of the ions on the MCP. The density plots show the experimental data (red) as well as best-fit results from calculations using the N2DEC code (see section 7.6.4), which are used to extract temperatures from the experimental points.

The results show a strong correlation between the sharpness of centrifugal separation and the temperature of the positron plasma. As the temperature of the plasma was increased by tuning the laser frequency, the sharpness of separation decreased: panel a) ($T = 7.1 \pm 0.5\text{K}$) in figure 7.8 shows a much steeper onset in the ion density as a function of radius than the hotter plasma in panel b) ($T = 10.1 \pm 0.2\text{K}$), indicating that they were more overlapped with the positrons that were positioned in the now hollow centre. At much hotter temperatures (panel c), $T = 370 \pm 100\text{K}$, there is barely any measureable separation between the species. The experimental data shows good agreement with the simulations outlined in section 7.5. The details of how the temperature was extracted from this data is discussed in section 7.6.4.

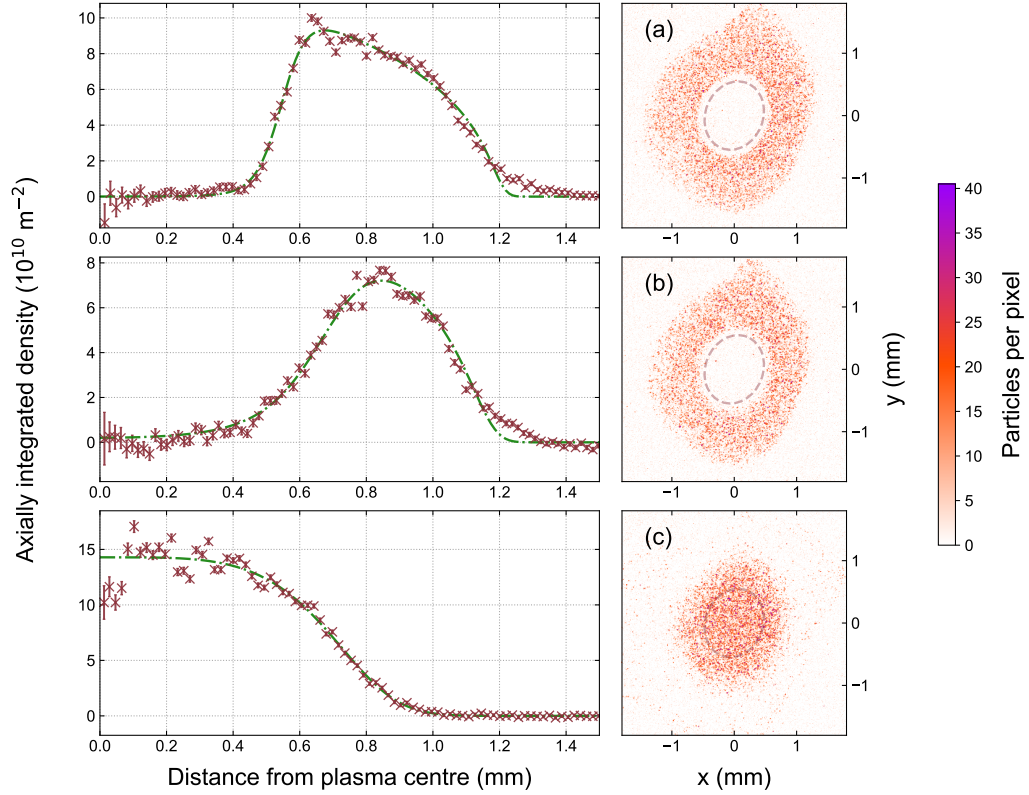


Figure 7.8: Axially-integrated radial density profiles of Be^+ from a cold, mixed Be^+ -positron plasma. $5\mu\text{s}$ before imaging the Be^+ ions, the positrons are ejected from the plasma with an ekick in the opposite direction. The left-hand plots show the axially integrated radial density profiles of the Be^+ ions, while the right-hand plots show the raw MCP image. The red points show the experimental density data, while the green dot-dashed line shows results of the best-fit distribution extracted from the N2DEC code described in section 7.5. The laser detunings used were a) -20Γ , b) -36Γ , and c) -128Γ . The measured axial positron temperatures were a) $7.1 \pm 0.5\text{K}$, b) $10.1 \pm 0.2\text{K}$, and c) $370 \pm 100\text{K}$. The fit results were a) $6.2 \pm 0.6\text{K}$, b) $19.2 \pm 2\text{K}$, and c) $253 \pm 54\text{K}$. The corresponding positron densities extracted from the N2DEC code were a, b) $6.2 \pm 0.1 \times 10^{13} \text{m}^{-3}$, and c) $1.2 \pm 0.2 \times 10^{14} \text{m}^{-3}$. The slightly elliptical nature of the MCP images is assumed to be due to distortions from electrostatic fields from the MCP, and is corrected for in the analysis of the radial profiles (see section 7.6.4).

It is critical that the Be^+ ions are imaged soon after the positrons are kicked away, as the hollow ring shape formed by the Be^+ ions is unstable if the central core of positrons is removed, and the ions will collapse inwards. The time taken to image the Be^+ ions must be significantly lower than the collapse timescale of the hollow Be^+ structure.

The collapse timescale was not studied in-depth with Be^+ ions, but it has been previously studied in ALPHA with antiproton-electron mixed plasmas [45]. This previous study was conceptually similar to the Be^+ -positron experiment described above, and started with a mixed antiproton-electron plasma that had centrifugally separated. The electrons were then removed, and the antiprotons were imaged after a variable delay to determine the timescale associated with the collapse of the hollow antiproton ring. The collapse timescale was measured to be around 5–10ms for antiprotons. The antiproton collapse proceeded via the $l = 1$ diocotron instability [86].

Although not studied in detail with Be^+ ions, the collapse timescale was briefly investigated. Figure 7.9 shows Be^+ ions imaged 1ms and $50\mu\text{s}$ after the positrons were ejected. After 1ms, the Be^+ ions appear to be part-way through collapsing, while after $50\mu\text{s}$, the hollow region is very clearly defined. This is very roughly in line with the more detailed measurements of the antiproton collapse timescale. It was therefore inferred that the $5\mu\text{s}$ delay used for the data collected in figure 7.8 was much shorter than the collapse timescale of the Be^+ ions.




Figure 7.9: MCP images of centrifugally-separated Be^+ ions after the positrons have been removed. There were 2×10^5 Be^+ ions and 1.9×10^6 positrons. The delay between the positron removal and imaging the ions was: Left: $50\mu\text{s}$, Right: 1ms. After the $50\mu\text{s}$ delay, there are no clear signs that the hollow plasma is collapsing. After the 1ms delay, the hollow Be^+ structure appears to be part-way through collapsing.

