Production of Slow Antihydrogen from Cold Antimatter Plasmas.

A thesis presented
by

Paul Kevin Oxley

to
The Department of Physics
in partial fulfillment of the requirements
for the degree of
Doctor of Philosophy
in the subject of

Physics

Harvard University
Cambridge, Massachusetts
October 2003

Thesis advisor: **Gerald Gabrielse** Author: **Paul Kevin Oxley**

Production of Slow Antihydrogen from Cold Antimatter Plasmas.

Abstract

Several major achievements are detailed in this thesis. Large numbers of antiprotons are captured and cooled to 4 Kelvin in a Penning trap. Antiprotons are provided by the Antiproton Decellerator at CERN, and the accumulation of cold, trapped antiprotons is found to increase in proportion to the number delivered by CERN. As many antiprotons as possible are required to achieve the goal of antihydrogen production.

The physical properties of positron and antiproton plasmas are measured for the first time in the ATRAP apparatus. These findings are essential for interpreting the results of antihydrogen experiments, for accurate calculations of recombination rates, and for the design of future experiments.

The interaction of antiprotons and positrons is induced by the application of radio-frequency drives. During the interaction large numbers of antihydrogen atoms are created and detected with a newly-devised background-free detection technique. A total of 718 atoms are directly detected in a 1 hour trial.

A series of experiments are performed to measure how tightly bound the atoms are, and a theoretical investigation is made of the dynamics of an atom under the experimental conditions. This infers a positron - antiproton separation of $\sim 0.2\ \mu$m, which indicates that the atoms are formed in a three body recombination process, and then collisionally de-excited while inside the positron plasma.

Contents

Publications

1. "Driven Production of Cold Antihydrogen and the First Measured Distribution of Antihydrogen States,"
   G. Gabrielse, N. S. Bowden, P. Oxley, A. Speck, C. H. Storry, J. N. Tan, M. Wessels, D. Grzonka, W. Oelert, G. Schepers, T. Sefzick, J. Walz, H. Pittner, T. W. Hänsch, E. A. Hessels, Phys. Rev. Lett. 89, 233401 on line (2002).

2. "Background-Free Observation of Cold Antihydrogen and a Field-Ionization Analysis of Its States,"
   G. Gabrielse, N. S. Bowden, P. Oxley, A. Speck, C. H. Storry, J. N. Tan, M. Wessels, D. Grzonka, W. Oelert, G. Schepers, T. Sefzick, J. Walz, H. Pittner, T. W. Hänsch, E. A. Hessels, Phys. Rev. Lett. 89, 213401 on line (2002).

3. "Stacking of Cold Antiprotons,"
   G. Gabrielse, N. S. Bowden, P. Oxley, A. Speck, C. H. Storry, J. N. Tan, M. Wessels, D. Grzonka, W. Oelert, G. Schepers, T. Sefzick, J. Walz, H. Pittner, T. W. Hänsch, E. A. Hessels, Phys. Lett. B 548, 140 (2002).

4. "First Positron Cooling of Antiprotons,"
   G. Gabrielse, J. Estrada, J. N. Tan, P. Yesley, N. S. Bowden, P. Oxley, T. Roach, C. H. Storry, M. Wessels, J. Tan, D. Grzonka, W. Oelert, G. Scheppers, T. Sefzick, W. Breunlich, M. Carngelli, H. Fuhrmann, R. King, R. Ursin, H. Zmeskal, H. Kalinowsky, C. Wesdorp, J. Walz, K. S. E. Eikema, T. W. Hänsch, Phys. Lett. B 507, 1 (2001).

5. "Aperture Scraping method for Characterizing Trapped Particle Distributions," P. Oxley, N. S. Bowden, R. Parrott, A. Speck, C. H. Storry, J. N. Tan, M. Wessels, G. Gabrielse, D. Grzonka, W. Oelert, G. Schepers, T. Sefzick, J. Walz, H. Pittner, T. W. Hänsch, E. A. Hessels, Anticipated, 2003.

6. "Positron Cooling of Antiprotons Requires a New Theoretical Cutoff," N. S. Bowden, P. Oxley, A. Speck, C. H. Storry, J. N. Tan, M. Wessels, G. Gabrielse, D. Grzonka, W. Oelert, G. Schepers, T. Sefzick, J. Walz, H. Pittner, T. W. Hänsch, and E. A. Hessels, Anticipated, 2003.

7. "Magnetic Field Dependence of Positron Accumulation via the Ionization of Highly Magnetized Rydberg Positronium," N. S. Bowden, P. Oxley, A. Speck, C. H. Storry, J. N. Tan, M. Wessels, and G. Gabrielse, Anticipated, 2003.

Acknowledgments

Firstly I would like to thank my advisor Jerry Gabrielse, the leader of our ATRAP collaboration, for his continual support and encouragement during my Ph.D., and for allowing me to be involved in such an interesting and exciting project.

The subject of this thesis - production of slow antihydrogen, is an endeavor undertaken by the ATRAP collaboration. The scientists in this collaboration come from four institutions: Harvard University Physics Department, IKP Forschungszentrum in Jülich, Germany, Max-Planck-Institut für Quantenoptik in Garching, Germany, and York University Physics Department in Toronto, Canada. From Harvard I have had the great pleasure of working with two superb postdocs: Cody Storry and Joseph Tan. They have taught me a tremendous amount about how world-class research is done. Cody's never-say-die attitude and his broad experimental experience helped me through many difficult projects, while Joseph's insight into the basic physics underlying our experiments was invaluable.

Graduate student Nathaniel Bowden was my partner in the majority of the work detailed in this thesis. His efforts were absolutely essential to the success of the project and his thesis should be read to complete the account of the experiments performed at CERN in 2001-2002. I am also very grateful for the support and hard work of three other graduate students: Andrew Speck, Melissa Wessels, and David Lesage. I would also like to thank three former graduate students: John Estrada, Peter Yesley, and David Hall who laid the groundwork for the experiments described in this thesis.

The part of the collaboration from Jülich is headed by Walter Oelert. Along with Georg Schepers and Dieter Grzonka he was responsible for creating and running the complex detector system needed to detect antiproton and positron annihilations. It is a testament to their skill and hardwork that in all my time at CERN problems with this system were extremely rare and always quickly fixed. Their good humor and supplies of German chocolate during the night shift were greatly appreciated.

The main role of the three scientists, Jochen Walz, Heiko Pittner, and Max Herrmann from Garching is to provide the laser systems needed to cool and interrogate trapped antihydrogen atoms. Since the atoms have only recently been produced they have not yet had the opportunity to fulfill this role, but instead have provided invaluable input into virtually all parts of the experiments during the last two years. I very much enjoyed working with Heiko on a number of projects and wish him the best for the future.

From York University Eric Hessels and Daniel Comeau were heavily involved with these, and other, experiments at CERN. Eric's clear thinking has guided me on very many occasions and it has been a great pleasure to work with him, while the experiments have benefited considerably from Daniel's willingness to do anything, anytime.

I would like to thank Jan Ragusa at Harvard for her patience and hardwork in sorting out the administrative procedures at both Harvard and CERN. Thank you very much Jan. Also at Harvard, Jim MacArthur provided electronics expertise of the highest quality, for which I am very grateful. I have also had the pleasure to know and become friends with a number of graduate students from the Gabrielse lab at Harvard: Brian Odom, Brian D'Urso, Tanya Zelevinsky, and Dan Farkas amongst them.

On the personal side my father has, as always, been entirely supportive during my Ph.D., as well as helping me understand the finer points of the induction motor and many other electrical systems. I appreciate that I would never have been able to come this far in my studies were it not for him. Finally, I have had tremendous encouragement from my wife Amanda. She has helped me celebrate when things were going well, and commiserate when they were not. I am glad to say that in the last few years we have done a lot of celebrating.

Dedicated to my parents.

List of Tables

List of Figures

Chapter 1

Introduction

This thesis reports the detection of slow antihydrogen atoms [1, 2]. These atoms, and those observed by another team [3], were produced during the positron cooling of antiprotons in a nested Penning trap, a technique which we proposed long ago [4], and were the first to demonstrate [5]. The atoms that our ATRAP collaboration observed are produced in large numbers and detected by field ionization - a technique which avoids all backgrounds. The first investigation of the internal structure of the atoms is made, and theoretical analysis of these experiments [6] determines a positron - antiproton separation of ~ 0.2 μm. This indicates that the atoms are initially formed in a three body recombination process, and then collisionally de-excited while traveling through the positron plasma. The antihydrogen detection experiments are discussed in Ch. 6, and the technique used to bring about the antiproton - positron interaction which leads to antihydrogen, is described in Ch. 5.

The production and subsequent investigation of the antihydrogen internal structure was made possible by a number of advances in antiproton and positron trapping, manipulation, and diagnostics. An intricate cryogenic Penning trap is used to capture and cool antiprotons and positrons, and then combine them to create antihydrogen. Detectors surrounding the trap detect the products of antiproton and positron annihilations, and form an integral part of the apparatus described in Ch. 2.

The simultaneous accumulation of 3×10^5 antiprotons and 1 million positrons for antihydrogen production, both cooled to 4.2 K, is now routine. The positron accumulation technique is the most efficient method of trapping large numbers of positrons directly into a cryogenic UHV environment [7]. The loading rate is as high as 3.5×10^4 positrons $\mathrm{hr}^{-1}\mathrm{mCi}^{-1}$, or 1 million positrons in 25 minutes with a 70 mCi source. From a pulse of 3×10^7 antiprotons delivered by the CERN particle accelerator, 2×10^4 are trapped in flight [8] and cooled to 4 K by Coulomb collisions with cold electrons [9]. Another antiproton pulse from CERN can then be trapped and cooled along with the first. This process, known as antiproton stacking [10, 11, 12], is the only way to accumulate more antiprotons for antihydrogen studies. The techniques used to repeatedly trap and cool antiprotons were developed by the TRAP collaboration [13, 14, 15, 16], and a review of the techniques is given in Ref. [10]. Nearly half a million cold antiprotons have been accumulated in 1 hour using the stacking technique, and given the linearity of the number of trapped, cold antiprotons with the number of pulses delivered by CERN, this record number could be increased.

A new diagnostic determines the physical parameters of the antiproton and positron plasmas confined in the Penning trap [17]. The knowledge of these parameters is essential for interpreting the results of antihydrogen experiments, for accurate calculations of recombination rates, and for the design of future experiments. This is the first determination of antiproton cloud shapes, and the only technique which can be applied to the moderately sized positron plasmas we employ. Numerical modelling of the antiproton plasmas reveals that they are not the spheroid shapes commonly assumed, and these studies are detailed in Ch. 4.

This work is done as part of the ATRAP collaboration. As a result I worked closely with a number of other scientists, most notably another graduate student, Nathaniel Bowden. We operated, maintained and upgraded the cryogenic Penning trap apparatus, devised and interpreted experiments, and took charge of the million-and-one other responsibilities involved with these challenging experiments. This thesis focusses upon parts of the experiments in which I played a leading role: in the studies of antiproton and positron plasmas detailed in Ch. 4, the driving of antiprotons in the nested well discussed in Ch. 5, and the antihydrogen production experiments and analysis included in Ch. 6. Nathaniel's thesis [18] details advances made in positron loading, studies of positron cooling of antiprotons, and production of antihydrogen during this cooling, and should be read to complete the story of the work at CERN in 2001-2002.

1.1 Theoretical Motivation

The theoretical motivation for the production of antihydrogen is to test the CPT theorem [19]. This states that the combined C, P, and T transformation is a basic invariance of nature. C is the charge conjugation transformation (replacing a particle by its antiparticle), P is the parity transformation (the inversion of spatial coordinates), and T is the time reversal symmetry (reversing the direction of motion).

CPT symmetry requires that a particle and antiparticle have identical masses and mean lifetimes, and equal magnitude, but opposite sign, values of electric charge and magnetic moment. As a consequence atoms and anti-atoms should have identical atomic structure - a test of this prediction being the ultimate goal of our ATRAP collaboration.

In 1956 it was noted [20] that the widely held belief that nature was invariant under the P transformation alone had not been tested experimentally. The following year an experiment by Wu *et al* [21] found that positrons emitted in beta decay of ^{60}Co were preferentially emitted in one direction relative to the cobalt nuclear spin. This violates P symmetry since the P transformation reverses the direction of the positron motion but leaves the (pseudovector) nuclear spin unchanged. The P transformed result is never observed in nature, and so the process is a violation of P symmetry.

After this experiment it was assumed that the combined CP transformation was a symmetry of nature, but in 1964 Christenson *et al* [22] found that K mesons in a CP eigenstate of +1 could decay into two pions with a CP eigenstate of -1, thus violating CP symmetry.

Current belief is that the CPT transformation is an invariant of nature. CPT symmetry is on a firmer footing that either P, or CP symmetry since it is a fundamental consequence of all local, Lorentz-invariant quantum field theories of point particles [25], such as the standard model and quantum electrodynamics (QED). However, a more fundamental theory, such as string theory, allows CPT violation, and even within the frameworks of the standard model and QED, CPT violating effects are proposed [23, 24] which are not currently excluded by experiment. The fundamental status of CPT invariance and the active theoretical work being done on the subject make experimental tests of CPT symmetry extremely important.

Another motivation for antihydrogen production is testing the weak equivalence principle for antimatter. Being electrically neutral, antihydrogen atoms avoid the main difficulty involved with gravity studies on charged antimatter particles, which is sensitivity to stray electric fields. Possible tests of the gravitational force on antihydrogen [26] require very cold atoms (~ 1 mK) but such a temperature could nearly be achieved by laser cooling the atoms to the Doppler cooling limit of 2.5 mK.

1.2 Current CPT Tests

Numerous experiments test CPT symmetry. Tests have been made on baryons, leptons, and mesons, and so far none of these have found a violation of CPT symmetry, to the accuracies denoted in Fig. 1.1.

Tests with baryons have come a long way in the last 15 years since antiprotons were first captured in a Penning trap by the TRAP collaboration [8]. A set of three measurements [14, 15, 16] comparing the charge to mass ratio of an antiproton and a proton culminated in a 90 ppt measurement in 1999 [16]. The lepton test is a comparison of the positron and electron magnetic moments [27], while the meson experiment deduced and compared the mass difference of neutral K mesons [28, 29].

The current best measurement of the atomic structure of hydrogen is the 1S-2S transition [30]. This interval has been measured to 46 Hz for an accuracy ($\Delta\nu/\nu$) of 1.8×10^{-14}. If this result can be replicated with antihydrogen confined in a magnetic trap then this would provide an extremely accurate test of CPT symmetry in a

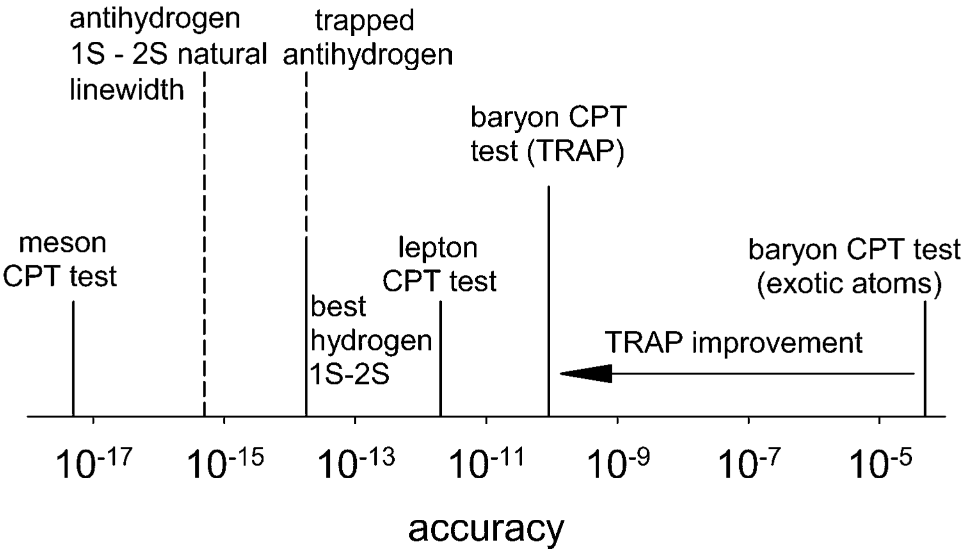

Figure 1.1: Antihydrogen spectroscopy to the level now attained in hydrogen is highly competitive with the most accurate CPT tests in the baryon, lepton and meson systems.

different system (a baryon-lepton system) to those which have been studied before. It has also been proposed [31] that certain 1S hyperfine Zeeman transitions are very sensitive to CPT violating effects, and a measurement of these transitions would provide an additional test of CPT symmetry.

1.3 Antihydrogen Production Methods

Before the production of slow antihydrogen described here, fast antihydrogen was produced by two collaborations [32, 33]. In these experiments a high energy beam of antiprotons within a storage ring strikes a gas target and in the collision process electron - positron pairs are produced. Occasionally a positron will form a bound state with a passing antiproton to make an antihydrogen atom. Unfortunately, the atoms created move with the speed of the incident antiproton, which in this case is relativistic - far too fast to be trapped for precise studies.

Antiprotons and positrons can be trapped and cooled to 0.36 meV (4.2 K) in a cryogenic Penning trap. This allows the possibility to create antihydrogen atoms with the same low energy. A number of different ways to compel a positron and an antiproton to form a bound state have been investigated theoretically [34, 35, 36, 37, 38, 39, 40], and some recombination experiments have been tried with matter particles [41, 42, 43, 44]. Here I will only summarize the two techniques easiest to implement in our apparatus.

Three body recombination occurs when two positrons and an antiproton collide to form an antihydrogen atom and the spare positron carries away the excess energy and momentum from the collision. The process is described by:

$$\overline{p}+\overline{e}+\overline{e}\rightarrow\overline{H}+\overline{e}, \tag{1.1}$$

and occurs with a rate *per antiproton* [34] of:

$$\Gamma_{tbr}(n_{\overline{e}},\ T)=4\times10^{-9}\frac{n_{\overline{e}}^{2}}{T^{\frac{9}{2}}}\ s^{-1}=600\ s^{-1}, \tag{1.2}$$

for a positron density, $n_{\overline{e}} = 10^7$ cm^{-3}, at a temperature, $T = 4.2$ K. This is an extremely high rate which for 2×10^5 antiprotons corresponds to 1 atom every 9 ns. The above rate pertains to a case with no magnetic field present; the application of a strong field reduces the rate by a factor of 10 [35]. The strong temperature dependence has yet to be experimentally verified at very low temperatures. Experiments on cold electron-ion plasmas at $T \sim 1$ K [45] are not fully understood, and suffered from imperfect knowledge of the electron temperature. Recent calculations (not including a strong magnetic field) suggest that the strong temperature scaling is valid down to at least $T \sim 10$ K [46], and that large deviations are not expected at 4 K.

The disadvantage of the three body recombination method is that the atoms are created in states bound by an energy only a few times kT below the ionization limit, since this is the average energy the additional positron can remove while conserving energy and momentum. This means that the atoms are in highly excited atomic states (principal quantum number > 100) which are easily ionized by the fields used to confine the positrons and antiprotons. Sufficient collisional de-excitation of the atom as it continues through the positron plasma [35, 37, 46] is needed to make it deeply enough bound to survive the trap fields.

The radiative recombination process:

$$\overline{p} + \overline{e} \rightarrow \overline{H} + \nu, \tag{1.3}$$

occurs at a much lower rate [34],

$$\Gamma_{rr}(n_{\overline{e}},\ T) = 6 \times 10^{-11} \frac{n_{\overline{e}}}{T^{\frac{1}{2}}}\ s^{-1} = 3 \times 10^{-4}\ s^{-1}, \tag{1.4}$$

for $n_{\overline{e}} = 10^7$ cm^{-3}, $T = 4.2$ K. This low rate is only 1 atom every 16 ms for 2×10^5 antiprotons. However, the cross section for recombination is inversely proportional to the principal quantum number of the recombined state [47], and so atoms are preferentially created in the ground state. They are therefore stable and ready to be confined in a magnetic trap for laser spectroscopy.

Chapter 2

Apparatus

The ingredients of antihydrogen – antiprotons and positrons, are confined and cooled in a cryogenic Penning trap. Positrons are detected non-destructively, via the image charge they induce in the electrodes of the trap, while antiprotons are detected by their annihilation products. The apparatus used is described in this chapter.

2.1 Penning Traps

Penning traps are devices for trapping charged particles. They have been used for decades to make precision measurements of particle properties, such as charge to mass ratios [16], and for studies of plasmas containing a single sign of charge [48, 49, 50, 51, 52]. A thorough review of Penning trap theory is given by Brown and Gabrielse [53].

For the experiments detailed in this thesis the Penning trap is cooled to 4.2 K and the resulting vacuum is less than 5×10^{-17} Torr, as measured in a similar apparatus [14]. Antimatter particles can be confined for months in such an environment without colliding with a background gas molecule and annihilating. Trapped particles come into thermal equilibrium with the 4.2 K environment resulting in a particle energy of only 0.36 meV. Positrons and electrons cool to this temperature in a fraction of a second by emission of synchrotron radiation, while the much heavier antiprotons cool via Coulomb collisions with a cold electron plasma. Such cold antiprotons and positrons are the required ingredients for making cold antihydrogen, and a cryogenic Penning trap is therefore an ideal apparatus for such an endeavor.

2.1.1 Penning Trap Theory

A perfect Penning trap consists of a uniform applied magnetic field (along a direction defined to be the z axis) and an electric potential which varies quadratically in the z and ρ directions ($\boldsymbol{\rho} = x\hat{\boldsymbol{x}} + y\hat{\boldsymbol{y}}$, is the radial coordinate in a cylindrical coordinate system). The magnetic and electric fields, and electric potential are given by,

$$\mathbf{B} = B_0\hat{\boldsymbol{z}}, \qquad \mathbf{E} = \frac{V_0}{d^2}(z\hat{\boldsymbol{z}} - \frac{\rho}{2}\hat{\boldsymbol{\rho}}), \qquad V = V_0\left(\frac{z^2 - \rho^2/2}{2d^2}\right), \tag{2.1}$$

where d is a characteristic dimension of the trap. Charged particles are confined radially by the magnetic field and axially by the electric field. This leads to three independent particle oscillations: an oscillation along the z axis (axial motion) and two oscillations in the $x - y$ plane (magnetron and modified cyclotron motions). The axial oscillation at frequency ω_z depends upon the electric field and the particle charge to mass ratio (Eq. 2.4 in Sec. 2.1.3). The modified cyclotron motion ($\omega_c^{'}$) is a fast oscillation about the magnetic field direction, modified slightly from the free space cyclotron motion (ω_c) by the trap electric field. The magnetron motion (ω_m) is a slow $\mathbf{E} \times \mathbf{B}$ drift oscillation. These frequencies are related by:

$$\omega_m \approx \frac{\omega_z^2}{2\omega_c^{'}} \tag{2.2}$$

$$\omega_c^{'} = \omega_c - \omega_m. \tag{2.3}$$

The motions and their corresponding oscillation frequencies for typical trap parameters are shown in Fig. 2.1.

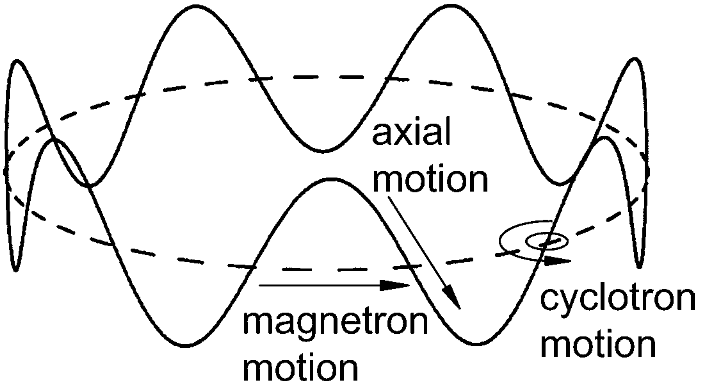

	e^+ and e^-	$\bar{p}$ and p
cyclotron oscillation, ν_c :	151 GHz	82.2 MHz
axial oscillation, ν_z :	30.4 MHz	1.59 MHz
magnetron oscillation, ν_m :	3.06 kHz	15.4 kHz

Figure 2.1: Particle motion in a Penning trap can be decomposed into three independent oscillations. For typical trap parameters: $B_0 = 5.4$ T, $V_0 = 10\ (-10)$ V, for electrons (positrons), and $V_0 = 50\ (-50)$ V, for antiprotons (protons), the frequencies are as shown.

2.1.2 Penning Trap Apparatus

The magnetic field is provided by an NMR grade superconducting solenoid. It can provide a field of up to 6 T (5.4 T is used in these experiments) and has a liquid-nitrogen-cooled 10 cm diameter bore, into which the trap apparatus is lowered. The bore is evacuated to $\approx 10^{-6}$ Torr to thermally isolate the separate vacuum container for the trap. The ideal quadratic potential is approximated to high accuracy by a careful choice of geometry and voltage applied to a set of five hollow cylindrical gold-plated copper electrodes. These are stacked one on top of the other and are separated by macor spacers which provide a gap of 0.18 mm between adjacent electrodes. The choice of a cylindrical electrode geometry allows particles to be introduced easily from both ends of the trap.

Two entirely separate Penning trap apparatuses are used, which, for the purposes of the experiments detailed in this thesis, are identical. They are referred to as hbar1 and hbar2. Here I describe the hbar1 apparatus (Fig. 2.2) which was most often used during these studies. A total of 29 cylindrical electrodes with a 1.2 cm inner diameter make up the whole hbar1 trap. The majority of the electrodes are needed for particle manipulation, antiproton and positron trapping and antihydrogen production experiments. The careful spacing of three sets of electrodes allow the potential to be made exceedingly quadratic in these regions which are used for particle detection (Sec. 2.1.3). The electrostatics of a cylindrical Penning trap are reviewed in detail in Ref. [54]. In addition, there are two tungsten crystals which are used to load positrons (TMOD and RMOD discussed in Sec. 3.2.2), two electrodes which can accommodate a high voltage for antiproton trapping (UPHV and DEG discussed in Sec. 3.3.4), and a set of three special electrodes which make up what is called the ball valve.

The ball valve separates the electrode stack into two parts: the upper part is used to load positrons while the lower part is for antiproton loading. The ball valve is a gold-plated copper sphere with a hole drilled through the center, sandwiched between two copper plates with similar holes. Many turns of copper wire are wound around and then glued onto the ball. A 200 mA current is passed through these wires creating a torque on the wire when it is situated in the 5.4 T field of the superconducting solenoid. Depending on the current direction the torque will either align the hole in the ball with the trap axis (the open position) or at 90° to the axis (the closed position). When the ball valve is open particles can be transferred between the positron and antiproton loading regions of the trap, and when closed the two parts of the trap are divorced from one another. In earlier experiments performed at CERN no ball valve was present. High energy antiprotons from the CERN accelerator could therefore strike the TMOD crystal while antiproton loading was underway. This removed adsorbed gas cryopumped onto the crystal surface, and arrested the positron loading [55] (more details of this effect are given in Sec. 3.2.2). The normal loading rate did not return until the apparatus was thermally cycled. The ball valve prevents antiprotons striking TMOD during antiproton loading.

The RMOD tungsten moderator and a field emission point (FEP described in Sec. 3.1.1) for loading electrons into the lower trap are also situated on the ball but are electrically isolated from each other and the ball itself. When the ball valve is in the closed position the FEP points down the axis of the lower trap and RMOD is aligned along the same axis but in the upper trap. A detailed explanation of the

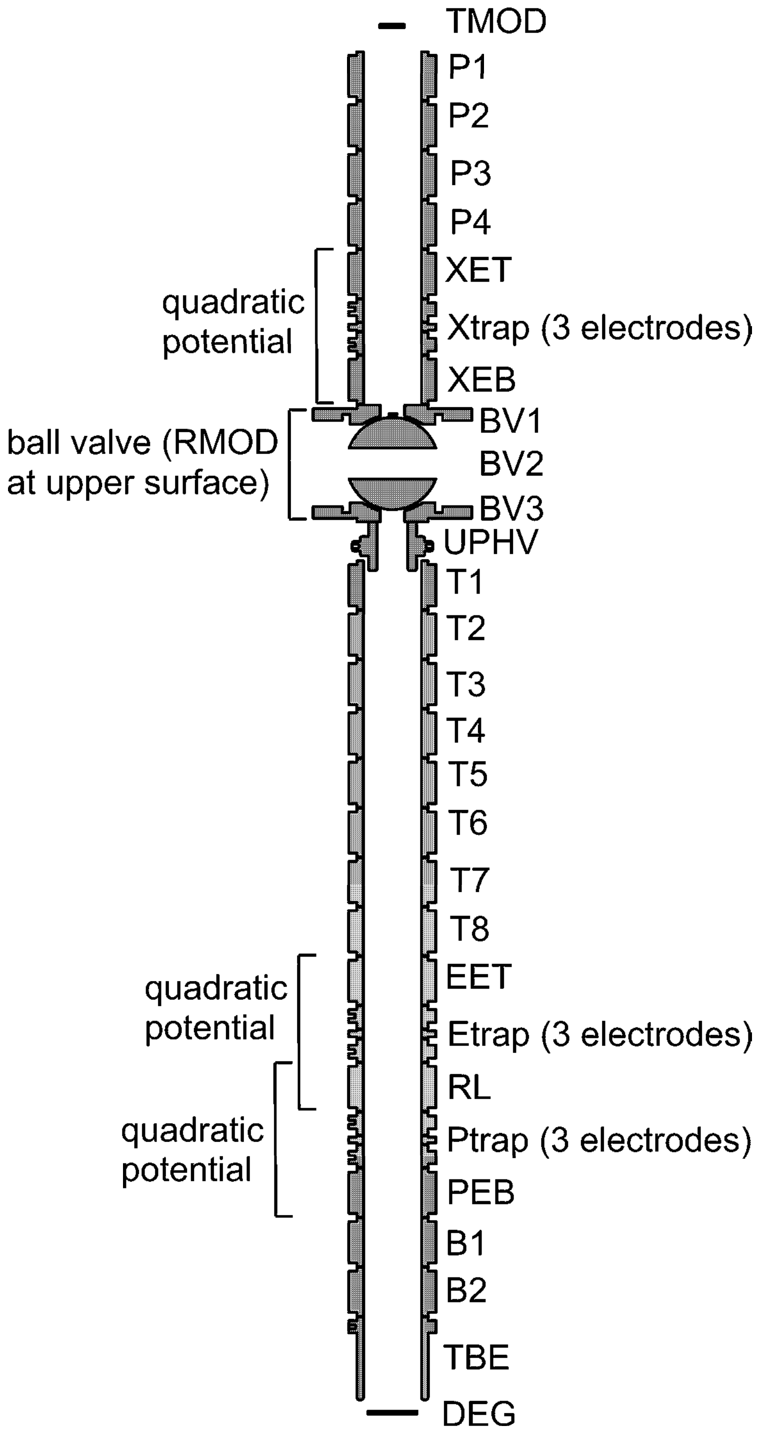

Figure 2.2: The electrodes of the hbar1 trap.

design, construction and operation of the ball valve is presented in [56].

