Title

The ALPHA antihydrogen trapping apparatus

Permalink

https://escholarship.org/uc/item/5nt9j6r3

Authors

Amole, C
Andresen, GB
Ashkezari, MD
et al.

Publication Date

2014

DOI

10.1016/j.nima.2013.09.043

Peer reviewed

The ALPHA antihydrogen trapping apparatus

C. Amole [a], G.B. Andresen [b], M.D. Ashkezari [c], M. Baquero-Ruiz [d], W. Bertsche [e,f,g],
P.D. Bowe [b], E. Butler [h,1], A. Capra [a], P.T. Carpenter [i], C.L. Cesar [j], S. Chapman [d], M. Charlton [e],
A. Deller [e], S. Eriksson [e], J. Escallier [k], J. Fajans [d], T. Friesen [l], M.C. Fujiwara [m], D.R. Gill [m],
A. Gutierrez [n], J.S. Hangst [b], W.N. Hardy [n], R.S. Hayano [o], M.E. Hayden [c], A.J. Humphries [e],
J.L. Hurt [i], R. Hydomako [l], C.A. Isaac [e], M.J. Jenkins [e], S. Jonsell [p], L.V. Jørgensen [e],
S.J. Kerrigan [e], L. Kurchaninov [m], N. Madsen [e], A. Marone [k], J.T.K. McKenna [q], S. Menary [a],
P. Nolan [q], K. Olchanski [m], A. Olin [m], B. Parker [k], A. Povilus [d], P. Pusa [q], F. Robicheaux [i],
E. Sarid [r], D. Seddon [q], S. Seif El Nasr [n], D.M. Silveira [j], C. So [d], J.W. Storey [m], R.I. Thompson [l],
J. Thornhill [q], D. Wells [q], D.P. van der Werf [e], J.S. Wurtele [d,t], Y. Yamazaki [u,s],
ALPHA Collaboration*

[a] Department of Physics and Astronomy, York University, Toronto ON Canada, M3J 1P3
[b] Department of Physics and Astronomy, Aarhus University, DK-8000 Aarhus C, Denmark
[c] Department of Physics, Simon Fraser University, Burnaby, BC Canada, V5A 1S6
[d] Department of Physics, University of California at Berkeley, Berkeley, CA 94720-7300, USA
[e] Department of Physics, College of Science, Swansea University, Swansea SA2 8PP, UK
[f] School of Physics and Astronomy, University of Manchester, Manchester M13 9PL, UK
[g] The Cockcroft Institute, Warrington WA4 4AD, UK
[h] Physics Department, CERN, CH-1211 Geneva 23, Switzerland
[i] Department of Physics, Auburn University, Auburn, AL 36849-5311, USA
[j] Instituto de Física, Universidade Federal do Rio de Janeiro, Rio de Janeiro 21941-972, Brazil
[k] Brookhaven National Laboratory, Upton, NY 11973, USA
[l] Department of Physics and Astronomy, University of Calgary, Calgary AB, Canada, T2N 1N4
[m] TRIUMF, 4004 Wesbrook Mall, Vancouver BC, Canada V6T 2A3
[n] Department of Physics and Astronomy, University of British Columbia, Vancouver BC, Canada V6T 1Z4
[o] Department of Physics, University of Tokyo, Tokyo 113-0033, Japan
[p] Department of Physics, Stockholm University, SE-10691 Stockholm, Sweden
[q] Department of Physics, University of Liverpool, Liverpool L69 7ZE, UK
[r] Department of Physics, NRCN-Nuclear Research Center Negev, Beer Sheva 84190, Israel
[s] Atomic Physics Laboratory, RIKEN Advanced Science Institute, Wako, Saitama 351-0198, Japan
[t] Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA
[u] Graduate School of Arts and Sciences, University of Tokyo, Tokyo 153-8902, Japan

ARTICLE INFO

ABSTRACT

The ALPHA collaboration, based at CERN, has recently succeeded in confining cold antihydrogen atoms in a magnetic minimum neutral atom trap and has performed the first study of a resonant transition of the anti-atoms. The ALPHA apparatus will be described herein, with emphasis on the structural aspects, diagnostic methods and techniques that have enabled antihydrogen trapping and experimentation to be achieved.

© 2013 Elsevier B.V. All rights reserved.

1. Introduction

Trapping antihydrogen ($\overline{\text{H}}$) atoms is an important milestone towards the goal of precision spectroscopic comparisons of the

* Comment to the author can be directed to alpha-contact@cern.ch.
[1] Present address: Centre for Cold Matter, Imperial College, London SW7 2BW, United Kingdom.

properties of antihydrogen and hydrogen. Such comparisons will allow direct tests of the CPT theorem, according to which no difference is expected between the energy levels of the atom and its antimatter counterpart. The ALPHA collaboration has recently trapped antihydrogen atoms. In the first report [1], 38 atoms were trapped for 0.17 s, soon followed by much longer (1000 s) confinement times [2] of a larger trapping sample. Atoms trapped for such long times are expected to only occupy their ground state [2]. Recently, ALPHA demonstrated the first resonant microwave interactions probing the hyperfine structure of the antihydrogen ground state [3]. It is anticipated that the most precise measurement of the energy levels of the antihydrogen atom will be achieved via two-photon excitation from the 1S to the 2S level. A comparable measurement with normal hydrogen atoms enabled the frequency of this transition to be determined with a relative precision of better than 10^{-14} [4].

The ALPHA experiment uses antiprotons ($\overline{p}$) supplied by the CERN Antiproton Decelerator (AD), a unique facility that provides bunches of the antiparticles at relatively low energy (5.3 MeV) [5]. Following further energy degradation, by about three orders of magnitude, a fraction of these antiprotons are stored in a charged-particle trap. Positrons (e^+) are routinely obtained from a ^{22}Na source, and a near-monoenergetic low energy beam can be formed to facilitate their capture using well-established techniques. ALPHA uses Penning-Malmberg [6] traps to separately confine the antiprotons and positrons. Only after these charged particles are manipulated and cooled, are they mixed to produce antihydrogen atoms.

The ALPHA experiment is the successor to the ATHENA experiment, which produced the first cold antihydrogen atoms in 2002 [7]. The ATHENA traps, however, were only for charged particles, and no means were provided to confine the neutral antihydrogen atoms. As a consequence, these anti-atoms escaped to annihilate on the wall of the trap shortly after they were formed. The main new feature introduced by ALPHA is an inhomogeneous magnetic field that can hold neutral atoms using the interaction of the magnetic field with the magnetic dipole moment of the atoms. This trap uses an octupole magnet to provide the radial confinement of the antihydrogen, and mirror magnets to provide the axial confinement (see Section 2.6).

Because the magnetic potential well is shallow for ground state atoms, only those that are very cold (kinetic energy equivalent to less than around 0.5 K) can be trapped. Thus, the antihydrogen must be produced from antiparticles cooled as far as possible in order to enhance the capture probability. Cooling is also important to enhance the production of the antihydrogen atoms, since the main production mechanism is considered to be three-body recombination (i.e. $\overline{p}+e^++e^+\rightarrow\overline{H}+e^+$) which depends very strongly on the positron temperature (see e.g. [8]).

The fundamental cooling mechanism employed in ALPHA is the emission of cyclotron radiation of the charged particles gyrating in the magnetic field. Because of their low mass, positrons can cool directly, in principle equilibrating with the cryogenic walls of the trap. The heavier antiprotons are confined together with electrons, and transfer kinetic energy to them through collisions. The electrons, in the same way as the positrons, radiate this energy away. In practice, we find that both the positrons and the antiprotons are hotter than the temperature of the wall of the trap, which is measured to be 8 K. We employed additional cooling techniques, such as evaporative cooling (see Section 4.3.1) in order to further cool the species.

In the first trapping experiments, trapping was demonstrated by allowing the anti-atoms to escape after turning the confining magnetic fields off. The anti-atoms could then hit the walls of the trap where they annihilated. Our main diagnostic device for the annihilation products (mainly pions from antiproton
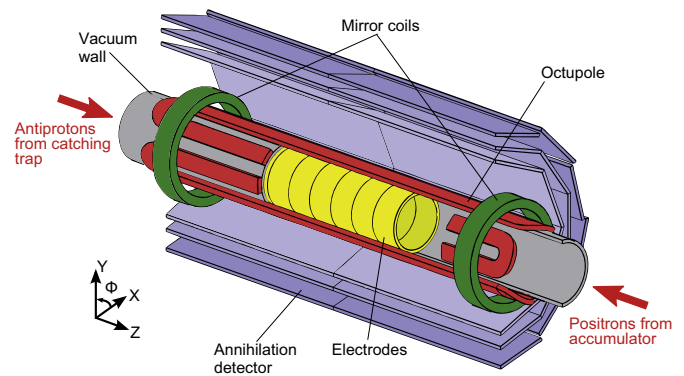
Fig. 1. Schematic illustration of the inner section of the ALPHA experiment showing the Penning-Malmberg trap electrodes, the neutral trap (comprising the octupole and mirror coils) and the silicon-based annihilation detector. The components are not drawn to scale.

annihilations) was a 60-module silicon detector that surrounded both the trap electrodes and the magnetic trap, as shown schematically in Fig. 1. This detector is described more fully in Section 3.1. A major effort was necessary to unambiguously identify the annihilation events as being caused by antihydrogen atoms by establishing their charge neutrality (to rule out trapped antiprotons) and to carefully distinguish them from cosmic rays events [1,9].

The success of the antihydrogen trapping endeavour was based on many innovations in methods of handling the antiprotons and positrons in order to mix them with the maximum spatial overlap, while keeping them as cold as possible. Amongst the techniques used to manipulate the charged particle plasmas, and to be described below (Section 4), were 'rotating walls' to control the size of the plasmas, and autoresonance excitation to enable the mixing of the antiprotons and the positrons without excessive heating.

The general structure of the paper is as follows. Section 2 contains a general description of the layout and vacuum and cryogenic infrastructure of ALPHA, together with details of the charged particle traps, the positron accumulator and magnetic minimum neutral atom trap. Section 3 comprises a summary of the detection systems used in the experiment, whilst Section 4 describes methods and processes, including details of charged particle manipulations, as well as the monitoring systems developed by ALPHA.

2. Apparatus structure, vacuum and cryogenics

2.1. Overview

The ALPHA experiment resembles, in several aspects, its predecessor, the ATHENA apparatus, which was used to produce antihydrogen from cold trapped plasmas in 2002 [7]. It features an open geometry, which allows particle insertion into the cold portion of the apparatus from the room temperature region (electrons) and also from adjacent machines (positrons and antiprotons). This geometry also allows easy particle extraction for the plasma diagnostic techniques available in the warm region (radial profile imaging and temperature measurement) and the introduction of microwave radiation into the trap [3].

From a functional point of view, the experimental system must produce a strong solenoidal magnetic field for stable charged-particle confinement in Penning-Malmberg traps and to supply the high gradient magnetic fields required for (neutral) antihydrogen trapping. In addition, it should also provide a cryogenic

environment to allow the non-neutral plasmas to reach very low temperatures and to facilitate the confinement of antihydrogen and its constituents for sufficiently long times (reducing the losses associated with annihilation with the residual gas molecules). Furthermore, space provision must be made for the annihilation vertex detector, used to characterise antihydrogen production and trapping.

A schematic drawing of the entire apparatus is shown in Fig. 2. The low-energy antiproton bunch arrives from the left (upstream) side, while positrons are transferred from the positron accumulator on the right (downstream) side. The apparatus is comprised of two independent cryostats: the large external cryostat houses a superconducting solenoid (see Section 2.7) which produces the field required for the operation of the Penning-Malmberg trap. The 260 mm-diameter room temperature bore of this solenoid hosts a second cryostat in which four superconducting magnets are sited – three for the neutral atom trap (see Section 2.6) and one to produce an extra solenoidal field in the antiproton catching trap region (see Section 2.4.1). The electrodes of the charged particle traps are contained within this latter cryostat. Around the internal cryostat (but still in the external cryostat bore) lies the cylindrically symmetric annihilation vertex detector.

2.2. Cryogenics

The inner cryostat uses liquid ^{4}He at 4.2 K as the coolant. It features a horizontal section, containing the four aforementioned magnets, the vacuum chamber with the Penning-Malmberg trap electrodes and a vertical section with the vapor-cooled leads for these magnets.

The cylindrical horizontal section features three concentric tubes which define three different regions: the trap vacuum chamber, the helium reservoir and the outer vacuum chamber. The innermost tube delimits a 47 mm-diameter cold bore where the electrodes of the Penning-Malmberg trap are inserted. The annular space between the inner tube and the middle tube houses the four superconducting magnets and is filled with liquid helium. The outermost volume (between the middle tube and the outer tube) is evacuated to thermally isolate the helium vessel from the room temperature surfaces. The vertex detector is slid onto the external surface (140 mm in diameter) of this outermost tube.

The vertical section of the cryostat is simpler: it is also cylindrical in shape, but it features only two concentric tubes, which delimit the helium vessel (containing the eight vapor-cooled leads for the four magnets) and the outer vacuum chamber. The horizontal and vertical sections of the outer vacuum chamber are directly connected by a large vacuum chamber, while the two sections of the helium vessel are connected through flexible metal hoses. These hoses also contain the superconducting leads that carry the various magnet currents and the signal wires for the quench protection system (see Section 4.6).

The ultimate temperature of the electrodes is limited by heat conducted by the leads carrying the excitation voltages and by black body radiation originating from warmer surfaces of the apparatus. The trap electrodes are in weak thermal contact with the cold surface of the trap vacuum chamber, which reduces the rate at which they cool down. To alleviate this condition, the leads are thermally anchored to a heat sink inserted in a pipe that runs through the liquid helium reservoir. Still, the trap takes longer to cool than the chamber walls. Temperature measurements on one of its electrodes show that the trap reaches ∼8 K a couple of hours after liquid helium starts accumulating in its vessel. Due to the lack of radial space and the desire to reduce the scattering material between the annihilation and detection points of the antiprotons (which would reduce the detector resolution), there are no actively cooled thermal radiation shields between the room-temperature surfaces and the cryogenic surfaces. To reduce the radiative power transferred to the helium bath, multilayer super insulation is wrapped around all cold surfaces.

2.3. Vacuum

Electromagnetic traps for particles (charged or neutral) rely on ultra-high vacuum (UHV) conditions to ensure a long trapping lifetime. In the case of antimatter particles, additional care has to be taken due to another possible loss channel: annihilation on background gas molecules. The upstream end of the trap vacuum chamber is closed by a 12.5 μm-thick stainless steel foil, which separates the UHV part of the apparatus (the trap region) from the vacuum chamber crossed by antiprotons when extracted from the AD, where the pressure is higher. This foil is also used as the first degrader for the antiproton bunch (see Section 2.4.1). The downstream end of the trap vacuum chamber has to be open to the room temperature portion of the apparatus (where the vacuum pumps, gauges and valves are located), to allow positron transfers, electron loading (see Section 2.8) and particle extraction for diagnostic purposes (see Sections 3.3. and 4.3). The surfaces of the trap vacuum chamber and the trap electrodes are held at
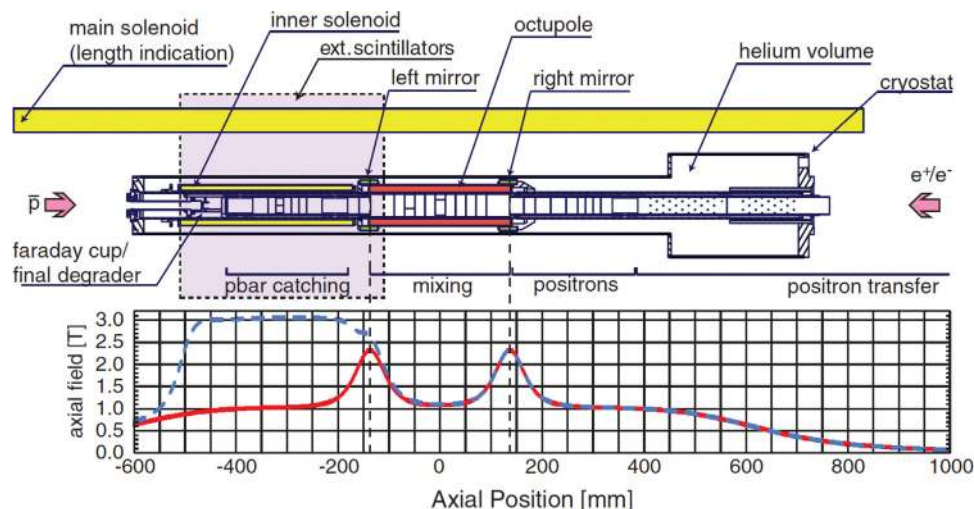

Fig. 2. Schematic illustration of the ALPHA apparatus. The antiprotons arrive in a burst from the AD, whilst the positrons are supplied from a separate accumulator system described in Section 2.5. The lower panel shows the variation of the axial magnetic field strength, as described more fully in Sections 2.6 and 2.7.

cryogenic temperatures, and as such they do not outgas; rather they act as powerful vacuum pumps. After the pump-down stage, outgassing from the warm surfaces becomes a major gas load for the vacuum system. The high vacuum pump used is a $150\,\mathrm{l\,s^{-1}}$ ion pump, in conjunction with a titanium-sublimation pump for pumping active gases, and pressures around 10^{-10} mbar are routinely achieved (measured at the room temperature region). The pressure in the trap region is expected to be lower than this and simple estimates from measurements of the antiproton lifetime [10] point to pressures in the 10^{-13}–10^{-14} mbar range.