7.6.4 Extracting Temperatures from Radial Profiles

In order to extract a temperature from the radial profiles in figure 7.8, calculations were performed using the N2DEC code described in section 7.5. The resulting calculations were then fitted to the experimental data to extract estimates of experimental values of the calculation's input parameters: the number and density of both species, and the temperature of the mixture.

The shape of the Be⁺ ions on the MCP is visibly elliptically distorted, which is thought to be due to stray potentials on the MCP housing (discussed in chapter 6.5.5). The distortion can be seen most easily in figure 7.8 a), and was corrected for in the analysis that was used to create the radial profiles shown in figure 7.8. In order to make this correction, an ellipse was fitted to the edge of the central hollow region in one of the coldest and most well-separated MCP images. The whole image then underwent a transformation which transformed the ellipse into a circle, and it is this transformed image that was then used to extract the radial profile. The hollow region was not re-fitted for each individual image, and instead the ellipse that was fitted to the image in figure 7.8 a) was used for the transformation of both b) and c). This was done as it becomes harder and harder to get a meaningful fit of the central hollow region as the temperature increases and the separation becomes more poorly defined, and because the ellipticity seemed to be very reproducible between images. This elliptical transformation therefore assumes that the distortion happens after the ions have been ejected and are travelling towards the MCP, and assumes that the distortion is the same from run-to-run.

There is a second distortion effect that is visible on the images at high radius, particularly on images 7.8 b) and c), which gives the plasma an asymmetrical and slightly avocado-like shape. This is consistent with the edge distortion effects from potentials on the MCP housing. This distortion was not accounted for in the analysis, and so it should be understood that these edge-distortion effects cause some artifacts in the radial profiles. By eye, these edge effects distort the plasma at radii of above around 1mm, and so data points above 1mm in figure 7.8 should be ignored.

7.6.5 Final Remarks About Positron Temperatures

The experimental measurements of positron temperatures presented in this chapter used two techniques: the plasma temperature measurement in sections 7.6.1, 7.6.2, and the centrifugal separation measurements in section 7.6.3. The good agreement between these two experimental measurement techniques and the simulations discussed in section 7.5 is generally reassuring, and in particular allays some concerns about systematic uncertainties with the plasma temperature measurements that were laid out in chapter 3.9.2.

One question that may remain is whether these measurements are of the parallel (T_{\parallel}) or per-

pendicular (T_{\perp}) temperatures. While the plasma temperature measurements and the centrifugal separation measurements are both measurements of T_{\parallel} , the long cooling times of several seconds that are used in the sympathetic cooling sequences ensure that there is ample time for T_{\parallel} and T_{\perp} to thermalise, even when the positrons are cold and strongly magnetised, i.e. when collision rates between the parallel and perpendicular degrees of freedom are low. Parallel-to-perpendicular collision rates can be calculated with equation 3.12, and a typical positron plasma with $T = 6\text{K}$, $n = 1.3 \times 10^{14}\text{m}^{-3}$ would have a collision rate of $\nu_{\perp\parallel} = 0.9\text{kHz}$ in a 1T magnetic field. In this parameter regime, T_{\parallel} and T_{\perp} will start to decouple when plasma temperatures drop below around 2 – 3K.

7.6.6 Improving Temperature Measurements

Performing both the plasma temperature diagnostic and the radial profile measurement on the same plasma is desirable, as this allows two independent methods of measuring the temperature to be performed back-to-back on the same plasma. This is a useful sanity check that both temperature measurements are producing reasonable temperatures, and allows for cross-calibration between the two methods.

To achieve this, a temperature measurement of the positrons could be performed to one MCP using the partial ejection method described in section 7.4, before the rest of the positrons were removed with an ekick. The remaining ions could then be quickly ejected to a second MCP at the other end of the trap to obtain the radial profile. The main problem that would come up when performing this kind of double temperature measurement would be that the timescale of a plasma temperature diagnostic measurement is usually around 10ms, which is comparable to the collapse timescale of the Be^+ ions. The relatively slow speed of the plasma temperature diagnostic is dependent on the limited speed at which the control hardware can change the potentials applied to the trap electrodes. With fast enough hardware, the lower limit of the speed at which a plasma temperature measurement can be performed is when an individual particle’s transit time through the plasma could cause a delay in the time it takes to escape [29]. Transit times are typically on the order of $1\mu\text{s}$ in ALPHA. In addition, this would require two MCPs – one at either end of the trap – to perform both measurements back-to-back. Although currently not set up, this would be a minor hardware upgrade.

7.7 Mixture Heating Rate

The heating rate of the mixed plasma after laser-cooling was measured by preparing a mixed plasma in the manner described above 7.6.2, turning the cooling laser off, waiting for a variable amount of time, and then performing a temperature measurement on the positrons. Figure 7.10 shows the results of this experiment, as well as a similar experiment where the laser was left on for the duration of the variable wait. The laser-off results (red points) show an initial heating rate of around 25K/s, which slows as the positrons reach an equilibrium temperature of around 125K. The laser-on results (blue points) show that if the laser is left on, the positrons can be maintained at their minimum temperature for > 20 s. The increase in temperature seen in the blue points after around 35s is possibly due to radial expansion eventually reducing the overlap between Be^+ ions and the cooling laser, or due to the Be^+ and positrons decoupling for similar reasons.




Figure 7.10: Time evolution of the positron temperature after laser cooling the Be^+ -positron mixture. The mixture contained 1.4×10^6 positrons and $4.7 \pm 0.1 \times 10^5$ Be^+ ions, and was cooled by chirping the laser from -110Γ to a final detuning of -7Γ over approximately 40s. Red points show the evolution when the laser was turned off after cooling, which resulted in the positrons heating at a rate of around 25K s^{-1} initially, before stabilising at around 125K. The blue points show experiments where the cooling laser was left on after the laser cooling sweep, at a fixed detuning of -7Γ . The temperature of the positrons remained steady at a value of $6.7 \pm 0.3\text{K}$ for over 20s in this case.

7.8 Mixture Expansion Rate

The radial expansion rate of the mixed plasma was measured by preparing a mixed plasma in the manner described in section 7.6.2 with 4×10^5 Be^+ ions and 2.6×10^6 positrons, compressing it, and then holding the plasma for a variable amount of time in a 1 T magnetic field. After the wait, the plasma was imaged with an MCP. This process was then repeated with a plasma containing only 4×10^5 Be^+ ions, and repeated again with a plasma containing only 2.6×10^6 positrons. The results of these experiments can be seen in figure 7.11. The positron-only plasma is stable in size over timescales of > 1 minute, but the mixed plasma and Be^+ -only plasma were found to expand much more rapidly, doubling in radius after approximately 25 seconds. When measuring the mixture radius, both species are imaged on the MCP, but the image is completely dominated by positrons due to their greater number and larger response per-particle. This means that the mixture radius is really a measurement of the positron radius in the mixture. It is quite clear that the expanding Be^+ ions are still able to expand unimpeded in the mixture, and are able to drive the positrons to expand at the same rate.