The stack of electrodes is clamped to a copper plate called the pinbase, situated directly above the electrode stack (Fig. 2.3). Into the pinbase are welded a total of 53 cryogenic electrical feedthrough pins which connect on the underside of the pinbase to gold-plated copper straps. The other end of each strap is soldered to an electrode. Surrounding the electrode stack is placed a copper can (the trapcan) which mates with an indium vacuum seal to the pinbase. The trapcan is evacuated and then sealed. There is a 10 μm thick titanium window in the center of the pinbase through which positrons pass to enter the trap. Roughly one third of the incident positrons are absorbed in this window [59]. To enter the trap antiprotons pass, without any significant absorbtion, though an identical window located on the bottom of the trapcan.

Bolted to the top side of the pinbase is a 13 cm tall three-legged copper structure called the tripod, the other end of which is bolted to the bottom of a liquid helium dewar. The three copper legs provide the thermal path to cool the trap to 4.2 K. All electrical circuits for d.c. biassing, d.c. filtering, and radio-frequency drive and detection are soldered to the feedthroughs on the top side of the pinbase. This dense array of circuits covers the entire pinbase area, apart from the central part where the window for positrons is located, and extends upward within the tripod volume. The entire trapcan, tripod, and helium dewar are thermally isolated from the 77 K magnet bore (and detectors of Sec. 2.2) by an aluminium shield.

Above the helium dewar is a thermal isolation stage constructed from G10. The midpoint of this stage contacts with the 77 K magnet bore, while the top of the stage

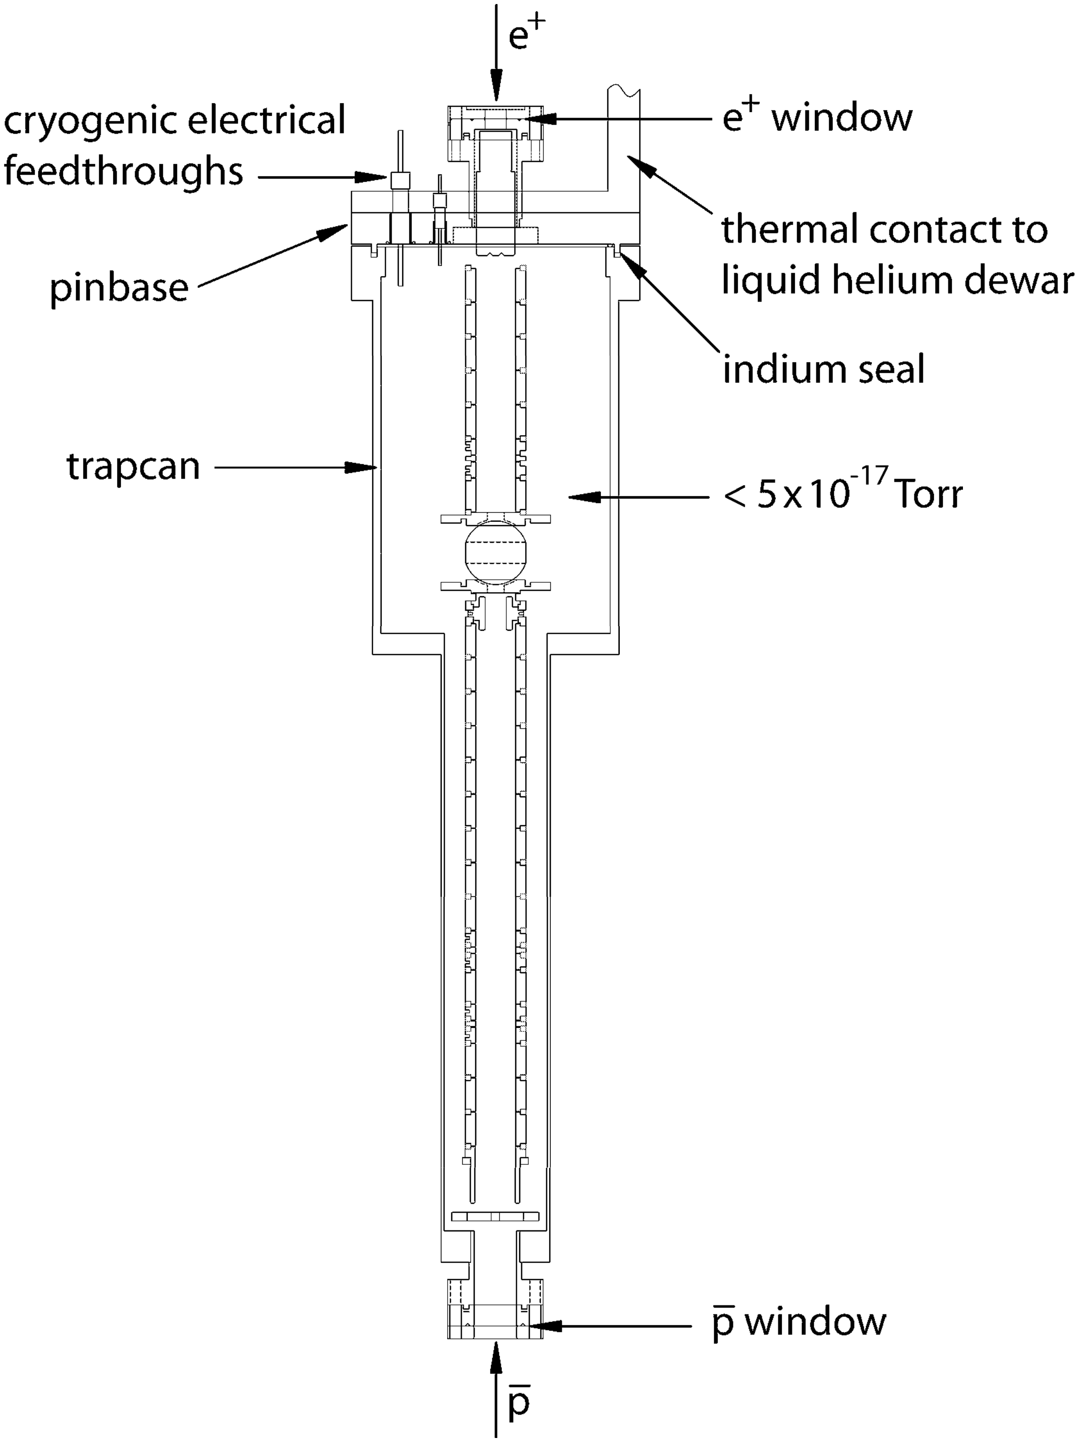

Figure 2.3: The hbar1 trap (for clarity the electrode support structure is not shown).

is at room temperature. More than 100 electrical leads from the cold electronics at the tripod region run up through sealed tubes in the dewar and holes in the isolation stage before emerging at the top of the apparatus from a number of room temperature feedthroughs. These feedthroughs separate the bore vacuum from atmospheric pressure outside the bore. There is also a continuous path along the axis of the dewar and isolation stage to allow access for the positron source described in Sec. 3.2.1. Figure 2.4 shows the apparatus (including parts of the detector system of Sec. 2.2).

2.1.3 Electronics and Non-Destructive Particle Detection

The electrical circuits used to manipulate and detect particles in the trap are described in detail by Estrada [57] and Yesley [56] and will be only briefly summarized here.

Voltage supplies designed by MacArthur [58] provide the d.c. potentials for each electrode. These are carried on double-shielded BNC cables to either an RC or an LC low pass filter before connecting to the room temperature feedthroughs. Constantan wires, 0.1 mm diameter, are soldered to the vacuum side of the feedthroughs, extend down to the tripod region and connect to 1 ms low pass RC filters. The small cross section and poor thermal conductivity of the constantan wires reduces the heat load on the 4.2 K apparatus. The far side of the 1 ms filters solder to the cryogenic pinbase feedthroughs which connect inside the trapcan to copper straps going down to the electrodes.

Radio-frequency signal lines are a.c. coupled at the room temperature feedthrough and proceed to the pinbase as one of a twisted pair of constantan wires. Twisting the

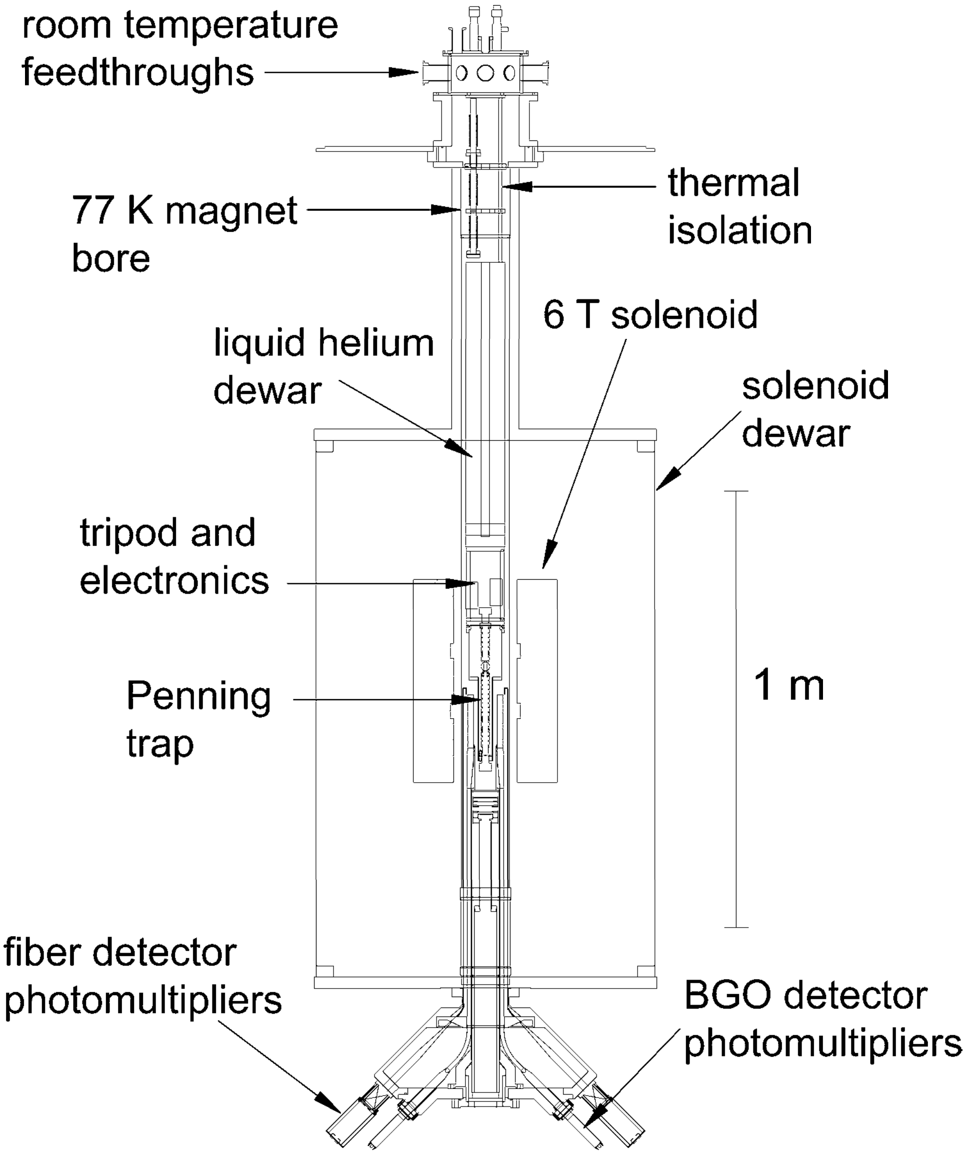

Figure 2.4: The hbar1 apparatus.

signal line around another wire, grounded at both ends, provides some attenuation to electrical pick-up and broadcasting of the signal. At the pinbase the signal line is a.c. coupled to a cryogenic feedthrough. These lines are used to drive and detect the axial and magnetron oscillations of the particles confined in the trap.

Many electrodes also have a "pulse line" connection between the room temperature feedthroughs and the pinbase. This is a small diameter (0.89 mm) stainless steel coaxial cable which allows the transmission of nanosecond pulses without substantial dispersion. At the pinbase a 50 Ω resistor to ground prevents unwanted reflections as the pulse leaves the coax and arrives at a cryogenic feedthrough. The setup and use of these pulse lines is explained in more detail in Sec. 4.1.1. One change to the apparatus described by Estrada and Yesley is an increase from 3 to 10 in the number of pulse lines. This provides for much more flexibility when manipulating particles. The high voltage for the degrader and upper high voltage electrodes use similar coaxial cables as they are high voltage compatible and a very fast change of the degrader potential is required for antiproton trapping. Figure 2.5 shows the typical electrode circuits.

The set of electrodes used to produce a quadratic potential variation along the z axis [54] consists of a "ring" electrode (1.75 mm long) in between two "compensation" electrodes (each 4.72 mm long). An "endcap" electrode (10.08 mm long) next to each compensation electrode completes the set. If a voltage, V_0, is applied to the ring electrode and $0.881V_0$ is applied to the two compensation electrodes, while the endcaps are held at 0 V, then the particles oscillate in the axial direction at a frequency ω_z given by:

$$\omega_z = \sqrt{\frac{C_2 e V_0}{m d^2}}. \tag{2.4}$$

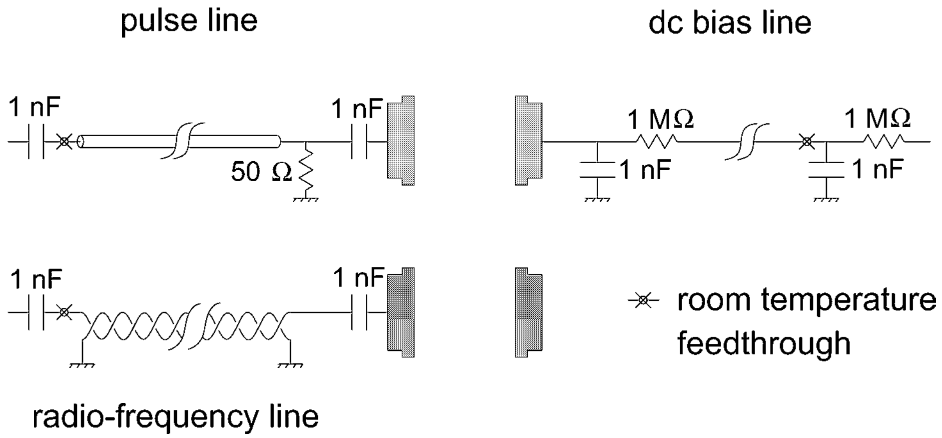

Figure 2.5: Typical circuits for d.c. biassing an electrode and for connecting either a pulse line OR a twisted pair for radio-frequency signals. The symbol used to represent a room temperature feedthrough is noted.

C_2 is a (unitless) geometry-dependent factor, d is the characteristic trap size ($C_2 = 0.545$ and $d = 5.1$ mm for the electrodes used here), and e and m are the particle charge and mass.

Attached across the ring and a compensation electrode is a parallel LC tuned circuit (Fig. 2.6(a)). The inductor has some unavoidable series resistive loss (r) which is a source of thermal noise (V_{therm}) at a temperature of 4.2 K - the temperature of the inductor. The signal voltage, V_{out}, taken from a tap on the inductor, modulates the gate of a cryogenic FET (not shown). The FET drain source current passes through a radio-frequency transformer circuit to match the output impedance of the FET (~ 5 kΩ) to a 50 Ω coaxial transmission line which carries the radio-frequency signal up to the room temperature feedthroughs. The signal is then amplified further before being measured on a spectrum analyzer (Fig. 2.6(b)).

When no particles are present the frequency response of V_{out} is simply the tuned circuit response driven by the thermal noise in the resistor, r. When particles are

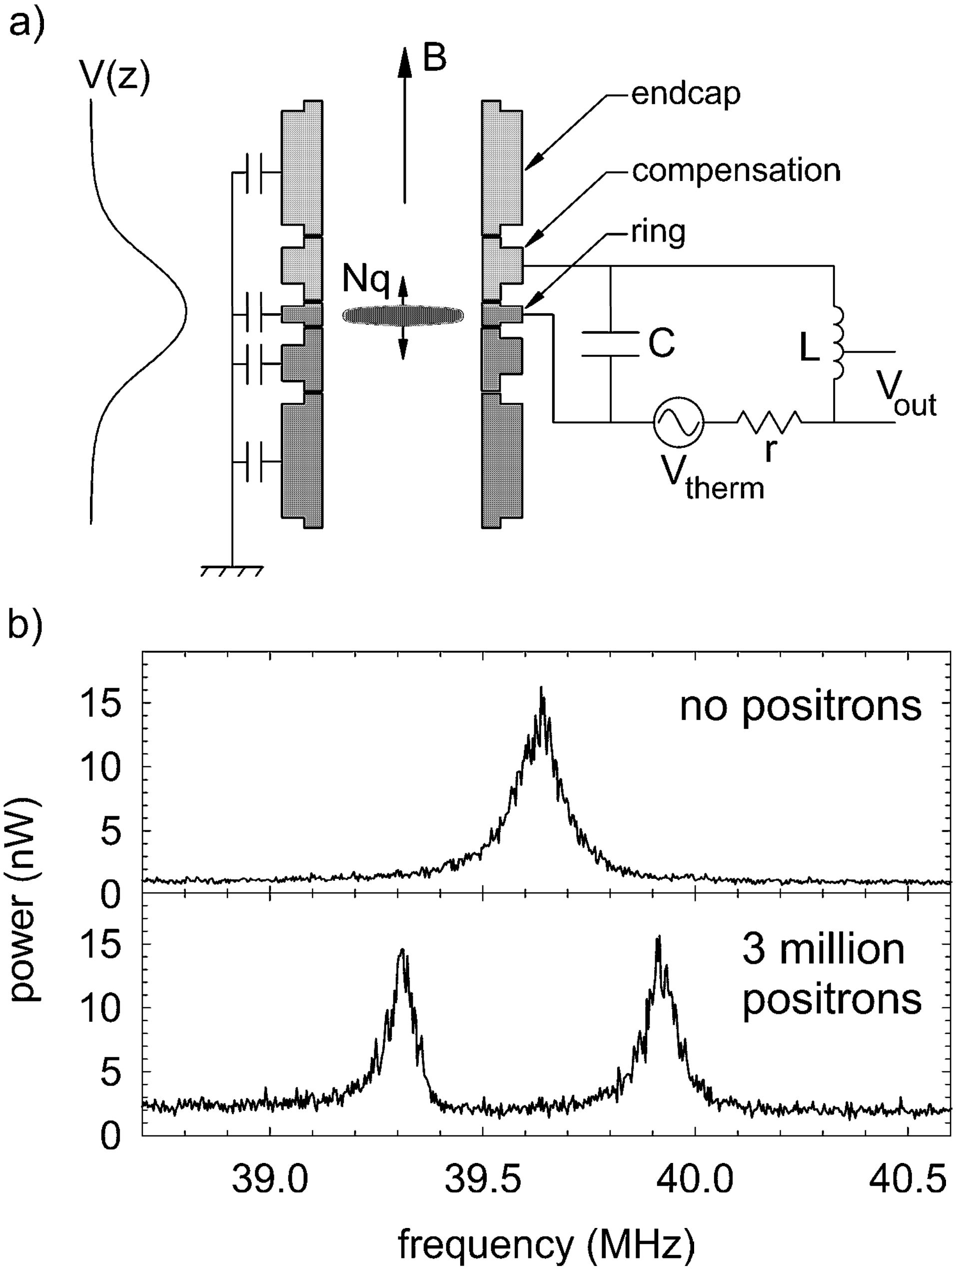

Figure 2.6: a) The electrodes and circuit used to detect the axial oscillation of the N charges confined in a highly quadratic potential. b) The frequency response measured with no particles and with 3 million positrons confined in the trap.

confined on the ring electrode the particle center of mass axial oscillation drives image charges in the electrodes. These image charges couple the particle oscillation to the tuned circuit oscillator. The result is a pair of coupled oscillators and two normal mode frequencies. The difference between the normal mode frequencies is a function of the coupling between the two oscillators and hence is a measure of the number of particles in the trap. A large number of trapped particles drives a large number of image charges which corresponds to a large coupling and a large difference between the normal mode frequencies. This is seen in Fig. 2.6(b) where first no particles and then 3 million positrons are confined in the trap. This particle detection technique works for any charged particle species and is sensitive enough to detect a single particle [59].

2.2 Detectors and Data Acquisition

2.2.1 Outer Scintillator Paddles

A set of 18 rectangular plastic scintillator "paddles" surround the magnet (Fig. 2.7). These are used for detecting charged pions from antiproton annihilations. When a pion deposits energy passing through a paddle a light pulse is produced which is detected by a photomultiplier attached to the paddle. To reduce the background count rate due to electrical noise it is required that two paddles in line with each other and the magnet center (where antiprotons from the Penning trap annihilate) register a pion within 50 ns of each other for an event to be counted. This is called a singles count. The maximum measurable count rate is 2×10^7 s^{-1} (1/(50 ns)) and the background count rate from cosmic rays is 60 s^{-1}. On average an antiproton annihilation results in 3.5 charged pions. This, coupled with the fact that the paddles only cover part of the solid angle and do not have a 100% inherent detection efficiency, gives them an effective antiproton detection efficiency of 47% [60].

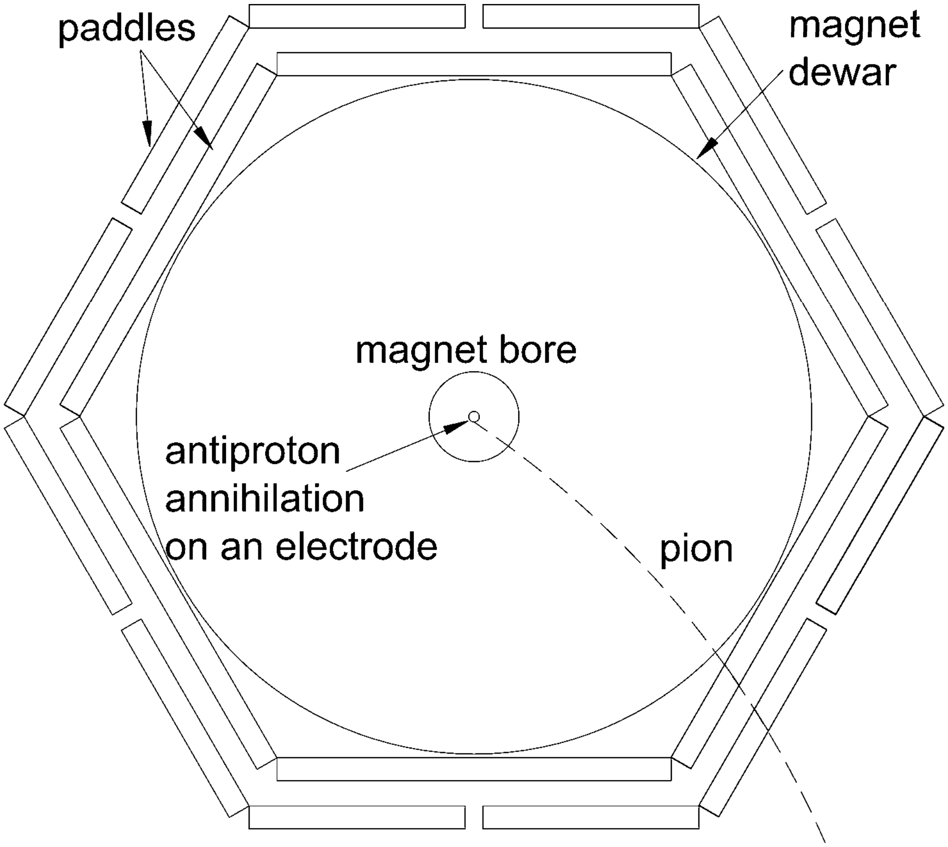

Figure 2.7: Plan view of the scintillator paddles for antiproton annihilation detection. The paddles surround the magnet and each is as tall as the magnet.

2.2.2 Fiber Detector

There are two detectors situated inside the 10 cm diameter magnet bore. Both of these are inserted into the bore from below the magnet and are cooled to 77 K by contact to the bore.

The first is the fiber detector consisting of three layers of 128 scintillating fibers. These are made of the same plastic scintillator used for the paddles and also detect antiproton annihilations. The fibers exit the bore at the bottom of the magnet where they couple to a number of photomultipliers. Since the fibers are much closer to the trap than the paddles they have a much higher detection efficiency, roughly 70%. However, given the relative positions of the fiber detector and the trap electrodes (Fig. 2.8), the solid angle efficiency depends upon where in the trap an antiproton annihilates. If only events which register in at least two of the three fiber layers within a 50 ns window are counted, the background count rate is 20 s^{-1}.

2.2.3 BGO Detector

Also situated in the magnet bore, but at a smaller diameter than the fiber detector, is a set of 12 BGO ($Bi_4Ge_3O_{12}$) crystals. These are 12 cm long and have a trapezoidal cross section so that they can be arranged in a circle with minimal gaps. When the trap is inserted into the magnet bore the lower stack of electrodes sits within the crystals.

A gamma ray passing through a BGO crystal can deposit some of its energy, producing a blue light pulse from the crystal. This pulse is transmitted by plexiglass lightguides to photomultipliers situated at the bottom of the magnet. The energy of the gamma ray will determine the amount of energy deposited in a crystal, which is the measured quantity for each BGO crystal. When a positron and an electron annihilate two back-to-back 511 keV gamma rays are produced and the signature of positron annihilations is therefore a peak in the BGO energy spectrum at 511 keV from two crystals opposite each other.

The light pulse in a crystal is typically integrated for 10 μs to measure the full energy deposited in the crystal. This gives the BGO detector a maximum count rate of 10^5 s^{-1} which is substantially less than that of the paddles and fiber detector. The BGO can itself be used to count antiproton annihilations as is explained in Sec. 5.4. The positions of both the BGO and the fiber detector, relative to the electrodes and to each other are shown in Fig. 2.8.

2.2.4 Trigger Counting and the Data Acquisition System

Very low background counting of antiproton annihilations is achieved if one requires a temporal coincidence of an event in two or more layers of the fiber detector *and* a singles count. This combination has a background of 1.2 s^{-1} which is entirely due to cosmic rays passing through both the outer paddles and the fiber detector. The combination of a fiber and a singles event is called a trigger count because it is used to trigger a more extensive data acquisition system. Whenever this data acquisition is triggered the status of all the detector elements is measured. That is, each of the 18 outer scintillators, all 384 (3*128) fibers, and the 12 BGO crystals. This process takes 1 ms to complete and so the data acquisition system is limited to operating at trigger count rates of 1 kHz or less.

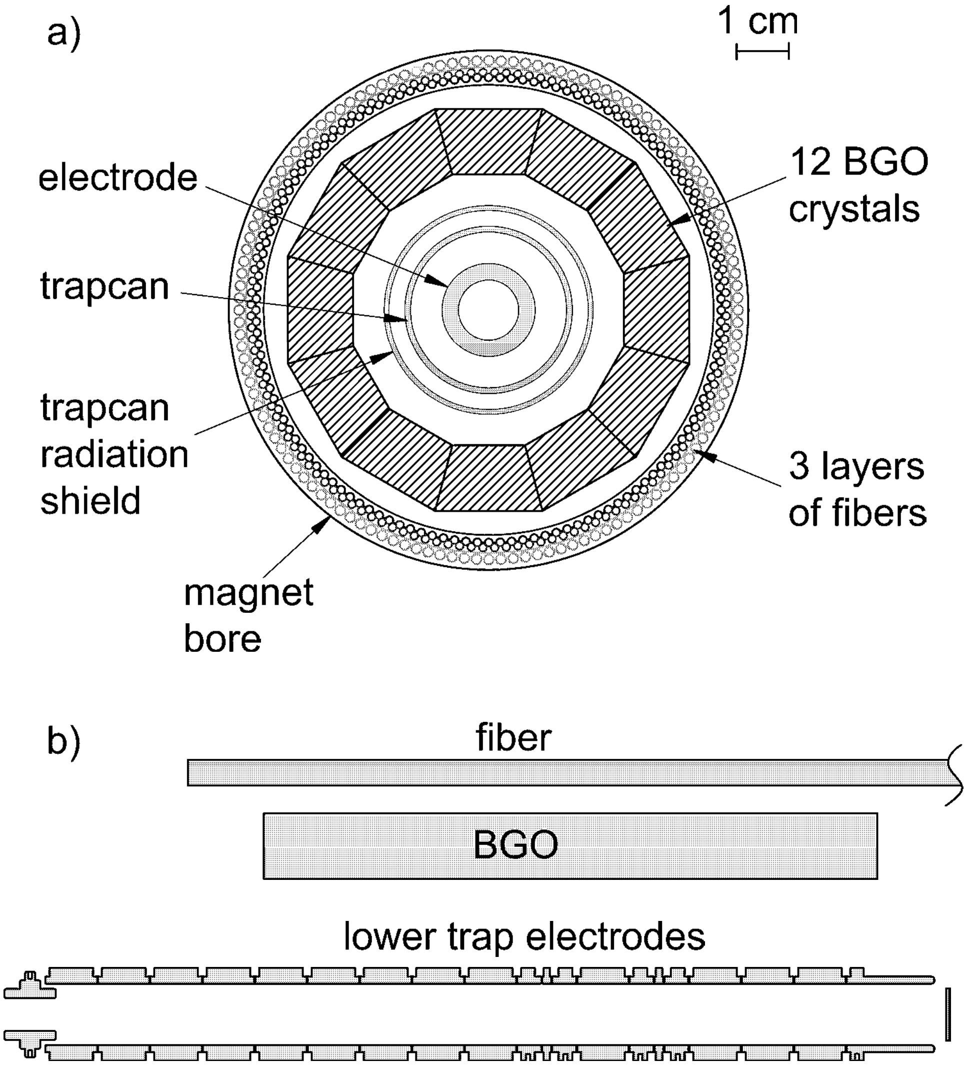

Figure 2.8: a) Top and b) side view of the fiber and BGO detectors surrounding the trapcan and electrodes.

Chapter 3

Particle Loading

In order to produce cold antihydrogen large quantities of cold antiprotons [10, 11] and positrons [7] are needed. Positrons are created in the radioactive decay of the unstable sodium isotope, ^{22}Na, while antiprotons are provided by the Antiproton Decelerator at CERN. Cold antiprotons in our Penning trap have an energy more than 10^{10} times lower than that of the antiprotons delivered by CERN. Cold electron clouds, loaded into the trap prior to the arrival of the antiprotons, play a vital role in this energy reduction and so large quantities of cold electrons are also required. This chapter describes the techniques used to capture and cool these three types of particles.