2.4. Central particle traps

Axial confinement of charged particles is accomplished by biasing the coaxial cylinders that form the electrode stack of the ALPHA Penning-Malmberg trap. These electrodes are all made of aluminium to minimise the scattering material through which pions from antiproton annihilations have to pass. To avoid oxidisation, and the resulting patch potentials, the electrodes are plated with a layer of gold. Below the 2 μm of gold, sub-layers of nickel and copper, with thicknesses of 0.1 and 1 μm respectively, are deposited to improve plating performance and avoid diffusion of the gold into the aluminium former. Static or varying voltages can be applied independently to each electrode, allowing the electric field in the trap to be shaped according to the desired operation.

The electrode stack consists of 34 elements organised in three distinct groupings, each with a particular use in mind. In Fig. 3 an overview is given where antiprotons arrive from the left. Proceeding from the left the first eleven electrodes are used to catch, cool and accumulate antiprotons (catching trap). The following group of thirteen elements, which are axially positioned on the centre of the magnetic minimum trap, is where antihydrogen is formed and is referred to as the mixing trap. The final ten electrodes, which are mainly used for re-capturing positrons and preparing them for antihydrogen synthesis, are referred to as the positron trap.

2.4.1. The catching trap

The catching trap is comprised of 11 electrodes (see Fig. 4). The first (HVA) and the ninth (HVB) electrodes are specially designed to support high voltage (HV) to catch the incoming antiproton bunches. The inner radius of the high voltage electrodes is 29.6 mm whereas for the other electrodes it is 33.6 mm. The non-HV electrodes are held electrically isolated from one another, and azimuthally fixed by 1.5 mm ruby spheres locked in
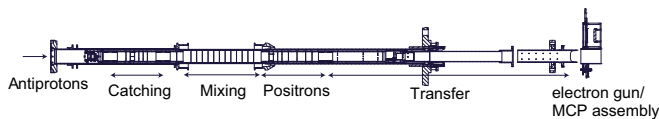
Fig. 3. The full electrode stack of the central ALPHA apparatus. Antiprotons arrive from the left, pass through a stainless steel vacuum window and are degraded in an aluminium foil before being captured in the catching region. Positrons enter from the right via a stack of transfer electrodes and are initially caught in the combined mixing and positron region. A vacuum manipulator to the right can be used to lower an electron gun or a micro-channel plate (MCP) detector onto the axis of the system (see text). A later version of this manipulator holder also housed elements to allow the insertion of microwaves into the device (see Fig. 12).
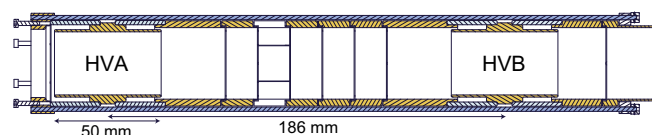
Fig. 4. Axial section of the catching trap electrodes with the two high voltage electrodes used for catching the energetic antiprotons highlighted. The six-fold segmented electrode is apparent.

appropriate spherical holes. One of the non-HV electrodes is azimuthally split into six segments in order to allow for the application of the rotating wall technique for plasma compression [11], as described in Section 4.2.2. The electrode stack is held together by stainless steel bars fixed to locking rings at each end of the system. The electrodes are connected to vacuum feedthroughs using Lakeshore KAP3 copper/kapton cables via screw points on each electrode. Before exiting the cold part of the system these cables are thermally anchored to a cold surface. High thermal conductance is ensured by allowing the signal to pass from the cable to a copper/kapton stripline, which can be easily anchored to the cold points.
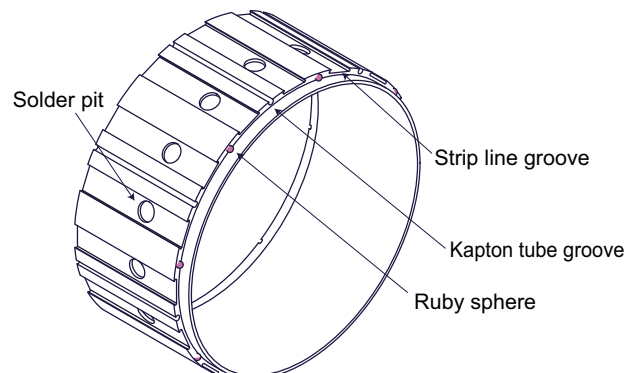
2.4.2. The mixing trap

The mixing trap is comprised of 13 electrodes. In order to maximise the depth of the magnetic minimum trap the trapped atom orbits should be allowed as close as possible to the wiring of the octupole. To allow this the mixing trap has been manufactured using a novel ultra-thin design, which does not rely on a support structure (see Fig. 5). This has the further advantage of minimising the scattering material encountered by particles produced from antiproton annihilations, in order to aid vertex reconstruction. The electrodes have a maximum wall thickness of 1.5 mm and the signal is fed along 4 mm wide, 300 μm thick, kapton/copper strip lines captured in groves on the external surface of the electrodes. The inner diameter of the trap is 44.5 mm. The axial separation and azimuthal alignment are ensured by ruby spheres of 1 mm diameter. The stack is held together by aluminium wires secured inside kapton tubing captured in 1/4 open cylindrical 0.92 mm diameter groves in the outer surface of the electrode. Electrical insulation from the vacuum wall (with an inner diameter of 48 mm) is ensured by a kapton film wrapped around the outer surface on insertion.

2.4.3. The central positron trap

When positrons are transferred from the positron accumulator the bunch length is about 500 mm. To capture and prepare these positrons the mixing trap is extended downstream by a positron trap which consists of 10 electrodes of the same construction as the catching trap, i.e. with an inner diameter of 33.6 mm. The last of these is of the same design as the HV electrodes of the catching trap. This stack also contains one electrode which is 4-way azimuthally segmented to allow for application of a rotating wall.

2.4.4. The transfer electrodes

Outside the trap system on the right a number of long cylinders with an inner diameter of 33.6 mm are mounted to ensure a well defined ground potential for the transfer of positrons, the loading

Fig. 5. Illustration of a single mixing trap electrode. The wall thickness is 1.5 mm, the inner diameter is 44.5 mm and the signal runs in strip-lines soldered to the relevant electrode.

of electrons as well as for the ejection of all particle species to the micro-channel plate detector assembly to the right of the system, as shown in Fig. 3.

2.5. The positron accumulator

A schematic illustration of the positron accumulator, together with a sketch of the on-axis static electrical potential used to trap the positrons, is shown in Fig. 6. The instrument is fed with low energy (typically around 80 eV) positrons from an accompanying beamline. The positrons are derived as β^+ particles from a ^{22}Na radioactive source, of peak activity around 2.8 GBq ($\sim$75 mCi), and are formed into a beam in vacuum with an efficiency of approximately 0.5% using a solid neon moderator [12]. The latter is the most efficient positron moderator to date and is used in the internal conical geometry [13]. The source and moderator are attached directly to a closed cycle helium cold head capable of maintaining a temperature of around 5 K. An exploded view of the source/moderator arrangement is given in Fig. 7 which shows how a 2 mm thick sapphire disk spacer is used to electrically isolate the cold head from the source and moderator, whilst maintaining good thermal contact. The beam energy is set by electrically biasing the source holder.

The source chamber, shown on the left of Fig. 6, is evacuated using an oil-free magnetically levitated turbomolecular pump, backed by a scroll pump. To grow a moderator, this chamber is isolated from the rest of the accumulator system using an in-line gate valve and neon is admitted into the region, typically at a pressure around 10^{-4} mbar for a period of about an hour. The growth in low energy positron intensity is monitored by a calibrated NaI scintillator/photomultiplier detector situated adjacent to the aforementioned gate valve. The positrons are magnetically guided to this valve. A solenoid furnishes the guiding field around a $\sim$0.85 m long, 19 mm internal diameter vacuum tube, which also serves as a pumping restriction when the accumulator is in use. This restriction is necessary to limit the flow of nitrogen gas, which is used to promote positron capture in the accumulator (as described below), into the moderator region, where it has a detrimental effect on the efficiency of beam production. When moderator growth is complete, the neon supply is switched off, and the beam is ready for use. From time-to-time, and depending upon experimental circumstances, the moderator needs to be replaced. This is done by raising the cold head temperature to 30–40 K using in-situ heaters to allow the neon to sublimate. Once this has occurred and the cold head has re-cooled to the required temperature, growth of a fresh moderator can begin. All of the

manipulations in the moderator growth sequence are remotely controlled using LabView software. Thus, a beam of around 5 million slow positrons per second (per GBq of source activity) is guided into the accumulator (see Fig. 6) which is a Surko-type three-stage buffer gas instrument [14]. The device is essentially a Penning-Malmberg trap, in which nitrogen gas molecules are introduced to cause the positrons to lose kinetic energy via collisions to promote trapping. The instrument is housed inside a 1.5 m long, 300 mm internal diameter, vacuum tube which is pumped on either side via chambers, each equipped with a 1200 l s^{-1} cryopump. The positrons are confined radially using a large bore solenoid which produces an on-axis central field of 0.14 T.

The first stage of the accumulator is a single electrode in the form of a 509.6 mm long, 12.7 mm internal diameter, tube. This tube is held at a voltage around 7–8 V lower than a 21 mm long, 12.7 mm internal diameter auxiliary electrode, which together form the trapping potential. The second stage is made up of a 539.8 mm long, 30.5 mm diameter, electrode with the final stage composed of four separate electrodes, each 154.0 mm long and of 200.7 mm internal diameter. One of these electrodes, as depicted schematically in Fig. 6, is segmented azimuthally into six pieces to allow, as explained in Section 4.4.1, the application of a rotating wall to control the radial dimension of the positron plasma.

Nitrogen gas is admitted to the centre of the first stage of the accumulator at a pressure of around 10^{-3} mbar. It has been found that the nitrogen molecule is the most efficient gas to promote capture of the positrons [14]. This is due to a prominent resonance
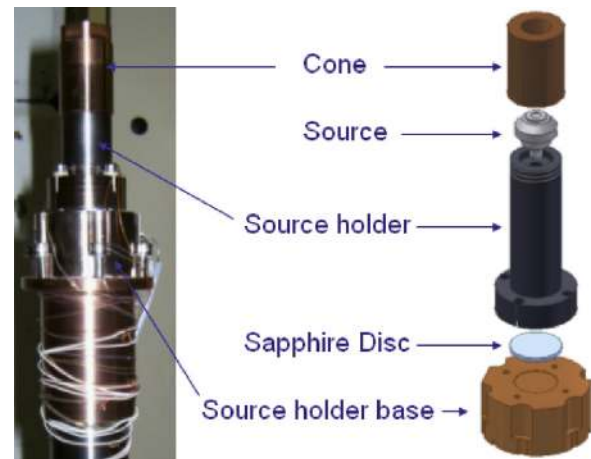
Fig. 7. Source and moderator assembly on the coldhead (see text).

Fig. 6. Schematic illustration of the positron beamline and the buffer gas accumulator. The lower panel is a representation of the axial electrical potential of the trap and shows how collisions, progressively in each stage, result in accumulation in the third stage. When the nitrogen line is closed, the gas is pumped out promptly in readiness for transfer of the positrons to the main ALPHA system.

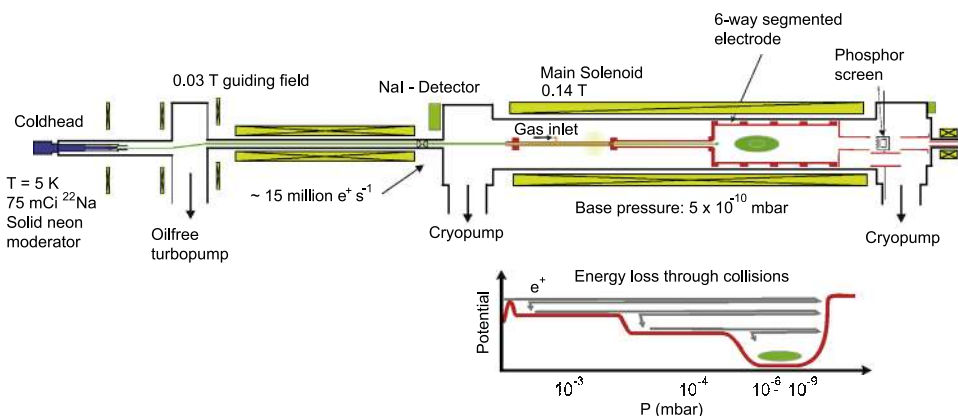in the positron-nitrogen electronic excitation cross-section [15] close to the threshold (at around 9 eV) for the reaction,

$$e^+ + N_2 \rightarrow e^+ + N_2^*, \tag{1}$$

which competes effectively with the usually dominant positron loss channel, positronium (Ps) formation, as

$$e^+ + N_2 \rightarrow Ps + N_2^+. \tag{2}$$

Typically around 30% of the positrons passing into the accumulator will be captured, mainly in the first stage, where the gas pressure is highest. Once a positron is captured, a further excitation of the gas will confine it to the second stage and then finally into the third stage, which has a gas pressure of around 10^{-6} mbar. Thus, positrons collect in this region, and over 250 million of them can be captured on a 3–4 min timescale for a source activity of 75 mCi: they are so numerous and dense that a plasma is formed after about 10 s of accumulation. This allows the application of a dipolar rotating wall electric field using the segmented electrode driven from a supply with signals with a phase difference of $\pi/3$ between adjacent electrode segments. Typically a fixed frequency of 600 kHz is used with an applied voltage amplitude of 0.4 V. The electric field, which rotates in the same sense as the natural plasma $\mathbf{E} \times \mathbf{B}$ drift, applies a torque which increases the plasma angular frequency such that it shrinks due to the conservation of angular momentum [16–18].

2.6. The magnetic neutral atom trap

A typical design of a magnetic trap consists of two so-called mirror coils (or pinch coils) for axial confinement and a multipole field for radial confinement. These fields are superposed on a solenoidal field necessary for non-neutral plasma confinement. In ALPHA, the axial well depth is usually deeper than the transverse one. The transverse (radial) well depth, ΔB, is then approximately given by Eq. (3) where B_w is the multipole field strength at the radius of the electrodes, and B_z is the z-directed field from the solenoid and mirror coils:

$$\Delta B = \sqrt{B_w^2 + B_z^2} - B_z. \tag{3}$$

Currently, antihydrogen is produced at relatively high temperatures with respect to the typical depth of the neutral atom trap, which is of the order of 0.8 T (i.e. ~ 0.54 K for ground state antihydrogen) [1,2]. Thus, in order to enhance the trapping probability, the trap depth should be as large as possible. However, this condition gives rise to large magnetic gradients which distort the field in the solenoid region, which may affect the stability of the positron and antiproton plasmas used in the formation process [19]. In the following we will describe how we have combined the condition of large well depth and stable plasmas. A more detailed description of the design can be found in Ref. [20].

The Ioffe–Pritchard configuration is a commonly used neutral atom trap. In this trap the radial multipole field, $B_s(r)$, is produced by a quadrupole (multipole order $s=2$). Fajans et al. [19], however, showed that the stability of non-neutral plasmas is severely impeded for $B_2(r_p)/B_z \geq 0.05$, where r_p is the plasma radius. Assuming that this inequality is valid for all s, we can graphically compare the fields of several multipoles, given by

$$B_s(r) = K_s r^{s-1}, \tag{4}$$

with a straight line given by $B_s(r_p)/B_z = 0.05$ and find the maximum stable plasma radius (see Fig. 8) for each s.

Typically, our plasma radii are smaller than $0.2 r_w$, so ALPHA opted for an octupole to radially confine the anti-atoms. Although higher order multipoles lead to even flatter radial fields in the plasma region [21], the slope at the wall is also steeper leading to a
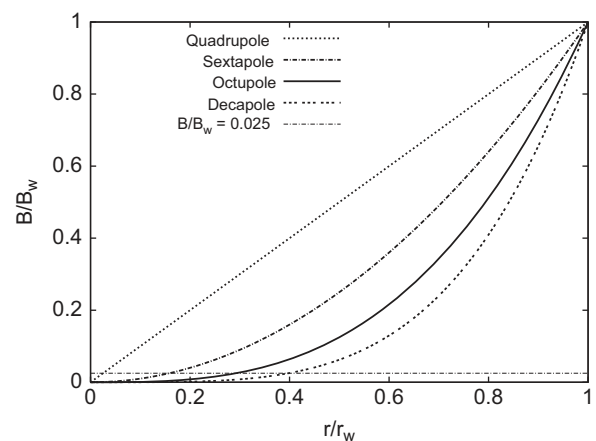
Fig. 8. Radial field dependence for $s=2$–5 (quadrupole through decapole) normalised to unity at the electrode wall radius. The significance of the horizontal line is explained in the text. However, the ratio of 0.025 is chosen for the figure since the multipole field at the trap wall is about twice that due to the solenoid: see also Ref. [20].


Fig. 9. First layer of the octupole [20].

reduced effective well depth due to the finite thickness of the electrodes and the vacuum tube.

The resolution of the silicon detector, placed around the dewar containing the magnets and traps, worsens with increasing amount of material between the position of antiproton annihilation and the silicon strips due to scattering of the annihilation products. Therefore, the wire chosen for winding the octupole and mirror coils (Supercon 56S53) contains the relatively low copper to superconductor (NbTi) ratio of 0.9. Seven strands of this wire have been woven into a cable with an overall diameter, including insulation, of 1.156 mm. The use of a larger diameter cable reduces the number of layers and turns, which makes for easier manufacturing, and also reduces the inductance of the coils and octupole, so increasing their maximum ramp (and fall) rates.

Normally a multipole is constructed from so called race-track coils. However, the windings of these coils give rise to an undesirable axial component at the ends. We opted therefore for a serpentine winding pattern (see Fig. 9). The eight layers of the octupole are azimuthally staggered at 45° with respect to one another, almost cancelling the magnetic field on axis around the position of the windings.