Figure 7.11: Time evolution of the radius of three different plasmas without laser cooling in a 1 T magnetic field. The red points show the evolution of a Be^+ -positron mixture, the blue points show the evolution of Be^+ ions only, and the green points show the evolution of positrons only.

7.9 Summary

The results presented in this chapter showed positron plasmas containing up to 2.6M positrons that were cooled to around 7K. This marks the first time that such a large number of positrons

has been cooled to such a low temperature. Since antihydrogen formation in ALPHA shows such a strong dependence on positron temperature, these sympathetically-cooled positron plasmas have the potential to provide serious increases in antihydrogen trapping rates in the future.

There is still some work to be done in incorporating this sympathetic cooling sequence into a full antihydrogen production sequence. The sympathetic cooling sequence developed in this chapter is currently too slow to be neatly slotted into current antihydrogen formation sequences, and the inability to perform simultaneous laser cooling and RW compression on Be^+ ions, along with the Be^+ ions' tendency to expand when not actively compressed, has led to sympathetically-cooled positron plasmas that are radially larger than the plasmas that are typically used for antihydrogen formation. The following chapter will discuss some of the future work that could go into introducing sympathetically-cooled positrons into an antihydrogen production sequence, as well as other possible upgrades to the beryllium experiment and some concluding remarks.

8/ Future Work and Conclusions

This chapter will discuss the conclusions of this thesis, and also outline some future work that could improve the beryllium experiment in ALPHA. A large part of the future work will focus on how the sympathetic cooling sequence could be integrated into the antihydrogen production sequence, but there will also be some discussion about potential hardware and software upgrades in the beryllium setup, and the possibility of using beryllium ions for magnetometry.

8.1 Towards Beryllium-Antihydrogen Compatibility

The previous chapter (7) discussed the development of a sequence to sympathetically cool positrons, and presented some of the results that came out of this sequence. This section will discuss the practicalities of integrating this sympathetic cooling sequence with the antihydrogen production sequence described in chapter 4. When thinking about integrating the sympathetic cooling sequence into an antihydrogen production sequence, it is best to consider the constraints imposed by the antihydrogen production sequence on the sympathetic cooling sequence. As mentioned in the previous chapter, the constraints are as follows:

- The number of positrons used should be between 1×10^6 and 3×10^6 ,
- The number of Be^+ ions used should be minimised,
- The sympathetically-cooled positrons should be produced in under 4 minutes,
- Plasma sizes should be kept small enough that the neutral trap does not produce any significant heating or expansion.

Each of these constraints will be discussed separately, with particular focus on how close the sympathetic cooling sequence is to fitting within these constraints, and how it may be altered to fit within them if necessary.

8.1.1 Number of Positrons

The first constraint describes the number of positrons that are currently used in antihydrogen formation sequences. The sympathetic cooling sequence discussed in the previous chapter was able to produce cold positron plasmas containing 0.5×10^6 , 1.4×10^6 , and 2.6×10^6 positrons. This demonstrated that the sympathetic cooling process can work over a large range of values of positron number, and that enough positrons can be sympathetically cooled to be used in antihydrogen formation.

8.1.2 Number of Be^+ ions

The second constraint arises because having too many Be^+ ions could potentially lead to antiprotons being captured by the Be^+ ions during antihydrogen formation, leading to antiproton loss from the trap. Although calculations indicate that this may not be a significant antiproton loss mechanism during a merge with positrons [82], the experiments performed in the previous chapter 7.6.2 show that once enough Be^+ ions have been added to the mixed Be^+ -positron plasma, adding more Be^+ ions does not measurably reduce the temperature of the positrons. Minimising the number of ions also generally helps keep the mixed plasma radially small, which is an important point that is discussed in section 8.2 below. A good target for the number of Be^+ ions used in antihydrogen formation is therefore just above the threshold discussed in chapter 7.6.2, which is roughly a factor of 10 fewer ions than positrons.

8.1.3 Cutting Out the Arrested Rotating Wall

The sympathetic cooling scheme outlined in the previous chapter is too slow to be integrated into current antihydrogen production sequences, and so alternative, faster preparation schemes should be considered.

The current sympathetic cooling sequence includes a lengthy arrested RW section, where the ions are compressed and cooled before being merged with a positron plasma, as in figure 7.1. The mixed plasma is then sympathetically compressed. If the uncompressed, hot ions could instead be immediately merged with a positron plasma, and then simply cooled and compressed as a mixture, this could cut out a large chunk of the preparation time.

Some more recent experiments were performed (without using the cooling laser) where the ions were loaded into the trap and immediately merged with a pre-prepared positron plasma. The ions could then be sympathetically cooled by the positrons¹, and sympathetically com-

¹This is essentially the reverse of the final result we want to achieve – here, positrons are being used to sympathetically cool Be^+ rather than the other way around. A true symbiotic relationship.




Figure 8.1: A sympathetically compressed Be⁺-positron plasma is shown on the left (blue). The image is dominated by positrons. In the right-hand image, the positrons have been removed after compression, leaving only the sympathetically compressed ions (orange). Sizes given are the calculated sizes in a 3T field. The plasma contained 10^5 Be⁺ and 10^6 positrons, and was driven with a 400kHz, 5V RW for 30s.

pressed with the positrons. This scheme is very similar to the preparation scheme currently used for antiprotons (discussed in chapter 4.1), except with a mixed Be⁺-positron plasma rather than an antiproton-electron plasma.

This much shorter preparation scheme has shown some promising results, and is likely the way forward in future sympathetic cooling experiments in ALPHA. A comparison of the mixed plasmas produced by the arrested RW sequence and this new, shorter sequence are summarised in table 8.1, and figure 8.1 shows a sympathetically compressed Be⁺-positron plasma. The main difference is the much hotter plasma produced by the shorter sequence, but even this hotter temperature should be able to be easily laser cooled.

	Number of e+	Number of Be ⁺	Radius (mm)	Temperature (K)
Arrested RW	0.1M - 2.6 M	100k - 500 k	~ 0.4mm	350K
Short Sequence	1M - 3M	~ 220k	~ 0.4mm	3000K

Table 8.1: Plasma parameters of a typical mixed plasma prepared with the arrested RW sequence, and with the shorter sympathetic compression and cooling sequence.