3.1 Electron Loading

3.1.1 FEP Loading

Electrons can be loaded from a field emission point (FEP) which is mounted on the ball valve, but is electrically isolated from the ball. When the ball valve is closed the FEP is pointing vertically down the axis of the lower trap. The FEP is a 0.5 mm diameter tungsten wire which is etched to a fine point at one end. The sharpness of this point is such that if negative a few hundred volts is applied to the FEP relative to the ball, the electric field at the point is high enough that electrons can tunnel out of the tungsten surface. Once emitted they are confined by the strong magnetic field and form a beam less than 40 μm in diameter which travels down the trap axis and strikes the degrader. A layer of atoms are necessarily frozen onto the degrader surface when the trap is initially cooled. These are liberated by the FEP electron beam, along with many secondary electrons. Collisions with the liberated gas can deposit a secondary electron into one or more potential wells, as shown in Fig. 3.1. Typically 5 million electrons are loaded in 10 s.

The figure also illustrates the "spilling" technique used to reduce the number of electrons in the wells after the FEP beam has been turned off. Clouds containing a large number of electrons are often found to be unstable. This is presumed to be due to collective oscillations triggered by trap imperfections or noise leaking into the trap. Also, during electron cooling of antiprotons, the energy added to the electrons by the energetic antiprotons may trigger instabilities. Large electron clouds approach the electrode where trap imperfections have the greatest effect, and in addition, the

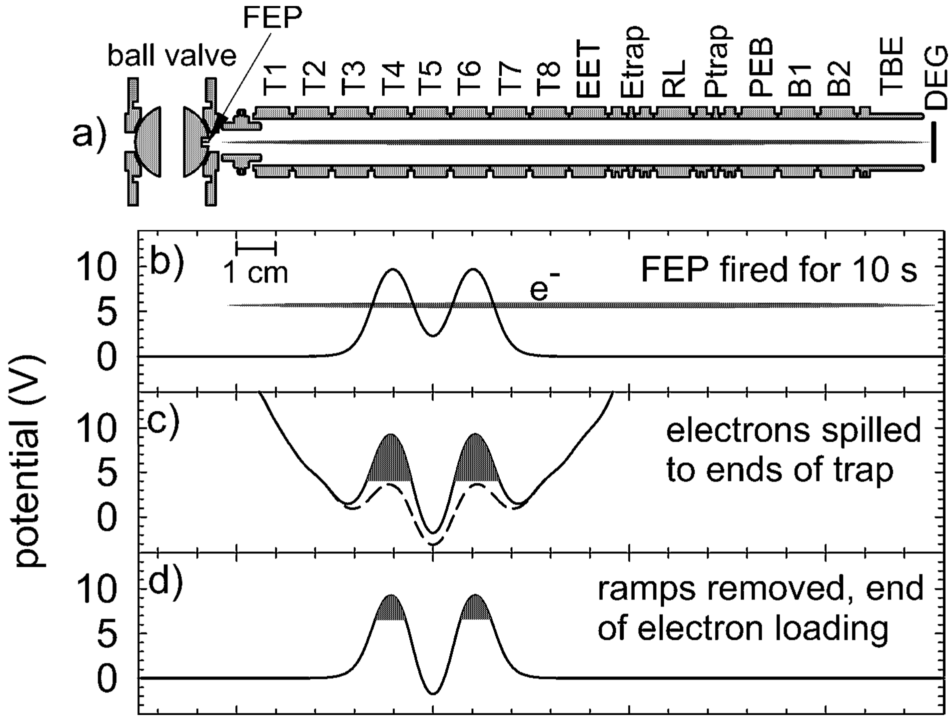

Figure 3.1: a) Electrodes, b) potentials while firing the FEP to load electrons into two wells, c) the number of electrons is reduced by making a ramp to the ends of the trap and shallowing the potentials to the dashed line for 2 s, and then returning to the solid line. d) Ramps removed and two moderate clouds of electrons remain, ready to cool antiprotons.

larger the electron cloud the greater the number of collective modes it can support. An unstable electron cloud has a detrimental effect on the stability of antiprotons loaded into the electron wells, and, especially with large antiproton clouds, can cause antiprotons to be lost from the trap.

3.1.2 Source Loading

A different electron loading technique uses the positron source to dislodge frozen gas and secondary electrons from the degrader. The loading rate with this method is much lower than when using the FEP, but this makes it more controllable. Overloading the well with electrons is not a concern and no spilling is necessary. The electron clouds are more stable and tend to be of a smaller diameter than those loaded from the FEP and then spilled. It typically takes 8 minutes to load 5 million electrons using this technique.

The procedure for electron loading via the positron source is explained in Fig. 3.2, with reference to the hbar2 trap. As many wells as possible are made in the trap but electrons load predominantly onto the TBE electrode since it is closest to the degrader. The degrader is biased slightly negative to repel secondary electrons, and the UPHV electrode is set to the same potential. This forces low energy electrons to repeatedly oscillate the length of the trap giving them a greater chance of a collision which can deposit them in a well.

It is likely that positive ions, created by bombarding the degrader with high energy positrons from the source, will be loaded onto the electrodes between the electron wells. When the electrons are grouped into one well (Fig. 3.2(c)) no well for

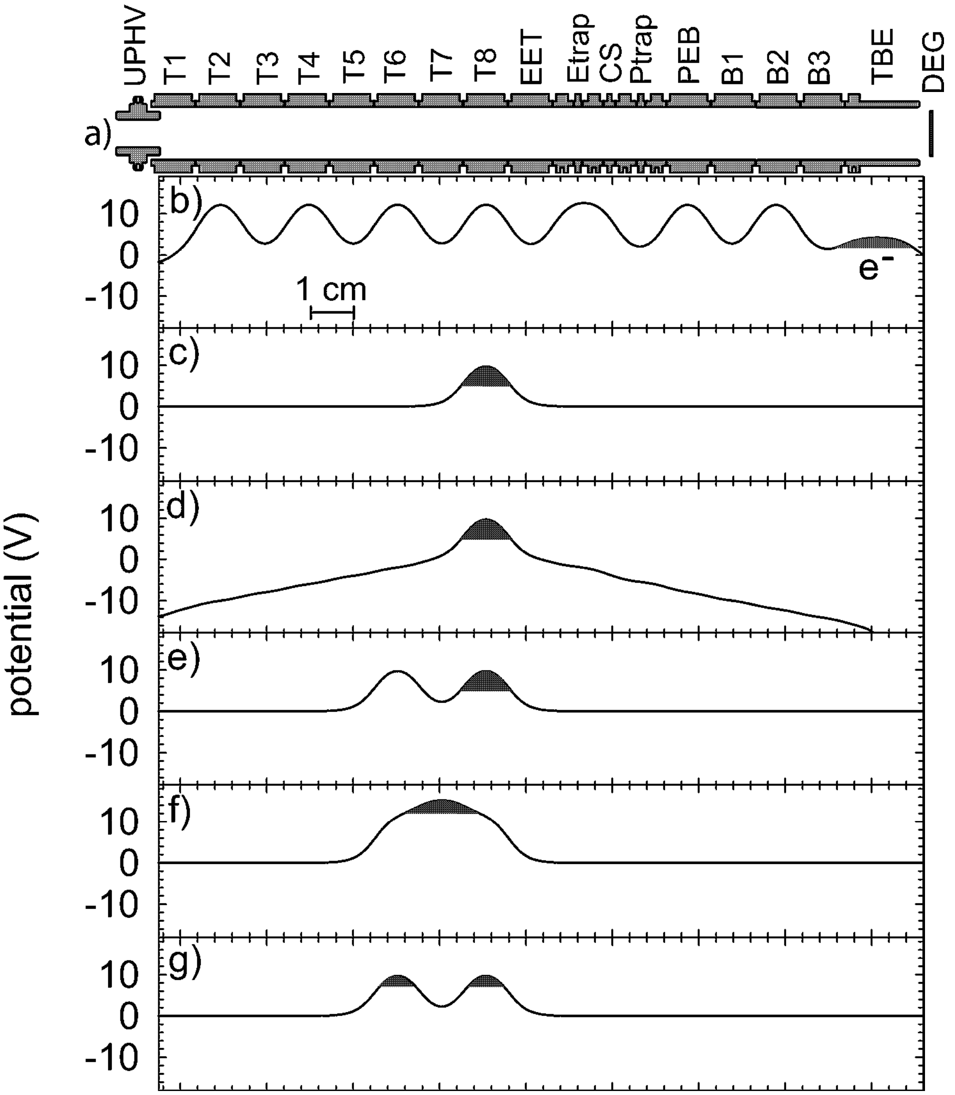

Figure 3.2: a) Electrodes of the hbar2 trap, b) potentials to load electrons onto the TBE electrode, c) all wells are joined to make a single well on T8, d) ramps are made to the ends of the trap to be certain to expel any ions in the trap, e) ramps removed and a potential well is made on T6 ready for dividing the electron cloud in two, f) potentials changed to collect all electrons on T7, before, g) the cloud is divided by returning the voltage on T7 to zero.

ions remains and they should therefore be lost from the trap. However, because of patch effects and small voltage offsets on the power supplies, there can exist small wells where ions can linger. A ramp of several volts per electrode ensures the comprehensive ejection of all ions (Fig. 3.2(d)).

Figure 3.2(d)-(g) shows how a cloud is divided. All the electrons are initially in a 12 V well on T8, and 12 V is applied to T6. The potential on electrode T7 is then set to 16 V forcing the particles onto T7. When T7 is returned to 0 V several seconds later the particles spill onto the T6 and T8 wells. This dividing technique is usually lossless and results in roughly equal numbers of particles in the two final wells. It can be repeated to divide a cloud into any number of smaller clouds and the technique works equally well for positrons and antiprotons.

Interestingly, if a cloud of electrons is loaded onto the TBE electrode, moved to, say, Etrap and then the loading potentials shown in Fig. 3.2(b) are reapplied, the majority of electrons will now load on Etrap and not TBE. This indicates that collisions with electrons already trapped in a well are most effective in capturing more electrons into that well. Figure 3.3 shows how the number of electrons loaded on Etrap grows, given that 1.6 million are initially confined there.

After the degrader is exposed to the positron source for $\sim$ 10 hours the electron loading rate diminishes substantially. This also occurs when using the FEP, but is less of a problem since the FEP is fired for only about 10 s each time. The loading rate reduces because, over time, the frozen gas layer on the degrader is removed by the impact of high energy positrons from the positron source, or electrons from the FEP. If there are no gas atoms to collide with then fewer secondary electrons will

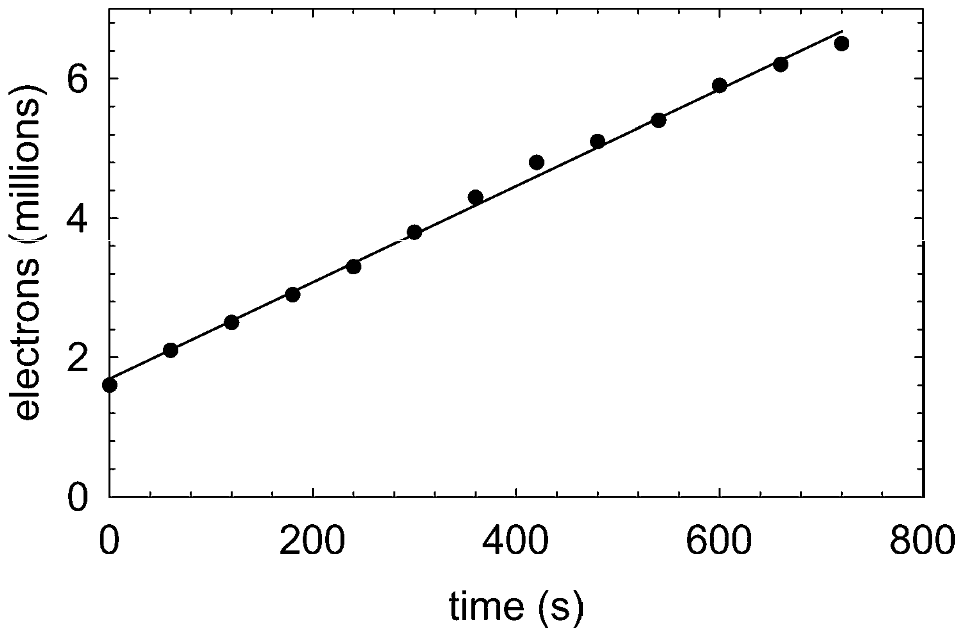

Figure 3.3: Electron loading directly onto Etrap from the radioactive source.

be trapped. However, the electron loading does recover after the trap is warmed to 300 K and then cooled back to 4.2 K and the gas layer is redeposited.

As mentioned in Sec. 2.1 electrons (and positrons) cool to 4.2 K by the emission of synchrotron radiation. In fact, only the cyclotron motion cools to this temperature, but collisions between electrons will transfer energy from the axial to the cyclotron oscillation, cooling the axial motion as well. Direct cooling of the axial motion can be achieved by coupling an electron cloud to a tuned circuit as shown in Fig. 2.6(a). The cloud will come into thermal equilibrium with the 4.2 K resistor, r.

3.1.3 Electron Loading Summary

Electrons are loaded by two different, but complimentary techniques. Loading from the FEP takes only a few seconds but can result in unstable clouds. These can have a detrimental effect on the stability of antiprotons cooled by the electrons, especially when loading large numbers of antiprotons. This electron loading technique is therefore best used for short experiments when small numbers of antiprotons are required. If larger numbers of antiprotons are needed then loading electrons from the radioactive source is more reliable. This takes 5-10 minutes. However, an experiment using many antiprotons takes $\sim$ 45 minutes and the electron loading time is therefore an acceptable fraction of the total.

3.2 Positron Loading

The trapping of positrons for precision experiments [61, 62, 63], plasma physics studies [64, 65, 66], and for antihydrogen production [67, 68, 69, 70, 71, 72] has a long history. The technique described here is the most efficient method of trapping large numbers of positrons directly into a cryogenic UHV environment [7]. The loading rate is as high as 3.5×10^4 positrons hr^{-1} mCi^{-1}, or 1 million positrons in 25 minutes with a 70 mCi source.

3.2.1 Radioactive Source Delivery

The ^{22}Na source used in these experiments was purchased in October 1999, and at that time had a strength of 150 mCi. The half life of ^{22}Na is 2.6 years and so by the summer of 2002, when these experiments were performed, its strength had reduced to 70 mCi. It remains, however, an extreme health hazard which has necessitated extensive thought and effort to build a system which will enable the source to be used safely. This is described in detail by Estrada [57] and I will only briefly describe it here.

The radioactive source is a paste of ^{22}Na, contained within a tungsten capsule mounted on the bottom of a rod 19 cm long and 1.54 cm diameter. The rod is made from the dense material Elkonite (90% tungsten, 10% copper) to absorb radiation emitted in the upwards direction. The source rod can be safely contained within a 1300 kg lead enclosure, situated above the Penning trap apparatus. A nearby experimenter then suffers no radiation dose above the background rate. While inside the enclosure the positron and antiproton annihilation detectors are shielded from the source, preventing them from exposure to ionizing radiation. The enclosure also supports a small liquid nitrogen dewar to cool the source. To load positrons a remotely controlled motor lowers the source on a long string until it rests at the pinbase, just above the trap vacuum enclosure. The pre-cooling of the source in the lead enclosure drastically reduces the heat load imparted to the trap when this is done. Once out of the lead enclosure access to the experiment is prohibited and the detector photomultipliers are turned off to prevent damage from over exposure to the source radiation. The tension of the string is continuously monitored and two cameras are used to observe the progress of the source rod as it is being lowered and raised. A switch indicates when the rod has reached the safety of the lead enclosure and a Geiger counter monitors the radiation at all times. These diagnostics have allowed years of safe positron experiments at CERN. Figure 3.4 illustrates the source delivery system.

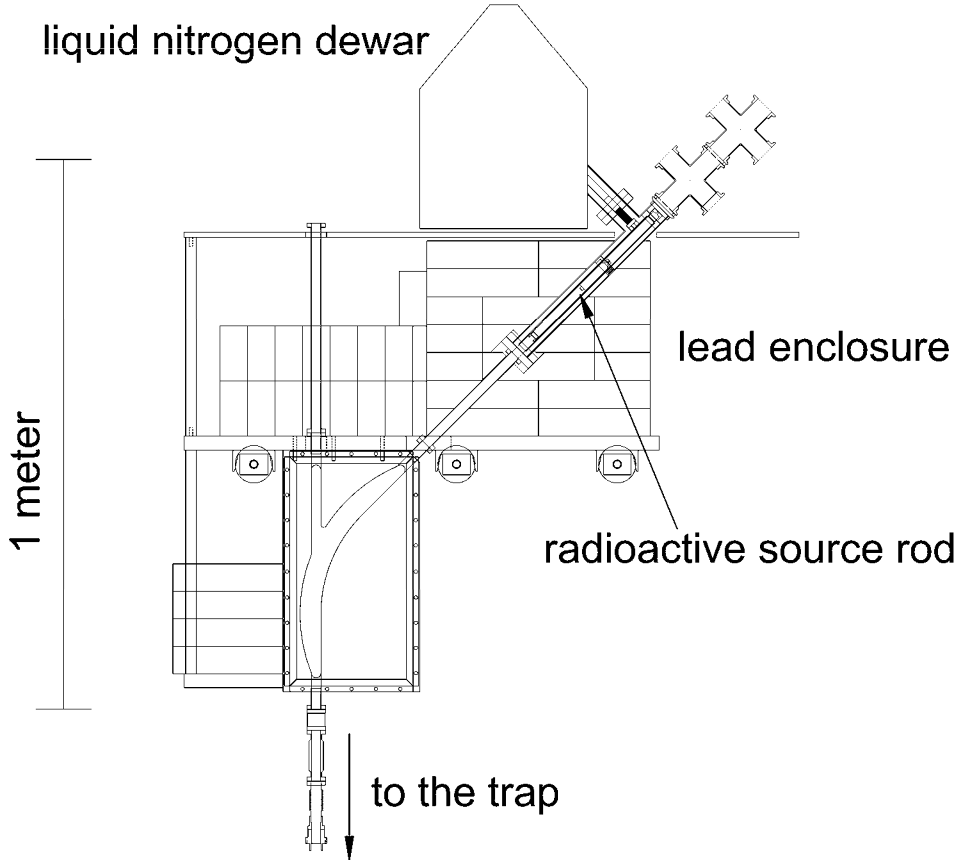

Figure 3.4: Positron source delivery system.

3.2.2 Rydberg Positronium Loading Technique

Positrons are emitted by the ^{22}Na source with a continuous range of energies between 0 and 545 keV. After emerging from the source capsule they are guided by the magnetic field and pass through the 10 μm thick titanium foil window in the pinbase which separates the trap vacuum enclosure from the vacuum of the magnet bore. They then encounter a 2 μm thick single crystal tungsten foil (TMOD). Such foils are used as transmission moderators for high energy positrons [73] and $\sim 0.05\%$ of the incident positrons emerge from the moderator with an energy of ~ 2 eV [7].

When exiting the foil a slow positron can be followed by a secondary electron and the two become bound together as a highly excited, or Rydberg, positronium atom. These atoms have a typical inter-particle separation of a few microns. Both particles are tightly confined in the radial direction by the strong magnetic field, while their axial positions can be influenced by applying an electric field at the moderator surface. This field is used to minimize the axial separation of the particles thereby increasing the likelihood of forming a bound state. The positronium atom is still guided by the magnetic field as it progresses down the trap where it is ionized by carefully constructed electric fields, which also trap the liberated positron. This is shown in Fig. 3.5. By reversing the potentials in the trapping region electrons are trapped at the same rate, confirming the role of positronium in the loading process.

Also shown in this figure is an additional moderator (RMOD) mounted on the ball valve but electrically isolated from the ball. This moderator is a thick tungsten crystal which works as a reflection moderator, i.e., slow positrons exit the crystal in the opposite direction to the incident positron beam. Since the majority of positrons pass

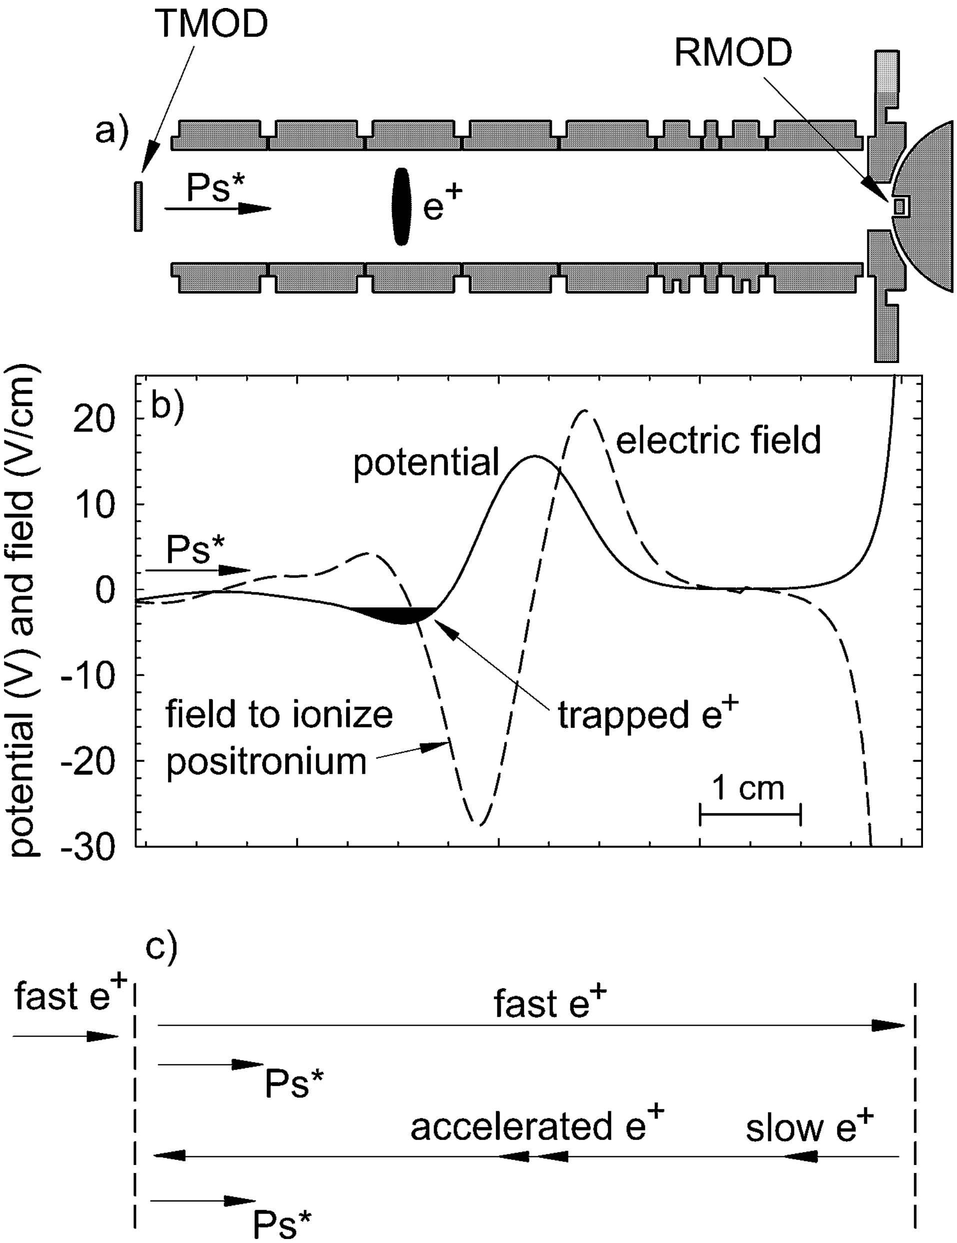

Figure 3.5: a) Electrodes, b) potentials used to trap positrons from positronium created at the transmission moderator, and c) schematic showing the positron currents which contribute to positronium formation.

straight through the transmission moderator they impinge upon RMOD. About 0.25% of these re-emerge as slow positrons and are accelerated back towards the transmission moderator by a large positive potential applied to RMOD. They diffuse a short distance into the surface of TMOD and can return back down the trap accompanied by an electron, as described above. The use of RMOD increases the positron loading rate by a factor of ~ 2.5. Figure 3.6 shows how the positron loading rate depends upon the voltage applied to TMOD, which is used to tune the positronium formation, and RMOD which accelerates positrons back towards TMOD. As can be seen the positron loading rate is strongly dependent on the voltage applied to the moderators. Accelerating positrons from RMOD back towards TMOD clearly increases the loading rate. This increase tails off as positrons are returned to TMOD with enough energy to implant deeply into the crystal, reducing the chance of re-emerging from the surface. The maximum loading rate occurs at a transmission moderator voltage of -1.25 V and a reflection moderator voltage of 500 V.

As with electron loading the positron loading rate diminishes over time as frozen gas atoms are removed from TMOD by the impact of high energy positrons. Unlike during electron loading, collisions with gas atoms do not play a role in positron loading. However, the gas can provide secondary electrons which are essential for positronium formation. The removal of the frozen gas layer on TMOD arrests the positron loading rate by roughly 6 positrons s^{-1} for each hour of exposure. Again, warming and then cooling the trap returns the loading rate to its initial high value.

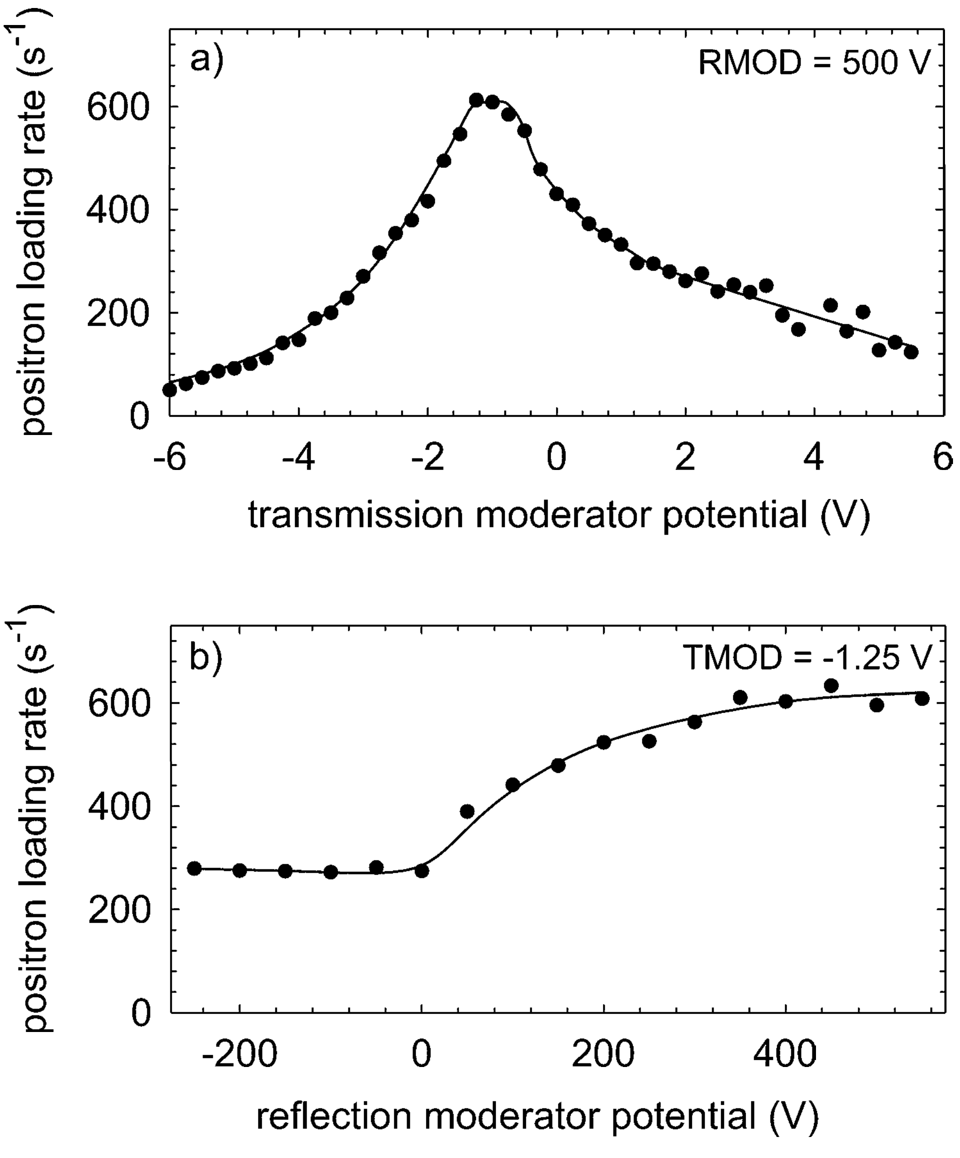

Figure 3.6: Loading rate variation with the potential applied to a) TMOD, and b) RMOD. The TMOD voltage tunes positronium formation and biasing RMOD provides additional positrons at TMOD for positronium formation.

3.3 Antiproton Loading

Antiproton capture in a Penning trap was first demonstrated in 1986 [8], a high voltage switching system used for this first experiment is explained in Ref. [13], while electron cooling of antiprotons is considered [74], and was demonstrated [9] in 1989. A thorough review of these experiments is given in [60]. The system used in the experiments detailed here was described by Estrada [57] and Yesley [56] based on the earlier design, and the techniques used to trap and cool large numbers of antiprotons are explained, investigated, and extended here.

3.3.1 AD Delivery of Antiprotons to ATRAP

Currently the only source of low energy antiprotons is at CERN in Switzerland. There are several stages involved in the eventual production of these particles. First protons produced by ionizing hydrogen gas are accelerated in a linear accelerator and injected into a synchrotron. There they are accelerated further, up to a momentum of 26 GeV/c. Approximately 10^{13} protons at this energy then strike an Iridium target and 3.57 GeV/c antiprotons produced in this process are guided to the Antiproton Decelerator (AD) shown in Fig. 3.7. Stochastic cooling [75] is used to reduce the antiproton bunch momentum spread and also the spatial emittance of the beam at 3.57 GeV/c. Radio-frequency deceleration to 2 GeV is then performed, followed by further stochastic cooling. While deceleration is underway the magnetic field in the bending magnets of the AD ring is ramped down to maintain the antiproton bunch trajectory in the center of the beamline. The antiprotons are then decelerated further, to 300 MeV/c, and electron cooling is employed to reduce the bunch spreading while at

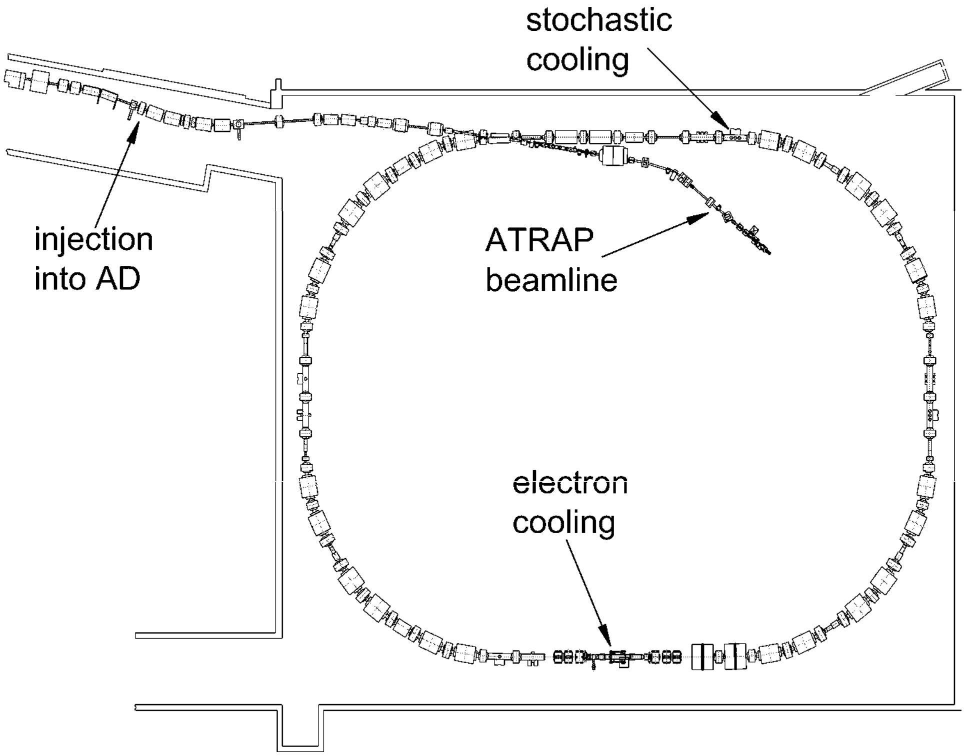

Figure 3.7: The AD ring showing the position of the cooling regions and the ejection line to the ATRAP experiment.

this energy. The final step is a deceleration to a momentum of 100 MeV/c followed by more electron cooling. At each cooling stage the number of antiprotons is monitored as a diagnostic for the cooling efficiency. About 100 s is required to complete all the deceleration and cooling steps which result in up to 3×10^7 antiprotons with an energy of 5.3 MeV compressed into an 80 ns bunch. After the final cooling stage the antiproton bunch is diverted to the ATRAP beamline before being bent upwards to enter the vertical ATRAP apparatus (Fig. 3.8).