The octupole layers (Fig. 9) are directly affixed to the vacuum tube using a special technique developed at the Brookhaven National Laboratory, wherein the wire is initially ultrasonically glued into place, and then permanently secured with G-10 spacers (in winding gaps), B-stage epoxy, and fibre overwraps [22]. This technique allows the wires to be placed in almost any pattern, with complete freedom to change the pattern between layers.

The final design for the neutral trap consists of an eight-layer octupole, and, as shown in Fig. 10, two mirror coils (each of thirty turns and four layers) located on top of the octupole. On one side, the small catching trap solenoid (and boost coil) is located, capable of raising the solenoidal field by 2 T, for improved antiproton trapping efficiency [23]. The whole assembly is immersed in a background solenoidal field of 1 T, provided by an external superconducting magnet. The nominal design currents through the
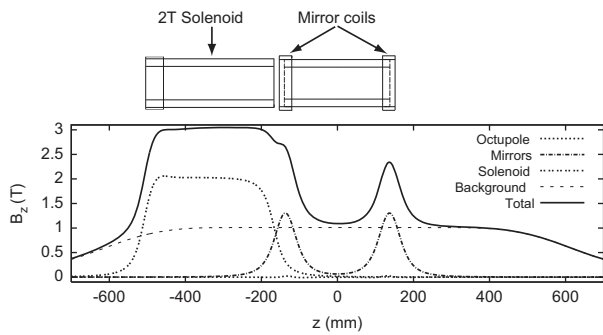
Fig. 10. The total axial magnetic field when the octupole and all the solenoids and coils are energised to their design currents. The 1 T background field is provided by a large bore external magnet [20].

octupole, mirror coils and small solenoid are 1100 A, 750 A and 230 A, respectively.

The radial field from the mirror coils either increases or reduces the radial field from the octupole, depending on the azimuthal angle, resulting in so-called 'holes' in the total field. These 'holes' can be clearly observed in Fig. 11 where the total field, calculated using the TOSCA/OPERA3D package [24], is plotted for a number of different radii and three different azimuthal angles. Moving the mirror coils outwards opens up holes in the magnetic well. The smallest field along the well boundary is at the saddles near the mirror coils; it is through these holes that antihydrogen may escape. The lowest saddle has a design value of 2.25 T. The lowest field within the magnet system has a corresponding value of 1.09 T giving the trap a nominal well depth of 1.16 T, corresponding to a temperature of 0.78 K for ground state antihydrogen. In practice, the coils are routinely run at a conservative ∼70% of maximum, producing a well depth of 0.54 K.

2.7. The external solenoid

The external solenoid magnet was constructed at the Kurchatov Institute in Moscow, Russia and used in UC Berkeley in the period 1995–2004 [25]. The magnet has a 260 mm horizontal room temperature bore into which our apparatus fits. It can produce a 3 T solenoidal field with a uniform field that is 600 mm long and 50 mm in radius, oriented along the solenoid axis. The axial field in the uniform region is constant to ± 0.25%, and azimuthally symmetric to an accuracy of 10^{-4}. During our experiments, this magnet was usually run at a field of 1 T (see Fig. 10).

The magnet assembly is placed onto an inside tube of a cylindrical helium vessel. The helium vessel is surrounded by a liquid nitrogen vessel, and there is an additional vapor-cooled radiation shield between the nitrogen and helium vessels. The outermost vessel of the magnet is corrugated stainless steel, the corrugation is to allow for thermal contraction. Inside this is a vacuum space with some multilayer superinsulation, then a liquid nitrogen vessel. The nitrogen vessel has a volume of 260 l, with inner and outer diameters of 915 and 1035 mm, respectively. The nitrogen vessel is suspended by eight bakelite supports which attach to bronze flanges on the two large end flanges. Inside the nitrogen jacket is more superinsulation, and then a radiation heat shield made of aluminium tube with copper flanges. Inside this shield there is vacuum space, and then the helium vessel. The helium surrounds the magnet windings and has a volume of 380 l and an outer diameter of 800 mm. The helium vessel is suspended by Kevlar strings, 4 mm in diameter, from the nitrogen vessel.

The magnet has a main coil which is wound in five sections on a fiberglass form and two steering coils wound on a second fiberglass form. The wire consists of NbTi filaments in a copper matrix with a lacquer insulating coating. The main coil uses about
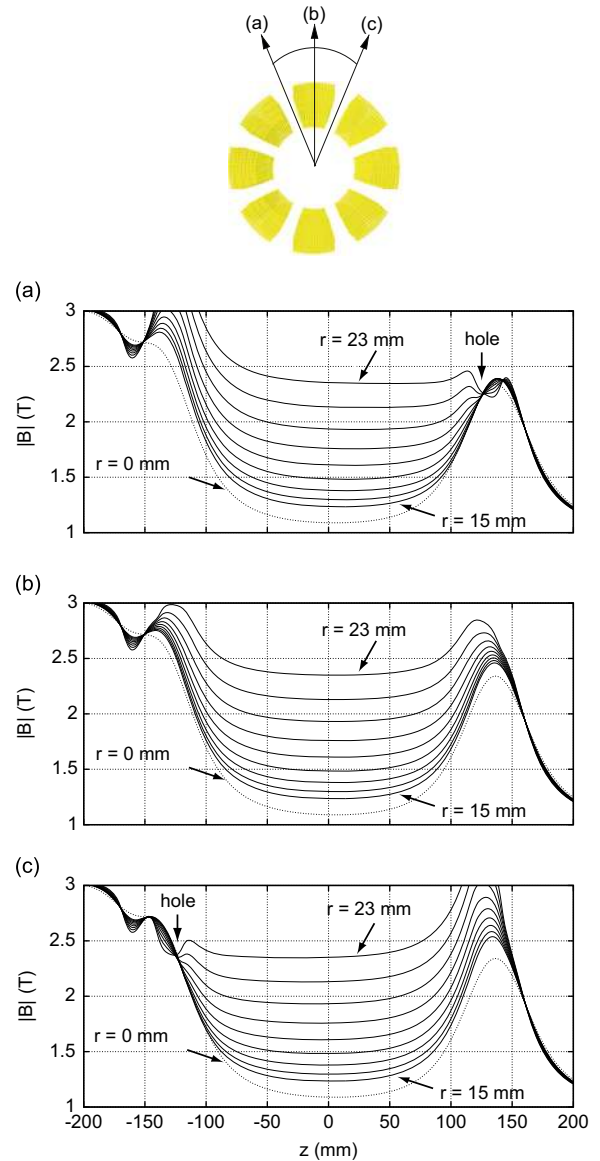
Fig. 11. The upper drawing is an axial view of a cross-section in the middle of the magnet system indicating the azimuthal angles for the planes for (a), (b) and (c). These three plots show the total field for a number of different radii as function of z. The radius for subsequent solid lines differs by 1 mm. (a) −22.5°; (b) 0°; (c) 22.5° [20].

30 km of wire. The room-temperature resistance is 2 kΩ. The inductance is 76 H. It takes 130 A to produce a field of 3 T. The average current density is $1.2 \times 10^4 \, Acm^{-2}$, and the stored energy is 650 kJ. Two transverse field steering coils, placed at an angle of ninety degrees with respect to each other, are capable of producing a field of 1.0% of the main field (i.e. 300 G.). These magnets can be used for alignment purposes, and have peak to peak variations of less than 1.0% in the uniform field region. The steering coils are wound in a saddle shape. They each have an inductance of 0.02 H and produce a field of $2.27 \, GA^{-1}$, with the maximum allowed current being 160 A.

In order to obtain high time stability of the main magnet field as well as to decrease liquid helium evaporation, the main magnet is supplied with a thermally activated persistent current switch, non-inductively wound from multi-filamentary 0.7 mm diameter NbTi wire in a resistive (CuMn alloy) matrix. Its normal resistance at 4.2 K is 23 Ω, and its inductance is approximately 0.07 mH. Prior to charging the main magnet, the persistent switch is activated by

an electrical heater. After charging is started, the heater can be switched off while the switch stays partially normal due to charging voltage if the Joule heating exceeds 0.7 W. Switching off the heater becomes possible due to properly chosen thermal insulation of the switch and helps to decrease evaporation of liquid helium during magnet charging. The switch is electrically shunted to prevent its overheating in the event that it quenches while the magnet remains superconducting. The magnet itself is passively quench protected by three electrical shunts. The steering coils are charged independently, and each of them is supplied with a persistent switch made of the same wire as the winding. In the persistent mode, the time stability is better than 4×10^{-7} per hour for the main magnet and better than 10^{-5} per hour for both steering coils.

2.8. The linear vacuum manipulator ('stick')

A vertically mounted translator, called 'the stick', is located between the positron accumulator and the mixing trap (see Fig. 3). It is used to align different devices with the plasma traps axis during the cycle of the experiment. To allow for the unobstructed passage of positrons from the accumulator into the trap it also includes a pass-through cylinder. For diagnostic purposes it has an imaging assembly, composed of an micro-channel plate (MCP) device and a phosphor screen, and also a microwave mirror and horn used for injecting microwave radiation to the trap. An earlier version also had a phosphor screen that faced the positron accumulator. An important device is the electron source, supplying the electrons that are used to cool the captured antiprotons. The positions of the various components are shown in Fig. 12. The microwave injection systems are described in Section 2.9. The MCP is described in Section 3.3.

The stick, which has a 15 cm travel, is located in the fringing field of the solenoid at a field of 0.024 T. The stick's base vertical
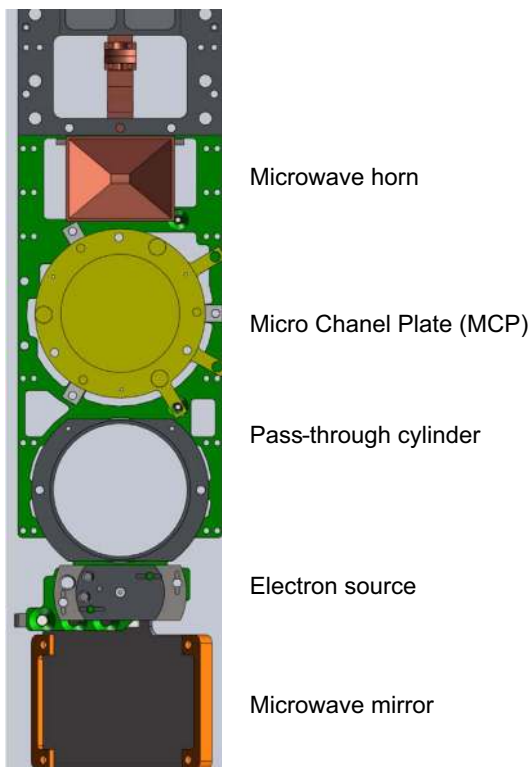

Fig. 12. The vacuum manipulator. From top to bottom, the components are the microwave horn, MCP, a pass-through cylinder to allow positrons to move into the trap, an electron source and a microwave mirror.

position is adjusted so that the electron gun is as close as possible to the trap magnetic axis. (The density and reproducibility of the resulting trapped plasmas are used to optimise the position). The typical time for the stick to move from one device to another is on the order of 10 s. Positional reproducibility is assured by the use of a stepper motor and confirmed by an optical encoder. Long term reproducibility is ensured by optical flags.

The electron source is a barium oxide emitter manufactured by Kimbal Physics. The gun is located in the 0.024 T field, and the emitted electrons, which typically have energies of several tens of volts, must propagate into the 1–3 T field found in the main trap. Hundreds of microamps of current are emitted by the gun, but only about 0.1–1 μA reaches the trap region, likely because of magnetic mirroring effects. The field increase has the beneficial effect of directing the beam towards the trap axis, partially compensating for misalignments between the trap mechanical and magnetic axes.

2.9. The microwave injection systems

Microwaves injected into the electrode stack are used as both a diagnostic tool and as a direct probe of trapped antihydrogen atoms. In both cases the frequencies employed lie in the Ka-band of the microwave spectrum, and are typically in the range 28–30 GHz. This corresponds to the range in which the electron cyclotron resonance frequency and the antihydrogen positron spin flip frequencies occur in a magnetic field of order 1 T.

Two independent systems are used for injecting microwaves from an external source into the trapping apparatus. The first consists of a vertical length of WR-28 waveguide that is attached to (and moves with) the vacuum manipulator, and which feeds a 20 dB gain rectangular microwave horn mounted above the MCP and pass-through cylinder electrode, as shown in Fig. 12. The axis of this internal feed horn is horizontal, and can be aligned with the axis of the electrode stack by adjusting the height of the vacuum manipulator. Once in place, the axial distance between the 3 cm-by-4 cm aperture of the horn and the first transfer electrode (Section 2.4.4) is 1 cm. Integrity of the vacuum at the point the waveguide enters the apparatus is maintained using a home-built non-magnetic hermetically sealed quartz window. The width and height of this 3.27 mm thick window are 5.46 and 2.67 mm, respectively, and its thickness corresponds to a half-wavelength at the frequencies of interest. It is glued (using Torr-Seal) into a closely fitting rectangular hole machined in a 3.27 mm thick stainless steel plate. A secondary seal is provided by a 0.127 mm thick Kapton film that is bonded to the face of an immediately adjacent waveguide flange. The measured microwave power transmission through the vacuum window and the ~75 cm length of waveguide leading to the internal feed horn is 0.79 at 28 GHz. The practical low frequency limit for microwave injection via the internal feel horn is approximately 22 GHz, which corresponds to the cutoff frequency for WR-28 waveguide.

The second microwave injection system consists of a pair of metal mirrors. One of these is a planar stainless steel mirror mounted at the bottom end of the vacuum manipulator, and which can thus be translated in the vertical direction and aligned with the bore of the electrode stack. It is oriented such that vectors normal to its surface lie in the horizontal plane and make a 45 degree angle with respect to the trap axis. The surface area of this internal reflector is 42 cm². The second mirror is a 5×10^2 cm² section of the inner surface of a prolate ellipsoid of revolution with semi-major and semi-minor axes of 19.5 and 17.0 cm, respectively. This convex surface is machined from a single block of aluminum and has a surface area that corresponds to approximately 1/8 that of the full spheroid. This external reflector is located outside the trapping apparatus within line-of-sight view of the internal

reflector via a 9 cm diameter glass viewport mounted in a CF100 vacuum flange. It is positioned so that both of its focal points lie in the same horizontal plane as the axis of the electrode stack. A second 20 dB gain rectangular microwave horn is held in place with its throat being coincident with the closer focal point (the external focal point) of the reflector. The position of the mirror/horn assembly is then adjusted so that the distant focal point (the internal focal point) lies in the vicinity of the internal reflector. Translations of the external reflector/feed horn assembly (and hence the internal focal point) over multiple wavelength distances are facilitated by an orthogonal pair of micrometer screw adjustments. Provision is also made to enable precise rotations of the entire external reflector/feed horn assembly about a vertical axis passing through the internal focal point. Empirical observations show that maximum microwave injection efficiency is obtained when the external reflector is placed so as to produce an image of the feed horn between the internal reflector and the first transfer electrode.

The propagation of microwaves down the electrode stack involves a complex geometry which, from a modelling perspective, amounts to a largely intractable problem. At frequencies of order 30 GHz, where the free space wavelength for electromagnetic radiation is of order 1 cm, the situation bears some resemblance to the propagation of microwaves down an over-moded circular waveguide into which many azimuthal breaks in current paths and several abrupt changes in diameter have been introduced. One anticipates that this structure will support a strongly frequency- and spatially dependent superposition of standing and travelling wave modes, with efficient and frequent conversion between modes. This picture is consistent with transmission experiments performed ex-situ, during which the microwave power transmitted through the entire electrode stack was measured, from a 20 dB transmitter horn to a 20 dB receiver horn. The transmission was observed to vary between -1 dB and -100 dB relative to the incident power over the range 22–40 GHz, with the average being of order -6 dB. Coarse scale (a few GHz) and fine scale (a few tens of MHz) changes in frequency both typically result in changes in transmission efficiency of order 5 dB. Extremely low transmission was observed at certain frequencies corresponding to nulls in the mode that the receiver horn accepts, but the total transmitted power has a weaker dependence on frequency.

Electron cyclotron resonance experiments performed in-situ (see Section 4.9.1) reveal similar behaviour. These experiments, which are performed on electrostatically confined electron clouds, measure the absolute time average microwave electric fields along the axis of the electrode stack, from which estimates of the injection efficiency (and microwave magnetic fields) can be obtained. In addition, they permit one to generate one-dimensional maps reflecting the time averaged microwave electric field along the trap axis with millimetre-scale resolution. These images are obtained through a combination of applying axial magnetic field gradients (analogous to conventional one-dimensional NMR imaging) and physically changing the position of the trapped electron cloud. As expected, these maps show standing wave-like structure in which both fine scale (~ 1 cm) and large scale (tens of cm) variations in amplitude are evident.

3. Detection systems

3.1. The Silicon vertex detector (SVD)

The antihydrogen atoms and the antiprotons are detected by means of a silicon vertex detector. The antiparticles annihilate predominantly on the trap walls, and the annihilation products of the antiprotons can be detected with the silicon vertex detector.
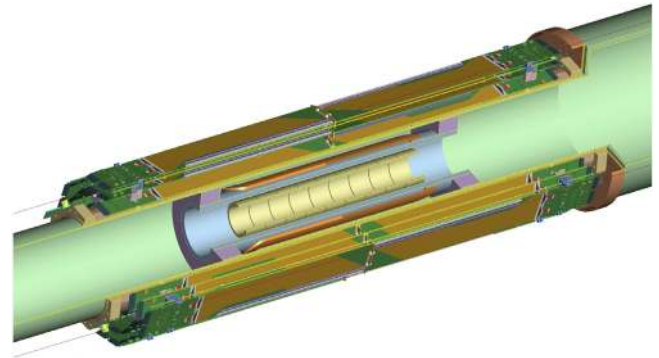
Fig. 13. The ALPHA SVD. The three-layer detector is placed between the beam pipe and the inner bore of the external solenoid (not shown in the figure). The detector is operated at atmospheric pressure and is kept at an operational temperature of 23 °C by cooled, dried and filtered air.