8.1.4 Further Reductions in Positron Temperature

The previous chapter demonstrated positron plasmas that had been sympathetically cooled to around 7K, and there is a good chance that further optimisations of the process could produce even colder positron plasmas. Adjusting the density of the positron plasma could reduce temperatures, as shown by simulations in chapter 7.5, but this may not be an easy parameter to optimise as the antihydrogen formation process is also dependent on positron density in a non-monotonic fashion (see chapter 4.6.4). The density of the Be^+ plasma could also prove to be a useful parameter to tune to optimise sympathetic cooling. The density of the Be^+ plasma was not tightly controlled in the experiments in the previous chapter as the plasma was first compressed and then cooled, allowing the plasma to expand during the cooling process. Section 8.4 discusses an upgrade to the experiment that could potentially allow laser cooling during RW compression of the plasma, which could allow for much tighter control of the density of the Be^+ during the laser cooling process. Increasing the power of the cooling laser could also lead to a reduction in positron temperature, as well as potentially reducing the time taken to laser cool the Be^+ ions at the start of the sequence. Finally, reducing external plasma heating sources, such as patch potentials, thermal radiation entering the trap, and other mechanisms discussed in chapter 3.4 has always been a focus in ALPHA, and should continue to be a focus in order to achieve the lowest possible positron temperatures going forward.

8.2 The Neutral Trap and Large Plasmas

As discussed in section 3.4, energising the neutral trap distorts the solenoidal magnetic field of the Penning-Malmberg trap. This causes magnetic field lines to guide particles into the trap walls above a critical radius, and can cause rapid plasma heating and expansion even in plasmas that are below this critical radius. As the current antihydrogen production sequence utilises antihydrogen stacking, it requires the neutral trap to be energised for the duration of the sequence. Keeping plasmas as small as possible is therefore very necessary to avoid rapid particle expansion and particle loss.

8.2.1 Sympathetic Cooling in the Neutral Trap

Some investigations into the effects of the neutral trap on the sympathetic cooling sequence were undertaken. In one experiment, 1.9×10^6 positrons and $\sim 5 \times 10^5$ Be^+ ions were prepared in a manner similar to the methods in the previous chapter 7, except the neutral trap's octupole magnet was ramped up to a current setpoint during the final mixture recompression stage. The octupole being ramped up does not affect the recompression, which occurs in the

positron catching region of the trap, but does affect the final cooling stage, which occurs in the central region of the trap². During the final laser cooling stage, the laser was chirped from an initial frequency of -129Γ to a final frequency of -51Γ . The temperature of the positrons was then measured as a function of octupole current, which is shown in figure 8.2. As the octupole current is increased, there is a drastic increase in the temperature of the positrons, indicating that the octupole’s field is causing heating of the mixed plasma. The positron temperature at zero octupole current (around 17K) is significantly hotter than positron temperatures achieved in the previous chapter (around 7K) because the laser was tuned much further from the resonance.




Figure 8.2: Figure created by Dan Maxwell [72]. Positron temperature as a function of octupole current. The mixture of 1.9×10^6 positrons and $\sim 5 \times 10^5$ Be^+ ions was prepared in a manner similar to the previous chapter. The mixture was then recompressed with the RW while the octupole was being ramped to the current setpoint. The mixed plasma was then shuffled back to the central region of the trap, which is influenced by the octupole magnet, to be laser cooled.

This result tells us that in order to be able to effectively sympathetically cool inside the octupole’s magnetic field, plasma radii need to be decreased further. To test this, a similar experiment was run with fewer positrons ($0.5M$) and fewer Be^+ ions (1×10^5), again with the octupole being ramped up during the mixture recompression stage. In this series of runs, the octupole was ramped to its nominal operating current of 900A. By reducing the number of Be^+ ions, the expansion rate of the plasma is decreased, and reducing the number of both species generally leads to a smaller plasma, which can potentially avoid becoming too large to be heated by the octupole’s fields. This set of experiments was brief, but was able to produce positrons with a

²Recall from chapter 4 that the neutral trap only covers the central region of ALPHA-2, and not the positron or antiproton catching regions of the trap.

temperature of $12.8 \pm 0.9\text{K}$ consistently over a series of 4 runs. This gives credence to the idea that sympathetically cooling a larger amount of positrons would be possible inside the octupole's magnetic field if the plasma radii can be kept small.

8.3 Improved Fluorescence Diagnostics

As discussed in earlier chapters, the current fluorescence detection hardware installed in ALPHA is severely limited. Since fluorescence detection is the de-facto method for imaging trapped, laser-cooled ions, it is important that improvements to the fluorescence diagnostics system are considered.

There are three main limitations of the current fluorescence detection setup: the small solid angle from the ions to the detector, the high level of background caused by light scattered inside the trap from the laser, and the limited sweep range of the AOM. The small solid angle and high background are both consequences of the trap geometry. The current fluorescence detection system uses a Silicon Photomultiplier (SiPM) which is positioned on one of the off-axis laser path windows, outside of the trap, as shown in figure 8.3. The large distance between the SiPM and the centre of the trap (1.6m), means that the solid angle subtended from the centre of the trap by the detector is around 2×10^{-6} steradians.

When the cooling laser is on, some photons from the main cooling beam are scattered by reflective surfaces inside the trap. Since the SiPM is very close to the axis of the laser, a large amount of these scattered photons are picked up by the SiPM, causing a large background signal on the SiPM.




Figure 8.3: The layout of the ALPHA-2 Penning-Malmberg trap, and the current position of the SiPM used for detecting fluorescence photons.

The first two problems could potentially be solved by installing detectors inside the trap

vacuum chamber in a position much closer to where the ions are trapped. One place that such detectors could be placed is at the radial step in the electrode stack (see figure 8.3). Since there is already a gap to let light out of the trap where the step occurs, this would avoid modification of the trap electrodes. This would place the detectors around 8 degrees off-axis, and around 250mm from the centre of the trap. The solid angle achievable with six $3\text{mm} \times 3\text{mm}$ detectors mounted at the radial step would be around $1 \times 10^{-3}\text{sr}$. Currently, the photodetectors that would be used in this setup are Hamamatsu SiPM S13370-3050CN.

Other similar experiments often used modified electrodes with gaps that allow fluorescence photons to escape the trap, where they can be detected. A popular choice is to use a segmented RW electrode with gaps between the segments, with photodetectors mounted on the outside of the electrodes. Other designs in the same vein could also work, such as drilling a small hole in an electrode to allow fluorescence light to escape to a detector. Such a hole could potentially be covered by a transparent, conductive material. Any kind of hole or gap would produce effective patch potentials on the electrodes, which would lead to plasma heating. As such, these kinds of solutions are probably not desirable in ALPHA, as minimising plasma temperatures for antihydrogen formation is very important. A second issue arises in ALPHA in particular, as the central electrodes are designed to be as thin as possible in the radial direction, to maximise the depth of the neutral trap. This makes the electrodes mechanically difficult to work with, and previous attempts at creating an azimuthally-segmented version of one of these thin electrodes, in order to use it as a rotating wall, have gone poorly.