3.3.2 Antiproton Beam Monitoring

Before the antiproton bunch enters the trap it passes through a parallel plate avalanche counter (PPAC) [76] which is used to monitor the position and intensity of the beam. This consists of two anode-cathode pairs orientated in the $x - y$ plane, perpendicular to the beam. Both cathodes are thin sheets of aluminized mylar. The two anodes are similar mylar sheets that have had some of the aluminum removed to leave five parallel conducting strips, one set is aligned in the x, and the other in the y direction. The strips are 2 mm wide with a gap of 0.5 mm between adjacent strips.

Between each anode-cathode pair 200 mbar of argon gas is flowed continuously and 200 V is maintained across each pair. When an antiproton passing through the PPAC strikes an argon atom it can ionize the atom. The resulting electron and ion are accelerated towards the anode and the cathode where they are collected. The electron will not move appreciably in the $x - y$ plane during this acceleration and so the anode strip which it strikes indicates the x or y position of the ionizing antiproton. The electron current on each strip flows through a series resistor and the voltage it develops

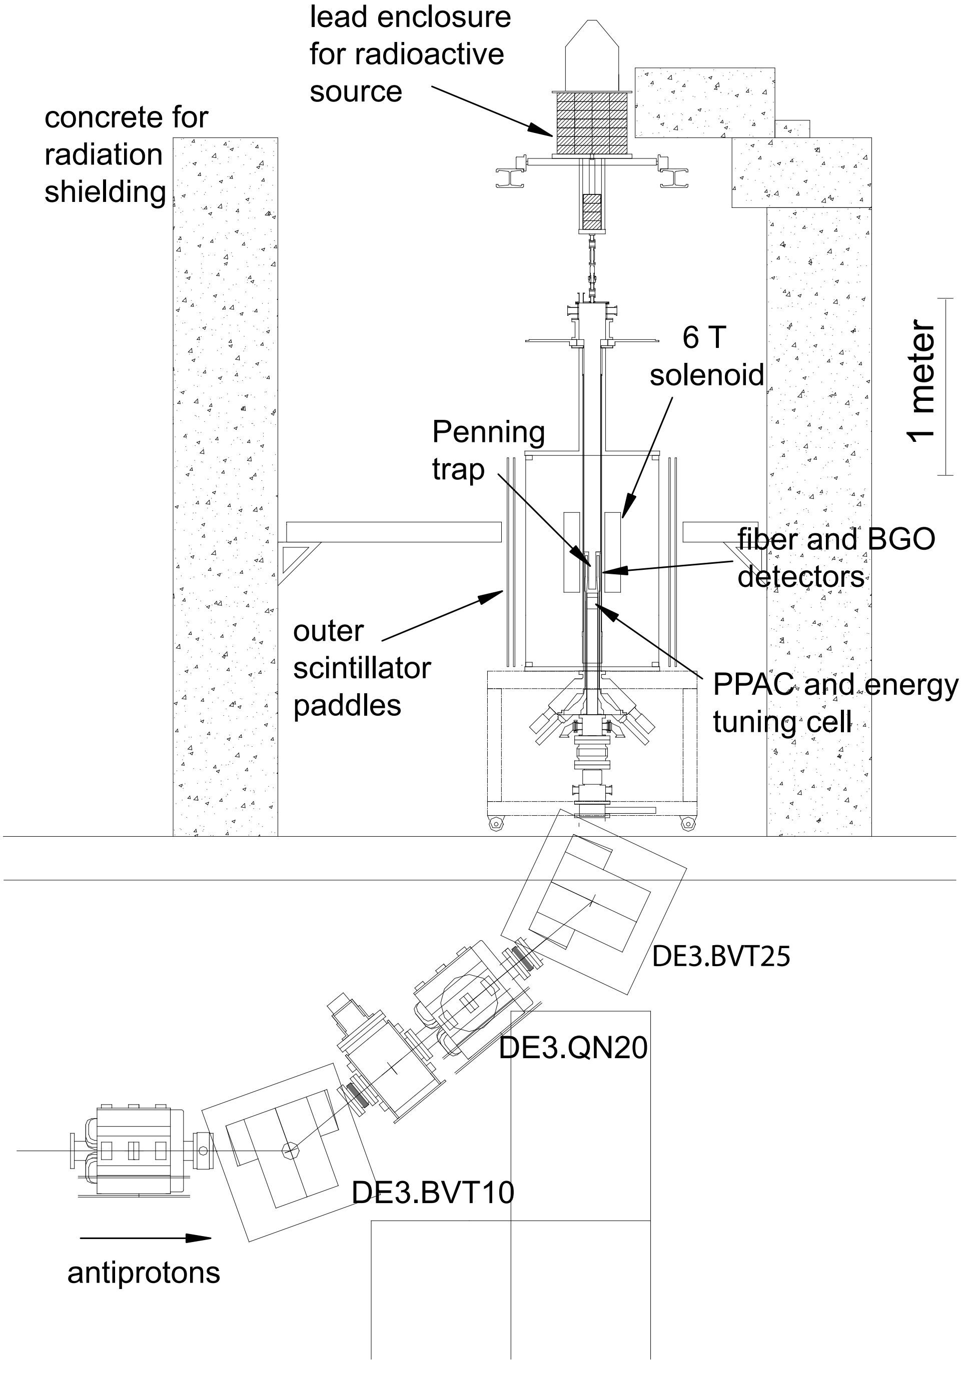

Figure 3.8: Side view of the entire ATRAP apparatus.

is amplified and then measured on a scope. Integrating this voltage determines the charge collected by each anode strip and allows the position of the antiproton beam to be determined.

The PPAC is mainly used as a diagnostic for positioning and focusing the antiproton beam. Steering the beam in the y direction is achieved by changing the current in the magnets DE3.BVT10 and DE3.BVT25 in Fig. 3.8, and in the x direction by tuning the DE0.DHZ45 magnet which is located further upstream. Focusing is performed using DE3.QN20 in the y direction and DE0.QN50 in the x direction. In this way the beam can be positioned and focused almost entirely onto the central 2 mm wide anode strip in each direction.

The total thickness of the PPAC is kept as low as possible to avoid blow up (straggling) of the antiproton beam. The SRIM software package [77] is used to calculate the energy loss of 5.3 Mev antiprotons passing through the PPAC material and argon gas. This shows that the energy loss is 0.48 MeV, or equivalently 33 μm of aluminium, which will produce little straggling with such energetic antiprotons.

Under typical operating conditions the voltage between the anode and cathode is too low to cause an electron avalanche. Instead, the PPAC operates in a linear mode, where each electron is unlikely to ionize further argon atoms. The integrated charge, which is measured by the scope, is therefore proportional to the number of antiprotons passing through the PPAC. However, if, as is often the case, the antiproton beam overlaps with a gap between two anode strips then the integrated electron charge will not fully represent the total number of antiprotons in the bunch. Because of this a different detector is normally used to monitor the intensity of the antiproton bunches arriving at the trap. This beam intensity monitor is a 7 cm square of plastic scintillator positioned on the concrete wall used for radiation shielding. In this position it sits ~ 2 m above the magnet and trap. As the bunch enters the trap, pions from antiproton annihilations strike the scintillator producing light which is detected by a photomultiplier. The output current from the photomultiplier is integrated to give a charge (Q_I) proportional to the number of incident ionizing particles and therefore to the number of antiprotons. This monitor provides a reliable indication of the relative strength of each antiproton bunch. Figure 3.9 shows one month of antiproton bunches measured by the beam intensity monitor. Also shown in the figure is the number of particles in the AD after the electron cooling at 100 MeV/c (N_{AD}), and the correlation between these two measurements.

The beam intensity steadily decreased through the course of the month due to a reduction in the number of antiprotons in the AD. The dashed line in Fig. 3.9(c) indicates an upper bound on the efficiency with which antiprotons can be ejected from the AD and steered successfully to the ATRAP apparatus. As can be seen from the clustered points at about 25 million antiprotons in the AD and -300 pC from the intensity monitor this ideal transfer is not usually reached. Plots such as these allow monitoring of the beam intensity, ejection and steering, and the AD staff can be informed if there is a problem.

3.3.3 Antiproton Energy Tuning

Antiprotons arrive from the AD with an energy of 5.3 MeV, far too high to be trapped with the modest potentials supported by our Penning trap. The antiprotons

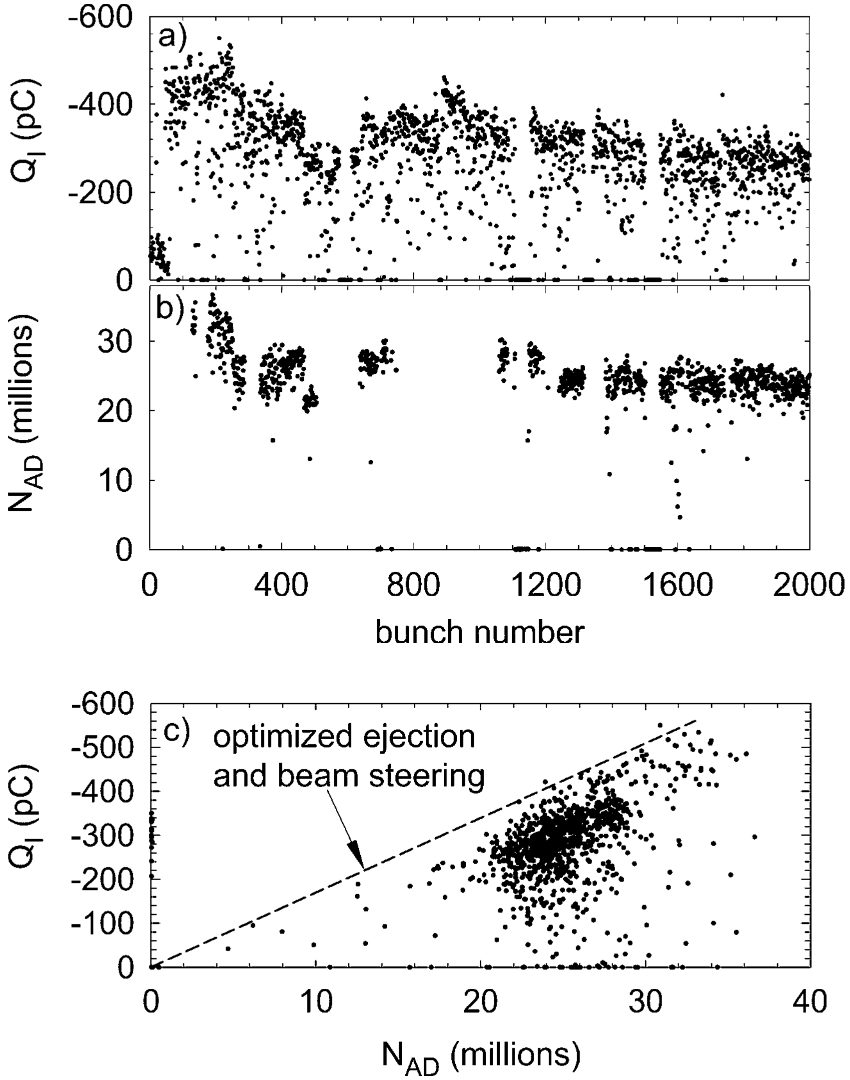

Figure 3.9: a) Variation of the beam intensity arriving at the ATRAP apparatus for each bunch delivered in October 2001, b) the number of antiprotons circulating in the AD just prior to ejection to the ATRAP beamline, and c) the correlation between these two measurements. The AD beam monitor was not always working and so some bunches were not measured in the AD.

are therefore slowed by passing through material on the way to the trap. Some of this material is necessary to monitor the beam, to keep the various vacuum spaces separate, and to provide thermal shielding, but the degrader and the energy tuning cell are used specifically for antiproton energy reduction. The degrader is a beryllium foil 125 μm thick, positioned directly underneath the electrode stack, while the energy tuning cell is a gas space filled with a variable mixture of sulfur hexafluorine (SF_6) and helium, and is located above the PPAC.

When the 15 mm long energy tuning cell is filled with helium the antiprotons are hardly slowed at all, but when only the much heavier SF_6 gas is present, substantial energy loss occurs. Controlling the relative amounts of the two gases in the cell therefore enables the energy of the emerging antiproton beam to be tuned to maximize the number of particles trapped. The total tuning range is 0.6 MeV, or equivalently 40 μm of aluminium.

The optimum SF_6:He ratio can, however, change from one cooldown to the next. This is due to inadequate evacuation of the magnet bore. When the trap is cooled any residual gas in the bore will freeze onto the cold surfaces through which the antiprotons pass. Different thicknesses of the frozen layer will occur for each cooldown, resulting in a shifting of the optimum SF_6:He ratio. After discovering this effect care is taken to fully evacuate the magnet bore before cooling the trap. In Fig. 3.10 the optimum ratio is determined by maximizing the number of antiprotons captured in the Penning trap.

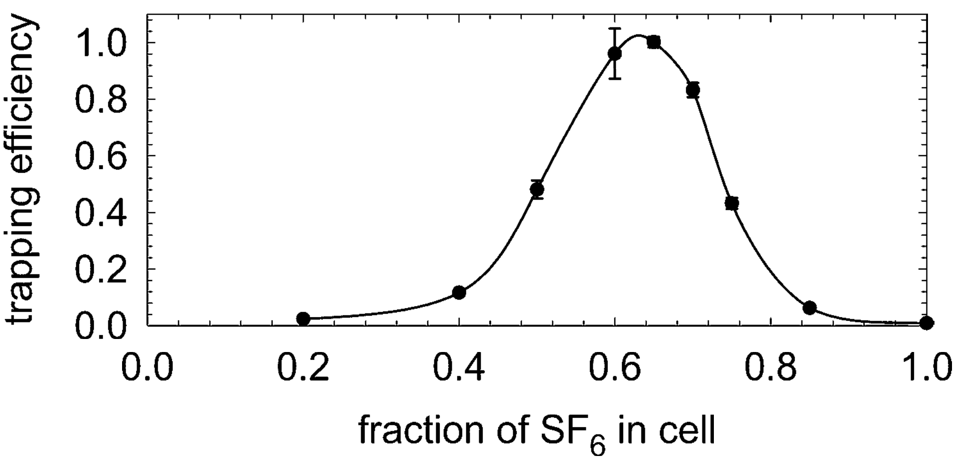

Figure 3.10: Varying the fraction of SF_6 in the energy tuning cell to optimize the number of trapped antiprotons. The efficiency is normalized to the strength of the antiproton bunches as measured by the beam intensity monitor.

3.3.4 Antiproton Trapping

After passing through the energy tuning cell an antiproton bunch enters the trap through the 10 μm thick window on the bottom of the trapcan. It then encounters the 125 μm thick beryllium degrader which provides the final energy reduction: from 3.5 MeV to $\sim$ 0 MeV. Prior to the arrival of the antiproton bunch a static -3 kV potential is applied to the upper high voltage electrode (Fig. 3.11). Antiprotons emerging from the degrader with an energy less than 3 keV will therefore be reflected at the top of the trap and return towards the degrader. Higher energy particles simply strike the ball valve and annihilate. Before the reflected antiprotons can return to the degrader its potential is switched to -3 kV in $\sim$ 50 ns, and the particles are trapped in a high voltage well.

Figure 3.12 shows the results of a numerical calculation of the time it takes an antiproton of a given energy to return to the degrader. The highest energy antiprotons return in 530 ns, while those with the average energy of 900 eV take 850 ns. A typical

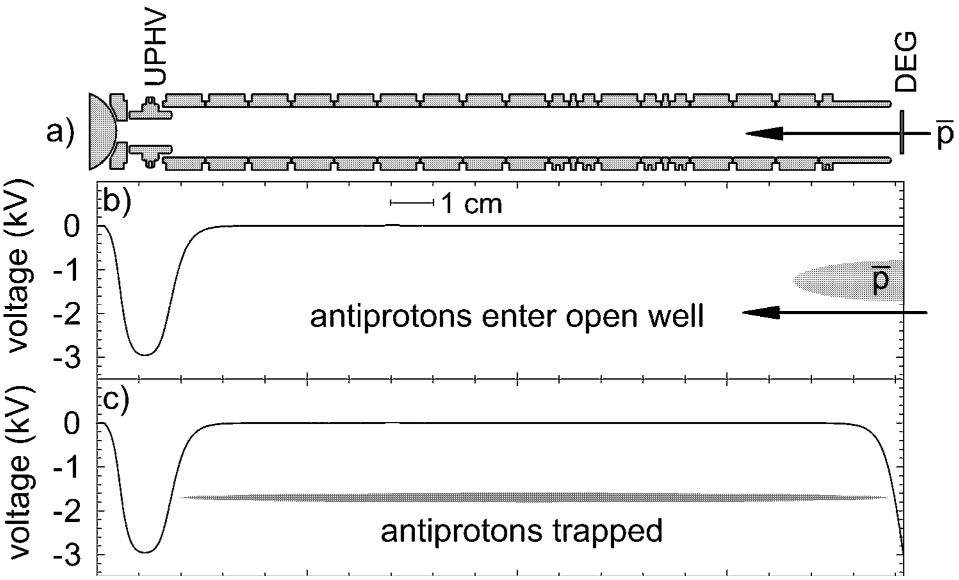

Figure 3.11: Antiproton capture in the high voltage well.

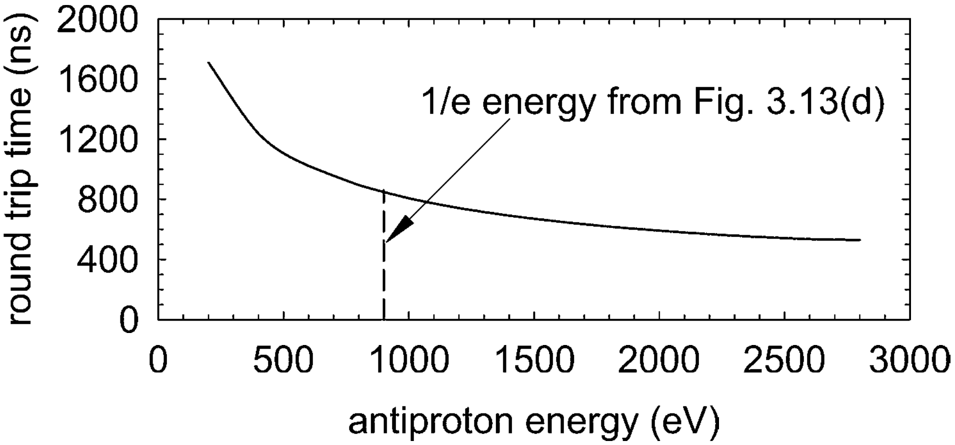

Figure 3.12: Calculated time taken for an antiproton traveling along the trap axis with a given energy to return to the degrader after reflecting from a -3 kV potential on the upper high voltage electrode.

antiproton bunch length, as measured by the PPAC detector, is 80 ns. The optimum time to switch the degrader voltage is just after all of the antiproton bunch has entered the trap and a switching time of 50 ns is clearly sufficient to prevent even 3 keV antiprotons at the front of the bunch from returning to the degrader before it is switched to -3 kV.

After 30 s the degrader voltage is ramped slowly to 0 V and the antiproton annihilations are detected as the particles strike the degrader. Monitoring the voltage at which annihilations occur enables the energy distribution (number of antiprotons per volt, or dN/dV) to be measured. This is shown in Fig. 3.13(d). Also seen in this figure (3.13(a) and (c)) is the variation of the trapping efficiency with the time at which the degrader voltage is switched to -3 kV and the high voltage well is established. If the degrader voltage is switched before the antiprotons arrive none will be caught. As the switching delay is increased a fraction of the antiproton bunch will enter the trap and be caught. This is evident in the the rapid increase in trapping efficiency at $t = 0$. This edge should therefore be the integral of the time profile of the antiproton bunch measured by the PPAC, and the comparison is made in Fig. 3.13(c). It is seen that the efficiency falls as the switch time is increased above $\sim$ 800 ns (the round-trip time of a 900 eV antiproton) in agreement with the numerical calculation shown in Fig. 3.12.

The correlation between the beam intensity and the number of trapped antiprotons can be measured. This is shown in Fig. 3.14(a) and indicates the efficiency with which a given antiproton bunch is trapped. This efficiency will be affected by small changes in the beam steering, the time at which the degrader voltage is switched, and the

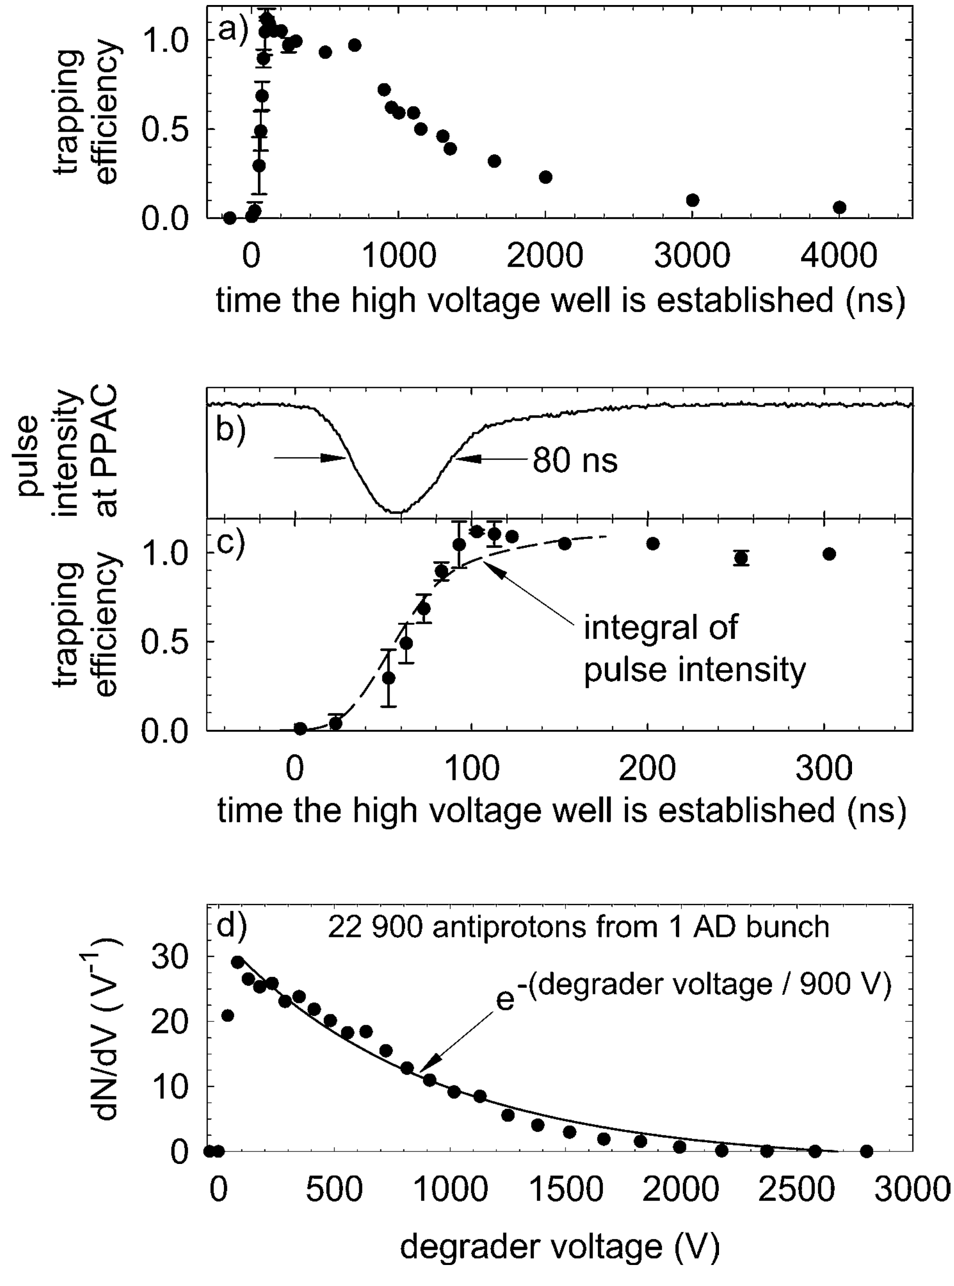

Figure 3.13: a) Variation of the trapping efficiency in the high voltage well with the time at which the potential on the degrader is switched, b) and c) the integral of the pulse intensity gives the shape of the initial rise in trapping efficiency, and d) after 30 s in the high voltage well the antiprotons are spilled and their energy distribution is measured.

energy tuning. For this particular data set there is seen to be two upper bounds to the efficiency with which antiprotons are trapped. The lowest of these two is from bunches trapped before 8^{th} October 2001, while the higher efficiency is for dates after 8^{th} October. The difference between the two is because on 9^{th} October the trap was warmed and then cooled back to 4 K but with more efficient pumping of gas from the magnet bore. This prevented an uneven layer of frozen gas being deposited on the bottom of the trap radiation shield, through which the antiproton beam must pass. An uneven layer of frozen gas present before 8^{th} October would mean that antiprotons at different $x - y$ positions pass through different amounts of matter, thus broadening the energy distribution and reducing the trapping efficiency. The effects of inadequate pumping have been seen on a number of occasions and care is now taken to avoid this problem

The data points in Fig. 3.14 showing a small number of antiprotons trapped are simply because normally one does not simply catch and spill the antiprotons from the high voltage well, but uses them for different experiments. Figure 3.14(b) indicates the efficiency with which antiprotons are captured. When fully optimized $\approx 2 \times 10^4$ antiprotons are trapped per bunch delivered by the AD, for an efficiency of 7.4×10^{-4}. This is a little higher than the efficiency of 5×10^{-4} achieved previously at CERN [55].

The number of antiprotons captured varies with the trapping voltage applied and with the time that the antiprotons remain in the high voltage well before being spilt to the degrader. Figure 3.15(a) indicates that the variation with trapping voltage is not linear up to the maximum trapping voltage of -3 kV, but instead saturates.

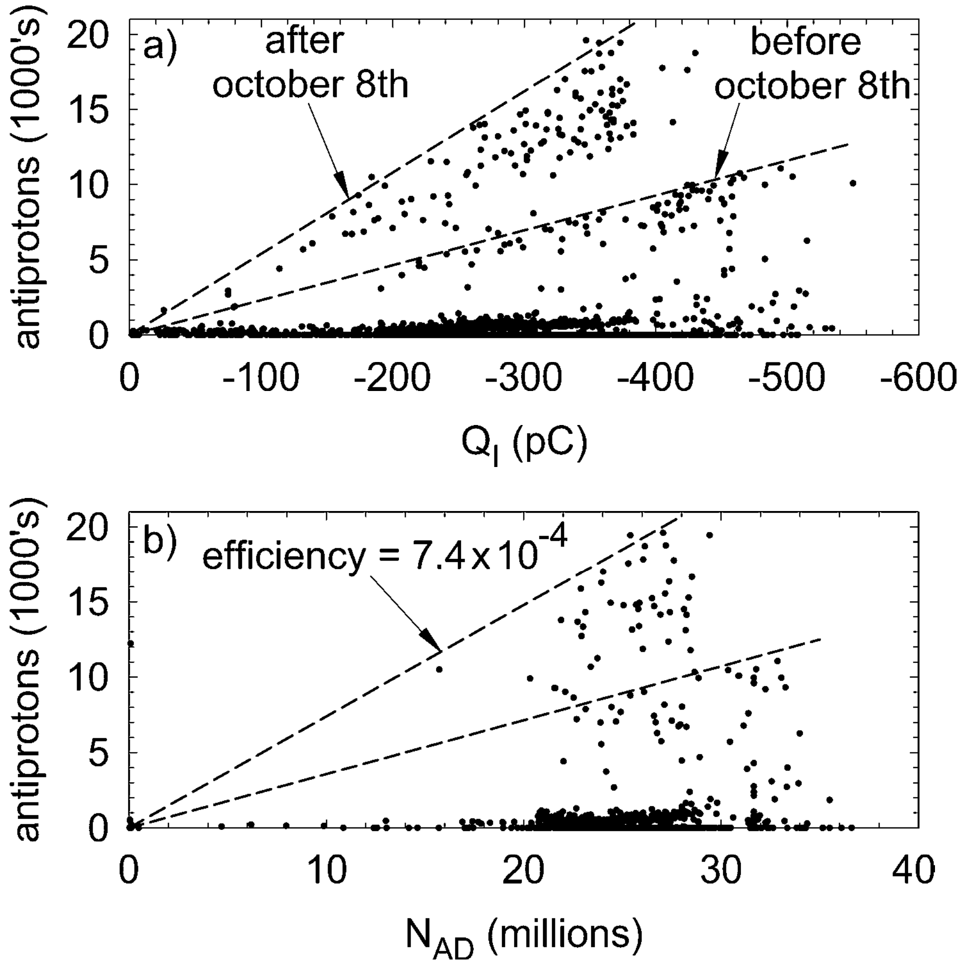

Figure 3.14: Variation of the number of trapped antiprotons with, a) the measured beam intensity, and b) the number of antiprotons circulating in the AD prior to ejection.