On average, three charged pions are emitted during antiproton annihilation. The pions deposit energy in the SVD. The spatial information of this energy deposition is then used to reconstruct the annihilation event and to determine the vertex location.

The SVD consists of 60 hybrids, assembled symmetrically in three layers in two opposite halves; see Fig. 13. The barrel-shaped detector symmetrically surrounds the neutral trap and it is placed between the outer wall of the outer vacuum chamber (OVC) and the inner bore of the solenoid. Each of the 60 hybrids has two double sided silicon sensors mounted on them. The silicon sensors are double sided p-on-n strip detectors having 128 strips perpendicular to the axial z-direction and 256 strips perpendicular to the ϕ-direction, the strip pitches being 875 µm and 227 µm, respectively. As can be seen in Fig. 13, there is a significant amount of scattering material between the annihilation region and the active silicon layers. The strip pitches were chosen with the help of pion scattering Monte Carlo simulations in order to achieve the optimal configuration and to avoid over/under-engineering. The active area of a single silicon sensor is 58.1×112 mm. The ϕ-strips of the two adjacent silicon sensors mounted on hybrids are wire bonded together and are read out as a single long strip.

The SVD has altogether 30 720 readout channels corresponding to 3.9 million 875×227 µm 'pixels'. Each hybrid is equipped with four 128-channel Va1Ta Application Specific Integrated Circuits (ASICs) [26]. The ASICs produce a fast trigger signal and a slow, typically 1 µs, shaping time analogue signal. The details of these hybrids and the detector mechanical mount are given in Ref. [27]. The SVD signals are guided through a Front end Repeater Card (FRC) array and the trigger signals are monitored by a Timing, Triggering and Control unit (TTC). If a positive trigger decision is made, a hold signal is sent to the ASICs after which they are read out to analogue to digital converters.

The primary use of the SVD is to locate single annihilation events within a given time window, e.g. to monitor antiprotons and antihydrogen atoms that escape the neutral trap. It is also used to diagnose particle losses during different cycles in the experiment. In addition, the trigger information can be used for counting purposes. The achieved vertex reconstruction resolution is 7–8 mm depending on the projection. The detector can be read out at 470 Hz and the overall annihilation reconstruction acceptance is $\sim 60\%$ taking into account the rejection of bad vertices and the trigger efficiency. The pion track identification and the vertex reconstruction are discussed in Sections 4.7 and 4.8.

3.2. The external scintillators

There are three sets of scintillator paddles arranged along the side of the ALPHA trap as shown schematically in Fig. 14. Each
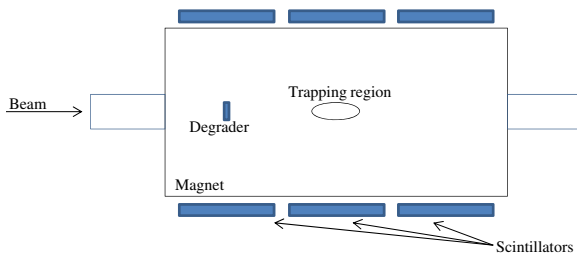
Fig. 14. Schematic showing the positioning of the external scintillators. The degrader is for antiproton slowing and its location is shown in Fig. 2.

paddle is composed of two plastic scintillators, 40 cm wide by 60 cm high, read out by magnetically shielded 5 cm diameter photomultipier tubes. The first paddle is located on either side of the beam degrader, the second paddle straddles the trapping region, and the third reaches beyond it. The purpose of these scintillators is to monitor antiproton annihilation rates at various stages in the trapping and manipulation sequences. The rates in each scintillator, the rate of coincidences in each paddle and the coincidence rate in each pair of paddles are accumulated. This information is used, for example, to cross-calibrate the functioning of various trapping sequences and to monitor such things as the quantity and quality of the antiproton beam delivered by the AD. For the latter, the rate in the scintillator paddles positioned at the degrader provides the necessary information.

3.3. The micro-channel plate (MCP)

A micro-channel plate/phosphor screen system is utilised as a plasma diagnostic device in ALPHA [28]. The MCP used was a type E050JP47 device manufactured by El-Mul Technologies [29]. The active face of the MCP is circular with a 41.5 mm active diameter and is covered with holes 12 μm in diameter spaced by 15 μm in a hexagonal array. The device has a gain of 8×10^5 at the maximum rated applied voltage of 1 kV. The gain behaviour of the MCP was investigated for each of the particle species used in ALPHA (antiprotons, positrons and electrons) over a range of operating parameters [28]. In the 2011 experimental season a new MCP device, manufactured by Photonis, replaced the El-Mul MCP device. In the Photonis MCP, the pore size is 10 μm in diameter with 12 μm center-to-center spacing. The device has a guaranteed minimum gain of a of 1×10^4 at the bias voltage of 1200 V.

The electron shower [30] produced by the MCP is ejected from the back of the plate and accelerated onto a phosphor screen where a signal is measured either through the induced current of the electrons captured by the phosphor screen or by imaging the subsequent emission of photons from the energy deposited using an external CCD camera. In this way, the system can be used either as a low-background charged particle detector or as a precision imaging device, to gain information on the spatial distributions of trapped plasmas. The phosphor type in the original El-Mul device was P47, in the Photonis device, P46.

To measure plasma properties, the trapping potentials are manipulated such that particles are allowed to escape from the trap and are accelerated toward the MCP. Over the majority of the extraction region, charged particles are highly magnetised and will follow magnetic field lines onto the MCP, forming a projected image that preserves information about the spatial profile of the particles in the trap. Since the MCP, located 1.3 m away from the trapping region, sits in the fringe region at a much lower magnetic field (24 mT) these projections can vary for different particle species. Leptons in our system remain tightly bound to the field lines over their entire trajectory, while antiprotons at the end of the trajectory are no longer well-confined and can deviate up to a

centimetre from the initial field lines by the time they reach the plate, depending on the magnetic configuration.

While the MCP is a common detector in systems with electron and ion beams, we have characterised the use of the detector with small numbers ($\leq 10^5$) of antiprotons [28]. Normally the electron cascades in each channel are the result of energy transfer from the kinetic energy of the impacting particle into the channel. However, antiprotons will annihilate with the MCP material creating pions and other secondary species. The byproducts of a nuclear-antiproton annihilation passing through the MCP material induce additional electron cascades in nearby channels, creating much higher charge amplification rates than matter impacts. In addition, we have measured that the impact energy of the antiprotons onto the MCP has much less influence on the observed amplification than with lepton impacts. Clear tracks from annihilation products can be easily observed in many of the images.

4. Methods and processes

The successful production and trapping of antihydrogen atoms necessitated the development of a large number of specialised processes and techniques. A brief description of some of them is given in the following sections. In Section 4.1 the sequencer that runs the experiment is described, along with the amplifiers used to feed the voltages onto the electrodes in the trap. In Section 4.2 we describe techniques that are used in the capture and manipulation of the antiprotons. Manipulation techniques included the transfer of antiprotons into the trap, their cooling with electrons and then the ejection of the electrons, evaporative cooling and compression. In Section 4.3 we describe our techniques for temperature diagnostics and evaporative cooling. In Section 4.4 we refer to the accumulation and manipulation of positrons. Autoresonance, described in Section 4.5, was used to transfer the antiprotons to where they can mix with the antiprotons. Essential to achieving high signal to noise was a fast ramp down of the atom trap, as described in Section 4.6. The reconstruction of events and trajectories of annihilation products from the signals of the silicon detector are described in Sections 4.7 and 4.8. Finally, we describe the microwave techniques used to probe the level splitting of the ground state of the antihydrogen atoms in Section 4.9 and describe how monitoring of environmental parameters was achieved (Section 4.10).

4.1. The sequencer and amplifiers

Timing in the ALPHA experiment is controlled by a digital and analog control timer which we refer to as a 'sequencer'. The full system consists of two major parts. The first handles the triggering of instruments and responds to triggers from external devices. The second portion of the sequencer controls the timing and amplitude of voltages applied to the ALPHA trap electrodes. Fig. 15 shows a block diagram of the system connectivity. The digital sequencer operates by processing a list of sequencer states. Each sequencer state consists of a time for the state to be executed, an array of digital line levels to be output, and an optional list of input triggers required to move on to the next state. The sequencer waits until all trigger conditions for a given state are met.
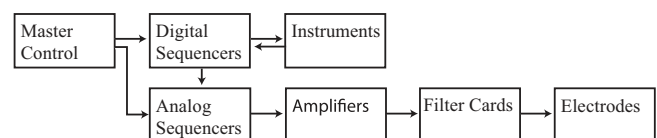
Fig. 15. Block diagram of the sequencer and amplifier setup.

The sequencer output is generated by a deterministically timed state machine, which keeps output generation jitter to the sub-nanosecond level (< 500 ps), and allows 12.5 ns resolution between states. Timing determinism only applies to states which do not wait for external triggers. The sequencer will respond to external triggers with a jitter of approximately 100 ns.

Each digital sequencer has 84 output channels, and 12 input channels. Most output channels behave as 'gates' in the sense that their levels are either HI or LO for the complete duration of any sequencer state, while 4 channels are special 'trigger' channels (level goes from LO to HI and then HI to LO when asserted in a state) that are exclusively used for updating the output of the analog sequencer. The input channels are independently config-ured to trigger on rising or lowering edges (signals transitioning from HI/LO to LO/HI), or respond to 'gate' states (trigger condition met when the input voltage is either HI or LO).

Timing and program logic is controlled by a National Instru-ments PXI-7811R card that is connected to a breakout board which buffers the output from the card using TTL-compatible logic.

The analog sequencer controls the voltages applied to the electrodes. It uses National Instruments PXI-6733 cards to output a pre-generated list of voltages, with output timing controlled by triggers from the digital sequencer. These National Instruments cards have 8 digital-to-analog-converter channels with 16-bit voltage resolution in a ± 10 V range, and can update all channels at a rate of approximately 10^6 s^{-1}.

The output from the DAC cards are fed into high voltage amplifiers, which drive the low pass filter channels that connect to trap electrodes. Fig. 16 shows how the analog channels are connected. We use two types of amplifiers in the traps. One type has an output range of ± 140 V and a bandwidth of ~ 1 MHz, while the other has a voltage range of ± 75 V, with a bandwidth of ~ 2 kHz. The filter channels serve the purpose of mixing high frequency and low frequency signals on the electrodes. The low

pass side of the channel is a passive RC filter with a cut-off frequency of ~ 25 kHz while the high pass side is similar, and cuts off at about 170 kHz.

4.2. Antiproton capture, cooling and manipulation

The Antiproton Decelerator (AD) delivers 100–200 ns long bunches of around 3×10^7 antiprotons at 5.3 MeV about every 100 s. These are further slowed by allowing them to pass through a degrading aluminium foil. The foil thickness has been optimised experimentally, guided by stopping power calculations using the SRIM code [31]. The antiprotons pass through a 12.5 μm stainless steel vacuum window before traversing the final degrader con-sisting of 185 μm of aluminium. This degrader also serves as a Faraday cup for measurements of charged particles.

In preparation to catch some of the energy degraded antipro-tons the HVB electrode (see Figs. 3 and 4) is raised to a given voltage, typically 4 kV. Subsequently HVA is turned on to the same voltage about 50 ns after the antiprotons arrive. The switching time of HVA is triggered by the AD ejection system, plus a delay. The delay is set experimentally by optimising the number of anti-protons caught. The catching efficiency is about 2.4‰ with a high voltage setting of 4 kV in a 3 T solenoidal field, resulting in around 10^5 captured antiprotons from the initial AD output. Fig. 17a shows the catching efficiency as a function of the high voltage applied to HVA for a fixed voltage on HVB.

The catching trap is located inside the main solenoid in a region where an additional inner 2 T solenoid (see Fig. 10 [20]) is installed. The inner solenoid allows the catching operation to be performed in a 3 T field, while having a reduced field in the mixing region to allow for a deeper neutral atom trap. The transverse energy spread of the antiprotons after passing through the degrader is quite large (see below) and the 3 T field reduces the cyclotron radii such that the trapping efficiency is increased. Fig. 17b shows the measured catching efficiency as a function of the total axial magnetic field in the catching trap. The slightly higher efficiency when using the external solenoid might be due to the small focusing effect of the fringe fields on the incoming antiproton beam.

4.2.1. Antiproton cooling, storage and stacking

To make trappable antihydrogen, it is necessary that the antiprotons have energies comparable to the depth of the mag-netic trap. To efficiently cool keV antiprotons trapped after
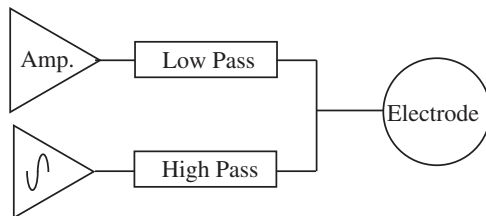

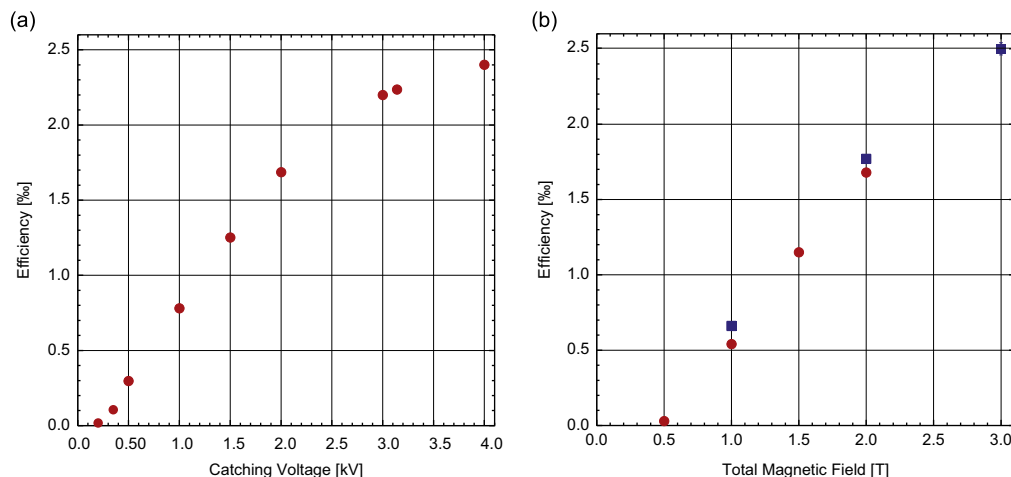Fig. 16. Block diagram of filter channel connectivity.

Fig. 17. Antiproton catching efficiencies. (a) Efficiency as a function of HVA gate voltage for $V_{HVB} = 4$ kV in a 3 T field and (b) efficiency as a function of magnetic field strength for the high voltages set to $V_{HV} = 4$ kV, ● inner solenoid only, ■ inner solenoid plus 1 T from external solenoid.

degrading, the catching trap is first loaded with an electron plasma before the antiprotons are caught. The electrons cool towards equilibrium with the surroundings through the emission of cyclotron radiation in the strong axial magnetic field (the trap system is held at 7–10 K by heat sinking to the cryostat of the superconducting neutral atom trap magnets). The antiprotons then cool through Coulomb collisions with the electrons. The cooling time for an unperturbed electron (of mass m_e and charge e) can be calculated from the Larmor formula for the power radiated by an accelerated charge [32] and is given by

$$\tau_e = \frac{3\pi\varepsilon_0 c^3 m_e^3}{e^4 B^2},\tag{5}$$

which for $B = 3$ T magnetic field gives 0.3 s. This cooling time is appropriate for collisionless plasmas, where only the motion transverse to the magnetic field is affected. Collisions lead also to the cooling of the longitudinal motion and therefore to a longer cooling time, by a factor of 1.5.

The pre-loaded electrons occupy a shallow well that is independent of the catching fields generated by the HV electrodes. Thus, antiprotons cooled by the electrons will fall into this well and be trapped independent of the HV electrode setting. This allows the study of the cooling process by measuring the number of antiprotons in two distinct groups. The first group of hot antiprotons are not cooled by the electrons and holding them requires the HV to be on. Thus, by lowering the HV and counting annihilations their number can be measured, independent of the second group. The latter can be counted by ejecting the particles from the shallow well after the hot antiprotons have been ejected.

The efficiency of the electron cooling depends on plasma density and on the radial overlap between the electron plasma and the captured antiprotons. The size of the antiproton cloud is given by a combination of the divergence and size of the incoming antiproton beam, any blowup that may occur during the capture and cooling process and the cyclotron radius of the antiprotons. The AD beam is focused to give a beam size (2σ radius) of 0.64 mm on the final aluminium degrader. From SRIM we estimate that the average transverse energy of antiprotons caught by our 4.0 keV potential well is ~ 2 keV, which corresponds to a cyclotron radius, in 3 T, of 2.0 mm.