The limited scan width of the acousto-optic modulator (AOM) is a limiting factor when attempting to use fluorescence signals to diagnose the temperature of the Be^+ ions. A temperature measurement is performed by using an AOM to quickly sweep the laser's frequency across the transition's resonance, and monitoring how the fluorescence signal evolves as a function of time. This allows a measurement of the lineshape of the transition, and the width of this transition can be used to estimate the temperature of the ions (section 6.5.2). The current AOM has a scan width of 80MHz ($\sim 4\Gamma$), which only allows measurement of temperatures below around 500mK. Ions with temperatures much higher than this will have much broader lineshapes, and producing a meaningful fit with the limited scan range of 80MHz will become impossible. Upgrading the hardware to have a larger scan range would help this greatly.

8.4 On-axis Cooling Laser

One major flaw in the current Be^+ laser cooling setup is the inability to simultaneously compress the ions with a RW and laser cool them. This is because the laser must enter the trap along one of the off-axis laser paths, which intersects the trap axis at the centre of the trap. This allows

for good laser overlap with trapped species in the centre of the trap only. The RW electrodes are not in the centre of the trap, and are instead close to either end of the trap, where there is effectively zero overlap between the trapped ions and the laser. The lack of ability to simultaneously compress and cool the Be^+ ions is one of the reasons that the sympathetic cooling sequence takes such a long time, and is the reason that the arrested rotating wall method was developed in chapter 7.2.1. Being able to simultaneously compress and cool the mixed plasma would also alleviate some of the radial expansion issues discussed in section 7.8: if the mixture is not being actively compressed, it tends to radially expand. Being able to compress and cool the mixture at the same time could potentially allow for a steady-state sympathetically cooled plasma. This would be much easier to incorporate into an antihydrogen trapping sequence, as the steady-state cooled plasma could simply be held for as long as desired before being used for antihydrogen formation.

A simple solution to this problem is to have a laser aligned along the axis of the trap, which would provide good laser overlap for the whole trap. This has already been implemented in the trap, although this was designed and installed after the majority of the work in the previous two chapters. This on-axis laser setup utilises two 45° mirrors mounted on the linear translators (“sticks”) at either end of the trap. The laser is aligned to enter the trap perpendicular to the trap axis, where it is reflected by the first mirror to pass straight along the trap axis. The second mirror guides the beam back out of the trap.

The main limitation of this setup is that it requires both sticks to be in the laser mirror position, which prevents any of the stick-based diagnostics (such as the MCP) from being used.

8.5 Fibre Delivery of the Cooling Beam

The cooling laser beam needs to pass through around 10 metres of air to get from the laser lab to the experiment. This is difficult to align, and prone to drifting in space, requiring daily re-adjustments of the laser alignment. One solution to this problem is the use of an optical fibre to carry laser light from the laser lab to the experiment. This would then only require careful alignment of the in-coupling optics in the laser lab, but would allow the out-coupling optics to be placed freely, much closer to the experiment. In this setup, the alignment of the laser into the fibre is decoupled from its alignment in the trap. The positioning of the laser through the trap could then be altered just by adjusting the out-coupling optics, which would not affect, or be affected by, the alignment of the laser in the laser lab before it reaches the in-coupling optics and the fibre.

Fibres for this kind of application are generally not commercially available. Fibres exposed to UV light will undergo a phenomenon known as UV solarisation, which causes UV fibres to

gradually increase in opacity for a certain wavelength as it transmits light of that wavelength. The physics causing this increase in opacity is colour-centre formation in the silica of the fibre core [87]. Recent work by Marciniak et al. [88] outlines a procedure for producing a solarisation-resistant optical fibre for UV, which involves high-pressure hydrogen loading and subsequent UV curing of a photonic crystal fibre. Some of these treated fibres have recently been acquired for the beryllium experiment at ALPHA, and will soon be installed in the experiment.

8.6 Beryllium for Magnetometry

Magnetic fields inside the ALPHA Penning-Malmberg trap are measured *in situ* by measuring the cyclotron frequency of an electron plasma within the trap. The frequency is measured by exciting the cyclotron mode with microwaves, and measuring the temperature of the plasma as a function of the microwave frequency. When the microwave frequency matches the cyclotron frequency of the electrons, the plasma is resonantly heated. The temperature of the plasma can be measured by measuring the amplitude of the quadrupole mode of the plasma [89] or, more recently, by using many small, reproducible plasmas that are destructively measured.

When using plasmas for temperature diagnostics, plasma dynamics and modes can affect the measurement, and make it difficult to get an accurate measurement of the true magnetic field in the trap. It is therefore perhaps desirable to use Be^+ ions for magnetometry, since we already have a setup for laser-cooling Be^+ ions in the trap. With an improved fluorescence detection system (section 8.3), a small sample of ions could be used to measure the magnetic field, which are not subject to plasma dynamics.

8.7 Conclusions

The matter-antimatter asymmetry, or the question of why our matter-dominated universe can even exist, is one of the largest unsolved questions in modern physics. Performing direct measurements on antimatter, and comparing them to measurements of their matter counterparts, is currently one of the best routes available to probing fundamental symmetries that may provide the answer to this question. The ALPHA collaboration has recently demonstrated the first ever direct laser cooling of anti-atoms [24], and this, combined with recent upgrades to the metrology system, means that the prospect of even more precise measurements of the 1S-2S transition in antihydrogen are on the horizon. This is happening in tandem with the recently-completed ALPHA-g, an addition to the experimental apparatus that is designed to measure the effect of gravity on antimatter. With this ambitious physics program ahead, increasing the trapping rate

of antihydrogen is especially important as it will maximise the physics output from the limited amount of antiprotons produced by CERN’s antiproton decelerator complex.

The main aim of this thesis was to develop a sequence to sympathetically cool positrons to temperatures significantly lower than those routinely used in ALPHA. These sympathetically-cooled positrons could be used in antihydrogen formation to increase the trapping rate of antihydrogen in ALPHA, therefore increasing the amount of antihydrogen that is available for experimentation.

Chapters 6 and 7 were the main experimental result chapters, and described the work that went into first laser cooling Be^+ ions in ALPHA, and then using these Be^+ ions to sympathetically cool positron plasmas. The results presented in chapter 7 showed the production of sympathetically cooled positron plasmas that are well-suited for antihydrogen production. Positron plasmas with temperatures as low as 7K were produced, compared to temperatures of around 17 – 30K in previous work. Increasing trapping rates of antihydrogen has been central to ALPHA’s work since the start: in 2010, ALPHA trapped only around 200 – 300 antihydrogen atoms in the entire year, compared to over 1000 in a single day in recent experiments. These cold positron plasmas may be key to further increases in antihydrogen trapping rates, and hence contribute to answering the question of why our universe can exist at all.