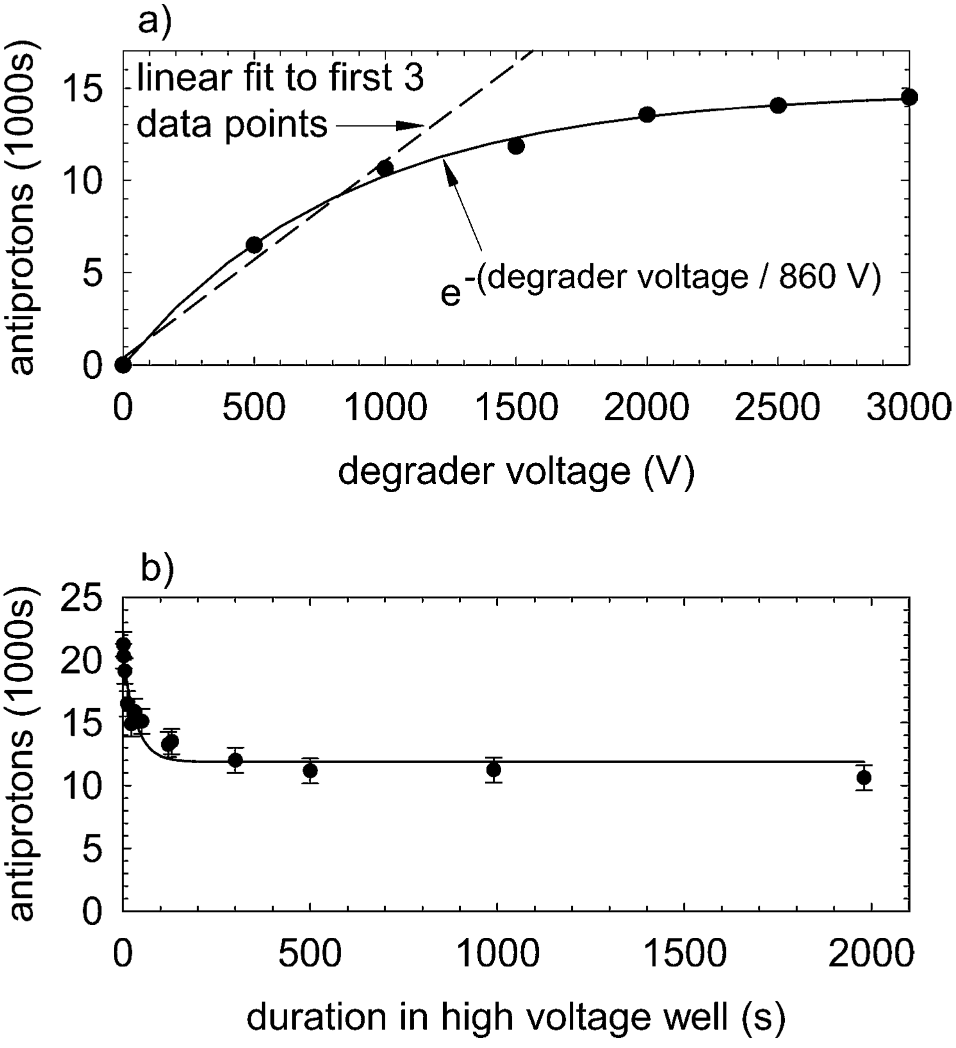

Figure 3.15: Variation of the number of trapped antiprotons with, a) the voltage applied to the degrader and the upper high voltage electrode, and b) the time the antiprotons are stored in this high voltage well before being released and counted.

The energy distribution of antiprotons emerging from the degrader is expected to be constant in the range 0-3 keV and the degrader voltage is switched rapidly enough to trap a 3 keV antiproton in the axial direction. The saturation of the trapping efficiency must therefore be due to antiprotons leaving the trap radially. This is consistent with the fact that few high energy particles are found in the energy spectrum when the antiprotons are released from a 3 kV long well as shown in Fig. 3.13(d). Presumably antiprotons with a lot of axial energy also have considerable transverse energy and so have large radius cyclotron orbits. They are therefore the most likely to strike an electrode and annihilate after a collision with another antiproton. A trap with a larger electrode radius should produce a trapping efficiency which remains linear up to higher voltages than the current trap. Also worth noting is that the 1/e energy given in Fig. 3.13(d) is roughly the same as that found in Fig. 3.15(a), as should be the case since varying the trapping voltage is just another way to measure dN/dV.

The initial reduction in the number of trapped antiprotons with hold time (see Fig. 3.15(b)) is due to collisions between antiprotons which nudge the outermost particles onto the electrodes. This loss indicates that the antiprotons, at least initially, cover the full diameter of the trap. After 100 s the number of trapped particles remains constant indicating that the trap vacuum is good, that there is no excessive noise to disturb the antiprotons, and that, basically, the trap is working to the high standards required of it.

3.4 Antiproton Cooling and Stacking

3.4.1 Electron Cooling of Antiprotons

Despite having an energy of order 10^3 times lower than that provided by the AD the antiprotons trapped in the high voltage well are still rather hot. If one were to define a temperature for them then it would be 10 million Kelvin. Electron cooling reduces this to 4.2 K, with little antiproton loss.

Electrons loaded in the trap prior to the arrival of the antiprotons quickly cool to 4.2 K. Typically 3 million electrons are loaded into a 12 V well (Fig. 3.16). When the antiprotons enter and are caught in the high voltage well they are forced to make multiple passes through the electrons. Coulomb collisions between the hot antiprotons and the cold electrons serve to reduce the antiproton energy until it is low enough for one final collision to deposit an antiproton in the electron well [9]. It takes ~ 50 s to achieve this for the majority of the antiprotons.

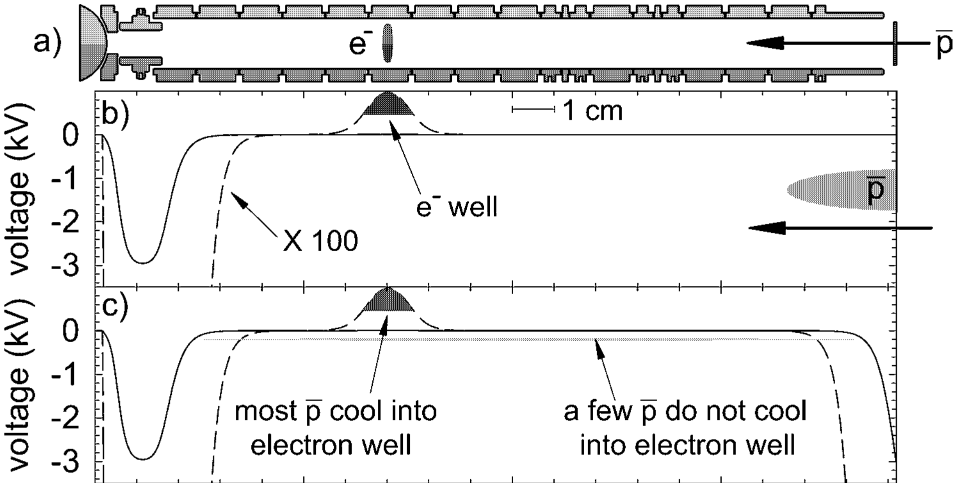

Figure 3.16: Electron cooling of antiprotons.

After the electron cooling time is over the electrons can be ejected. The technique to remove the electrons makes use of the fact that antiprotons move much slower in response to an applied field than do electrons. To remove the electrons a potential ramp is made towards the degrader and then the well containing both electrons and antiprotons is opened for 100 ns. The electrons stream out to the degrader during this time while the antiprotons hardly move, and so remain trapped when the well is restored. In some antihydrogen experiments it is better to allow the electrons to remain as a cooling source (see Ch. 5), and in others it is better to eject the electrons (see Ch. 6).

Figure 3.17 shows the energy distribution of a single trapped antiproton bunch after electron cooling for 80 s. The width of the distribution is narrow, but still appreciably larger than that of a 4 K thermal distribution. However, numerical calculations confirm that the width is consistent with the space charge potential of 25 000 antiprotons.

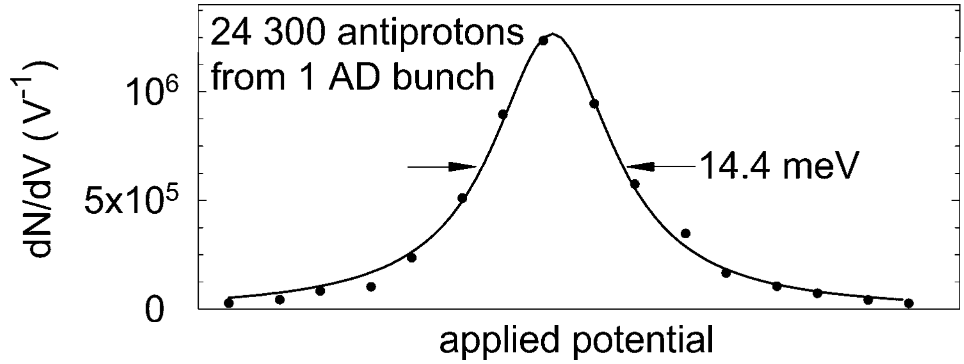

Figure 3.17: Electron cooled antiprotons show a very narrow energy distribution determined by the particles' space charge. The absolute energy scale is not given due to small uncalibrated offsets.

It is illustrative to observe the energy distribution of antiprotons released from the high voltage well after different electron cooling times. These are the antiprotons which have not been cooled sufficiently to become trapped in the electron well. The observed spectrum shifts down in energy as the cooling time is increased and also fewer antiprotons are released at the end of the cooling time as more are cooled into the electron well. Figure 3.18 shows the results. Even after 85 s of cooling some antiprotons have still not cooled into the electron well. These are presumed to be antiprotons loaded far off axis, or with large cyclotron radii, which therefore do not have a significant overlap with the electron cloud.

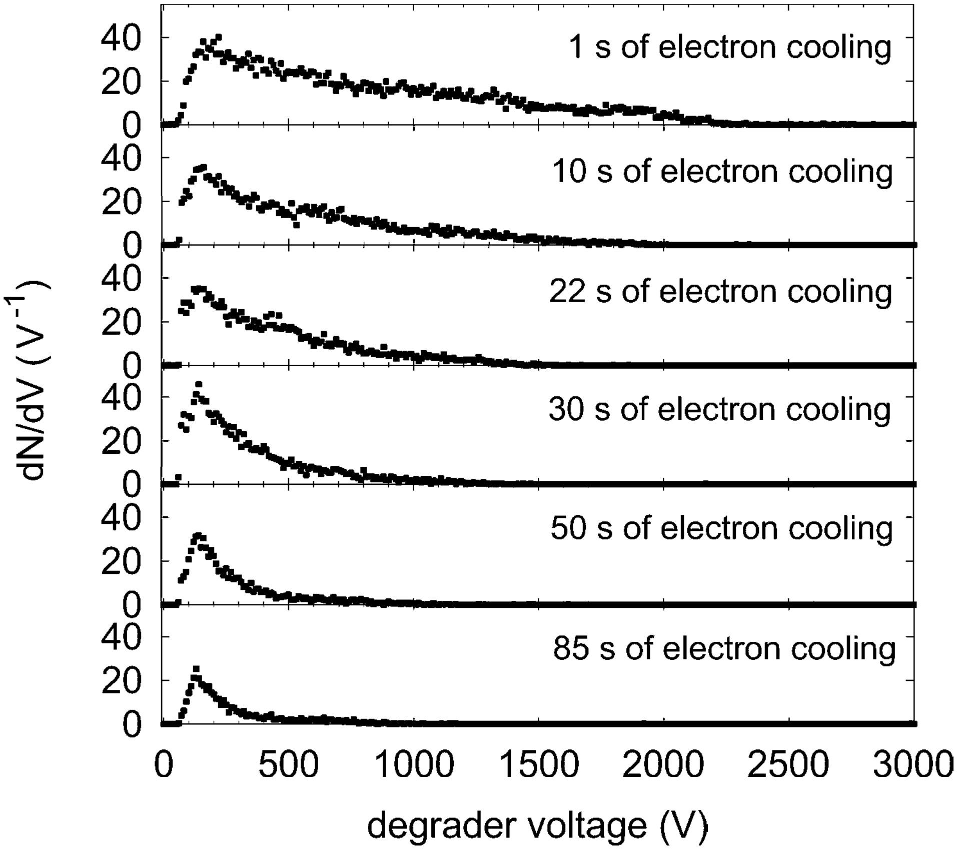

Figure 3.18: As the electron cooling time increases so the antiproton energy distribution shifts to lower energies and fewer antiprotons are released, as more are cooled into the electron well.

During electron cooling the antiproton loss from the trap is continuously monitored. This loss, shown in Fig. 3.19, is virtually the same as when no electrons are present. Since the electron cloud does not cover the entire 12 mm diameter of the electrode but, at least initially, the antiprotons do, there will be a number of antiprotons at large radii which do not interact with the electrons. Antiproton loss which is insensitive to the presence of electrons is therefore likely due to these far off-axis antiprotons confirming the previous assertion.

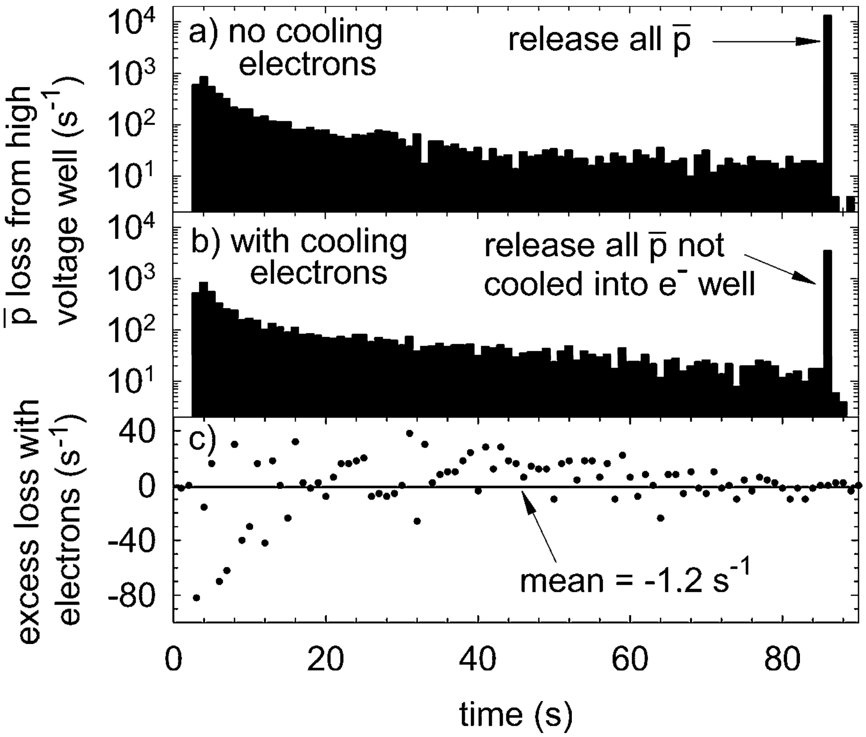

Figure 3.19: Antiproton loss from the high voltage well is insensitive to the presence of electrons.

The antiprotons entering the trap follow one of three paths: a fraction f_c are cooled into the electron well, a fraction f_l are lost from the trap during cooling, and a fraction f_u are uncooled and released from the trap when the degrader voltage is lowered after the electron cooling time is over. The annihilation detectors are turned off for 10 s when the antiproton bunch arrives at the trap so that they do not saturate and it is therefore not possible to measure the particle loss during this time. A complete account of the relative fractions, f_c, f_l, and f_u is shown in Fig. 3.20. As the

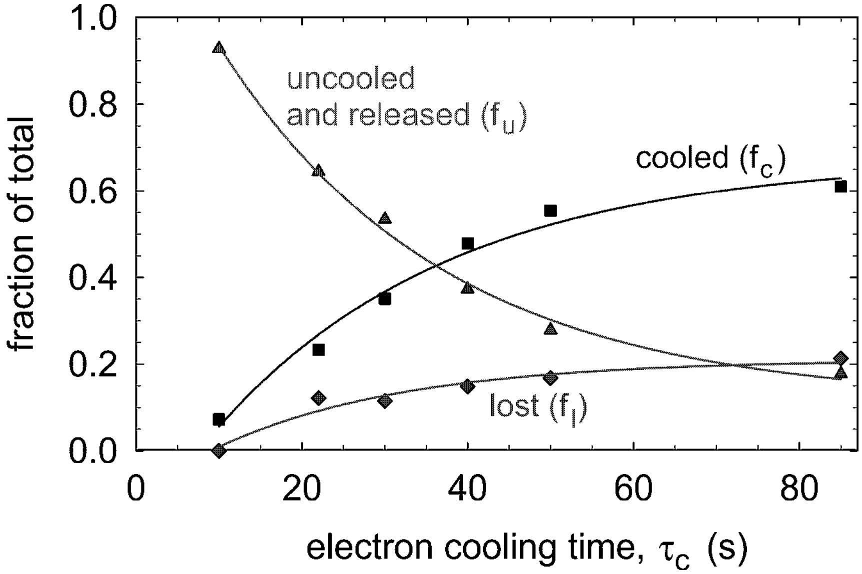

Figure 3.20: Antiprotons are either cooled into the electron well, lost from the trap during cooling, or uncooled and released from the trap at the end of the cooling time. The relative fractions change with the duration of the electron cooling.

cooling time is increased the fraction of cooled antiprotons increases and the uncooled fraction decreases. The accumulated loss also increases. Increasing the cooling time from 50 s to 85 s does not appreciably increase the fraction of cold antiprotons but simply allows more time for loss from the trap and a corresponding reduction in the fraction released at the end of the cooling. After 85 s the relative fractions of cooled, uncooled and lost antiprotons have all roughly reached a constant level. The total number of antiprotons used in each experiment remained approximately constant for the duration of this study. The cooling efficiency at time τ_c is defined as the ratio of the fraction of antiprotons cooled into the electron well at time τ_c, to the fraction of antiprotons that would be trapped in the high voltage well after a time τ_c, with no electrons present. Since the loss from the trap is insensitive to the presence of electrons the cooling efficiency, ϵ_c, is simply the ratio:

$$\epsilon_c = \frac{f_c}{f_c + f_u}, \tag{3.1}$$

and is shown in Fig. 3.21. The cooling efficiency is seen to have an exponential

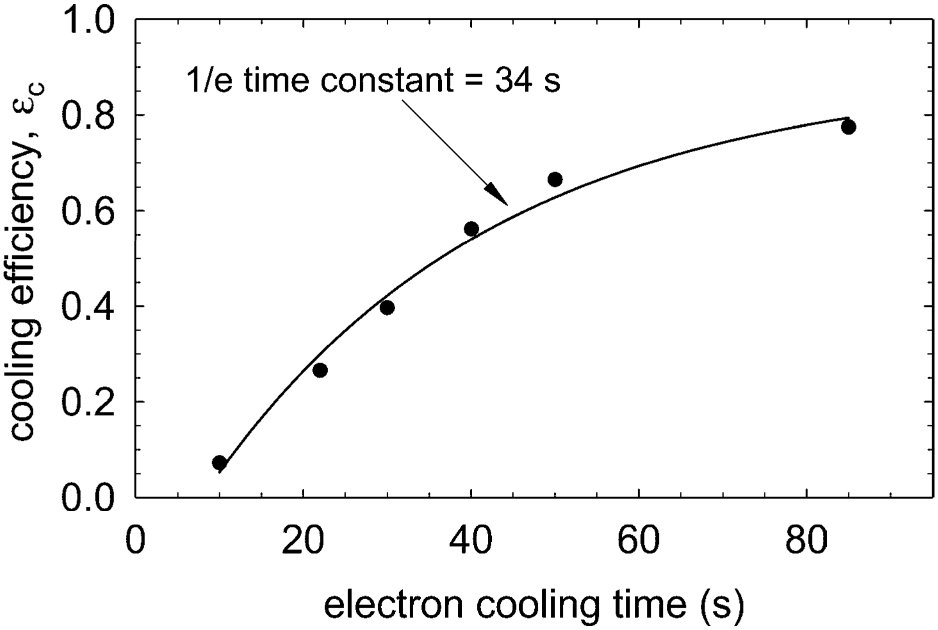

Figure 3.21: Electron cooling becomes more efficient as the cooling time increases.

time constant of 34 s, and an efficiency of 80% after 85 s. Electron cooling is often a delicate balance between having enough electrons for effective cooling but not so many that instabilities in the electron cloud can perturb the antiprotons. As such, higher efficiencies (up to nearly 100%) have been achieved, but were not evident at the time of this systematic study.

3.4.2 Antiproton Stacking

As many cold antiprotons as possible are needed for antihydrogen experiments and so the ability to load and cool multiple antiproton bunches is crucial [9, 10, 11, 12]. After a single bunch has been cooled into an electron well as described above the degrader potential can be lowered to admit another bunch of antiprotons. Any antiprotons not cooled into the electron well will be lost at this time. This procedure can be repeated many times and is know as antiproton stacking. It is the only way to capture more than 2×10^4 antiprotons for use in antihydrogen experiments. The energy distribution of a total of nine antiproton bunches stacked one on top of the other is shown in Fig. 3.22. Again, the width of the distribution is determined by the antiproton space charge.

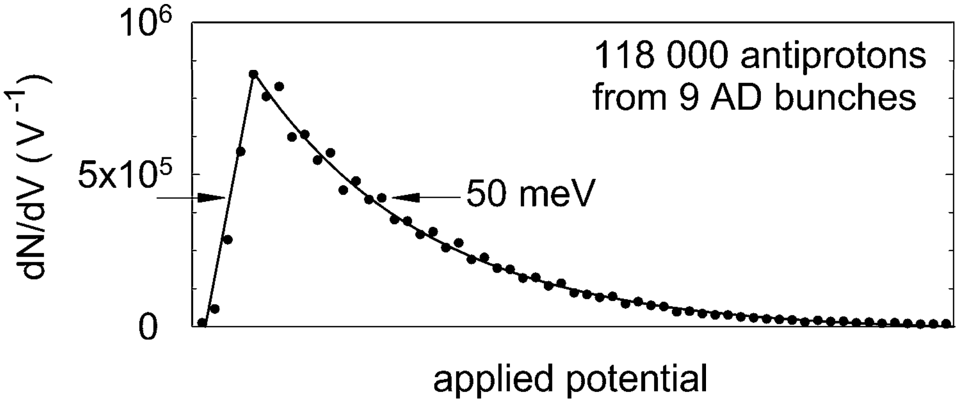

Figure 3.22: The energy distribution of nine antiproton bunches stacked together. The width of the distribution is determined by the particles' space charge.

Figure 3.23 shows a series of experiments where increasing numbers of bunches are stacked in the trap which is pre-loaded with 4 million electrons loaded from the radioactive source. After the required number of bunches have been stacked the electron well is spilt towards the degrader and the number of cold antiprotons is measured. A total of 4.4×10^5 antiprotons are trapped and cooled from 32 bunches,

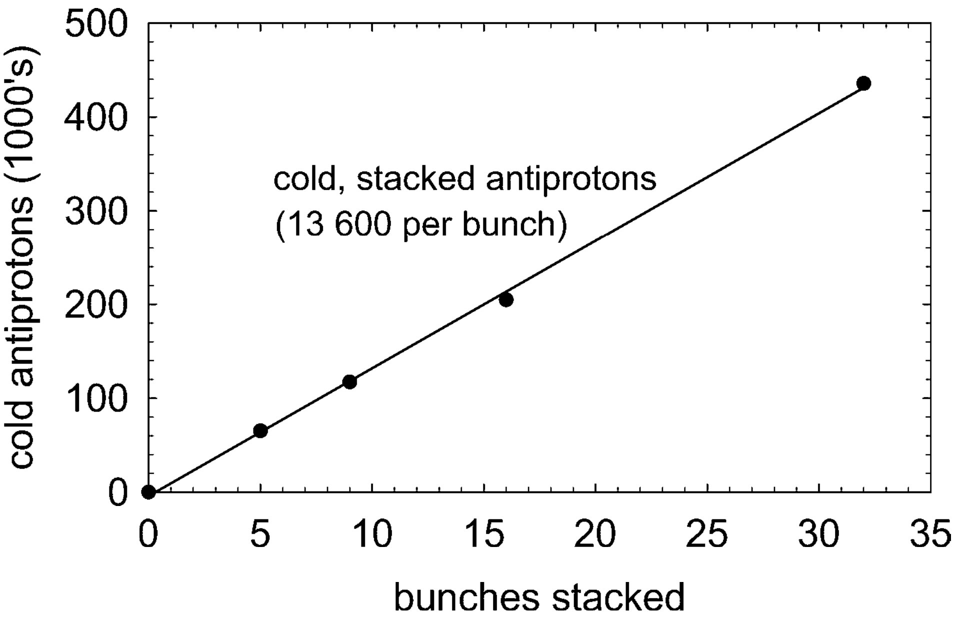

Figure 3.23: The number of cold antiprotons increases linearly with the number of bunches stacked in the trap.

an average of 13 600 per bunch. This is the largest number of cold antiprotons accumulated using the stacking technique. The AD was not operating at maximum intensity during this study and therefore larger antiproton numbers per bunch can be expected. Importantly, the number of cooled antiprotons increases linearly with the number of bunches delivered, indicating that more antiprotons can be trapped.

A careful account is kept of the number of antiprotons which are cooled, uncooled, and lost from the trap during stacking and is shown in Fig. 3.24. The fractional loss from the trap during cooling remains constant, however many bunches are stacked. This is because a constant fraction of each bunch are loaded far off axis and are collisionally ejected from the trap. However, the fraction of uncooled antiprotons reduces from 25% of the total after 5 bunches to 20% after 32 bunches. The fraction cooled

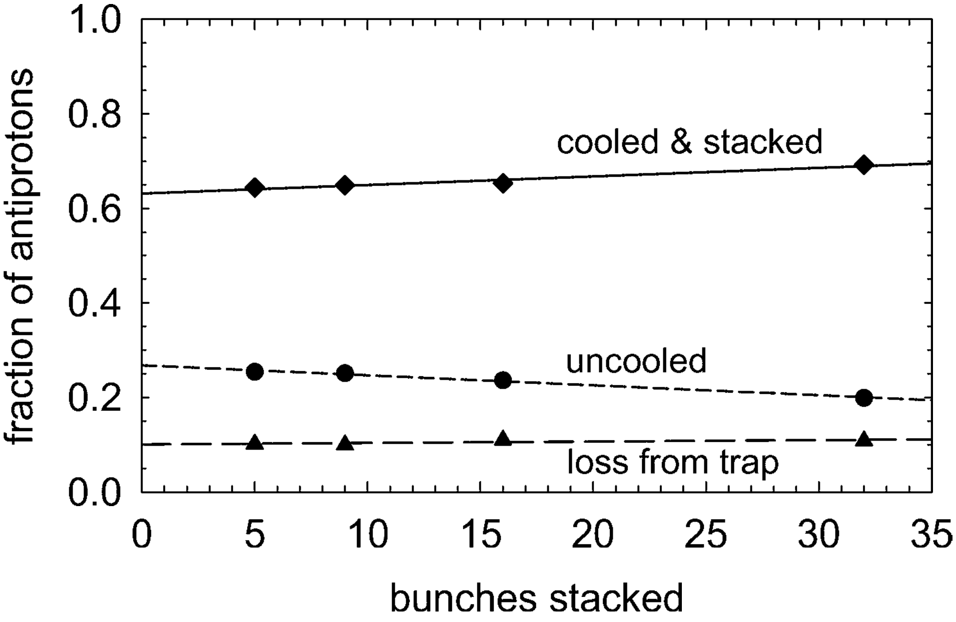

Figure 3.24: Analysis of the electron cooling and stacking capabilities.

rises from 65% to 70%. The same trend is observed in all the stacking studies made and is attributed to an enhanced cooling of hot antiprotons by antiprotons from previous bunches which have already cooled into the electron well. As more bunches are stacked this form of cooling becomes more efficient and the average cooling efficiency increases, as observed. Being of the same mass, antiproton - antiproton collisions can be highly efficient at transferring energy and clearly have a noticeable effect.

This effect is also observed by measuring the number of uncooled antiprotons released from the trap while a long stack is in progress. The reduction in this number due to increasing antiproton - antiproton cooling is clearly seen in Fig. 3.25.

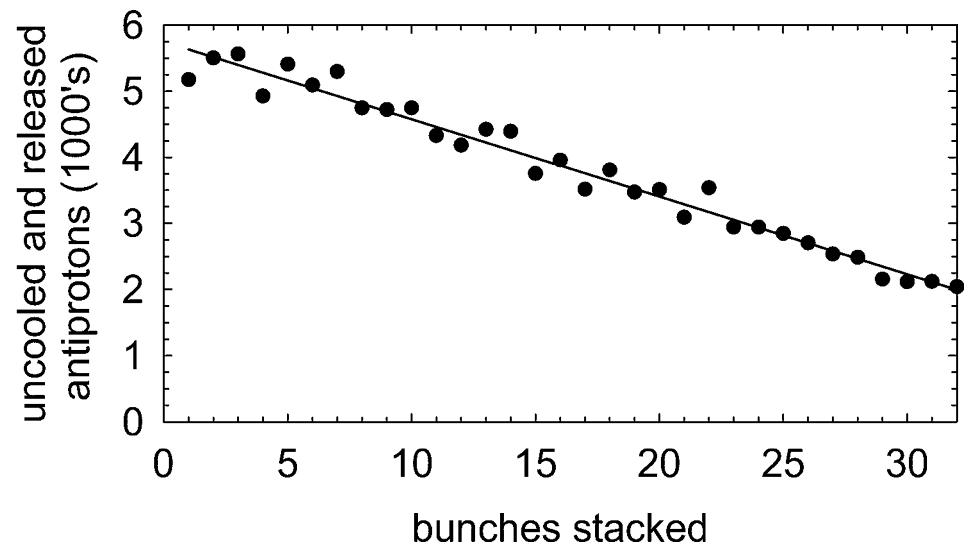

Figure 3.25: The number of uncooled antiprotons decreases linearly with the number of cooled antiprotons already in the trap.

3.5 Conclusions

Two methods are used to load electrons: the standard technique of loading from a field emission point and a new method using the positron source. This second technique is preferred for cooling large numbers of antiprotons since the electron clouds loaded in this way are more stable. The relative slowness of this method means that for short experiments with few antiprotons, electrons loaded with the FEP are more suitable. Robust and efficient positron loading is demonstrated and clouds of more than 1 million particles are routinely trapped. A loading rate of 3.5×10^4 positrons hr^{-1} mCi^{-1} is the most efficient method of trapping large numbers of positrons directly into a cryogenic UHV environment. Antiproton trapping, cooling to 4 K, and stacking are demonstrated. A fraction 7.4×10^{-4} of the high energy antiprotons delivered by the AD are trapped, the majority of which are electron-cooled to 4 K in 85 s. A careful analysis is made of the losses during the stacking process and it is observed that the accumulation of cold antiprotons increases linearly with the number of high energy antiprotons delivered to the trap. This stacking technique is vital to antihydrogen experiments where large numbers of cold antiprotons are required.