Fig. 18 shows an example of antiproton cooling with a plasma of 20 million electrons having a radius of 0.55 mm. The cooling time constant is about 10 s, and after 80 s no further antiprotons are cooled. The cooling efficiency, i.e. the fraction of caught antiprotons that are subsequently cooled by the electrons, is around 65% when all parameters are optimised. This number can be increased somewhat by using larger electron plasmas, but forming these requires extra time for little gain, and removing the additional electrons
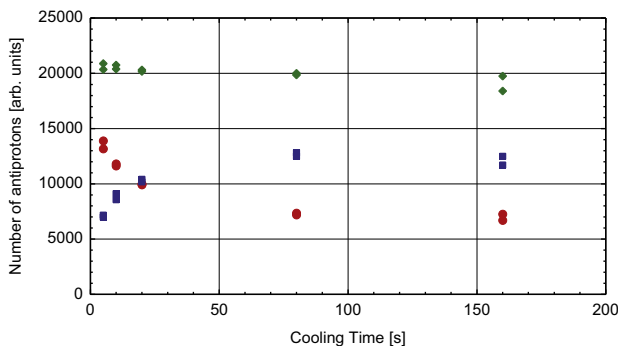

Fig. 18. Electron cooling of antiprotons. The electron plasma was comprised of 20 million particles with a radius of 0.55 mm. The uncooled antiprotons are those released when the high voltage is switched off, leaving only those which have fallen into the trap well containing the cold electrons. ● Uncooled, ■ Cooled, ◆ Total Number.

results in warmer and less-dense antiproton clouds. Thus, 65% is a typical value for the current ALPHA setup.

The lifetime of cold antiprotons is mainly determined by annihilation on residual gas in the system. Depending on vacuum conditions, lifetimes between 2 and 10 h are observed. Using the cross-section for annihilation of antiprotons on H_2 molecules [33] that are assumed to dominate the residual gas we find an estimated H_2 density of $(0.9–4.4) \times 10^4$ cm^{-3} which at 10 K is equivalent to a pressure of $(1–6) \times 10^{-14}$ mbar.

As these lifetimes are long compared to the average time between AD shots, as well as the cooling time of antiprotons in the trap, it is possible to stack large numbers of antiprotons. This is done by lowering the HVA electrode without releasing the cooled antiprotons held with the electron plasma. Up to twenty shots of antiprotons have been stacked with the expected linear increase in the number trapped. However, in standard operating conditions the optimal number of antiprotons is limited by an important experimental constraint. Antiprotons tend to heat the positrons upon injection into the positron plasma in a manner observed to scale with the number injected. It is therefore desirable to operate with a relatively low number of antiprotons and, to date, ALPHA has rarely used more than two stacks of antiprotons for an experiment.

4.2.2. Antiproton compression

For forming cold, trappable antihydrogen it is desirable to have radially small antiproton plasmas for a number of reasons. The transverse magnetic fields of the magnetic minimum trap distort the magnetic field lines such as to introduce a dynamic aperture for the charged particles in the traps. The dynamic aperture takes the form of a critical radius, beyond which field-lines are distorted to intersect with the walls of the trap [34,35]. Furthermore, plasmas with a radial extent comparable to the critical radius can be perturbed and heated by the field inhomogeneities associated with the transverse multipole [19]. These effects mean that it is advisable to maintain ensembles of particles significantly smaller than the critical radius in the neutral atom trap.

In equilibrium, a trapped plasma will rotate around the axis due to the crossed magnetic field and electric fields in the system [36]. The frequency of this rotation depends on the density of the plasma. For antihydrogen formation involving merged plasmas, which is the principal method employed in the ALPHA apparatus, positrons are more numerous and form the densest plasma in question. Antiprotons, which are introduced as a perturbation to the positron plasma, will experience the same fields as the positrons and thus rotate around the axis in the same manner. To be able to capture the nascent antihydrogen atoms they must have a kinetic energy lower than the depth of the trap. Any momentum due to the equilibrium rotation will increase the minimum kinetic energy antihydrogen atoms have, and this rotation is thus a potential limitation to trapping efficiency. As an example, for a constant density (denoted by n_e) cylinder of electrons the equilibrium rotation in a homogeneous magnetic field is [36]

$$\omega_{re}^{\pm} = \frac{1}{2}\Omega_e \left[1 \pm \sqrt{1 - \frac{2\omega_{pe}^2}{\Omega_e^2}} \right],\tag{6}$$

where ω_{re}^- (ω_{re}^+) corresponds to a slow (fast) rotation of the plasma column. The parameters are the electron cyclotron frequency $\Omega_e = eB/m_e$ and plasma frequency, $\omega_{pe}^2 = n_e e^2/\varepsilon_0 m_e$.

In the antihydrogen experiments plasmas rotate at the slow frequency ω_{re}^- which for typical parameters of $B = 1$ T and $n_e = 5 \times 10^7$ cm^{-3} gives an angular rotation rate $\omega_{re}^- = 4.5 \times 10^5$ s^{-1}. For a plasma of radius 1 mm, this would imply a velocity as high as 450 ms^{-1}, for both electrons and antiprotons. An antiproton of

this velocity would have a kinetic energy of 1 meV, which is about 20 times the neutral atom trap depth of 50 μeV (~ 0.5 K). Thus, it is important to keep both the antiproton and the positron plasmas small. The radial extent of a plasma can be controlled using the so-called rotating wall technique [37] where the plasma is exposed to a transverse, rotating electric field. With a sufficiently strong field, and with the frequency of the rotation matched to the plasma rotation frequency, efficient exchange of angular momentum can take place. Thus, by adjusting the frequency of the rotating wall, the plasma rotation frequency can be changed, which changes the plasma radial extent and density [38].

To compress a plasma in practice it is often sufficient to activate rotating wall fields at some frequency higher than the initial rotation frequency of the plasma. Adjusting the amplitude of the rotating wall will change the rate at which the compression occurs [11]. In the ALPHA apparatus the rotating wall is used in this manner to compress both positrons and electrons. The rotating wall causes the plasma to heat but these light species are radiatively self-cooled. Applying the same methods to an antiproton plasma would leave us with hot antiprotons. Further to this we have not been able to compress pure antiproton plasmas, possibly due to their lower density than in experiments where such compression was reported [39]. We therefore compress the antiprotons indirectly by compressing the mixed antiproton/electron plasma. When the compression rate is not too high, the system is observed to equilibrate such that the antiproton radial distribution follows that of the electrons, thus resulting in compressed antiproton plasmas [11]. After compression, the antiprotons re-cool through interactions with the self-cooling electrons. Typical temperatures reached at this stage are a few hundred Kelvin.

With the high densities that are involved ($\sim 10^8$ cm^{-3}), centrifugal separation of the two species is observed [40], causing the antiprotons to be transported to larger radii than the electrons. Ejecting the electrons thus results in a hollow plasma, which quickly fills in Ref. [41]. However, centrifugal separation introduces a lower limit on the final antiproton plasma size, as the canonical angular momentum is conserved when the antiproton plasma fills in. This means that the resulting antiproton plasma will in general be larger in size than the electron plasma.

4.2.3. Antiproton transfer and ejection of electrons

Once the antiprotons have been prepared, i.e. cooled and compressed, they are transferred to the mixing trap by adiabatically changing the potentials such that the plasma is moved one electrode at a time along the axis of the apparatus. In order to be able to cool them further, the electrons are transferred along with the antiprotons during this process. The antiproton losses during this stage are negligible.

To avoid competing processes limiting the effectiveness of antihydrogen formation when the antiprotons and positrons are merged, the electrons are ejected from the plasma before the antiprotons are brought into contact with the positrons. This is accomplished by applying voltage pulses of ~ 100 ns to an electrode constituting one side of a well confining the mixture. Such short pulses briefly open the confining potential on one side, which allows the less massive electrons to escape, while the slower antiprotons are only slightly perturbed. A number of such pulses are usually needed to eject all electrons. By tuning the amplitude of each pulse and leaving the remaining mixture to re-cool, the final temperature of the antiprotons is optimised to be as low as possible. Typical values are in the range 200–400 K.

4.3. Temperature diagnostics

We diagnose the temperature of our plasmas by directly sampling their energy distribution. This involves measuring the charge that escapes when reducing the potential difference confining the plasma axially. Leptons are generally extracted to the MCP, which is biased to maximise charge gain instead of our ability to image the plasma. We then measure the signal of the amplified charge impacting the phosphor screen. Antiprotons can also be extracted to the MCP. However, we have the additional option of simply extracting the antiprotons onto either end of the trap and observing the subsequent annihilation events with the scintillator paddles (Section 3.2). The latter approach has the advantage of not requiring the MCP be moved into place or energised and is our usual choice. If the plasma is in equilibrium, it has been demonstrated [42] that the number N of particles with charge q that escape the trap is approximately related to the change in the confining potential, V_c, by

$$\frac{\partial \ln(N)}{\partial V_c} \approx -\frac{|q|}{k_B T}. \tag{7}$$

There are two major additional factors for which we routinely correct. First, the approximation leading to Eq. (7) is only valid for the very first, highest energy particles to escape from the well, as it neglects the change in space charge due to escaping particles. Both the number of particles available for sampling and the energy range over which they will escape are expected to be proportional to the plasma temperature. Thus, for colder plasmas this approximation can become limited by either the resolution of voltages we apply to the trapping electrodes or by the smallest amount of charge we can effectively measure. We can extend our diagnostic capability deeper into the plasma distribution by applying a numeric model for the change in plasma space charge during extraction.

Second, the potential manipulations required to extract the plasmas necessarily change the shape of the electrostatic wells. Generally, this lengthens the plasma and results in a decrease in the plasma temperature. It also causes the plasma space charge to decrease, as the plasma fills a larger physical volume. This reduces the number of particles that escape for a given change in the confining potential and makes the plasma temperature appear higher. We can numerically calculate the expected influence of these two effects and correct the measured plasma temperature for them. Comparisons to particle-in-cell simulations [43] indicate that, after incorporating all of our corrections, we can determine the temperature of our plasmas to within 20%.

4.3.1. Evaporative cooling of antiprotons

The ultimate temperature of the antiprotons when cooled by electrons is given by the temperature of the electrons. The coldest possible temperature of the electrons is given by the temperature of the electrodes, which is 7–10 K. While temperatures in the vicinity of 20 K have been obtained in a 3 T field with about 10^5 electrons, similar temperatures are not routinely achieved when using 10^6 electrons in the 1 T field in the mixing region. We speculate that radiative heating from higher temperature surfaces, electronic noise and magnetic field inhomogeneities contribute to raising this lower limit. We have not yet identified a way to passively contain large numbers of electrons cooled to the electrode temperature. Additionally, the process of ejecting the electrons from the plasma after the cooling may also contribute thermal energy to the antiprotons.

In order to obtain antiprotons at lower temperatures, we have applied the technique of evaporative cooling, which is well established in the field of cold atom physics (see e.g. [44]), to our plasmas. In essence, the process operates by selectively removing the more energetic antiprotons from the trapped ensemble by lowering the depth of the electrical well confining them. In this manner the average kinetic energy of those that remain trapped is lower and they reach a new equilibrium distribution, via collisions, at a reduced temperature. We provide

a brief overview of this work here. More detailed descriptions have appeared elsewhere [43,45].

The experiments were performed in the mixing trap portion of the apparatus (see Fig. 3 and Section 2.4.2) in which the electrode temperature was typically around 8 K. At the start of the procedure typical antiproton clouds contained around 45 000 particles and had a radius of 0.6 mm and a density of $7.5 \times 10^6 \text{ cm}^{-3}$. The initial temperature of the plasma (as determined in a manner similar to that described in Section 4.3) was around 1040 K, with the antiprotons held in a potential well with an on-axis depth of ~ 1500 mV. To achieve cooling, one side of the confining well was lowered by linearly ramping down the voltage applied to one of the trap electrodes to a pre-determined value. The remaining antiprotons were allowed to re-equilibrate for 10 s before being ejected to measure their temperature and number. The shallowest well investigated had a depth of only (10 ± 4) mV.

Fig. 19 shows temperature fits for six well depths, corrected (downwards by about 16% [43,45]) using a particle-in-cell simulation for effects such as the time-dependent vacuum potentials and the self-fields of the clouds. The lowest temperature achieved in this study was found to be (9 ± 4) K, with around (6 ± 1)% of the antiprotons remaining. The antiproton temperature and the fraction remaining are shown versus the on-axis well depth in Fig. 20a and b respectively. A rate equation model describing the time evolution (as the well is lowered) of the trapped particle temperature, T, and number N was developed [43,45], based upon an earlier cold atom study [46]. These two equations can be written as

$$\frac{dN}{dt} = -\frac{N}{\tau_{ev}} - \gamma N, \qquad (8)$$

and

$$\frac{dT}{dt} = -\alpha \frac{T}{\tau_{ev}} + P. \qquad (9)$$

Here τ_{ev} is the timescale governing evaporation, α the excess energy removed per particle, $\gamma = 1 \times 10^{-4} \text{ s}^{-1}$ per antiproton is a loss term which accounts for annihilation of the antiparticles on residual gas in the trap and P (of order $-dN/dt \times 5$ mK) is a heating term that prevents predicted temperatures from falling below the measured limits. As can be seen from Fig. 20 there is good accord between the model and the experimental data. Further details concerning the relationship between α and the height of the potential barrier and other parameters and on the dependence of
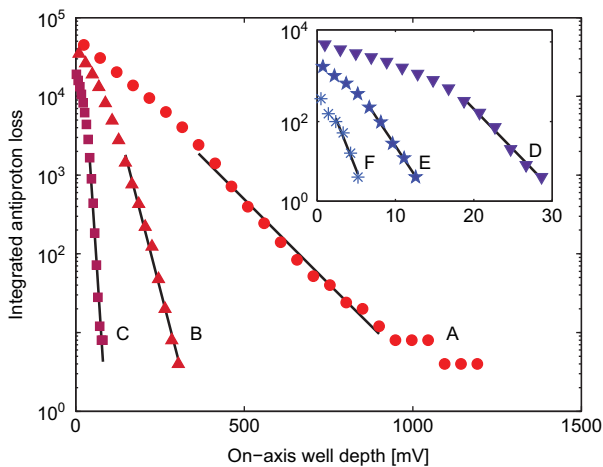

Fig. 20. Antiproton temperature (a) and remnant antiproton fraction (b) versus well depth, with the lines showing results of a model calculation (see text).

τ_{ev} and the relaxation time for antiproton-antiproton collisions are given elsewhere [43,45].

Note that in antihydrogen formation experiments, typical final temperatures were around 40 K. Furthermore, ALPHA has developed evaporative cooling of the much denser positron plasmas [45] used in this work, with similar final temperatures achieved before antiproton-positron mixing was performed.

4.4. Positron accumulation and manipulation

The use of nitrogen buffer gas to achieve positron trapping results in a constant accumulation rate, R, if interactions of the injected positrons with those already trapped can be ignored. As such, the number of trapped positrons at a time t is given by

$$N_{e+}(t) = N_{e+}(\infty)(1 - e^{-t/\tau}). \qquad (10)$$

Here τ is the lifetime of the positrons and $N_{e+}(\infty)$ is the limiting number of accumulated positrons given by,

$$N_{e+}(\infty) = R\tau = I_0 \varepsilon \tau, \qquad (11)$$

where R is the product of the positron beam intensity and the overall capture efficiency into the trap, I_0 and ε respectively. Accumulation curves of the form of Eq. (10) can be used as trap diagnostics, yielding values for the lifetime τ at the ambient nitrogen gas pressure, P. τ is affected by the pressure of the buffer gas, loss as a result of cross-field transport to the walls of the accumulator and also, perhaps, the loss due to annihilation on any background gases. The latter can be eliminated by good vacuum practice, including keeping contamination from hydrocarbons to

Fig. 19. The integrated antiproton loss over time as the well is lowered, plotted against well depth. The curves are labelled in decreasing temperature order as: A; 1040, B; 325, C; 57, D; 23, E; 19 and F; 9 K. The lines are the exponential fits used to derive the final temperatures [43,45].
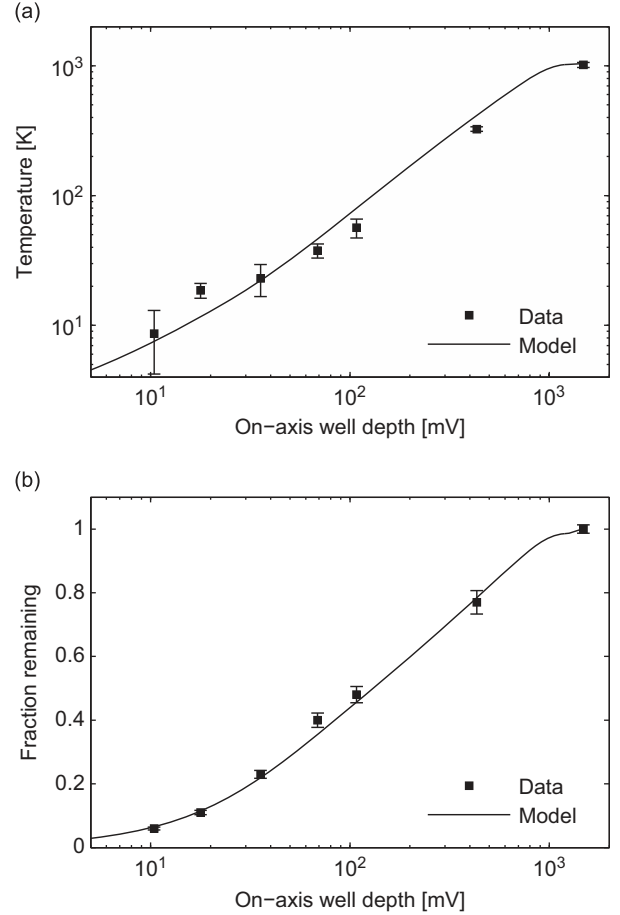
a minimum, as these molecules can have high positron annihilation cross-sections [47]. Assuming the effect of background gas can be ignored, τ is simply related to P as

$$1/\tau = CP. \tag{12}$$

Here C is a constant determined by the positron-nitrogen annihilation cross-section and the magnitude of the cross-field transport in the gas induced via collisions. We note that the latter is effectively stemmed by the application of the rotating wall technique (see below). Plots of $1/\tau$ versus P and the behaviour of $N_{e+}(\infty)$ reveal aspects of the performance of the accumulator. From Eqs. (10) and (11), $N_{e+}(\infty)$ can be written in terms of the branching ratio, f_{ex}, for positron-nitrogen collisions which involve excitation (see Eq. (1)) relative to other process, principally positronium formation (Eq. (2)), which remove positrons from the trap and beam, as

$$N_{e+}(\infty) = f_{ex}I_0(1 - e^{-DP})/CP, \tag{13}$$

where D is a constant related to the total scattering cross-section for positron-N_2 collisions. At high pressures (where $DP \gg 1$), this function behaves as $1/P$, as all the positrons interact in the first stage of the accumulator, whilst it tends to the constant value $f_{ex}I_0D/B$ as $P \rightarrow 0$. Strictly the P in the exponential relates to the pressure in the first stage, whilst that in Eq. (12) is for the third stage. However, these quantities are proportional to one another and C and D account for the constant pressure differential between the stages.