8.7.1 Final Remarks

Although it initially seemed like a straightforward goal, the production of sympathetically-cooled positron plasmas in ALPHA was a very time-consuming task, and was only made possible by the tireless efforts of the rest of the Beryllium Team, and the rest of the group at ALPHA. I really hope that some time can be set aside in future antiproton runs to integrate sympathetic cooling into full antihydrogen production runs, and that large gains in the antihydrogen trapping rate are seen as a result.

Bibliography

- [1] Dirac, P. A. M. The quantum theory of the electron. *Proc. R. Soc. Lond. A*, **117**:610–624 (1928). doi:10.1098/rspa.1928.0023.
- [2] Chamberlain, O., Segrè, E., Wiegand, C., & Ypsilantis, T. Observation of Antiprotons. *Phys. Rev.*, **100**:947–950 (1955). doi:10.1103/PhysRev.100.947.
- [3] Cork, B., Lambertson, G. R., Piccioni, O., & Wenzel, W. A. Antineutrons Produced from Antiprotons in Charge-Exchange Collisions. *Phys. Rev.*, **104**:1193–1197 (1956). doi:10.1103/PhysRev.104.1193.
- [4] Garwin, R. L., Lederman, L. M., & Weinrich, M. Observations of the Failure of Conservation of Parity and Charge Conjugation in Meson Decays: the Magnetic Moment of the Free Muon. *Phys. Rev.*, **105**:1415–1417 (1957). doi:10.1103/PhysRev.105.1415.
- [5] Landau, L. On the conservation laws for weak interactions. *Nucl. Phys.*, **3**:127–131 (1957). doi:10.1016/0029-5582(57)90061-5.
- [6] Charlton, M., Eriksson, S., & Shore, G. M. Testing Fundamental Physics in Antihydrogen Experiments (2020). arXiv:2002.09348.
- [7] Andresen, G. B. *et al.* Evaporative Cooling of Antiprotons to Cryogenic Temperatures. *Phys. Rev. Lett.*, **105**(1) (2010). doi:10.1103/physrevlett.105.013003.
- [8] Gabrielse, G. *et al.* Adiabatic Cooling of Antiprotons. *Phys. Rev. Lett.*, **106**:073002 (2011). doi:10.1103/PhysRevLett.106.073002.
- [9] Hori, M. & Walz, J. Physics at CERN’s Antiproton Decelerator. *Progress in Particle and Nuclear Physics*, **72**:206–253 (2013). doi:10.1016/j.pnnp.2013.02.004.
- [10] Ahmadi, M. *et al.* Characterization of the 1S-2S transition in antihydrogen. *Nature*, **557**(7703):71–75 (2018). doi:10.1038/s41586-018-0017-2.

- [11] Storry, C. ATRAP antihydrogen experiments and update. *Hyperfine Interactions*, **172**:91–96 (2006). doi:10.1007/s10751-007-9528-1.
- [12] Juhász, B. & Widmann, E. Planned measurement of the ground-state hyperfine splitting of antihydrogen. *Hyperfine Interactions*, **193**(1):305 (2009). doi:10.1007/s10751-009-0016-7.
- [13] Amole, C. *et al.* Description and first application of a new technique to measure the gravitational mass of antihydrogen. *Nature Communications*, **4**(1):1785 (2013). doi:10.1038/ncomms2787.
- [14] Kellerbauer, A. *et al.* The AEGIS experiment at CERN. *Hyperfine Interactions*, **209** (2012). doi:10.1007/s10751-012-0583-x.
- [15] Pérez, P. *et al.* The GBAR antimatter gravity experiment. *Hyperfine Interactions*, **233**(1):21–27 (2015). doi:10.1007/s10751-015-1154-8.
- [16] Ulmer, S. *et al.* Challenging the standard model by high-precision comparisons of the fundamental properties of protons and antiprotons. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, **376**(2116) (2018). doi:10.1098/rsta.2017.0275.
- [17] Wineland, D. J., Drullinger, R. E., & Walls, F. L. Radiation-Pressure Cooling of Bound Resonant Absorbers. *Phys. Rev. Lett.*, **40**:1639–1642 (1978). doi:10.1103/PhysRevLett.40.1639.
- [18] Shuman, E. S., Barry, J. F., & DeMille, D. Laser cooling of a diatomic molecule. *Nature*, **467**(7317):820–823 (2010). doi:10.1038/nature09443.
- [19] Larson, D. J. *et al.* Sympathetic cooling of trapped ions: A laser-cooled two-species non-neutral ion plasma. *Phys. Rev. Lett.*, **57**:70–73 (1986). doi:10.1103/PhysRevLett.57.70.
- [20] Ahmadi, M. *et al.* Antihydrogen accumulation for fundamental symmetry tests. *Nature Communications*, **8**(1):681 (2017). doi:10.1038/s41467-017-00760-9.
- [21] Jelenković, B. M. *et al.* Sympathetically cooled and compressed positron plasma. *Phys. Rev. A*, **67**:063406 (2003). doi:10.1103/PhysRevA.67.063406.
- [22] Ahmadi, M. *et al.* Observation of the 1S-2P Lyman- α transition in antihydrogen. *Nature*, **561**(7722):211–215 (2018). doi:10.1038/s41586-018-0435-1.

- [23] Ahmadi, M. *et al.* Investigation of the fine structure of antihydrogen. *Nature*, **578(7795)**:375–380 (2020). doi:10.1038/s41586-020-2006-5.
- [24] Baker, C. J. *et al.* Laser cooling of antihydrogen atoms. *Nature*, **592(7852)**:35–42 (2021). doi:10.1038/s41586-021-03289-6.
- [25] Baker, C. J. *et al.* Sympathetic cooling of positrons to cryogenic temperatures for antihydrogen production. *Nature Communications*, **12(1)**:6139 (2021). doi:10.1038/s41467-021-26086-1.
- [26] Johnson, M. A. *Design and Commissioning of Beamlines for the ALPHA Antihydrogen Experiment*. Ph.D. thesis (2019).
- [27] Tranquille, G. *et al.* Commissioning the ELENA Beam Diagnostics Systems at CERN. In *9th International Particle Accelerator Conference* (2018). doi:10.18429/JACoW-IPAC2018-WEPAF084.
- [28] Jørgensen, L. *et al.* A Positron Accumulator for Antihydrogen Synthesis. *Materials Science Forum*, **363**:634–636 (2001). doi:10.4028/www.scientific.net/MSF.363-365.634.
- [29] Hunter, E. D. *et al.* Plasma temperature measurement with a silicon photomultiplier (SiPM). *Review of Scientific Instruments*, **91(10)**:103502 (2020). doi:10.1063/5.0006672.
- [30] Bendiscioli, G. & Kharzeev, D. Antinucleon-nucleon and antinucleon-nucleus interaction. A review of experimental data. *La Rivista del Nuovo Cimento (1978-1999)*, **17(6)**:1 (2007). doi:10.1007/BF02724447.
- [31] Andresen, G. *et al.* Antihydrogen annihilation reconstruction with the ALPHA silicon detector. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **684**:73–81 (2012). doi:10.1016/j.nima.2012.04.082.
- [32] Bertsche, W. A., Butler, E., Charlton, M., & Madsen, N. Physics with antihydrogen. *Journal of Physics B: Atomic, Molecular and Optical Physics*, **48(23)**:232001 (2015).
- [33] Knoop, M., Madsen, N., & Thompson, R. C. *Trapped Charged Particles* (2016). ISBN 978-1-78634-011-5.
- [34] Glinsky, M. E. *et al.* Collisional equipartition rate for a magnetized pure electron plasma. *Physics of Fluids B: Plasma Physics*, **4(5)**:1156–1166 (1992). doi:10.1063/1.860124.