Chapter 4

Particle Cloud Shapes

Knowledge of the physical properties of antiproton and positron plasmas confined in our Penning trap is important for accurate calculations of recombination rates, for the interpretation of antihydrogen experiments, and to guide the design of future experiments. In the experiments described below, both positron and antiproton clouds are made to pass from one part of a Penning trap to another through the ball valve aperture. The efficiency with which the plasmas are transferred provides information about the radial extent of the plasma. A numerical code written by Spencer *et al* [78] solves Poisson's equation self-consistently for charged particles confined in a Penning trap, and is used to accurately determine the plasma cloud shapes from the transfer efficiency.

4.1 Pulse Transfer of Particles

It is well known that particles can be moved adiabatically from one standard Penning trap electrode to another [59]. For example, positrons held in a -10 V well on one electrode (T1, say) can be moved to the adjacent electrode (T2) by applying -15 V to T2 and then 0 V to T1. The voltage on T2 is then set to -10 V and the positrons have been successfully transferred one electrode distance. Typically this procedure takes several seconds. By repeating this process with electrodes T2 and T3 the particles are moved one further electrode, and so on. In this way particles can be moved anywhere within the trap without loss.

It has also been observed [9] that loss will occur if particles are held in a well formed by applying the same potential to two adjacent electrodes. In such “long” wells, as they are called, there is a significant volume within the well where the confining electric fields are small. Field inhomogeneities due to patch effects and electrode machining imperfections can then be of the same order as the confining fields which disturbs the normally excellent radial confinement. The magnetron motion of the particles is perturbed, leading to a radial expansion of the cloud and subsequent particle loss.

A similar problem arises if particles are confined in a potential well created by the ball valve electrode. The ball valve is longer and has a smaller diameter than a single standard electrode and hence also generates small confining fields. Transferring positrons through the ball valve using the standard adiabatic technique was unsuccessful [56] due to the loss when particles are confined within the ball valve. Successful transfer of positrons was achieved by speeding up the adiabatic technique to roughly 30 ms for the whole transfer. The quicker the transfer the less loss was observed [56], the limiting factors were then the filters on the DC lines and the power supplies.

The natural progression is to reduce still further the time particles spend within the ball valve. This can be done by "pulsing" them rapidly through. The pulsing technique is, however, qualitatively different to the adiabatic method since the particles are not confined on adjacent electrodes in turn, but travel rapidly through a field-free region from a well in the upper trap to one in the lower trap. In this way they are never confined on the ball valve at all, but simply pass through it in only about 20 ns, for positrons, and 1 μs for antiprotons. Figure 4.1 shows the transfer of positrons from electrode P2 to electrode T7 using this decidedly non-adiabatic technique. Initially, positrons are held on electrode P2 by applying 6 V to P1 and P3, and 3 V to P2 (similarly, T6, T7, and T8 have 6 V, 3 V, and 6 V applied to them). The voltage on P3 is then rapidly pulsed to 0 V allowing the positrons to escape to the right, accelerated by the sloping potential upon which they find themselves. At the same time the voltage on T6 is also changed to 0 V. The positrons move freely into the lower trap until they start to slow down as they climb the potential ramp provided by T7 and T8. When they have slowed to rest on this ramp the voltage on T6 returns to its initial value and the positrons are caught on T7. The voltage on P2 also returns to 6 V at this time. The transfer time, t_t, is 180 ns and so this procedure is accomplished by applying a -6 V pulse of duration 180 ns to the electrodes P3 and T6.

An additional benefit derived from this transfer technique is that any positive ions

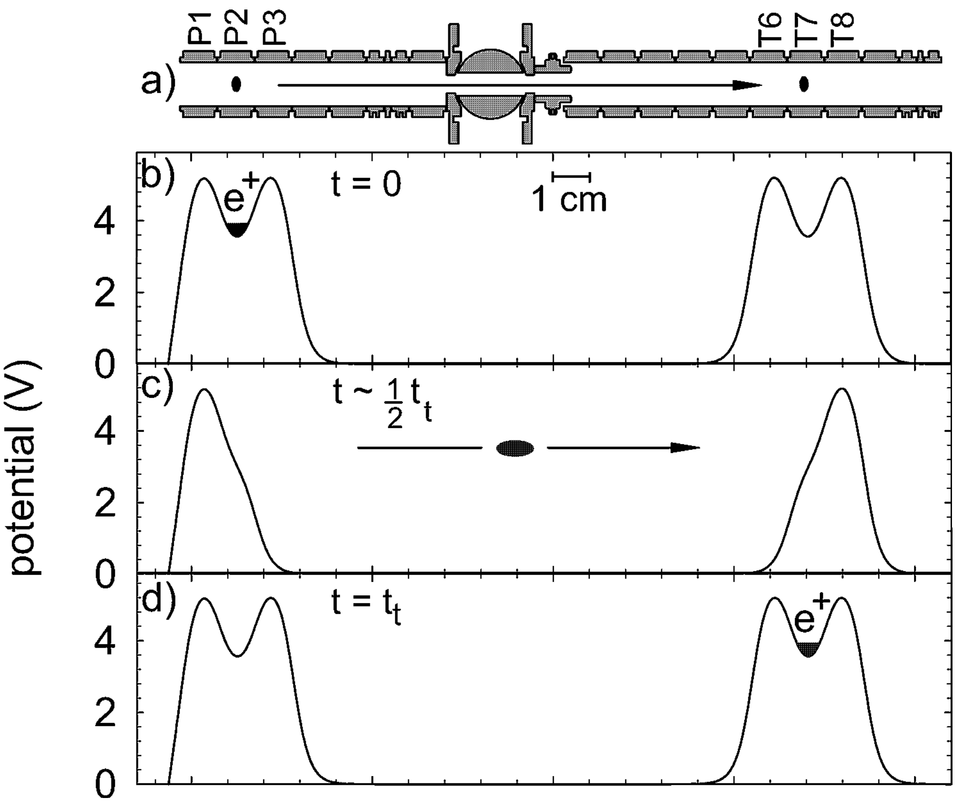

Figure 4.1: a) Electrodes, b) potentials before pulsing, c) pulsed potentials, and d) potentials return to the initial values.

which are trapped with the positrons will have a different t_t and therefore will not be caught in the lower trap. In fact, due to their great mass, the ions will not even have left P2 in 180 ns. They can then be safely ejected from P2 onto the transmission moderator. The problem of ions within a positron cloud is a great concern as they can collide with antiprotons during an antihydrogen experiment leading to antiproton loss. In some experiments (although not those described in Ch. 6) there is the possibility that these annihilations could be confused with annihilations from neutral antihydrogen escaping the trap, resulting in an incorrect claim of antihydrogen production.

4.1.1 Pulsing Hardware

A great deal of care is taken to use voltage pulsers which are low enough noise and have the necessary pulse amplitude and voltage rise and fall times. In addition, the voltage pulses must arrive at the electrodes with minimal attenuation and relative delay, and without suffering from reflections caused by impedance mismatching between different parts of the transmission line. The hardware setup is shown in Fig. 4.2. An externally triggerable SRS DG535 pulse generator is used to provide the initial

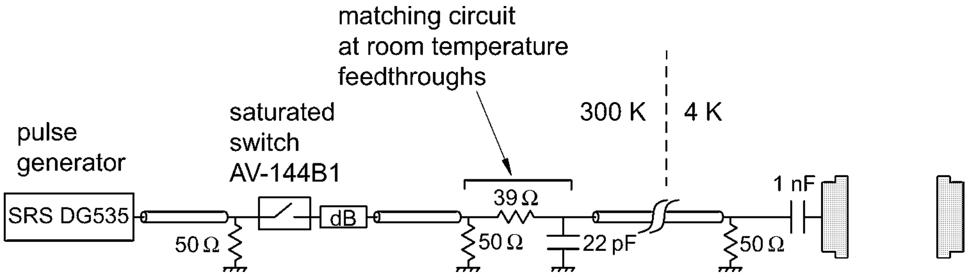

Figure 4.2: Setup used to provide fast voltage pulses to an electrode.

voltage pulse. These generators have rise and fall times of 2-3 ns and timing jitter of 60 ps. They are, however, limited to a pulse amplitude of only 4 V across 50 Ω. This pulse is therefore used to trigger a saturated switch (AV-144B1), the output of which has a fixed amplitude of +20 V and rise/fall times of 2-3 ns. Attenuators are placed on the output of the switch to vary the amplitude of the pulse. If a negative voltage pulse is required a toroid is added at the switch output. After the attenuator a short BNC cable carries the pulse to a matching circuit at the room temperature feedthrough. This reduces reflections as the pulse enters a 50 Ω microcoax transmission line traveling down to the pinbase. There, a 50 Ω resistor to ground provides further impedance matching before the pulse is a.c. coupled to the electrode.

Pulses applied to different electrodes can be arranged to have a relative delay of as little as 500 ps and rise/fall times of 3-4 ns at the pinbase. The maximum pulse amplitude at the pinbase is 8.5 V when the matching circuit at the room temperature feedthrough is used, and 12 V without this circuit. The pulse at the electrode itself will have very similar properties.

4.1.2 Pulse Transfer Capabilities

The transit time for pulsing a particle from one position (A) to another (B) in the trap is calculated using a Runge Kutta numerical integration technique (if $A = B$ the particles are re-caught on the starting electrode). The transit time calculation is tested by experiment as follows. A number of similar particle clouds are loaded and pulsed from A to B in the manner described above. Using slightly different pulse durations for each cloud allows the transfer efficiency as a function of the pulse duration to be measured. The calculated transit time is usually about 5 ns too short. This is because the effect of space charge is neglected in the calculation. The outer particles in the cloud screen the inner ones from the accelerating trap fields, thus slowing the launch of the inner particles, and giving a longer transit time for the cloud as a whole [18].

Figure 4.3 shows an example of the capabilities of pulsing and re-catching particles. In this particular example the particles are launched and caught on the electrode which is pulsed, rather than the adjacent electrode as shown in Fig. 4.1. The positrons start on P3 and are launched towards the ball valve when it is pulsed from -3 V to +3 V (from the solid to the dashed lines). They are reflected by a static potential ramp on Xtrap and XEB and return to P3 at the same time as the pulse returns the potential to -3 V. In this way the particles are re-caught on P3. As can be seen (Fig. 4.3(c)) it is possible to pulse and catch particles without loss. Although this example shows pulsing and catching positrons, the procedure works just as well for antiprotons, the only difference being the transit time is a factor $\sqrt{m_{\overline{p}}/m_{\overline{e}}} \approx 40$ times longer.

For the experiment of Fig. 4.3 the positron round-trip time is 114 ns. If the 6 V pulse is applied to P3 for a duration of 228 ns then the positrons will make two round-trips before being re-caught, a duration of 342 ns re-catches after three round-trips, and so on. Figure 4.4 shows the results when the positrons are allowed to make 1, 51 and 100 round-trips. The data indicates that the re-catch efficiency reduces and the timing curve broadens for a larger number of round-trips. However, one can still catch 80% of the positrons after 100 round-trips.

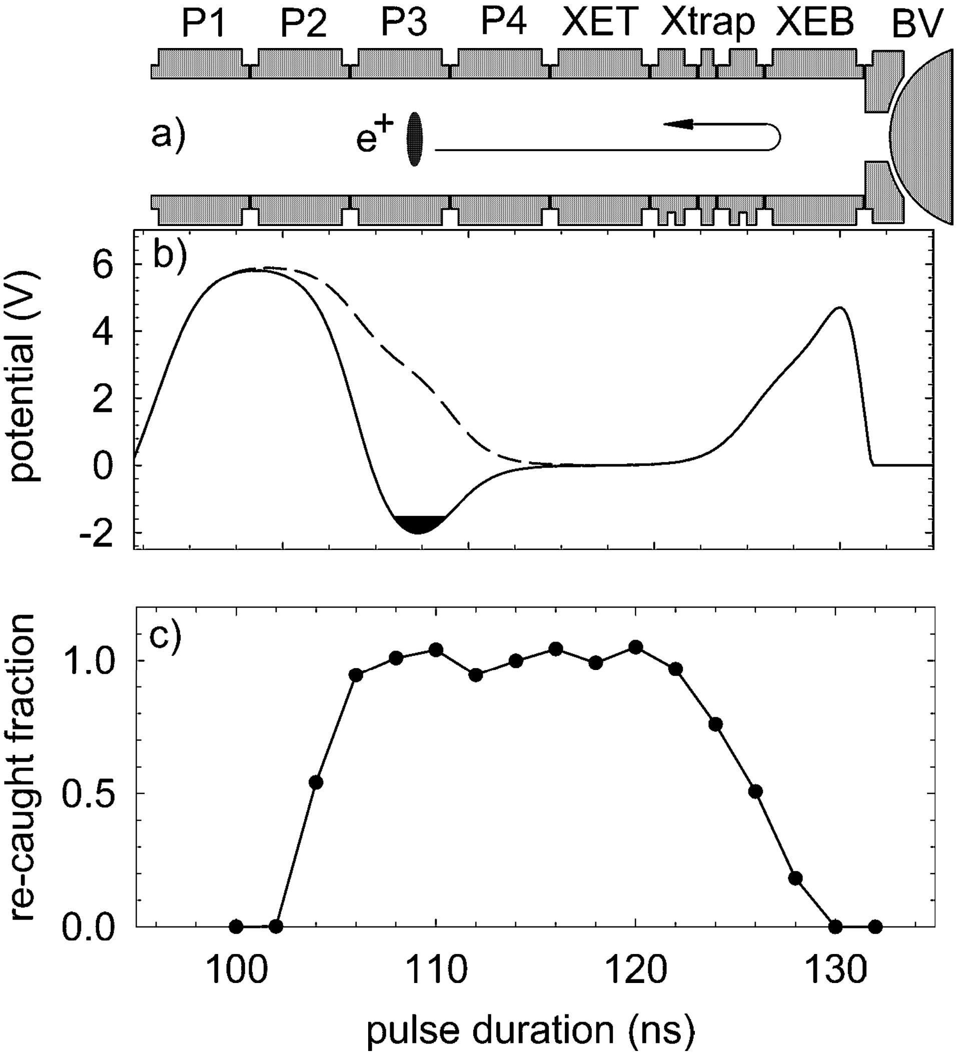

Figure 4.3: a) Electrodes, b) potentials used to pulse and re-catch positrons in the upper trap (dashed line is the pulsed potential), and c) the re-catch efficiency as a function of the pulse duration.

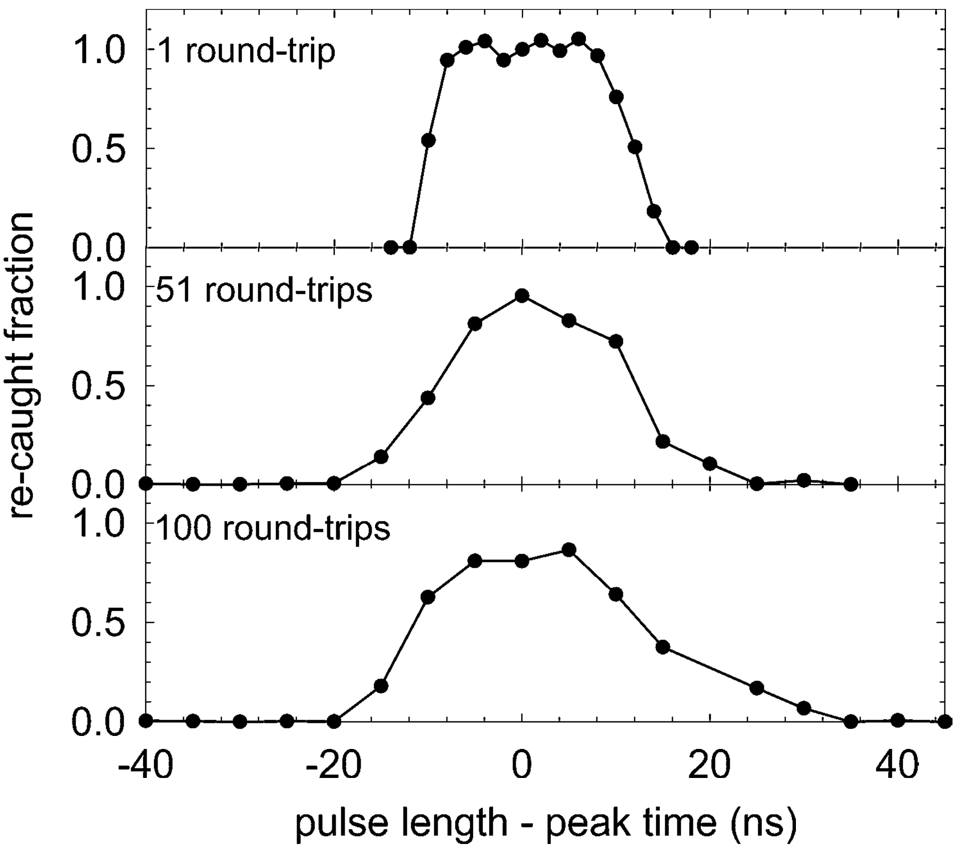

Figure 4.4: Re-catching positrons after multiple bounces. Each re-catch efficiency curve is centered on 0 ns to allow comparison.

Pulsing and catching particles within either the upper or the lower traps can always be achieved with 100% efficiency, but transferring between the traps is only 100% efficient when the number of particles is very small ($< 5 \times 10^4$ for positrons). Positrons at radii greater than 2.5 mm are striking the ball valve and so not passing through to the other trap. This effect is currently limiting the number positrons that can be transferred, as can be seen in Fig. 4.5. This figure illustrates that increasing the number of positrons does not proportionately increase the number transferred through the ball valve. There must be a corresponding increase in the radial dimension of the cloud which prevents proportionality. The exponential fit to the data places a limit on the maximum number of particles that can be transferred, that limit being about

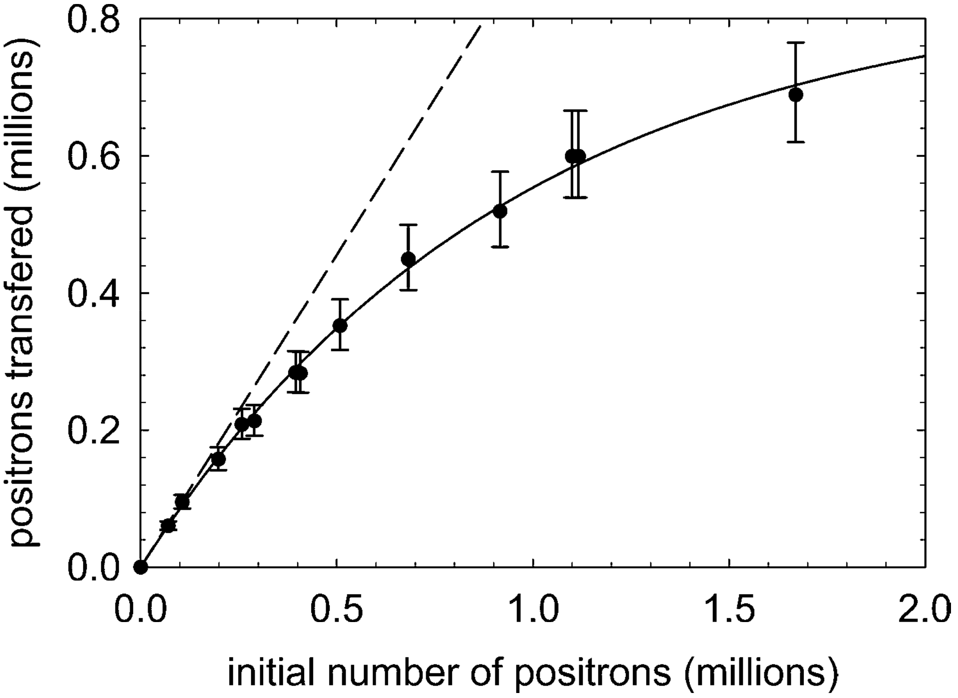

Figure 4.5: Number of positrons transferred to the lower trap as a function of the initial number in the upper trap. The dashed line is a linear fit to the first four data points.

0.85 million.

This clearly undesirable feature of the ball valve can, however, be used to advantage. It enables, for the first time, the physical parameters of the particle clouds to be measured. By physical parameters I mean the density, diameter, axial extent and angular momentum of the clouds. This will require an understanding of how the clouds can be treated as (non-neutral) plasmas of interacting particles.

4.2 Plasmas in Ideal Penning Trap Fields

Plasmas are collections of charged particles which exhibit collective effects, for example, density and shape perturbations which propagate through the plasma as waves. A particle cloud with density n and temperature T is considered a plasma when the dimensions of the cloud (d_c) are much larger than both the inter-particle spacing (a_s) and the Debye wavelength, λ_D, where,

$$\lambda_D = \sqrt{\frac{\epsilon_0 kT}{nq^2}}. \tag{4.1}$$

In this case a mean field approach can be used to describe the system. The Debye length is the distance over which the electric field of any individual charge is screened out by other charges in the plasma. This is easiest to understand in terms of a neutral plasma with equal numbers of positive and negative charges. Here, a positive charge will attract negative charges to surround it, and over a distance of λ_D these negative charges will screen the effect of the positive charge on the other, more distant, particles in the plasma.

Experiments detailed later in this chapter and summarized in Table 4.1 show that clouds of 1 million positrons and 1.5×10^5 antiprotons satisfy the requirements to be treated as plasmas. A comprehensive review of the properties of non-neutral plasmas confined in Penning traps is given by Dubin and O'Neil [79].

Plasma	λ_D (μm)	a_s (μm)	d_c (μm)
1×10^6 positrons	45	25	4000
1.5×10^5 antiprotons	70	40	440

Table 4.1: Length scales for plasmas of 1 million positrons and 1.5×10^5 antiprotons.

In an ideal Penning trap the electrostatic potentials vary quadratically in the axial (z) and radial (ρ) directions as seen in Eq. 2.1. If the plasma remains distant from the electrodes producing this potential the presence of image charges induced in the electrodes can be ignored. If, in addition, the plasma temperature is low enough (and hence λ_D small enough) the equilibrium plasma shape is a constant density spheroid which rotates as a rigid body with frequency ω_r. For a single particle ω_r is equal to the magnetron frequency but this is not the case for many particles since space charge modifies the rotation. The rotation, and the corresponding $q\mathbf{v} \times \mathbf{B}$ force, provides radial confinement. The spheroidal plasma has a diameter $2a$ and an axial extent $2b$ as shown in Fig. 4.6.

In the interior of an isolated charged spheroid the electric potential varies quadratically with the z and ρ coordinates. When a spheroidal cloud is confined by Penning trap fields the interior potential retains its quadratic dependence on ρ but is independent of z [80],

$$V(\rho) = \frac{m\omega_r}{2q}(\omega_c - \omega_r)\rho^2. \tag{4.2}$$

The z independence shows that there is no resistance to particle motion along the magnetic field lines within a cold plasma. Inserting Eq. 4.2 into Poisson's equation

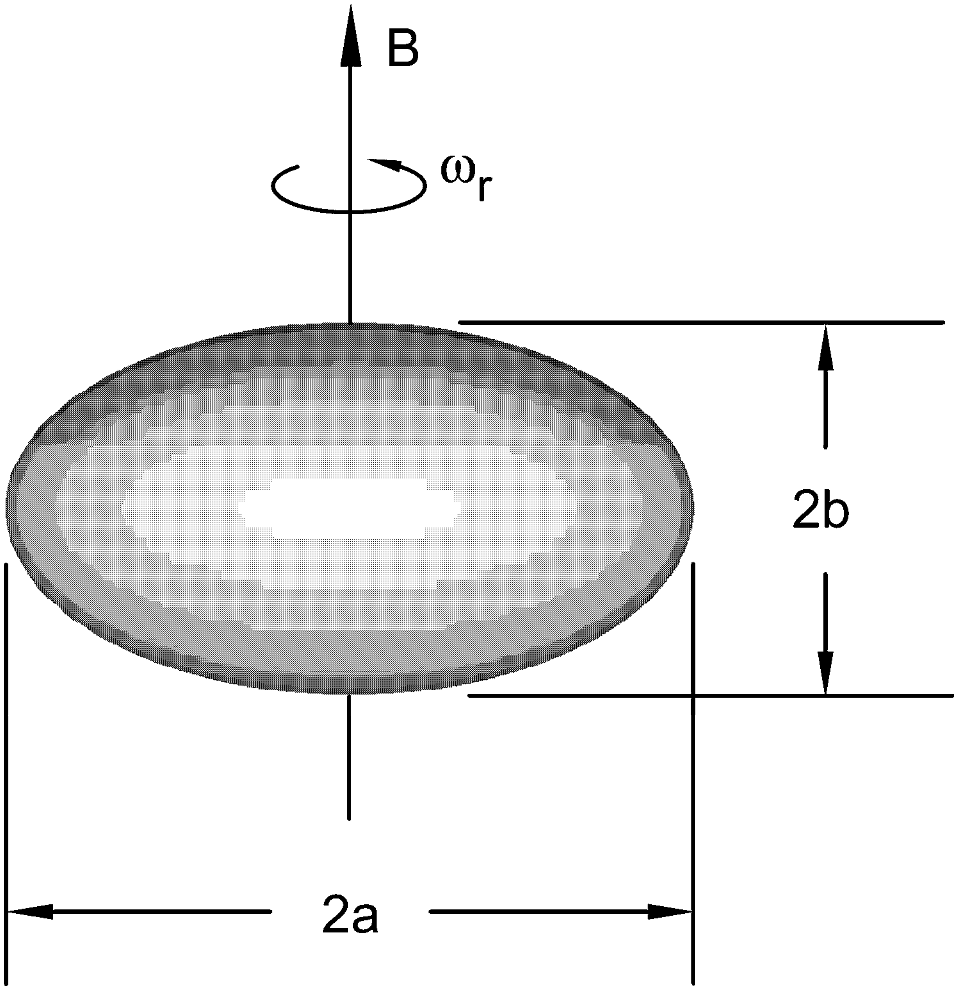

Figure 4.6: A cold spheroidal plasma confined in an ideal Penning trap.

results in a relationship between the density and rotation frequency of the plasma,

$$\omega_p^2 = \frac{nq^2}{\epsilon_0 m} = 2\omega_r(\omega_c - \omega_r), \tag{4.3}$$

where ω_p is the plasma frequency (the frequency at which density perturbations travel through the plasma). One can relate the plasma's axial oscillation frequency, the plasma frequency, and the plasma aspect ratio, α, by [81],

$$\frac{\omega_z^2}{\omega_p^2} = \frac{1}{\alpha^2 - 1}\left[\frac{\alpha}{\sqrt{1-\alpha^2}} tan^{-1}\left(\frac{\sqrt{1-\alpha^2}}{\alpha}\right) - 1\right], \quad \alpha = \frac{b}{a}. \tag{4.4}$$

Along with the straightforward relationship between density, volume, and particle number, N,

$$N = \frac{4\pi}{3} a^2 b n, \tag{4.5}$$

Eqs. 4.3 and 4.4 relate all the physical parameters of a cloud. The frequency ω_z is easily measured by non-destructive radio-frequency techniques or calculated (Eq. 2.4). For positrons N is also found non-destructively, while for antiprotons it has to be found by destructive techniques, as explained below. Once N and ω_z are known the above equations allow the determination of *all* cloud parameters if any *one* other parameter is known. Measuring the efficiency with which particles are pulsed through the ball valve aperture allows this to be done.

4.2.1 Spheroidal Clouds Cropped by an Aperture

If the radius of the ball valve aperture is denoted by a_{bv} and the particle clouds are assumed to be constant density spheroids, then the number of particles which are successfully transferred to the lower trap, N_f, is given by,

$$N_f = n\frac{4\pi b}{3}\left(a^2 + (a_{bv}^2 - a^2)\sqrt{1 - \frac{a_{bv}^2}{a^2}}\right). \tag{4.6}$$

This is simply the number of particles in the initial spheroid which are at a radius less than a_{bv}. An assumption implicit in this relationship is that the particles do not move radially during the transfer. The diffusion coefficient for particle transport across a strong magnetic field has been measured for magnesium ions [82]. Scaling this measurement to apply to positrons and antiprotons in the current experiment leads to a diffusion timescale $\sim 1 - 10$ s for both species. This is much longer than even the antiproton transfer time of 10^{-5} s and so no radial expansion of the cloud is expected during transfer.

The procedure for determining the positron cloud shapes is to load and count positrons in the upper trap, then pulse them to the lower trap and count again. Equations 4.3 to 4.6 can then be used to determine the values of n, a, b, and ω_r for the initial cloud. Positrons are held on electrode P2 in a well made by biasing P2 to $+V$ and biasing P1 and P3 to $+2 * V$. They are pulsed to the lower trap with a $-2V$ pulse on P3, where they are caught in an identical well structure made by electrodes T6, T7 and T8. This is shown in Fig. 4.1 with $V = 3$ V. A total of three different values of V are used to find how the cloud shapes vary with well depth. These values are 1.5, 3, and 4.25 V, and the well structures are referred to as 3-1.5-3 V, 6-3-6 V and 8.5-4.25-8.5 V well structures. Deeper wells are not used as they cannot be pulsed open using the limited amplitude saturated switches, while much shallower wells cannot hold large numbers of particles. The transit times for the three well structures are different and so a preliminary experiment is performed to determine these times experimentally. The potentials can be seen in Fig. 4.9(b).