4.4.1. Positron compression in the accumulator

A rotating wall is applied for the entire ALPHA positron accumulation cycle to counteract plasma expansion and loss of particles. The effect on the central density of the positron plasma is seen in Fig. 21, which shows the response to various amplitudes of the rotating voltage, at a selection of frequencies. The central density is a measure of the compression of the plasma, it increases when the diameter of the positron plasma is reduced. This parameter can be measured in dedicated experiments by ejecting the positrons onto a phosphor screen (shown on the right edge of Fig. 6), which can be lowered to intercept the particles. An external CCD camera records an image of the ejected flux, from which the physical dimensions of the plasma may be deduced. The rotating wall field also heats the plasma, and in this case cooling is provided by the nitrogen gas itself, despite the fact that nitrogen is known to be a poor positron cooler [17,48]. As is apparent from
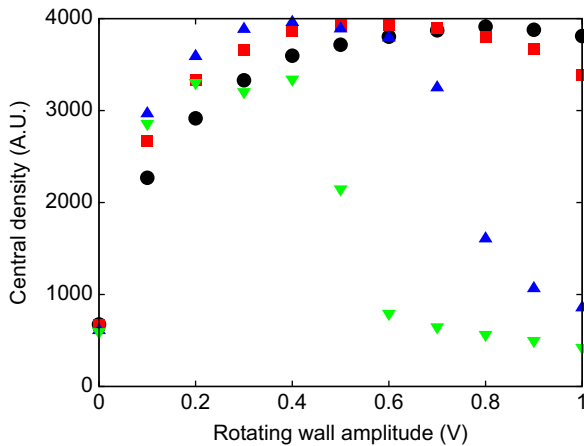
Fig. 21, only low amplitude voltages of the rotating wall could be applied. Higher amplitudes are detrimental particularly at the higher frequencies.

4.4.2. Positron transfer

Once the desired number of positrons has been accumulated they can be ejected from the accumulator and transferred to the ALPHA main magnet system, where the antiproton and mixing traps are located. On completion of accumulation, the buffer gas line is closed and the nitrogen gas is pumped out. After about 25 s, when the pressure in the third stage of the accumulator falls below about 10^{-8} mbar, a valve to the main antihydrogen experiment is opened, the confining potential at the end of the accumulator is lowered and the particles are released. To isolate as far as possible the relatively poor positron accumulator vacuum from that in the cryogenic antiproton/antihydrogen region, the positrons are transported via a 100 mm long tube incorporating a narrow pumping restriction. To prevent loss of the positrons, a 1 T magnetic field is pulsed on for 1 s along the length of the tube.

The positrons are transported from the accumulator to the cryogenic region of the apparatus at a kinetic energy of around 80 eV. The bunch is immediately captured by applying a trapping voltage after approximately 1 μs to an electrode in the main ALPHA trapping region. The positrons are then allowed to cool over a period of around 60 s, held in a trap which is configured by applying appropriate potentials to a pair of electrodes. Once held in the high magnetic field, the positrons cool by the emission of synchrotron radiation. Typically about 50% of the positrons are transferred from the accumulator to the high field trap, resulting in around 100 million positrons, at a density of about 10^8 cm^{-3} (for a primary ^{22}Na source activity of 2.8 GBq), being available every few minutes for antihydrogen production. Once this procedure is complete, it can be repeated if desired.

4.5. Autoresonance mixing for antihydrogen formation

When a swept frequency drive is applied to a nonlinear oscillator, whose frequency varies with its amplitude, the response can become phase-locked to the drive. Under these circumstances, the oscillator's amplitude can be controlled by adjusting the frequency of the drive. This phenomenon, known as autoresonance, occurs in a wide variety of driven nonlinear oscillators from plasma modes [49] to orbital dynamics [50]. We recently demonstrated that autoresonance can be used to controllably excite a cold, dense antiproton plasma [51]. We use autoresonance as an excitation technique for injecting antiprotons into positrons because the excitation of the antiprotons can be made largely independent of the initial conditions of the antiprotons and the positrons. This feature makes the technique robust against small fluctuations in our initial plasmas.

Autoresonance control of an oscillator requires an anharmonic potential that gives a monotonic relationship between the amplitude and response frequency. In the case of antiprotons electrostatically confined next to a positron plasma, the oscillation in question is the longitudinal bounce frequency ω_b of an antiproton with axial velocity $v_z(z)$ and energy $U = m_p v_z^2/2 - e\Phi(z)$. τ_b, the time it takes the antiproton to traverse the well is given by:

$$\tau_b(U) = \frac{\pi}{\omega_b(U)} = \int_{z_l}^{z_r} \frac{dz}{|v_z|}. \tag{14}$$

Here, z_l and z_r are the left and right turning points. Note that the presence of positrons significantly alters the electrostatic potential, $\Phi(z)$. In order to accurately calculate the full potential, the positron plasma must be characterised by measuring its density profile and particle number [28], and its temperature [42], and then

Fig. 21. Behaviour of the central density (see text) of the positron plasma in the accumulator versus the amplitude of the applied rotating wall voltage at various frequencies as: ● 600 kHz; ■ 700 kHz; ▲ 800 kHz and ▼ 900 kHz.

self-consistently solving the Poisson-Boltzmann equation to determine the charge distribution along the axial coordinate z [52]. Given that there are typically far fewer antiprotons than positrons in the experiments in ALPHA, we neglect the influence of antiprotons on the potential. Fig. 22 shows both the self-consistent and vacuum potentials used during a typical antihydrogen production sequence, as well as the bounce-frequency curves for the antiprotons. As a rule, the frequency decreases with amplitude for antiprotons in our nested well.

Autoresonant injection takes place by applying a swept-frequency drive of the form $\overline{\varepsilon} \sin \omega t$ to one of the confining electrodes thus creating a nearly uniform oscillating electric field across the antiproton plasma. The drive is started above the linear, small-amplitude resonance frequency, and swept down to a frequency corresponding to the desired final energy. The sweep-rate α and drive amplitude $\overline{\varepsilon}$ must be chosen so as to exceed the chirp-rate/amplitude threshold in order for the antiproton oscillation to phase-lock to the drive [50,51,53].

During the injection process, there are several important energy scales. ΔU_i is the potential energy confining the positrons: this energy scale needs to be overcome in order to to inject antiprotons into the positron plasma. In a standard injection scheme, antiprotons are ejected from a nearby electrostatic well at an energy aimed to overcome this barrier [7,54]. Autoresonant injection overcomes this energy scale by exciting the particles up in energy using the swept-frequency drive to a low frequency. Naively, one might assume this scheme would not work, because the energy at which antiprotons would pass into positrons corresponds with a discontinuity at $\omega_b = 0$ rad s^{-1}. However, we experimentally observe that a large fraction of antiprotons are injected before the drive reaches this point. We also measure that the temperature of positrons after antiproton injection increases as we drive to low frequency.

As a consequence of this, we use an optimisation strategy to tune the autoresonance injection final frequency. For a given set of antiproton and positron conditions, we conduct experiments in which we progressively lower the final drive frequency while measuring the number of antihydrogen atoms produced as a function of the temperature of positrons after the injection. We broadly find that the frequency that corresponds to the highest

number of produced antihydrogen atoms also results with the lowest positron temperature after injection.

While the exact dynamics of the autoresonant injection technique are not well-understood at the moment, it presents a number of advantages over the standard side-injection technique. First, the autoresonance technique is more stable against variation in the characteristics of the positron plasma than the standard scheme. In a standard injection, the antiproton injection energy must be matched to ΔU_i, the positron space charge potential in the final well. This quantity scales linearly with any shot-to-shot changes in positron number. If the space charge decreases, antiprotons will not have enough energy to mix with the positrons. If it increases, antiprotons will have too much energy and will fail to form trappable antihydrogen. In autoresonant injection, the final drive frequency is fixed. This final frequency corresponds to an energy difference ΔU_f between the space-charge level positrons and the final oscillator amplitude. Fig. 22 marks these quantities in their respective plots. Because of the nonlinear nature of the amplitude-frequency relationship, ΔU_f is a small, weak function of the positron number. Fig. 23 shows how these two quantities change as a function of relative positron number.
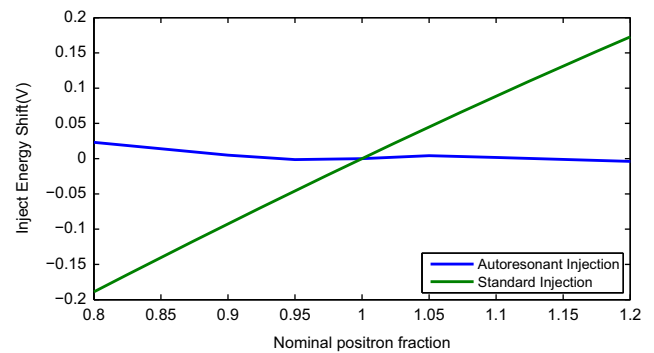
Fig. 23. Energy shift for standard and autoresonant injection schemes. The abscissa is fractional variation from a nominal $N_0 = 2 \times 10^6$ positrons. The energy shift shown is $\Delta U_i(N_0) - \Delta U_i$ for the standard injection and $\Delta U_f - \Delta U_f(N_0)$ for autoresonant injection. Antiprotons with negative energy shifts in the standard injection scheme do not pass through the positrons.
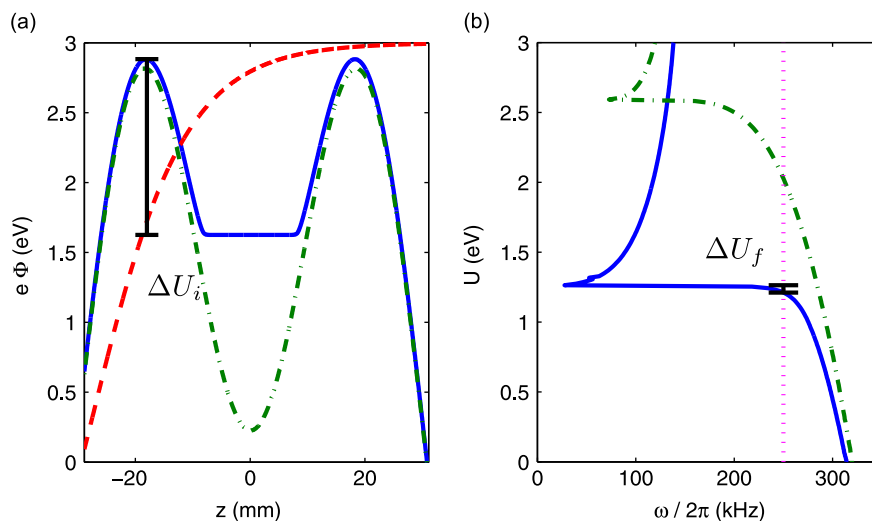
Fig. 22. (a) Potentials of interest for a typical autoresonant drive injection. Green (dot-dash) curve: (fixed) on-axis vacuum potential which confines antiprotons and positrons. Red (dashed) curve: ($100 \times$) potential for an autoresonance drive potential superimposed onto the fixed trap potential. Blue (solid) curve: typical self-consistent potential, with 2×10^6 positrons. Black bar: marks ΔU_i, the positron confinement energy and (b) calculation of antiproton bounce frequency $\omega_b/2\pi$ for the vacuum (green dot-dash) and self-consistent (blue solid) potentials in (a) as a function of total antiproton energy in the confinement region (zero energy is defined to be antiprotons at rest in the left-hand well). Purple (dashed) line marks a typical autoresonant drive end frequency. ΔU_f marks the energy difference between the space-charge confinement and the nominal energy of an antiproton at the end of the autoresonance drive (see text). (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this article.)

Collisions during any injection scheme rearrange the parallel and perpendicular velocity distributions for the antiprotons. This has two negative consequences for forming trappable anti-atoms. The first is simply that collisions can remove antiprotons into the side wells. These antiprotons are then unavailable for antihydrogen formation unless they are re-injected into the positron plasma. The second problem is that collisions, in adding energy to the perpendicular degrees of freedom, will reduce the likelihood of forming antihydrogen atoms that are cold enough to be trapped.

We assess the rate at which parallel energy is scattered into perpendicular degrees of freedom by using autoresonance to drive particles to a fixed energy and then measuring the longitudinal energy distribution at various times after the drive ends. Fig. 24 shows that the time scale for energy redistribution in our typical antiproton plasmas is a few tens of milliseconds, implying that formation should occur faster than this time scale.

An advantage of autoresonant excitation is that antiprotons remain cold in a side well up until the moment the drive is applied. The drive itself is typically around 1 ms in duration, which is far shorter than collisional time scales of the antiprotons. As the energy of antiprotons passing into positrons is quite low, we expect there is less need for collisions with positrons to bring these two species to match their velocities and initiate formation. We observe that formation using an autoresonant drive is immediate and rapid, in contrast to standard side-injection schemes in which antihydrogen formation may take several seconds to begin.

4.6. Atom trap shutdown and quench protection

In order to maximise the probability of detecting rare occurrences of the release of trapped antihydrogen against the cosmic ray background, the ALPHA device was designed with a novel mechanism for de-energising the atom trap magnets. The high currents (up to 800 A in the mirror coils and up to 1100 A in the octupole) can be rapidly diverted to a set of energy extraction resistors using an insulated gate bipolar transistor (IGBT) switch for each circuit. The decay time constant for the currents, about 9 ms (see Fig. 25), is limited by the manufacturer (BNL) specified voltage limit (400 V) across the coils. Thus, the time window in which we search for antihydrogen atoms being released from the trap is chosen to be 30 ms; slightly more than three e-folding times for the currents. The size of the background due to cosmic rays is proportional to the length of this window. Given that comparable systems often require many seconds to de-energise, the ALPHA design has proven to be a distinct advantage. The shut-off circuit also features in the quench protection system (QPS) for the magnets.

The ALPHA-designed and built QPS features the use of field programmable gate array (FPGA) modules to constantly monitor
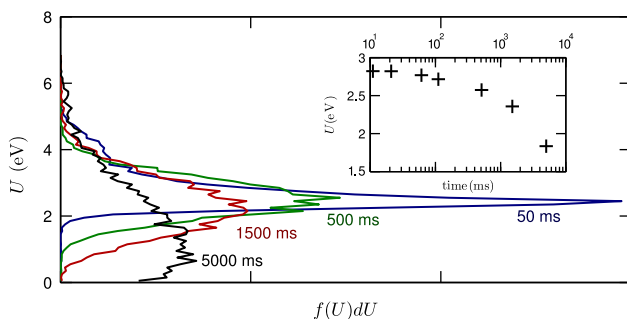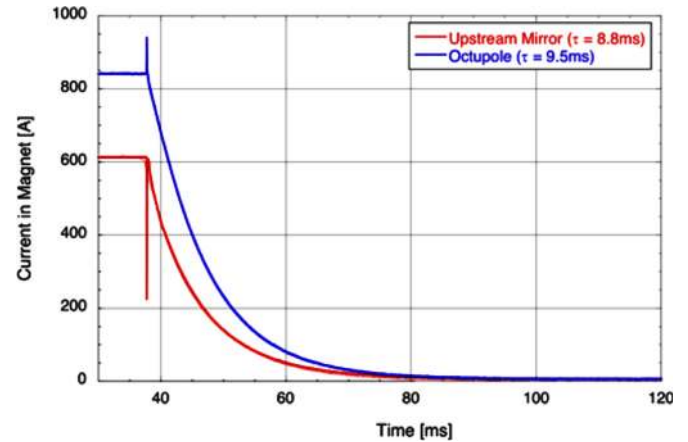
Fig. 25. Magnet quench shutdown comparison of the octupole (τ=9.5 ms) and the upstream mirror (τ=8.8 ms).

voltage taps on the magnet windings and leads for indications of quenches. Tap signals are filtered by LHC prototype signal conditioning modules provided by CERN. A quench indication results in triggering of the IGBT switch, as well as firing of a silicon controlled rectifier crowbar to short-circuit the power supply. The ALPHA QPS is based on a Brookhaven design by Ganetis [55]. The magnet control and interlock system also monitors voltage drops on the vapor-cooled leads for each magnet, temperatures on the water-cooled power semiconductors, and the helium level in the cryostat. The power supplies for the superconducting magnets are switching-type supplies constructed by PowerTen. The QPS also protects the inner solenoid magnet, which is not part of the atom trap.