- [35] Baker, C. J. *et al.* The Effects of Patch Potentials in Penning-Malmberg Traps. Unpublished.
- [36] Butler, E. *Antihydrogen formation, dynamics and trapping*. Ph.D. thesis (2011).
- [37] Huang, X. *et al.* Steady-State Confinement of Non-neutral Plasmas by Rotating Electric Fields. *Phys. Rev. Lett.*, **78**:875–878 (1997). doi:10.1103/PhysRevLett.78.875.
- [38] Anderegg, F., Hollmann, E. M., & Driscoll, C. F. Rotating Field Confinement of Pure Electron Plasmas Using Trivelpiece-Gould Modes. *Phys. Rev. Lett.*, **81**:4875–4878 (1998). doi:10.1103/PhysRevLett.81.4875.
- [39] Danielson, J. R. & Surko, C. M. Torque-Balanced High-Density Steady States of Single-Component Plasmas. *Phys. Rev. Lett.*, **94**:035001 (2005). doi:10.1103/PhysRevLett.94.035001.
- [40] Andresen, G. B. *et al.* Compression of Antiproton Clouds for Antihydrogen Trapping. *Phys. Rev. Lett.*, **100**:203401 (2008). doi:10.1103/PhysRevLett.100.203401.
- [41] Anderson, M. H. *et al.* Observation of Bose-Einstein Condensation in a Dilute Atomic Vapor. *Science*, **269**(5221):198–201 (1995). doi:10.1126/science.269.5221.198.
- [42] Davis, K. B. *et al.* Bose-Einstein Condensation in a Gas of Sodium Atoms. *Phys. Rev. Lett.*, **75**:3969–3973 (1995). doi:10.1103/PhysRevLett.75.3969.
- [43] Ahmadi, M. *et al.* Enhanced Control and Reproducibility of Non-Neutral Plasmas. *Phys. Rev. Lett.*, **120**:025001 (2018). doi:10.1103/PhysRevLett.120.025001.
- [44] Carruth, C. *Methods for plasma stabilization and control to improve antihydrogen production*. Ph.D. thesis (2018).
- [45] Andresen, G. B. *et al.* Centrifugal Separation and Equilibration Dynamics in an Electron-Antiproton Plasma. *Phys. Rev. Lett.*, **106**:145001 (2011). doi:10.1103/PhysRevLett.106.145001.
- [46] Evans, L. T. *Phenomenology of creation of antihydrogen and measurement of antihydrogen properties*. Ph.D. thesis (2016).
- [47] Eggleston, D. L. *et al.* Parallel energy analyzer for pure electron plasma devices. *Physics of Fluids B: Plasma Physics*, **4**(10):3432–3439 (1992). doi:10.1063/1.860399.

- [48] LHC Project Schedule. <https://project-hl-lhc-industry.web.cern.ch/content/project-schedule>. Accessed 2021-08-16.
- [49] Andresen, G. B. *et al.* Trapped antihydrogen. *Nature*, **468(7324)**:673–676 (2010). doi:10.1038/nature09610.
- [50] Jonsell, S. & Charlton, M. On the formation of trappable antihydrogen. *New Journal of Physics*, **20** (2018). doi:10.1088/1367-2630/aabc71.
- [51] Robicheaux, F. Simulations of antihydrogen formation. *Phys. Rev. A*, **70**:022510 (2004). doi:10.1103/PhysRevA.70.022510.
- [52] Jonsell, S., van der Werf, D. P., Charlton, M., & Robicheaux, F. Simulation of the formation of antihydrogen in a nested Penning trap: effect of positron density. *Journal of Physics B: Atomic, Molecular and Optical Physics*, **42(21)**:215002 (2009). doi:10.1088/0953-4075/42/21/215002.
- [53] Capra, A. Lifetime of magnetically trapped antihydrogen in ALPHA. *Hyperfine Interactions*, **240(1)** (2019). doi:10.1007/s10751-018-1526-y.
- [54] Glinsky, M. E. & O’Neil, T. M. Guiding center atoms: Three-body recombination in a strongly magnetized plasma. *Physics of Fluids B*, **3(5)**:1279–1293 (1991). doi:10.1063/1.859820.
- [55] Robicheaux, F. Atomic processes in antihydrogen experiments: a theoretical and computational perspective. *Journal of Physics B: Atomic, Molecular and Optical Physics*, **41(19)**:192001 (2008). doi:10.1088/0953-4075/41/19/192001.
- [56] Fujiwara, M. C. *et al.* Temporally Controlled Modulation of Antihydrogen Production and the Temperature Scaling of Antiproton-Positron Recombination. *Phys. Rev. Lett.*, **101**:053401 (2008). doi:10.1103/PhysRevLett.101.053401.
- [57] Jonsell, S., Charlton, M., & van der Werf, D. P. The role of antihydrogen formation in the radial transport of antiprotons in positron plasmas. *Journal of Physics B: Atomic, Molecular and Optical Physics*, **49(13)**:134004 (2016). doi:10.1088/0953-4075/49/13/134004.
- [58] Dubin, D. H. E. & O’Neil, T. M. Trapped nonneutral plasmas, liquids, and crystals (the thermal equilibrium states). *Rev. Mod. Phys.*, **71**:87–172 (1999). doi:10.1103/RevModPhys.71.87.