As a first step, however, it is necessary that the radio-frequency detection of positrons in the upper and in the lowers traps give identical measurements of particle number for identical clouds. This is achieved by loading a cloud which is small enough to pass through the ball valve without loss and comparing the particle numbers measured in each trap. A single cloud of ≈ 18 000 positrons is loaded into the upper trap and counted. It is then pulsed between the two traps a total of 15 times, measuring the particle number after each pulse. The round-trip transfer efficiencies can then be calculated. If these ratios are 1 then the transfer is lossless, the number of positrons constant, and the calibration between the traps can be found. Figure 4.7 shows the results of this test. The transfer is not entirely lossless, the data showing

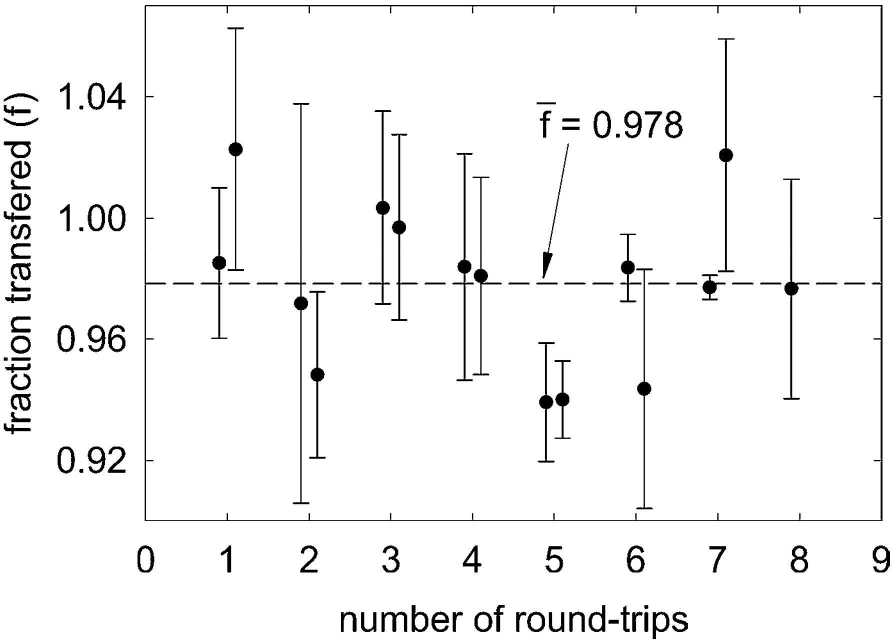

Figure 4.7: Re-catching a small positron cloud after multiple transfers through the ball valve.

an efficiency of 0.989 ($= \sqrt{0.978}$) per transfer, averaged over the 15 transfers. This enables the two traps to be calibrated to ≈ 1%. The actual number of positrons counted in each trap determines that a correction factor of 1.15 is needed to give identical number measurements in both traps for identical clouds:

$$N_{e^+}(lower\ trap) = 1.15 * N_{e^+}(upper\ trap). \tag{4.7}$$

Difficulties measuring the number of antiprotons in the trap using the standard radio-frequency techniques mean that the transfer efficiency for antiprotons is calculated in a different manner. First, a specific number of antiproton shots are stacked in the lower trap. The antiprotons are then simply ejected and counted using the annihilation detectors. Next, the same number of shots are stacked, the antiprotons are pulsed from the lower trap through the ball valve, where they reflect from a potential ramp on electrodes P4 and XET, and are re-caught in the lower trap (Fig 4.9(c)). The antiprotons are then ejected and counted. The transfer efficiency is the ratio of the two ejected clouds, normalized to the strength of the antiproton shots. It is necessary to eject both antiproton clouds from the lower trap since the trigger counting has different, and relatively uncalibrated, solid-angle efficiencies in the upper and the lower traps.

To investigate the effect of passing the antiprotons twice through the ball valve, rather than just catching them on the other side, an additional set of experiments are performed. In these experiments the antiprotons are re-caught after a total of four passes (three reflections) through the ball valve. The results are shown in Fig. 4.8. It can be seen that there is some loss, in addition to the initial clipping of the antiproton cloud, due to further passes through the ball valve. This is presumably due to a small shape change of the cloud upon reflection from the potential barrier in the upper trap which allows antiprotons that *just* passed through the ball valve on one pass to be clipped on the next pass. The timescale on which the cloud shape relaxes from

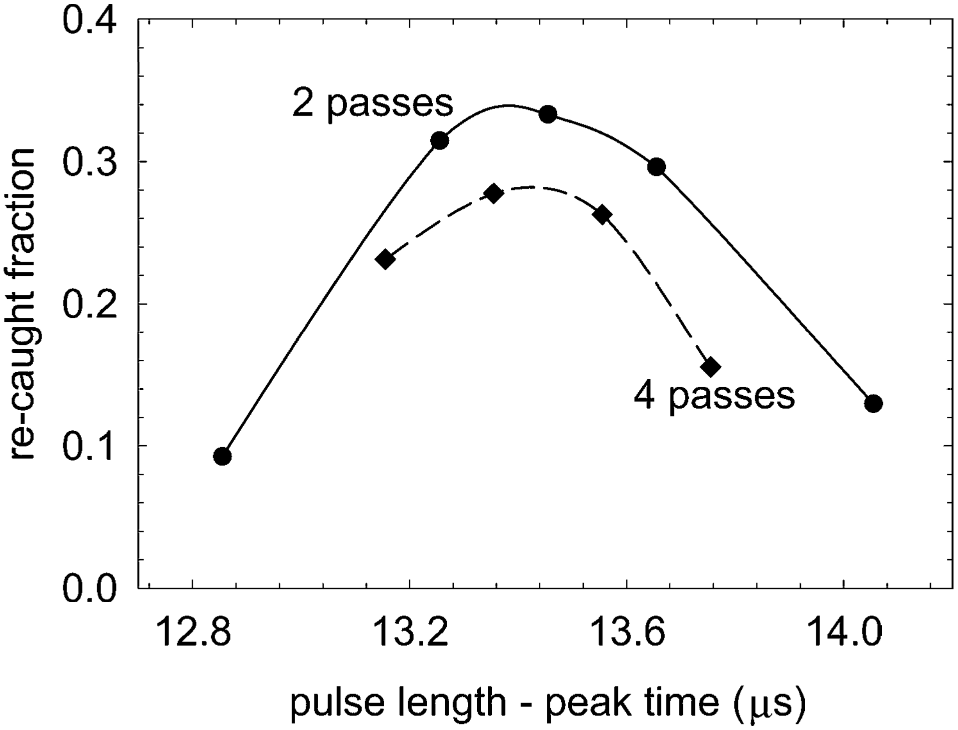

Figure 4.8: Re-catching antiprotons after two and four passes (one and two round-trips) through the ball valve.

a non-equilibrium shape, such as that created by pulsing through the ball valve, is determined by the plasma frequency. This is compared to the time it takes for a cloud to reflect from the potential barrier. If the reflection time is very much less than $1/\omega_p$ then the cloud will not have time to react and the plasma shape will remain unchanged. For these experiments the reflection time and $1/\omega_p$ are both of order 1 μs, so it is not surprising that some further clipping occurs. Assuming that, after the first pass, a constant fraction of antiprotons are lost per pass through the ball valve, then the results in Fig. 4.8 determine this fraction to be 0.095. The subsequent analysis takes this into account.

The antiproton cloud shapes are only measured using one well depth, $V = -3$ V, and are reflected from a barrier in the upper trap by biasing P4 to -6 V and XET to -3 V. Figure 4.9 illustrates the potentials used for the studies of both the positron and the antiproton cloud shapes.

The transfer efficiency for positron clouds containing between 5×10^4 and 1.5×10^6 positrons, in each of the three well depths is measured. As is the efficiency for antiproton clouds of between 10^4 and 1.4×10^5 particles. Analysis of this data using Eqs. 4.3 to 4.7 shows that the cloud diameters can be as large as 9 mm for positrons and 11 mm for antiprotons. Such a large diameter is cause to question the validity of the assumptions that the particles are confined in exactly quadratic potentials imposed by distant electrodes. The electrodes (inner diameter of 12 mm) are clearly not distant, nor are the potentials purely quadratic as far off axis as the radius of the particles clouds. This simple analysis, having given an estimate of the cloud parameters, now has to be refined in order to more accurately describe the particle

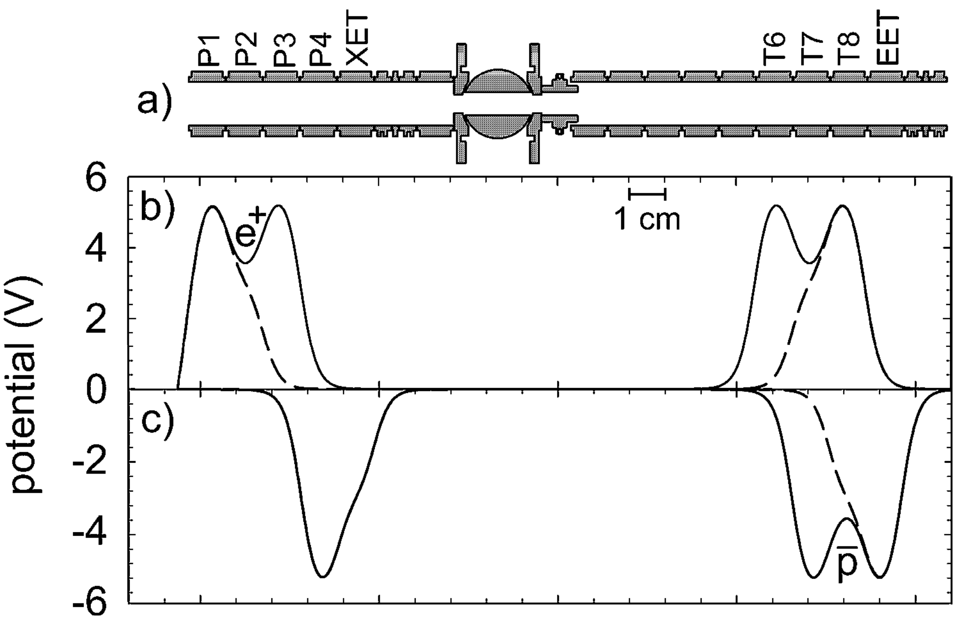

Figure 4.9: a) Electrodes, b) potentials used for pulsing positrons, and c) antiprotons through the ball valve. The dashed lines are the pulsed potentials.

clouds.

4.3 Refined Analysis of Cropped Particle Clouds

4.3.1 Numerical Code to Analyze Cropped Particle Clouds

A numerical code called Equilsor2 has been written by Spencer *et al* [78] which solves Poisson's equation self-consistently for charged particles confined in electrostatic fields produced by cylindrically symmetric electrodes. Parrott [83] adapted this code for our purposes. The essential new features are that the electric potential is not now assumed to be perfectly quadratic everywhere but is calculated numerically for a given set of electrode voltages, and the presence of image charges is also included.

The code can be run in a number of different modes, two of which are used in the following calculations. The first, called "thermalN" mode, has the number of particles and the radial extent of the plasma as two fixed input parameters. The density, aspect ratio, and angular momentum of the cloud are found self-consistently, while the particle number and radius in the code output are the same as the input values. The second mode is called "thrmcnsv" which fixes the particle number and the momentum of inertia, I, of the cloud, (I is proportional to the angular momentum) and solves for the density, radius and aspect ratio.

An input "*.run" file is used to define the electrode potentials, the number of radial and axial grid points used in the calculation, the spatial extent of the computation region, the input density, radial and axial extents of the cloud, and also the temperature of the plasma. For both species a temperature of 4.2 K is appropriate. To fix N is not quite as simple as defining it directly in the run file, since n, a, and b are the input variables. The code is written such that one utilizes the relationship between density, volume, and number for a spheroid in order to define N,

$$N = \frac{4\pi}{3} a^2 b n. \tag{4.8}$$

From the above equation it is seen that fixing N and a only fixes the product $n*b$. To test the code it was run in thermalN mode with a number of different values of n and b, while keeping $n * b$ fixed. Reassuringly, the code output was identical for all these combinations. The code was also tested by comparing its output for small values of N and a to the predictions using Eqs. 4.3 to 4.5. The results were the same, as should be the case for small N and a where image charges and non-ideal Penning fields are negligible. For the calculation of the positron cloud parameters radial grid points are separated by 50 μm and axial grid points by 75 μm. The antiproton clouds have the same number of radial grid points but twice as many axial points, since these clouds are rather thin and so require higher axial resolution. Doubling the number of grid points does not alter the code output for either species, neither does varying the size of the computation region.

4.3.2 Positron Clouds

The thermalN mode is used to determine the equilibrium positron cloud shapes. The particle number is measured using the non-destructive radio-frequency techniques described in Ch. 2, and the radial extent is used as a variable input parameter. The procedure is as follows: for a particular known particle number guess a possible radial extent. Run the code using these two values as input parameters. Repeat for a number of other possible radii. In this way one builds up a set of output clouds, each with a different density, radius, aspect ratio and angular momentum, but with the same particle number. Then, for each cloud, integrate the number of particles which are at a radius less than that of the ball valve. This is the number of particles which would pass through the ball valve and hence determines what the transfer efficiency would be for each particular output cloud. These efficiencies can then be compared to the experimentally measured efficiency in order to determine the actual radius of the cloud. This is illustrated in Fig. 4.10 for a cloud of 1.4 million positrons in a 6-3-6 V well. In this case, the experimental transfer efficiency, $f_{expl} = 0.436$, gives a cloud radius of, $a_{expl} = 4.17$ mm. Running the code once more with an input radius of 4.17 mm gives the correct cloud shape for 1.4 million positrons in a 6-3-6 V well,

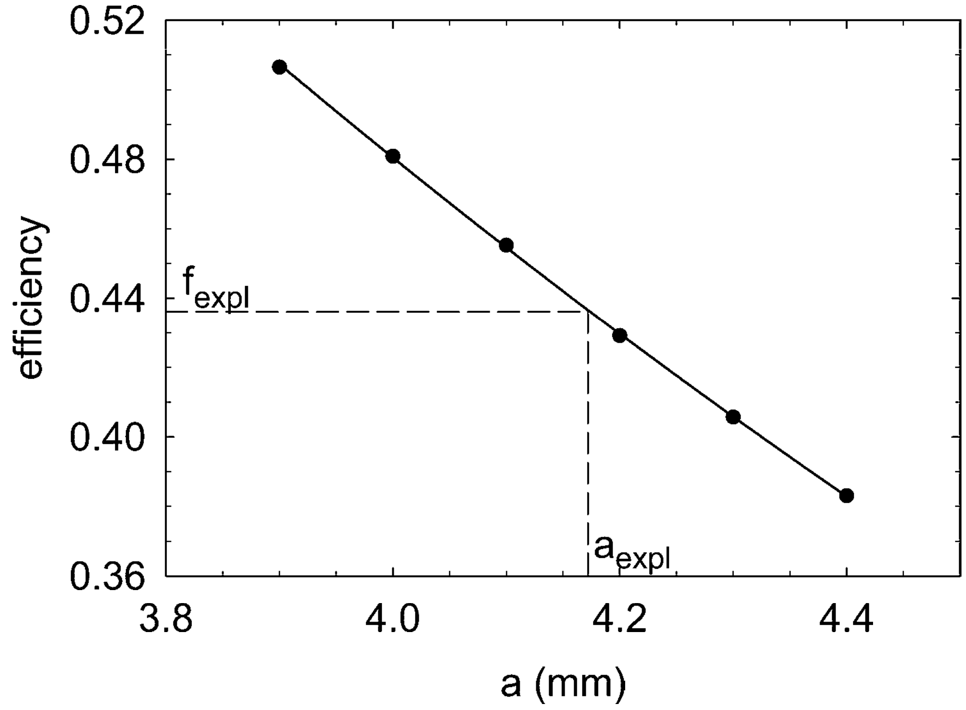

Figure 4.10: Numerical calculation of the transfer efficiency for a cloud of 1.4 million positrons with a given radius. Each point on the graph represents a numerical calculation of the transfer efficiency through the ball valve.

which is shown in Fig. 4.11.

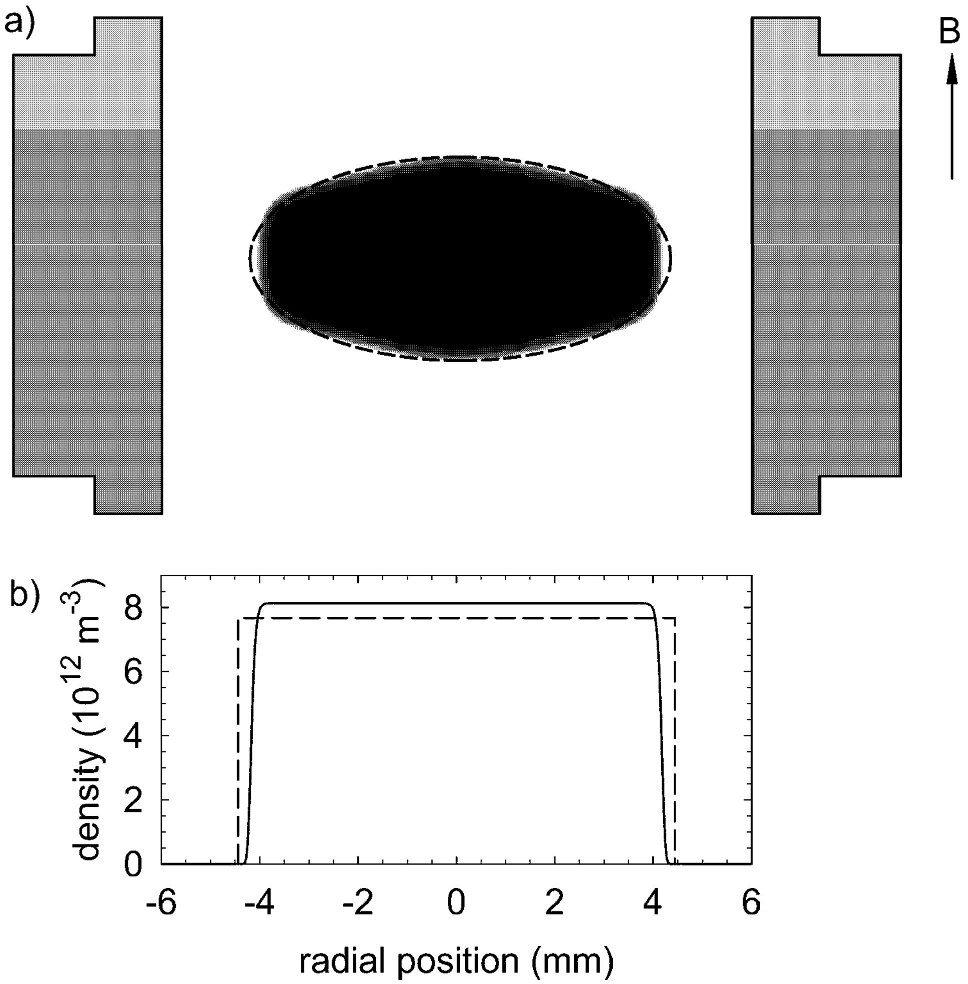

Figure 4.11: a) Contour plot showing the density of 1.4 million positrons held in a 6-3-6 V well and the cloud shape calculated from Eqs. 4.3 to 4.7 (dashed line), b) midplane density profile from the numerical calculation (solid line) and Eqs. 4.3 to 4.7 (dashed line).

As can be seen the actual positron cloud shape is slightly different to that which is calculated using Eqs. 4.3 to 4.7. The main difference is that the outer edge of the plasma is flattened in comparison with the spheroidal cloud. This is because the confining axial electric field this far off axis is weak so the cloud shape tends to follow the rather "flat" potential variation close to the electrode. Image charges induced on the electrode surface will provide an attractive force on the trapped positrons which is stronger for particles closer to the electrode. This will have the effect of attracting the closer particles closer still, and elongating the plasma radially. Clearly the effect of the flat confining potentials dominates the space charge effect. In Fig. 4.11(b) the density variation along the midplane of the cloud (the $z = 0$ plane) is constant with an abrupt fall to zero at the edge. The density falls to zero over a distance of order the Debye length, which is 50 μm in this case. The distribution for a cold spheroidal plasma is shown as the dashed line. In the $T = 0$ limit the Debye length is zero and hence the density discontinuously becomes zero at the plasma edge.

In addition to the self-consistent solution of Poisson's equation for the positron cloud shown in Fig. 4.11, other clouds with different numbers of positrons confined in three different depth wells are pulsed through the ball valve and analyzed numerically. An error bar of ±10% on the transfer efficiency is used in the data to account for the uncertainty in the measured particle numbers. The results are summarized in Figs. 4.12, 4.13, and 4.14 and the fits to the data guide the eye.

For any one well depth, as the number of positrons increases, so do the dimensions of the cloud due to the repulsive Coulomb interaction between the particles. The density also increases. The variation of the cloud parameters with well depth is understood in the following way. First, consider the angular momentum of the cloud, L. There are two contributions, one from the particles and one from the fields:

$$L = m\omega_r \sum_{i=1}^{N} \rho_i^2 + \frac{qB}{2} \sum_{i=1}^{N} \rho_i^2. \tag{4.9}$$

The first term is the mechanical angular momentum of the particles' rotation about the trap axis at frequency ω_r. The second term is the angular momentum in the

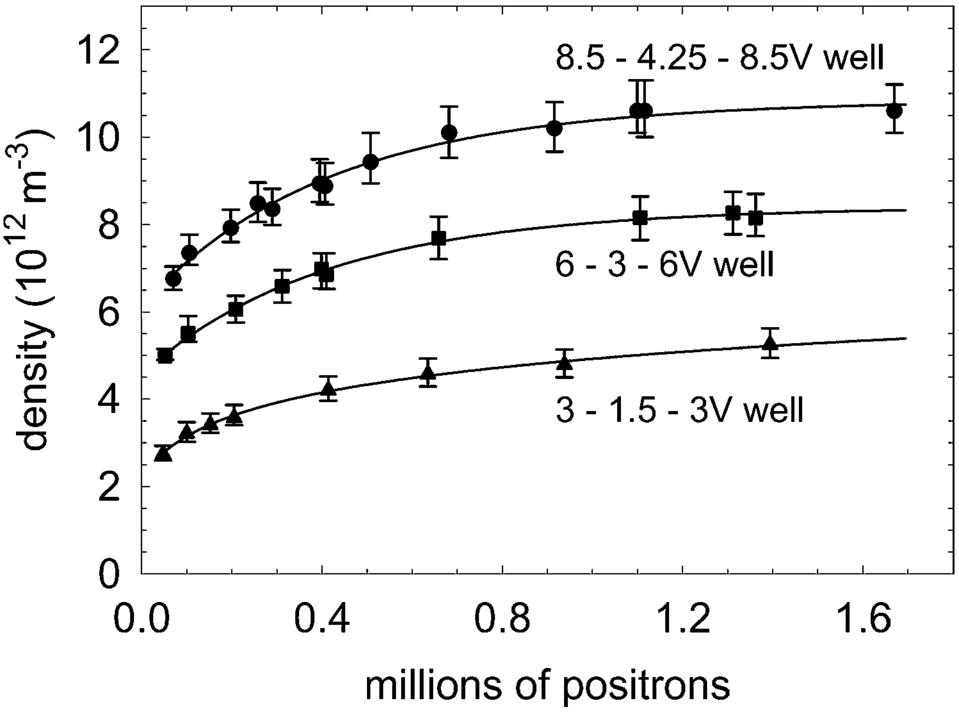

Figure 4.12: Variation of positron density with particle number for three different well depths.

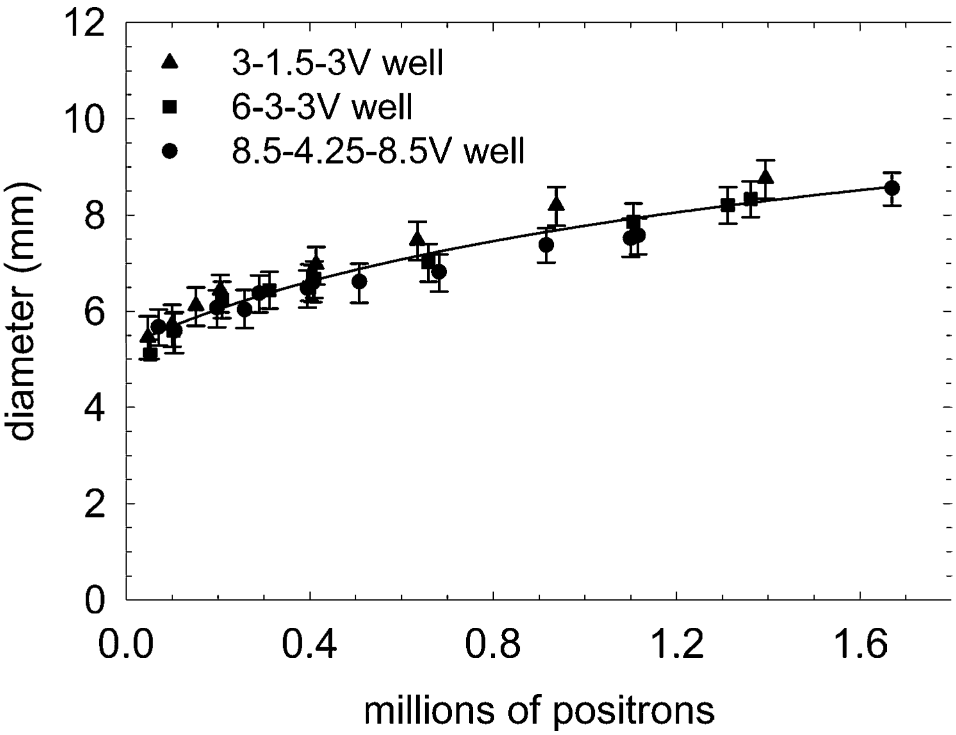

Figure 4.13: Variation of positron cloud diameter with particle number for three different well depths. Only one curve is fit to the entire data set.

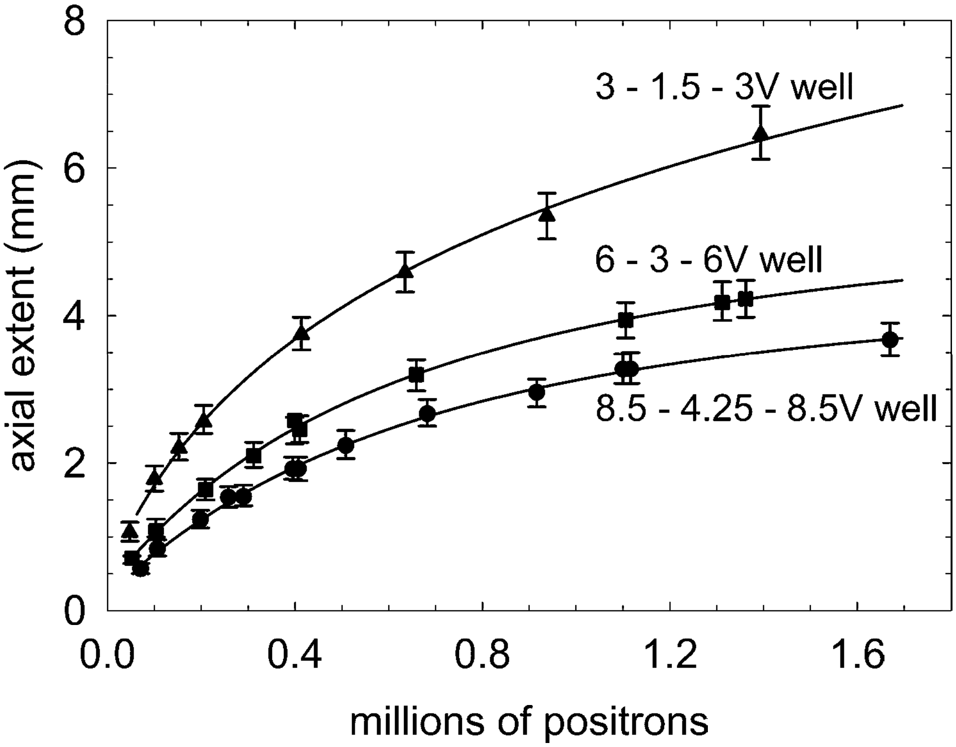

Figure 4.14: Variation of positron cloud axial extent with particle number for three different well depths.

electromagnetic fields, and is greater than the first term by the factor,

$$\frac{\omega_c}{2\omega_r} \gg 1. \tag{4.10}$$

Hence

$$L \approx \frac{qB}{2}\sum_{i=1}^{N}\rho_i^2 = \frac{qB}{2}N\rho_{rms}^2. \tag{4.11}$$

The self-consistent code defines the diameter of a cloud as twice the distance from the cloud center to the point where the density is half the central density. Diameter is therefore not quite equal to twice ρ_{rms}. In the absence of machining errors, trap misalignment, background gas atoms etc., the angular momentum of the cloud does not change in time – if some particles drift outwards, others must move inwards. The angular momentum also stays constant if the well depth is changed, due to the cylindrical symmetry of the electric field. Therefore, when the well depth is increased ρ_{rms} is unchanged but the axial extent of the plasma will reduce due to the larger axial electric field. The density will then increase. These effects are seen in the Figs. 4.12, 4.13, and 4.14.

In the limit of an infinitesimally thin "pancake" cloud the rotation frequency is equal to the single particle magnetron frequency, ω_m. Then, from Eqs. 4.3, 2.2, and 2.3,

$$n = \frac{2\epsilon_0 m\omega_m}{q^2}\omega_c' \approx \frac{\epsilon_0 m}{q^2}\omega_z^2. \tag{4.12}$$

From Figs. 4.13 and 4.14 one can see that as the number of particles approaches zero the axial extent also approaches zero, while the diameter remains finite. Therefore the aspect ratio, $\alpha \to 0$, and the cloud is pancake shaped. Extrapolating the fits of the density data to zero particles therefore gives three values of n which can be compared with Eq. 4.12. Figure 4.15 shows that the measured values for the minimum density of a plasma in the three well depths agree with the theoretical prediction.

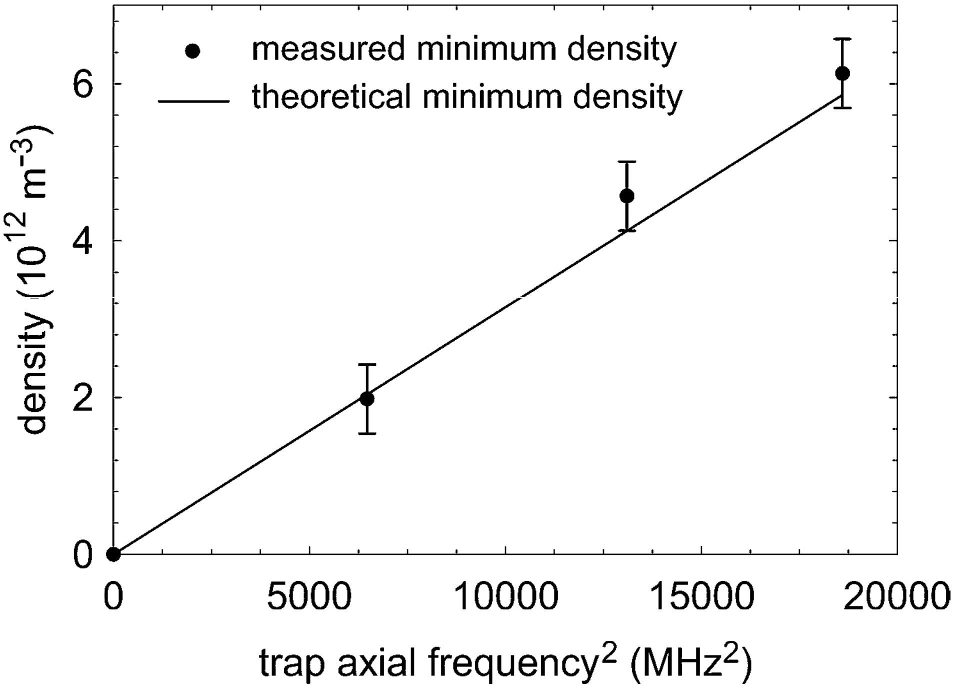

Figure 4.15: The experimentally determined minimum cloud densities agree well with the theoretical prediction, confirming our understanding and calculations of thin plasmas.