4.7. Si detector track reconstruction and particle detection efficiency

Signals from the detector system (see Section 3.1) are converted into useable particle track data as follows. Passage of a high-energy charged particle leaves charge deposits in the silicon wafers, which are detected by the on-board electronics in each detector module. Simultaneous signals from one or more p-side strips and one or more n-side strips in the same detector module are defined as a hit. Identification of the specific strips carrying signals, combined with precise knowledge of the position of each module in space, allows for a full three-dimensional identification of the location of each hit. Based on the physical size of the strips, the azimuthal or $R-\phi$ resolution for a single strip signal is 65 µm, while in the axial or z-direction it is 253 µm. On occasion, signals appear in neighbouring strips, corresponding to particles passing through two or more strips in a single wafer. In this case, a weighted average of the strip position and signal strength is used to localise the particle hit, often leading to a smaller uncertainty on the hit location than for the single strip case. Events in which two or more non-adjacent strips on one side of a wafer produce simultaneous signals lead to a situation known as ghost hits, due to the lack of clarity on which p-side signal pairs up with each n-side signal. Separation of real hits from ghost hits is achieved through the track reconstruction process.

As the number of hits resulting from a single antiproton annihilation is small, typically 6–15 corresponding to 2–5 charged particles generated in the annihilation, a brute force approach can be taken to the conversion of hit data into reconstructed charged particle tracks. That is, our analysis approach examines all 3-hit sets, known as track candidates, with one hit in each layer (inner, middle, and outer). The first step eliminates combinations in which hits in neighbouring layers are too spatially separated to correspond to the same particle track. Data from Monte Carlo

Fig. 24. Plot of the time-evolution of the longitudinal distribution $f(U)dU$ for antiproton plasmas with $\sim$ 50 000 antiprotons excited with an autoresonant drive. Inset: Mean longitudinal energy of the distributions as a function of the time between the autoresonant drive and the energy measurement dump.

simulations of our system using GEANT3 with the ROOT Virtual Monte Carlo (VMC) interface were used to determine effective thresholds for this step (see Fig. 26a and b), which eliminated track candidates with axial hit separations exceeding 6 cm or azimuthal hit separations exceeding 20°. All track candidates surviving these cuts are then fit to a helix, based on the fact that the detector is situated in a uniform magnetic field generated by the external solenoid. The resulting helices are then pruned with track candidates that have poor helical fits ($\chi^2 > 15$) or those tracks that do not pass within 1.5 cm of the electrode wall being rejected. As before, a thorough Monte Carlo analysis was used to set these thresholds (see Fig. 26c and d). All track candidates that pass these final cuts are promoted to particle tracks and carried forward to the vertex analysis set, to be described in the following section.

Given this ability to reconstruct tracks, one can now use the detector to monitor cosmic rays to measure the hit detection efficiency of our system. With no positrons or antiprotons in our system, the detector will continue to detect charged cosmic rays passing through its viewing volume. Typically a cosmic ray passing through the full detector and setting off a trigger will produce 5 or 6 hits in the detector (3 hits on one side of the detector and 2 or 3 hits on the opposite side), with 5 hits corresponding to the case where one strip has failed to produce the signal that it should have. A detailed analysis of the frequency of 5 hit passes and identification of the missing hit in the 5 hit line (cosmic rays travel in an effectively straight line through our apparatus) enabled the determination of a hit detection efficiency for each side of each module, as shown in Fig. 27. On average, the p-side efficiency was ∼99.5% and the n-side was ∼98.5%, which when multiplied together with the overall geometric efficiency of our detector (90%) yields an 88% hit detection

efficiency for the entire system. As both cosmic ray particles and annihilation particles act as minimally ionising particles, we conclude that the same hit efficiency applies to antiproton annihilations.

4.8. Event reconstruction and cosmic background suppression

The principal sources of tracks generated through the processes described above are the charged-particles produced in antiproton annihilations, either from antihydrogen atoms or bare antiprotons annihilating on matter, and charged cosmic ray particles. A key goal during the design and fabrication of the detector system was to ensure the ability to distinguish between annihilation events and cosmic ray trajectories. Fig. 28 shows examples of reconstructed particle trajectories for a cosmic ray event and for an annihilation event. It illustrates the primary topological features that are used to separate annihilations from cosmic rays. Cosmic rays typically produce two collinear tracks resulting from a single particle which enters from one side of the apparatus, passing out through the other. A detectable annihilation event produces two or more, non-collinear tracks, one for each charged particle that is generated in the annihilation. The annihilation events we typically observe involve between 2 and 5 charged particle tracks. The first step in distinguishing between cosmic background and annihilation signals is to reconstruct the vertices for each detected event, based on the tracks generated by the procedure described in the previous section. Vertices are then determined as follows:

1. For each event, the number of tracks (N_{tracks}) is identified; if $N_{tracks} = 1$, the event is discarded.
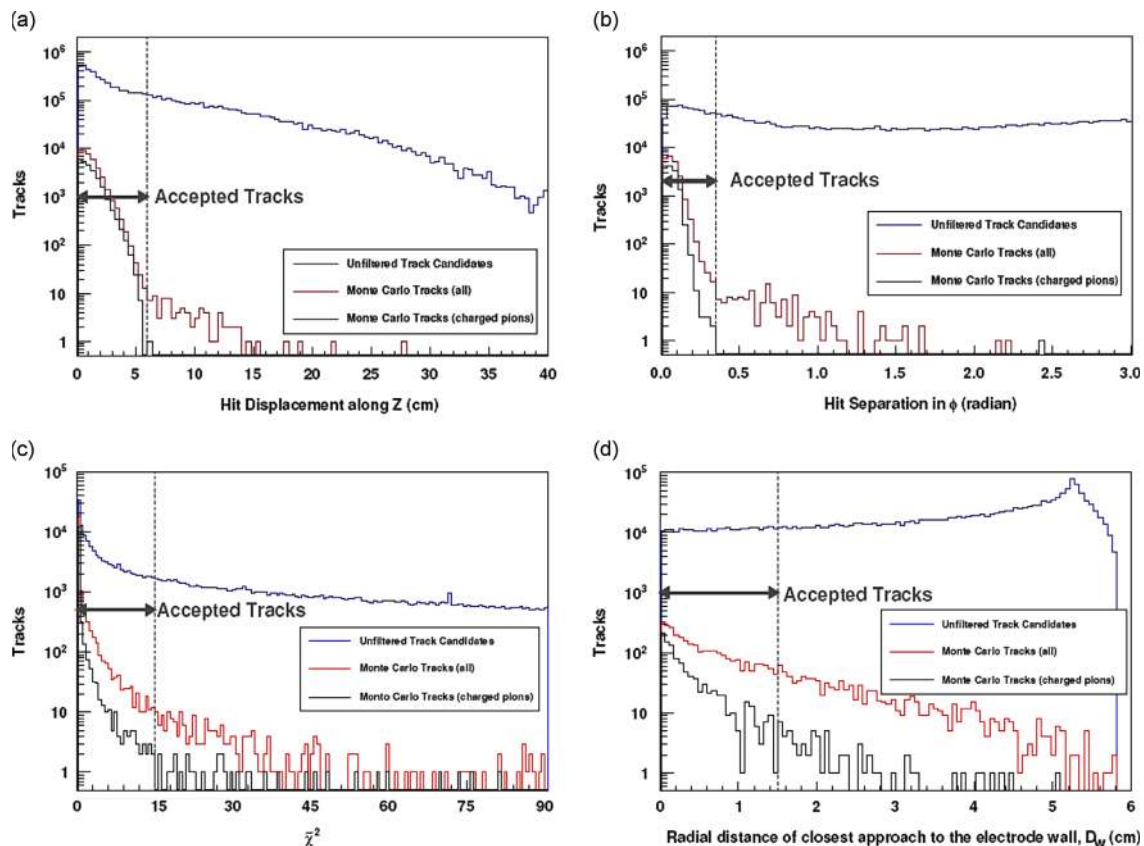
Fig. 26. Illustration of Monte Carlo data used to set the track identification criteria in the ALPHA Apparatus. The provided histograms demonstrate (a) axial hit displacement, (b) azimuthal hit separation, (c) goodness of helical fit, and (d) closest approach to trap electrodes. The blue line illustrates an experimental set of unfiltered track candidates, the red and black lines are the results of Monte Carlo simulations separated into all track candidates and known charged particle tracks, respectively. The vertical dashed lines indicate the selected thresholds. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this article.)
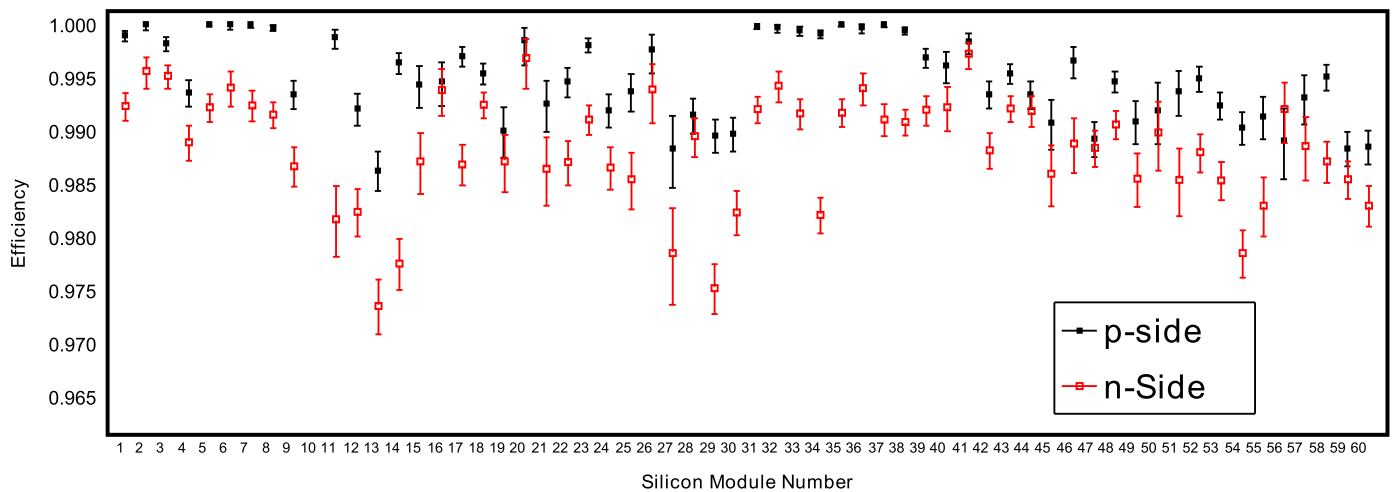
Fig. 27. Illustration of the hit detection efficiency for the p-side (blue) and n-side (red) of each detector module (numbered 1–60), based on an analysis of 5 and 6 hit cosmic ray data. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this article.)
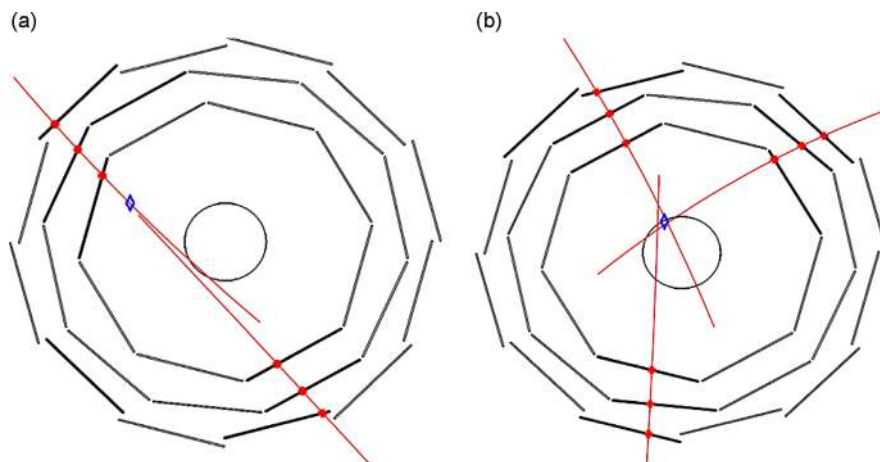
Fig. 28. Axial projection views of track reconstructions from experimental data for (a) an example of a cosmic ray event and (b) an example of an antiproton annihilation. The red dots indicate the location of detector hits, the red lines show the reconstructed trajectories, the blue diamonds show the location of the reconstructed particle event vertex and the inner circle represents the inner diameter of the Penning trap electrodes, which also forms the wall of the neutral atom trap. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this article.)

2. For events with $N_{tracks} \geq 2$, the point of closest approach is determined and the average distance of closest approach (D) for all tracks is calculated; if the event produces only two detected tracks, a single vertex location and a single value for D are determined.

3. For events with $N_{tracks} > 2$, we check to see if the vertex is improved by neglecting one of the tracks. Step 2 is then repeated as many times as necessary (ignoring one of the tracks) until all cases have been considered. The minimum value of D generated by this process is then compared to the value of D for the full set of tracks. If D is not significantly reduced, the result from step 2 is used to determine the vertex location and value for D. If D is reduced by 60% or more the corresponding track is deleted and the new, reduced, set of tracks is accepted as a candidate representation for the outcome of the event. If the number of tracks in the new set still involves more than 2 tracks, then Step 3 is repeated until the result converges, or the set is reduced to two tracks.

This technique yields a finite resolution for vertex location determination. Based on Monte Carlo simulations, the average resolution of our system was determined to be (0.67 ± 0.04) cm in the axial (z) direction, (0.68 ± 0.04) cm in the radial (R) direction and

(0.82 ± 0.04) cm in the azimuthal direction at the trap wall. Once the tracks and vertex locations have been established, one must distinguish between annihilations and cosmic rays, preferably in a manner that is not subject to experimenter bias. With our apparatus this was possible through the use of a blind analysis technique. That is, our parameters for cuts were set to distinguish between cosmic and annihilation signals without inspecting the data from antihydrogen trapping experiments, but rather through the use of training data sets. Specifically, experiments were carried out that examined the signals produced by antihydrogen annihilations and the signals produced by cosmic rays separately, and used these to inform the definition of our cuts, then applied the cuts to the trapping data. The training data sets we used consisted of a) a total of 165 520 events collected during 335 s of antiproton-positron mixing with all trapping fields engaged (the annihilation event training set) and b) a total of 109824 events collected during about 3 h of operation with the trapping fields engaged but no positrons or antiprotons present (the cosmic ray training set). We believe the annihilation training set comprises 98% annihilation event signatures, with the remaining 2% of the events being caused by cosmic rays. This is based on our usual observed cosmic ray detection rate of 10–11 events per second. We expect the cosmic ray training set comprises 100% cosmic ray signatures or noise events.

Extensive investigation of these training data sets enabled us to reduce the selection criteria to three measures: Number of tracks, radial location of the vertex relative to the trap axis, and the residual from a linear fit to the event tracks (δ). The number of tracks was selected as a measure because it effectively splits the data into one of two pools, the $N_{tracks}=2$ pool which is predominantly associated with the cosmic background signal and thus requires strict cuts, and the $N_{tracks}>2$ pool which is predominantly annihilations and thus requires less strict cuts. Following a complete analysis of our training data sets, the following cuts were selected. For $N_{tracks}=2$ events, a vertex radius of $<4\,cm$ and a linear residual of $>2\,cm^2$ was defined as an annihilation event. For $N_{tracks}>2$ events, a vertex radius of $<4\,cm$ and a linear residual of $>0.05\,cm^2$ were required.

Applying these cuts to the cosmic background data set resulted in a $(99.54 \pm 0.02)\%$ rejection rate. This leaves a $(0.46 \pm 0.02)\%$ cosmic acceptance rate, corresponding to an acceptable 0.047 ± 0.002 events per second background signal rate. On the other hand, $(64.4 \pm 0.1)\%$ of the events in the annihilation training set returned a vertex and passed our cuts, which when combined with our 90% trigger efficiency yields an overall annihilation event efficiency of $(58 \pm 7)\%$. Full details on the ALPHA detector system and its analysis and reconstruction systems are provided in a separate publication focused on this topic [9].

4.8.1. Alternative annihilation event acceptance criteria

The level of cosmic ray background suppression obtained using the annihilation event acceptance criteria described above is sufficient for monitoring the release of small numbers of atoms over short observation times, such as the 30 ms post-quench time windows used in our early reports of antihydrogen trapping [1,2]. A very different situation is encountered if one is interested in detecting small numbers of annihilation events over much longer observation times, such as the (multiple) 15 s time windows used in our demonstration of antihydrogen PSR transitions [3]. Observing rare annihilation events over a time window of tens of seconds per run required a reduction of the background level by an additional order of magnitude while retaining a large fraction of antihydrogen annihilations and thus a high signal efficiency. To achieve this, we used a bagged decision tree in the random forest approach [56–58] trained to separate antiproton annihilations on the trap walls from cosmic ray events [3]. The random forest multivariate classifier was chosen for its stability with higher dimensionality, training stability, and insensitivity to input variables with weak discriminating power.

The classifier combines several input variables describing the position of the reconstructed vertices, the hit and track multiplicities, and topological variables defined so as to characterise the shape of the event and its orientation relative to the apparatus. A complete listing and definition of the variables we have employed for this purpose can be found elsewhere [3]. The acceptance criteria for annihilation events are determined via a blind procedure by optimising a sensitivity figure of merit [59]: the optimisation relies on control sample datasets comprising of antiproton annihilations observed during antiproton-positron mixing and cosmic ray events from dedicated off-beam runs. When the criteria so determined are subsequently applied to the analysis of data from experiments in which magnetically trapped antihydrogen atoms are induced to undergo spin flip transitions, we find that the false positive cosmic ray background acceptance rate is reduced by an order of magnitude relative to our standard acceptance criteria.