- [59] Radics, B., Murtagh, D. J., Yamazaki, Y., & Robicheaux, F. Scaling behavior of the ground-state antihydrogen yield as a function of positron density and temperature from classical-trajectory Monte Carlo simulations. *Phys. Rev. A*, **90**:032704 (2014). doi:10.1103/PhysRevA.90.032704.
- [60] Ahmadi, M. *et al.* Observation of the 1S–2S transition in trapped antihydrogen. *Nature*, **541(7638)**:506–510 (2017). doi:10.1038/nature21040.
- [61] Nichols, E. F. & Hull, G. F. The Pressure Due to Radiation. (Second Paper.). *Phys. Rev. (Series I)*, **17**:26–50 (1903). doi:10.1103/PhysRevSeriesI.17.26.
- [62] Hänsch, T. & Schawlow, A. Cooling of gases by laser radiation. *Optics Communications*, **13(1)**:68–69 (1975). doi:https://doi.org/10.1016/0030-4018(75)90159-5.
- [63] Brewer, S. *et al.* Al+27 Quantum-Logic Clock with a Systematic Uncertainty below 10–18. *Phys. Rev. Lett.*, **123(3)** (2019). doi:10.1103/physrevlett.123.033201.
- [64] Monroe, C. *et al.* Demonstration of a Fundamental Quantum Logic Gate. *Phys. Rev. Lett.*, **75**:4714–4717 (1995). doi:10.1103/PhysRevLett.75.4714.
- [65] Foot, C. J. *Atomic Physics (Oxford Master Series in Atomic, Optical and Laser Physics)*. Oxford University Press, USA, 1 edition (2005). ISBN 0198506961.
- [66] Adams, C. & Riis, E. Laser cooling and trapping of neutral atoms. *Progress in Quantum Electronics*, **21(1)**:1–79 (1997). doi:https://doi.org/10.1016/S0079-6727(96)00006-7.
- [67] Wikipedia user: DePiep. Periodic table (demo 1).svg. [https://en.wikipedia.org/wiki/File:Talk,_Periodic_table_\(demo_1\).svg](https://en.wikipedia.org/wiki/File:Talk,_Periodic_table_(demo_1).svg) (2013). Accessed 2021-02-16. Licensed under Creative Commons Attribution-ShareAlike 3.0.
- [68] Brewer, L. R. *et al.* Static properties of a non-neutral ${}^9\text{Be}^+$ -ion plasma. *Phys. Rev. A*, **38**:859–873 (1988). doi:10.1103/PhysRevA.38.859.
- [69] Shiga, N., Itano, W. M., & Bollinger, J. J. Diamagnetic correction to the ${}^9\text{Be}^+$ ground-state hyperfine constant. *Phys. Rev. A*, **84**:012510 (2011). doi:10.1103/PhysRevA.84.012510.
- [70] Kondev, F. *et al.* The NUBASE2020 evaluation of nuclear physics properties. *Chinese Physics C*, **45(3)**:030001 (2021). doi:10.1088/1674-1137/abddae.
- [71] Bollinger, J. J., Wells, J. S., Wineland, D. J., & Itano, W. M. Hyperfine structure of the $2p^2P_{\frac{1}{2}}$ state in ${}^9\text{Be}^+$. *Phys. Rev. A*, **31**:2711–2714 (1985). doi:10.1103/PhysRevA.31.2711.

- [72] Maxwell, D. Private communications.
- [73] Sawyer, B. C., Britton, J. W., & Bollinger, J. J. Spin dephasing as a probe of mode temperature, motional state distributions, and heating rates in a two-dimensional ion crystal. *Phys. Rev. A*, **89**:033408 (2014). doi:10.1103/PhysRevA.89.033408.
- [74] Thompson, R. C., Donnellan, S., Crick, D. R., & Segal, D. M. Applications of laser cooled ions in a Penning trap. *Journal of Physics B: Atomic, Molecular and Optical Physics*, **42**(15):154003 (2009). doi:10.1088/0953-4075/42/15/154003.
- [75] Andelkovic, Z. *et al.* Laser cooling of externally produced Mg ions in a Penning trap for sympathetic cooling of highly charged ions. *Phys. Rev. A*, **87**:033423 (2013). doi:10.1103/PhysRevA.87.033423.
- [76] Sameed, M. *Laser-Ablated Beryllium Ions for Cold Antihydrogen in ALPHA*. Ph.D. thesis (2017).
- [77] Sameed, M., Maxwell, D., & Madsen, N. Ion generation and loading of a Penning trap using pulsed laser ablation. *New Journal of Physics*, **22**(1):013009 (2020). doi:10.1088/1367-2630/ab6066.
- [78] Bäuerle, D. *Laser Processing and Chemistry*. Springer (2011). ISBN 978-3-642-17613-5.
- [79] <https://www.quantel-laser.com/en/products/item/ultra-50-100-mj--134.html>. Accessed 2021-03-09.
- [80] https://www.newport.com/medias/sys_master/images/images/h24/h0e/8797301637150/Variable-Attenuator-for-Lasers-App-Note-26.pdf. Accessed 2021-03-11.
- [81] Amoruso, S. *et al.* Kinetic energy distribution of ions in the laser ablation of copper targets. *Applied Surface Science*, **127-129**:953–958 (1998). doi:https://doi.org/10.1016/S0169-4332(97)00772-1.
- [82] Madsen, N., Robicheaux, F., & Jonsell, S. Antihydrogen trapping assisted by sympathetically cooled positrons. *New Journal of Physics*, **16**(6):063046 (2014).
- [83] Kabantsev, A. A., Yu, J. H., Lynch, R. B., & Driscoll, C. F. Trapped particles and asymmetry-induced transport. *Physics of Plasmas*, **10**(5):1628–1635 (2003). doi:10.1063/1.1564089.

- [84] Donley, E. A. *et al.* Double-pass acousto-optic modulator system. *Review of Scientific Instruments*, **76(6)**:063112 (2005). doi:10.1063/1.1930095.
- [85] <https://www.comsol.com/>. Accessed 2022-03-05.
- [86] Driscoll, C. F. Observation of an unstable $l=1$ diocotron mode on a hollow electron column. *Phys. Rev. Lett.*, **64**:645–648 (1990). doi:10.1103/PhysRevLett.64.645.
- [87] Skuja, L., Hosono, H., & Hirano, M. Laser-induced color centers in silica. *Proceedings of SPIE - The International Society for Optical Engineering*, **4347** (2001). doi:10.1117/12.425020.
- [88] Marciniak, C. D., Ball, H. B., Hung, A. T.-H., & Biercuk, M. J. Towards fully commercial, UV-compatible fiber patch cords. *Opt. Express*, **25(14)**:15643–15661 (2017). doi:10.1364/OE.25.015643.
- [89] Amole, C. *et al.* In situ electromagnetic field diagnostics with an electron plasma in a Penning–Malmberg trap. *New Journal of Physics*, **16(1)**:013037 (2014). doi:10.1088/1367-2630/16/1/013037.