4.9. Microwave diagnostics

As time progresses, the focus of trapped antihydrogen research will shift to experiments that probe anti-atomic properties. For example, the first such experiment to have been conducted involved the use of microwave radiation to induce antihydrogen positron spin flip transitions, driving ground state anti-atoms from trapped low-field seeking hyperfine levels to untrapped high-field seeking hyperfine levels [3]. This conversion ultimately results in antihydrogen annihilation events as anti-atoms that are ejected from the trap collide with electrodes. The spin flip transitions are resonant quantum processes, requiring application of microwaves at precise frequencies that match the intervals between anti-atomic hyperfine levels (which are in turn set by the magnetic field in which the anti-atoms are located as the transition occurs).

The effort to probe anti-atomic properties will almost certainly be accompanied by the introduction of new diagnostic methods and techniques. A detailed description of these tools at this point in time would clearly be speculative, and well beyond the scope of this manuscript. Nevertheless, a few of the experimental tools and methods that facilitated our recent demonstration of resonant quantum transitions in trapped antihydrogen are likely to play an important role in future studies. A brief summary is thus given below. In particular, we focus on the use of electron cyclotron resonance methods to characterise and calibrate the static magnetic fields and microwave radiation fields that are required for antihydrogen spin flip experiments.

Under the experimental conditions explored to date, the minimum magnetic field in the ALPHA atom trap is set to be of order 1 T. This places the ground state antihydrogen positron spin resonance (PSR) transition frequencies and the electron cyclotron resonance (ECR) frequency in the range 28–30 GHz. This in turn corresponds to the low-end of the Ka-band of the microwave spectrum (26.5–40 GHz).
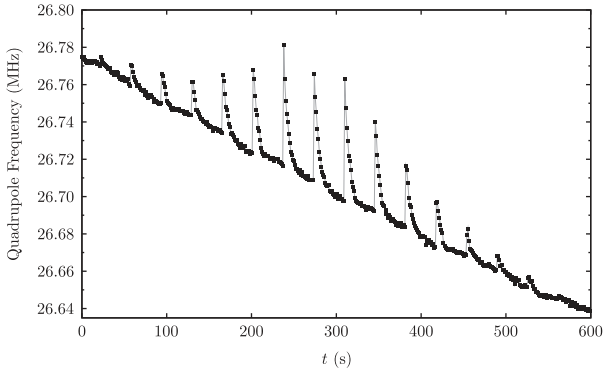
4.9.1. Electron cyclotron resonance

Electron cyclotron resonance experiments are conducted by first loading an electron cloud in the mixing region of the trap. Typical clouds vary between 1 and 4 cm in length, 1.5 and 3 mm in diameter, and contain 8–70 million particles. Pulses of microwave radiation are then injected into the electrode stack, either via the external feed horn and reflectors or via the internal feed horn (Section 2.9). The electric fields associated with these pulses heat the electron cloud when the frequency of the applied radiation matches the electron cyclotron resonance frequency. Between each pulse, we allow the electron cloud to cool to its equilibrium temperature via cyclotron radiation. Typical experiments involve the application of one microwave pulse every 30 s [3,60].

Temperature increases during ECR experiments are inferred by monitoring changes in the frequency of the axisymmetric quadrupolar vibrational mode of the trapped electron cloud. This frequency increases in a manner that is approximately linear with changes in temperature [61]. Vibrational motion of the trapped cloud is induced by applying a high frequency (HF) oscillating potential (typically 26.5 MHz) to an adjacent electrode. This motion, and hence the frequency of the quadrupolar vibrational mode, is then inferred by monitoring the potential induced on the other adjacent electrode. For a given set of conditions, the frequency response of the cloud is calibrated against measurements of its temperature performed using the MCP (Section 3.3).

4.9.2. Calibration and monitoring of static magnetic fields

Calibrations of the static magnetic field B along the trap axis are made by injecting a slow train of 4 μs duration microwave pulses, each at a different frequency. The maximum heating response, which is typically resolved at the level of a few MHz, is then interpreted as the electron cyclotron frequency; i.e. $\Omega_e = eB/m_e$. The evolution of the quadrupole mode frequency during a resonance scan in a uniform magnetic field is shown in Fig. 29.

Fig. 29. Real-time readout of the quadrupole mode frequency of an electron plasma as a train of microwave pulses are applied, scanning across the cyclotron resonance in 1 MHz steps. The quadrupole mode frequency is measured roughly once per second. The overall linear decay of the quadrupole frequency is consistent with the slow expansion of the electron plasma.

Fig. 30. Map of the cyclotron heating as a function of resonance position along the trap axis for a microwave frequency of 28.375 GHz. The diamonds indicate sampled points and the dotted connecting line serves to guide the eye. Temperature increase is measured using the plasma quadrupole mode frequency jump. The vanishing response at either end of the map indicates the extent of the plasma.

This simple picture is complicated by the finite extent of the electron cloud used to probe the field, which introduces sensitivity to inhomogeneous microwave excitation fields and inhomogeneous magnetic fields (such as the intentionally imposed trapping fields). Nevertheless, cyclotron frequency measurements enable us to track the stability of the magnetic field at the centre of the trapping apparatus and reproduce the same experimental conditions at the ± 70 ppm level over extended periods of time.

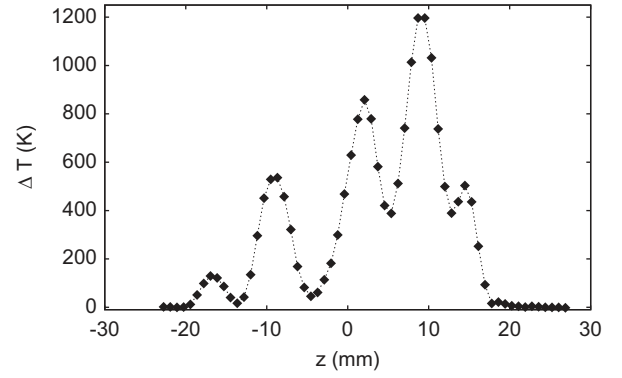### 4.9.3. Calibration of microwave fields

Calibrations of the microwave electric fields that are produced in the electrode stack are performed by injecting 80 ns microwave heating pulses at the ECR frequency. The duration τ_d of these rectangular pulses is short compared to characteristic damping/collisional times in the electron cloud, which leads to temperature increases ΔT that scale as $\tau_d^2 P_m$, where P_m is the injected microwave power. Equivalently, if E_+ is the amplitude of the component of the electric field co-rotating with the cyclotron motion, one finds

$$\Delta T = \frac{e^2 \tau_d^2 E_+^2}{12 m_e k_B}, \tag{15}$$

where k_B) is the Boltzmann constant. Hence, by measuring ΔT one can infer the microwave electric field in situ. This information can then be used to derive estimates of the microwave magnetic fields that are produced in the electrode stack and to infer the overall microwave injection efficiency [3].

4.9.4. Imaging of microwave field distributions

Information regarding the uniformity of microwave fields along the axis of the electrode stack in the mixing region can be obtained by performing ECR heating experiments on long (typically 4 cm) electron clouds while applying a magnetic field gradient. These gradients are generated by partially energising the internal (3 T) solenoid and taking advantage of the fact that its fringing field overlaps the mixing trap. Magnetic field gradients are typically selected so as to be of order 0.093 Tm^{-1} at the centre of the trapped cloud. Long duration (typically 4 μs) pulses of microwave radiation are then applied at the frequency of interest to excite electrons located in a narrow axial slice of the trapped cloud. By varying the static magnetic field in increments, the location of this resonant slice can be scanned from one end of the electron cloud to the other. This information can then be used to generate a plot of temperature increase as a function of axial position (see Fig. 30). Since the density of the electron cloud is approximately uniform over its length, variations in the

temperature response map are primarily due to variations in the time average microwave electric field. One-dimensional maps spanning the entire mixing region can then be generated by concatenating maps for electron clouds held in the centre of each electrode.

4.10. Environmental monitoring

Many parameters of the full ALPHA apparatus are important for both the smooth operation of the experiment but also for determining systematic effects in the experimental results. Anti-hydrogen trapping depends critically on the temperature of the synthesised anti-atoms and is very sensitive to environmental factors such as residual gas pressure and temperatures. It is therefore important that these factors be both controlled and monitored on a continuous basis. Under normal operating conditions, the rate at which the environmental parameters such as temperature and pressure change is relatively slow. Thus, a monitoring system has been established to record and store all of these parameters, with polling intervals between a few seconds for some variables and a few minutes for others.

Our system uses industry-standard Beckhoff BC9000 fieldbus units. These are stand-alone devices with a modular approach to data acquisition and instrument control. Each BC9000 controller can address up to 128 units on its bus. We rarely use more than about 20 at a time. These units are tailored for simple voltage or current monitoring or sourcing, or for interacting with thermocouples, digital signals or even serial ports. The controllers are connected to the intranet of the laboratory and can be interrogated by any computer that addresses them correctly. This setup allows different systems that need to access the associated information to do so directly rather than by having to contact a central server.

For continuous recording and storage, a personal computer based LabView program is used. It polls the various devices at fixed intervals, and then forwards the readings to the MIDAS system. Data acquisition in ALPHA is implemented using MIDAS (Maximum Integrated Data Acquisition System) system [62]. This DAQ accepts and controls input streams for detector data (Section 3) passed through a VME bus, environmental data and MCP images passed through a Labview interface, sequencer data controlling the electrode voltages, magnet currents, and hardware signals, including synchronising changes in these parameters (Section 4.1), and microwave frequency scan information (Section 4.9.1). MIDAS supports multiple electronic logbooks in which information relevant to each run is automatically logged and linked to operator comments and subsequent analyses.

It provides easy access to the historical environmental parameters and data rates, enabling detailed quality of data checks, and generates warnings and protects sensitive equipment from large environmental changes such as power outages and warm-ups.

5. Concluding remarks

We have given a comprehensive description of the ALPHA antihydrogen apparatus that has been used to trap for the first time samples of antihydrogen atoms [1]. The techniques we developed enabled lengthening the trapping periods to 1000 s or more [2]. Recently, we demonstrated the first resonant microwave interactions probing the hyperfine structure of the antihydrogen ground state [3]. We are now in the process of upgrading our capabilities through the construction of a new trapping system, ALPHA-II. The new trap will allow laser access for optical interaction with the trapped antihydrogen atoms. It will enable us to achieve higher precision spectroscopic measurements, including laser measurements that will probe the 1S-2S transition.

Acknowledgments

References

[1] G.B. Andresen, et al., ALPHA Collaboration, Nature 468 (2010) 673.
[2] G.B. Andresen, et al., ALPHA Collaboration, Nat. Phys. 7 (2011) 558.
[3] C. Amole, et al., ALPHA Collaboration, Nature 483 (2012) 439.
[4] G. Parthey, et al., Phys. Rev. Lett. 107 (2011) 203001.
[5] S. Maury, Hyperfine Interact. 109 (1997) 43.
[6] J.H. Malmberg, C.F. Driscoll, Phys. Rev. Lett. 44 (1980) 654.
[7] M. Amoretti, et al., ATHENA Collaboration, Nature 419 (2002) 456.
[8] M.H. Holzscheiter, M. Charlton, M.M. Nieto, Phys. Rep. 402 (2004) 1.
[9] G.B. Andresen, et al., ALPHA Collaboration, Nucl. Inst. Meth. Phys. Res. A 684 (2012) 73.
[10] X. Fei, Trapping low energy antiprotons in an ion trap (Ph.D. thesis), Harvard University 1990.
[11] G.B. Andresen, et al., ALPHA Collaboration, Phys. Rev. Lett. 100 (2008) 203401.
[12] G. Parthey, et al., ALPHA Collaboration, Phys. Rev. Lett. 49 (1986) 1121.
[13] R. Khatri, M. Charlton, P. Sferlazzo, K.G. Lynn, A.P. Mills, L.O. Roellig, Appl. Phys. Lett. 57 (1990) 2374.
[14] T.J. Murphy, C.M. Surko, Phys. Rev. A 46 (1992) 5696.
[15] J.P. Sullivan, J.P. Marler, S.J. Gilbert, S.J. Buckman, C.M. Surko, Phys. Rev. Lett. 87 (2001) 073201.
[16] F. Anderegg, E.M. Hollmann, C.F. Driscoll, Phys. Rev. Lett. 81 (1998) 4875.
[17] R.G. Greaves, C.M. Surko, Phys. Rev. Lett. 85 (2000) 1883.
[18] L.V. Jørgensen, et al., ATHENA Collaboration, Phys. Rev. Lett. 95 (2005) 025002.
[19] J. Fajans, W. Bertsche, K. Burke, S.F. Chapman, D.P. van der Werf, Phys. Rev. Lett. 95 (2005) 155001.
[20] W. Bertsche, et al., ATHENA Collaboration, Nuc. Inst. Meth. Phys. Res. A 566 (2006) 746.
[21] J. Fajans, A. Schmidt, Nucl. Inst. Meth. Phys. Res. A 521 (2004) 318.
[22] ⟨www.bnl.gov/magnets/BEPCII/BEPCII.asp⟩.
[23] G.B. Andresen, et al., ALPHA Collaboration, J. Phys. B: At. Mol. Opt. Phys. 41 (2008) 011001.
[24] ⟨http://www.vectorfields.com/⟩.
[25] A. Dudarev, V.E. Keilin, N.P. Kopeikin, I.O. Shugaev, A.V. Stepanenko, V.V. Stepanov, J. Fajans, D. Durkin, MT-15 (Magnet Technology Conference, Beijing), 1997.
[26] The VA1TA chip, User manual, IDEAS, ⟨http://www.ideas.no⟩.
[27] G.B. Andresen, et al., ALPHA Collaboration, J. Inst. 7 (2012) C01051.
[28] G.B. Andresen, et al., ALPHA Collaboration, Rev. Sci. Instrum. 80 (2009) 123701.
[29] ⟨http://www.el-mul.com⟩.
[30] J.L. Wiza, Nucl. Instum. Meth. 162 (1979) 587.
[31] J.F. Ziegler, U. Littmark, J.P. Biersack, The Stopping and Ranges of Ions in Matter, Pergamon, New York, NY, 1985.
[32] J.D. Jackson, Classical Electrodynamics, 3rd ed, Wiley, New York, NY, 1999.
[33] J.S. Cohen, Phys. Rev. A 56 (1997) 3583.
[34] G.B. Andresen, et al., ALPHA Collaboration, Phys. Plasmas 15 (2008) 032107.
[35] J. Fajans, N. Madsen, F. Robicheaux, Phys. Plasmas 15 (2008) 032108.
[36] R.C. Davidson, An Introduction to the Physics of Nonneutral Plasmas, Addison-Wesley, 1990.
[37] X.-P. Huang, F. Anderegg, E.M. Hollman, C.F. Driscoll, T.M. O'Neil, Phys. Rev. Lett. 78 (1997) 875.
[38] J.R. Danielson, C.M. Surko, Phys. Rev. Lett. 94 (2005) 035001.
[39] N. Kuroda, H.A. Torii, M. Shibata, Y. Nagata, D. Barna, M. Hori, D. Horvàth, A. Mohri, J. Eades, K. Komaki, et al., Phys. Rev. Lett. 100 (2008) 203402.
[40] T.M. O'Neil, Phys. Fluids 24 (1981) 1447.
[41] G.B. Andresen, et al., ALPHA Collaboration, Phys. Rev. Lett. 106 (2011) 145001.
[42] D.L. Eggleston, C.F. Driscoll, B.R. Beck, A.W. Hyatt, J.H. Malmberg, Phys. Fluids B 4 (1992) 3432.
[43] G.B. Andresen, et al., ALPHA Collaboration, Phys. Rev. Lett. 105 (2010) 013003.
[44] H.F. Hess, Phys. Rev. B 34 (1986) 3476.
[45] E. Butler, Antihydrogen formation, dynamics and trapping (Ph.D. thesis) Swansea University, 2011.
[46] W. Ketterle, N.V. Druten, Adv. At. Mol. Opt. Phys. 37 (1996) 181.
[47] K. Iwata, R.G. Greaves, T.J. Murphy, M.D. Tinkle, C.M. Surko, Phys. Rev. A 51 (1995) 473.
[48] M. Charlton, J.W. Humberston, Positron Physics, Cambridge University Press, 2001.
[49] W. Bertsche, J. Fajans, L. Friedland, Phys. Rev. Lett. 91 (2003) 265003.
[50] J. Fajans, L. Friedland, Am. J. Phys. 69 (2001) 1096.
[51] G.B. Andresen, et al., ALPHA Collaboration, Phys. Rev. Lett. 106 (2011) 025002.
[52] S.A. Prasad, T.M. O'Neil, Phys. Fluids 22 (1979) 278.
[53] J. Fajans, E. Gilson, L. Friedland, Phys. Rev. Lett. 82 (1999) 4444.
[54] G. Gabrielse, et al., ATRAP Collaboration, Phys. Rev. Lett. 89 (2002) 213401.
[55] G. Ganetis, Private Communication, Brookhaven National Laboratory, 2003.
[56] L. Breiman, Mach. Learn. 45 (2001).
[57] I. Narsky, Preprint at ⟨arxiv.org/abs/physics/0507143⟩, 2005.
[58] I. Narsky, Preprint at ⟨arxiv.org/abs/physics/0507157⟩, 2005.
[59] G. Punzi, Proceedings of the Conference on Statistical Problems in Particle Physics, Astrophysics and Cosmology PHYSTAT2003, 2003 p. 79.
[60] M. Ashkezari, et al., ALPHA Collaboration, Hyperfine Int. 212 (2012) 81.
[61] M.D. Tinkle, R.G. Greaves, C.M. Surko, R.L. Spencer, G.W. Mason, Phys. Rev. Lett. 72 (1994) 352.
[62] S. Ritt, P. Amaudruz, K. Olchanski, The midas data acquisition system, Midas online: ⟨https://midas.psi.ch/⟩, 1997.