4.3.3 Antiproton Clouds

The procedure for determining the antiproton cloud shapes given an experimental transfer efficiency through the ball valve is exactly the same as that used for the positron clouds. The density distribution and cloud shape for 1.4×10^5 electron-cooled antiprotons (a 12 bunch stack) is shown in Fig. 4.16. The antiproton cloud is not at all spheroidal. The outer edge of the plasma is flattened because the confining axial potential is weak so close to the electrode. This also allows the cloud to spread

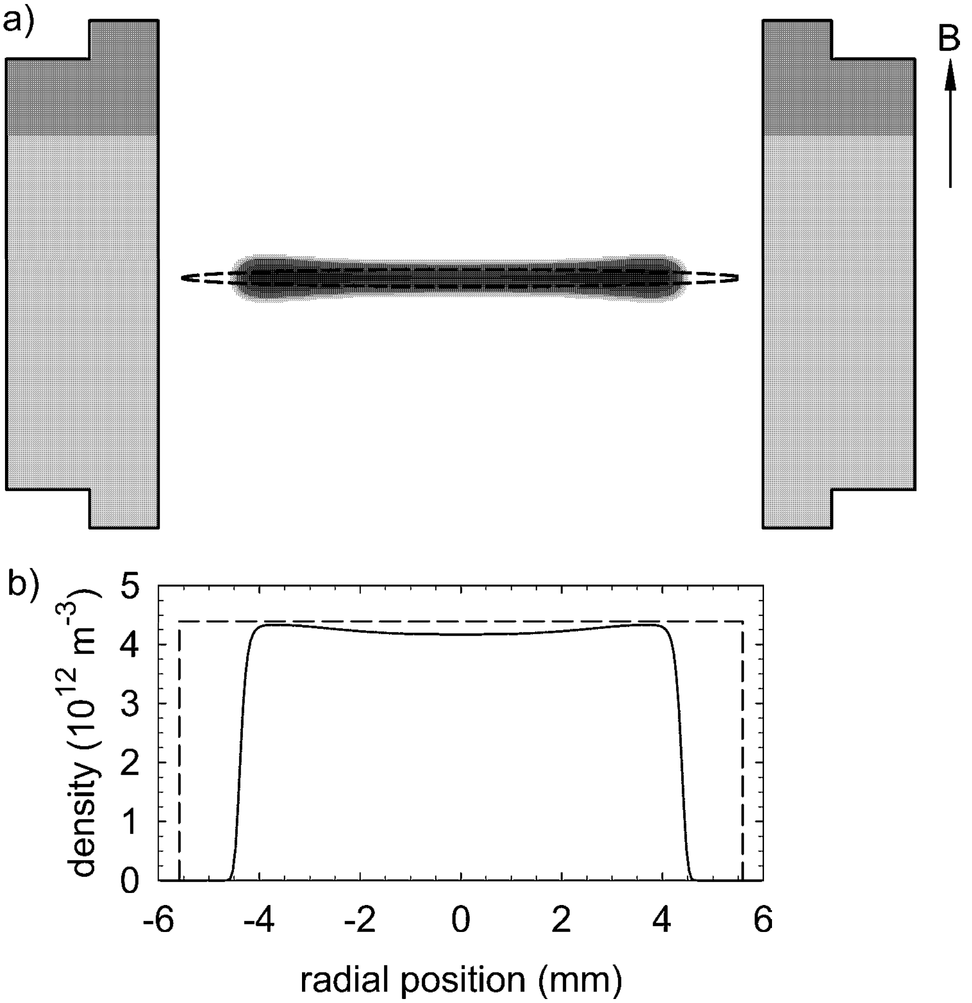

Figure 4.16: a) Contour plot showing the density of 1.4×10^5 antiprotons held in a 6-3-6 V well and the cloud shape calculated from Eqs. 4.3 to 4.6 (dashed line), b) midplane density profile from the numerical calculation (solid line) and from Eqs. 4.3 to 4.6 (dashed line).

out giving it a slightly "dog bone" shape - a shape similar to that found theoretically by Prasad and O'Neil [84]. It is unlikely that the image charges have a substantial effect since the number of antiprotons in the cloud is rather small, and hence the number of image charges is also rather small.

The particular cloud shown in Fig. 4.16 was cooled by 4.5 million electrons loaded from the radioactive source. Measurements of the transfer efficiency through the ball valve have been made with electrons loaded from the source. These show that the electrons have densities, diameters and axial extents which are virtually the same as those of the positrons in Figs. 4.12 to 4.14. Extrapolating the diameter fit in Fig. 4.13 to 4.5 million particles gives an electron cloud diameter of 10.1 +1.9 -1.8 mm. This is compared to the antiproton cloud diameter of 8.8 ± 0.2 mm. When high energy antiprotons first enter the trap after being slowed in the degrader material, they will cover the entire diameter of the electrodes. However, only those particles which are at a radius less than that of the electron cloud become cooled into the electron well and the calculated electron and antiproton diameters are consistent with this assertion. Once the antiproton cloud diameter is fixed by the size of the electron plasma all other parameters follow from the number of antiprotons in the cloud.

Several other clouds, containing different numbers of antiprotons were pulsed through the ball valve and analyzed using the Equilsor2 code. These antiprotons were cooled by electrons loaded from the field emission point and the results are summarized in Figs. 4.17, 4.18, and 4.19. The error bars reflect a ±10% uncertainty in the transfer efficiency due, primarily, to loading and stacking repeatability, and uncertainties in normalizing the antiproton shot strength. The data is fit to a linear

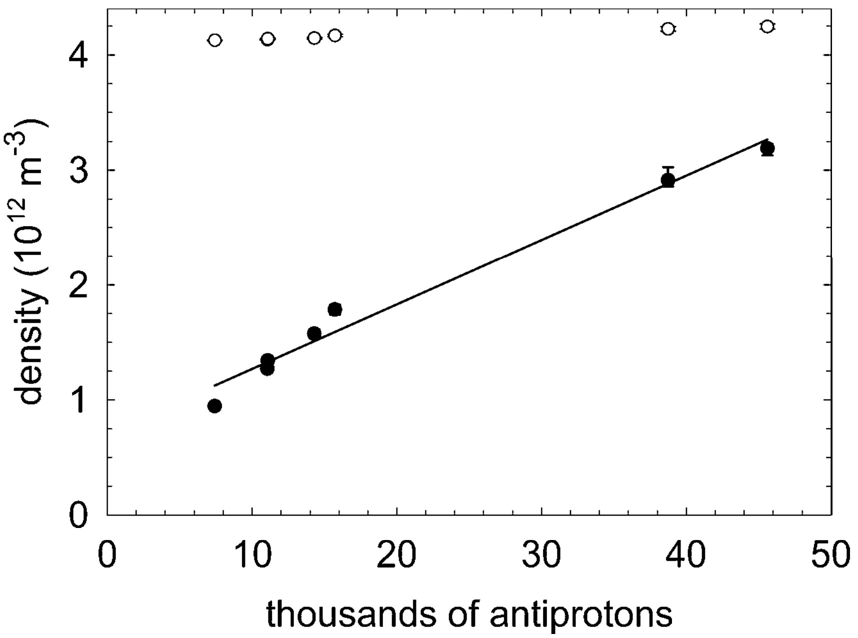

Figure 4.17: Variation of antiproton mean density with particle number. Filled points are results from the Equilsor2 code, hollow points are calculated from Eqs. 4.3 to 4.6.

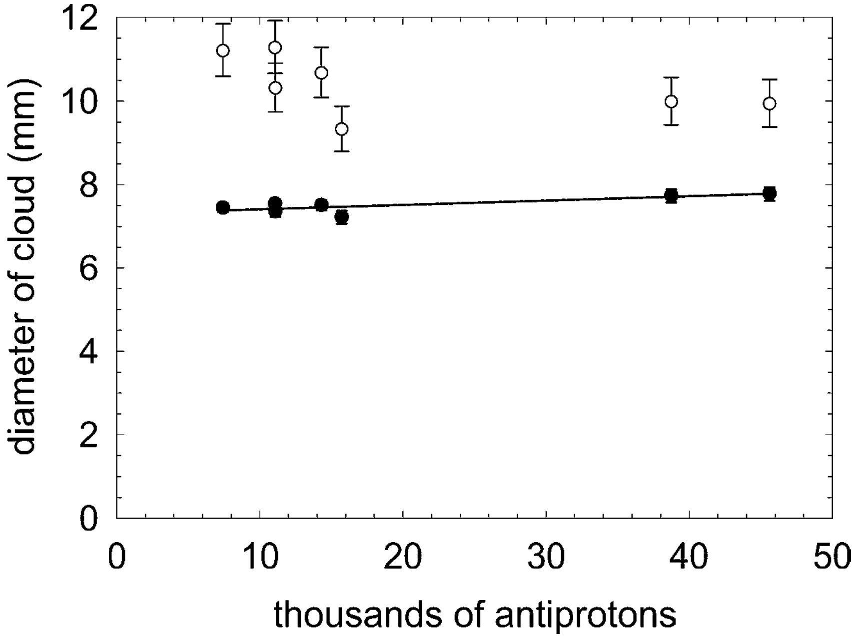

Figure 4.18: Variation of antiproton cloud diameter with particle number. Filled points are results from the Equilsor2 code, hollow points are calculated from Eqs. 4.3 to 4.6.

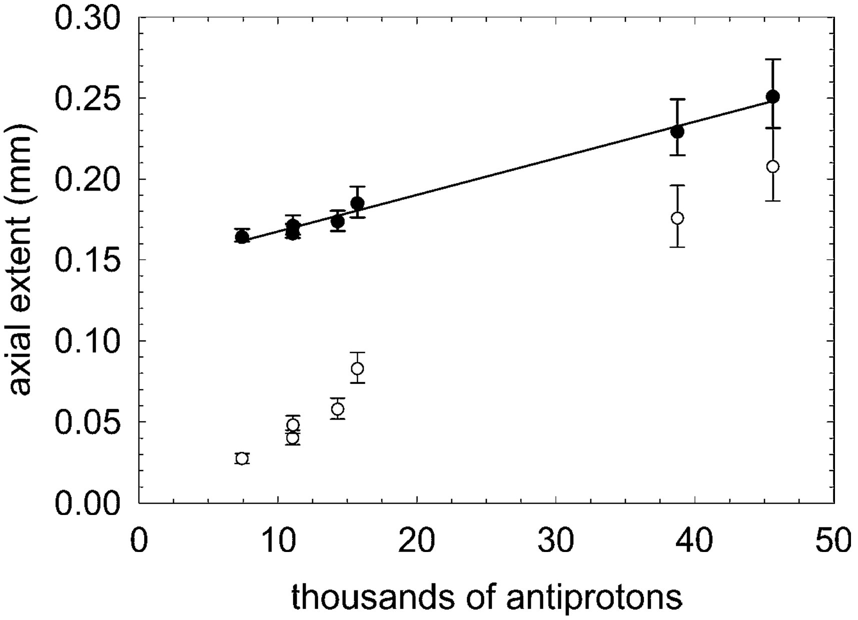

Figure 4.19: Variation of antiproton cloud axial extent with particle number. Filled points are results from the Equilsor2 code, hollow points are calculated from Eqs. 4.3 to 4.6.

variation with particle number. Also shown in these figures are the cloud parameters determined by the analytical approach of Eqs. 4.3 to 4.6 which assume the antiprotons are cold and confined by exact quadratic potentials. The ±10% uncertainty in the transfer efficiency can be seen in Fig. 4.18 to give a significantly larger uncertainty in the diameters of clouds analyzed according to Eqs. 4.3 to 4.6, when compared to the numerical analysis. This is because the antiproton clouds analyzed numerically have more particles concentrated at larger radii (see Fig. 4.16 and Fig. 4.20), while a spheroid has more particles at small radii. A 10% change in the transfer efficiency is therefore accounted for by a smaller change in the diameter of the numerically analyzed clouds than the spheroidal clouds.

One can see that even the smallest clouds have large diameters ($\approx$ 7.5 mm), supporting the claim that the electron cloud size determines the antiproton cloud size. As the number of antiprotons in the cloud increases so the axial extent of the cloud also increases due to space charge repulsion between particles. The density also grows with particle number. There is a significant difference between the results determined by the numerical code and by the simplifying of a cold plasma confined by perfect quadratic potentials. This is mainly because the clouds are always loaded with a large diameter and so experience non-ideal Penning fields.

One should be cautious drawing too many conclusions from this data since for the smallest clouds the axial extent of the plasma actually falls slightly below the Debye length and becomes comparable to the inter-particle spacing. The Equilsor2 code does not require the plasma size to be much greater than the Debye length, but it does assume that the cloud can be treated using a mean field approach [85]. When a cloud is only a few particles thick it is difficult to imagine how this is possible. The thermal energy of the particles will, however, smear the particle distribution, returning it towards a situation where a mean field approach is imaginable. Similar numerical calculations used to calculate the frequency of drumhead modes in equally thin plasmas have been able to reproduce the experimentally determined frequencies to within $\sim 10\%$ [86]. This gives confidence that the numerical calculation of antiproton cloud shapes presented here is satisfactory.

A feature of the antiproton clouds is that the midplane density is lower on axis than off-axis, as seen in Fig. 4.20. Due to the thermal energy of the antiprotons the thickness of a cloud cannot fall much below λ_D [87]. For fewer and fewer particles this lower limit on the axial extent means that the central density must decrease. A

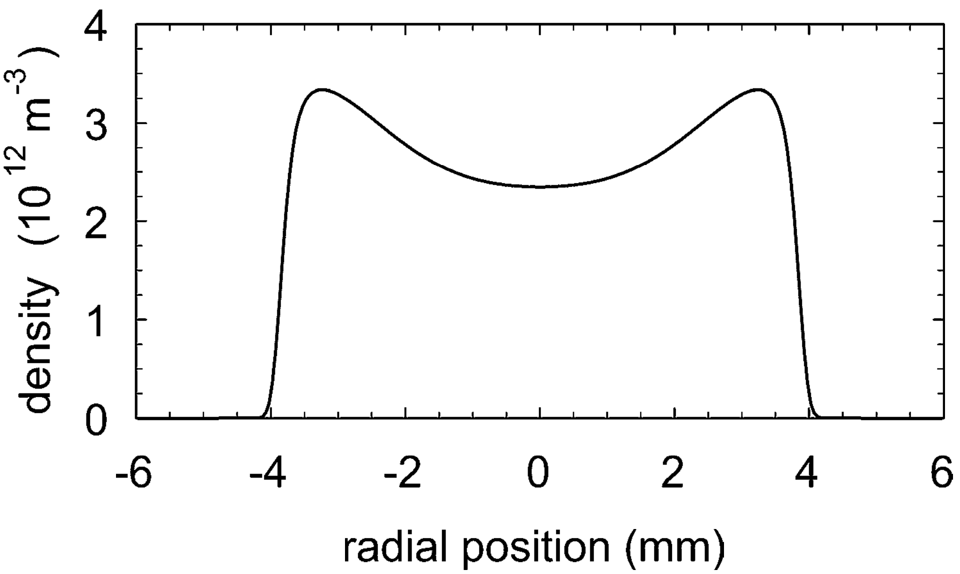

Figure 4.20: Variation of density across the midplane of a cloud of 4×10^4 antiprotons.

lower density at the center of thin clouds is also seen in other simulations [86].

4.3.4 One Million Positrons

Due to the limited amplitude of the voltage pulsers the deepest well that can be pulsed open, and in which the positron physical parameters can be measured, is produced by applying 8.5 V, 4.25 V and 8.5 V to three adjacent electrodes. Unfortunately, for most antihydrogen experiments (see Sec. 6.1.2 and Ref. [18]) the positrons are confined in much deeper wells. However, as noted above, the angular momentum of a cloud is independent of the depth of the potential well in which it is confined. This is confirmed in Fig. 4.21 where the angular momentum of positron clouds in all three well depths are shown to lie on a single smooth curve.

The Equilsor2 code can be run in thrmcnsv mode which keeps the cloud angular momentum fixed. It can therefore calculate the parameters of positron clouds confined

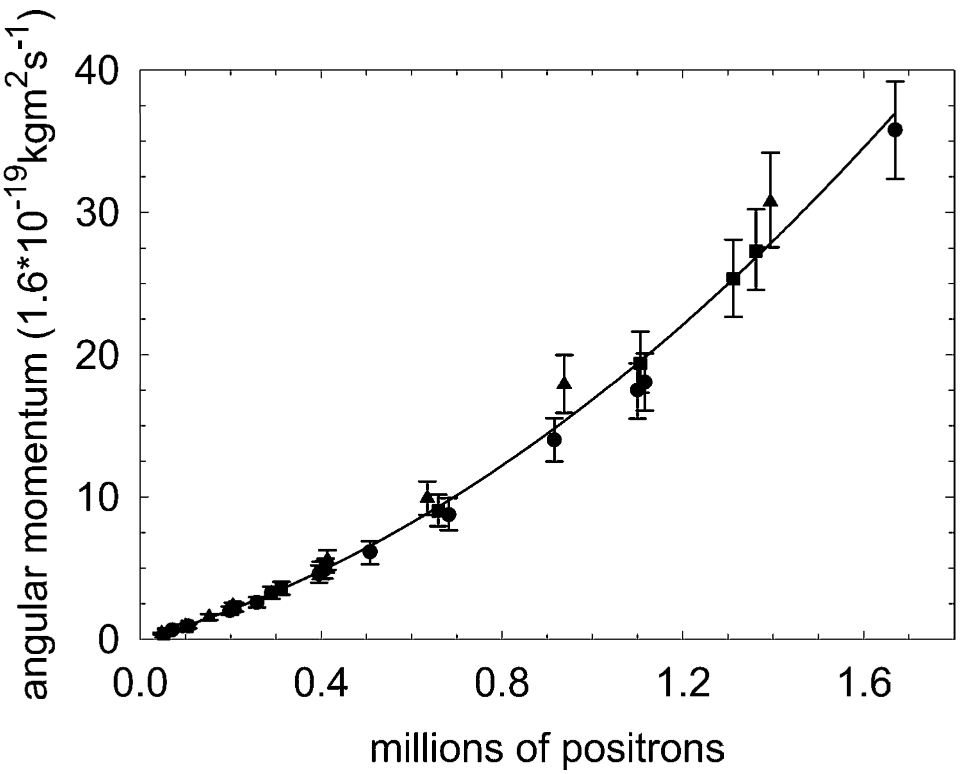

Figure 4.21: Angular momentum dependence on positron number for all three well depths.

in deeper wells, such as those used in the antihydrogen experiments. Plotted in Fig. 4.21 is the angular momentum, while the thrmcnsv mode uses values of the moment of inertia, I. The two are related by,

$$L = \frac{qB}{2} I. \tag{4.13}$$

The fit in Fig. 4.21 gives the value of I, which for a given number of positrons, will not change as the well depth is changed. The code can be run in thrmcnsv mode to find the cloud parameters in other well depths if one uses input parameters in the *.run file given by,

$$\begin{aligned} I = \frac{2}{5} N a^2 &= -0.01946 + 3.395 \times 10^{-6} N + 2.976 \times 10^{-12} N^2 \\ \textit{and} \quad N &= \frac{4\pi}{3} a^2 b n. \end{aligned} \tag{4.14}$$

Using the above equations and running the code for 1 million positrons gives the dependencies shown in Fig. 4.22.

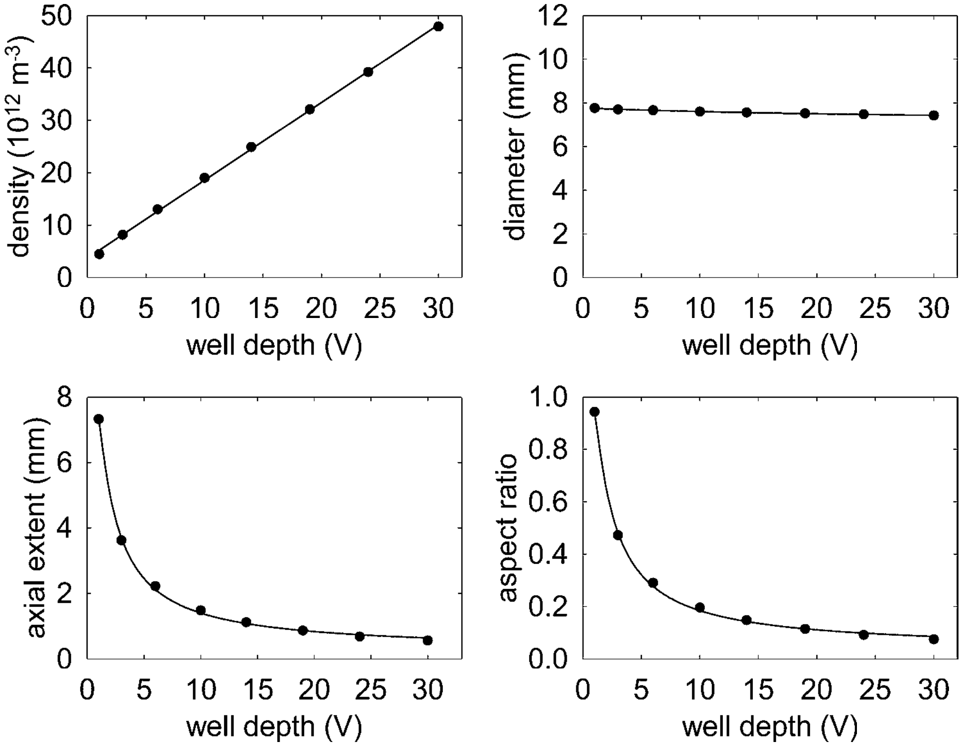

Figure 4.22: Variation with well depth of density, diameter, axial extent and aspect ratio for 1 million positrons. Particles are confined on a single electrode held at a voltage equal to the well depth, all other electrodes are grounded.

Chapter 5

Driving Antiprotons in a Nested Penning Trap

5.1 The Nested Penning Trap

In order to make antihydrogen a method is needed that will allow the oppositely charged positrons and antiprotons to interact with each other. This is not straightforward, since an electrostatic well to confine, say, positrons, will be anti-confining for antiprotons. A solution to this problem is a nested well configuration [4], illustrated in Fig. 5.1. Here the positrons are held in an inverted central well while the antiprotons are confined in an outer well with enough energy to move through the positrons and form antihydrogen. The nested Penning trap configuration is easy to implement and can be used for experiments favorable to antihydrogen production by both radiative recombination and three body recombination.

The axial confinement of oppositely charged particles is assured in the nested well

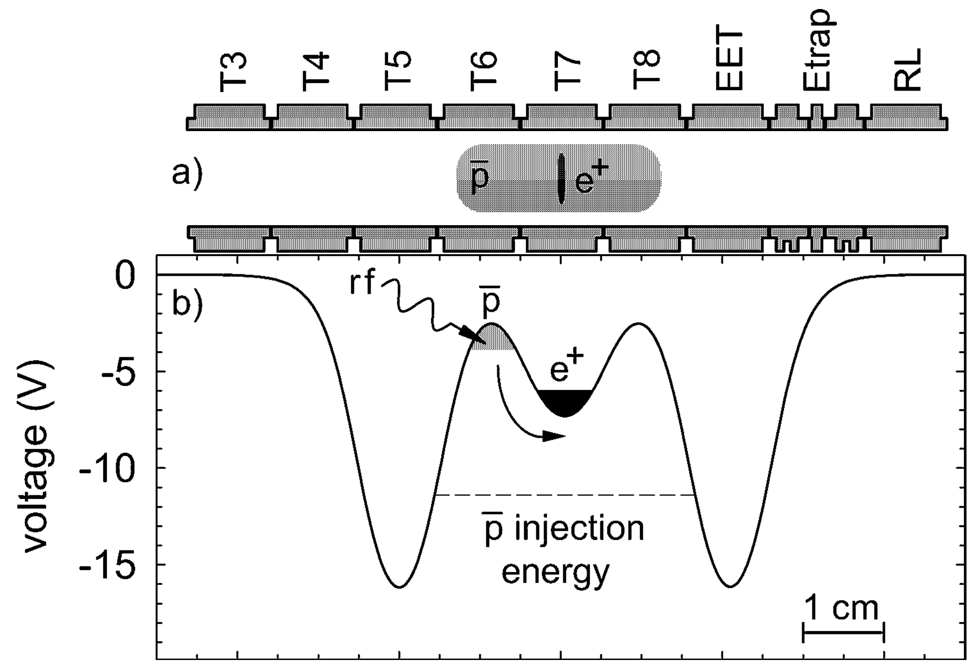

Figure 5.1: The nested Penning trap. a) Electrodes, b) nested well potentials, and, schematically, two techniques used to make the antiprotons interact with the positrons.

structure, but the question remains as to how to provide the antiprotons with sufficient energy to pass over the positron well in the center. There are two techniques to do this. The first is to inject the antiprotons into the nested well at an energy greater than the potential of the central positron well. Once injected, the antiprotons oscillate back and forth in the well formed by potentials on electrodes T5 and EET. They lose energy through Coulomb collisions with the positrons (which cool by radiating to the 4 K surroundings) until they have only just enough energy to pass over the top of the positron well. At this point the probability of recombination is the highest due to the low antiproton - positron relative velocity. Any antihydrogen produced will be immediately lost from the trap as the neutral atom is not confined by the Penning fields. This interaction technique is reported in detail elsewhere [5, 18].

The second technique is to start with the antiprotons in the side of the nested well (T6 in Fig. 5.1). Since the antiprotons are at a temperature of 4 K (0.36 meV) after electron cooling, they have insufficient energy to interact with the positrons sitting in the adjacent well several volts deep. However, a radio-frequency drive applied to the T6 electrode can resonantly increase the energy of the antiprotons' axial motion, allowing them to climb over the central well and interact with the positrons. This is also shown in Fig. 5.1, and the technique is the topic of this chapter.

5.2 Driving Antiprotons in a Nested Penning Trap

5.2.1 Calculation of Bounce Frequencies

In order to drive the axial motion of antiprotons in the sides of the nested well the frequency of this motion must first be determined. This is calculated using a Runge Kutta numerical integration technique. The calculation is performed for particles on the trap axis and 4 mm off axis (the latter is approximately the radial extent of the antiproton clouds). In this way one can determine an appropriate drive frequency. Figure 5.2 shows the results of such a calculation for typical nested well potentials used in the experiments. The potentials are also shown in the figure.

It can be seen that the oscillation frequency changes discontinuously as the antiprotons are excited out of the side well and begin to oscillate in the wider region of the nested well. Antiprotons 4 mm off axis have some double-valued oscillation frequencies. This is because particles with higher energies in the side well have a larger distance to travel, so some frequencies result from two different antiproton energies.

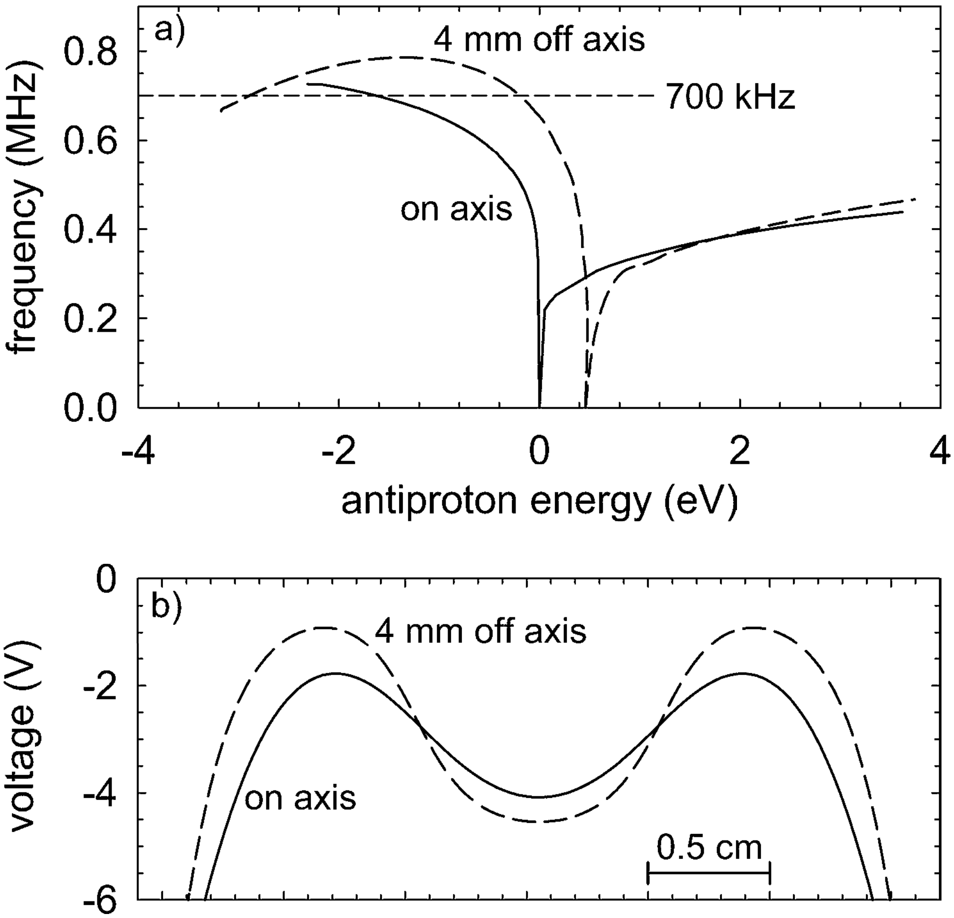

Figure 5.2: a) The axial oscillation frequency of antiprotons in the side of the nested well depends upon the antiproton energy and radial position. The energy is calculated with respect to the potential energy on axis at the center of the nested well. b) Nested well potentials on, and 4 mm off axis.

As noted in Fig. 5.2(a) a drive frequency of 700 kHz will be resonant with antiprotons oscillating close to the bottom of the nested side well.

The radio-frequency drive is applied to the electrode upon which the antiprotons are confined. If the equilibrium position of the antiprotons were at the axial midplane of the electrode, the drive would just modulate the well depth slightly and not drive the particle's axial motion. However, the nested well potentials are such that antiprotons do not sit at this midplane but are offset slightly. This allows them to be driven effectively.

5.2.2 Driving Hardware

Considering the variation of axial frequency with antiproton energy in Fig. 5.2(a) two driving techniques seem viable. The first is to use a drive at a single frequency close to resonance with antiprotons at the bottom of the side well, and rely on collisions between antiprotons to transfer the drive energy to non-resonant particles, thereby increasing the axial energy of all antiprotons. A typical single frequency drive for antiprotons is 700 kHz, provided by a PTS 250 synthesizer. The second technique is to apply a drive at frequencies resonant with all antiprotons in the nested side well. A drive at many frequencies, i.e., a noise drive, is available from an SRS DG535. This can behave as an arbitrary waveform generator, which is programmed to give a very tightly spaced comb of lines (200 Hz apart) in the bandwidth 500-1000 kHz, thus approximating a noise drive in this range.

It is of course important that the Jülich detector system is informed when the drive is on and antihydrogen could be produced, and when it is off and antihydrogen cannot be produced. To this end a hardware switch is set up as shown in Fig. 5.3. During an experiment the synthesizer is remotely switched either to the trap via a variable attenuator, or to the detector system via electronics which provide a high d.c. voltage when the drive is not applied to the trap, and zero voltage when the drive is applied to the trap. The isolator is required to keep separate the electrical grounds of the detector system and the trap. At the trap the drive is transmitted to the electrode via a high frequency microcoax transmission line (a pulse line).

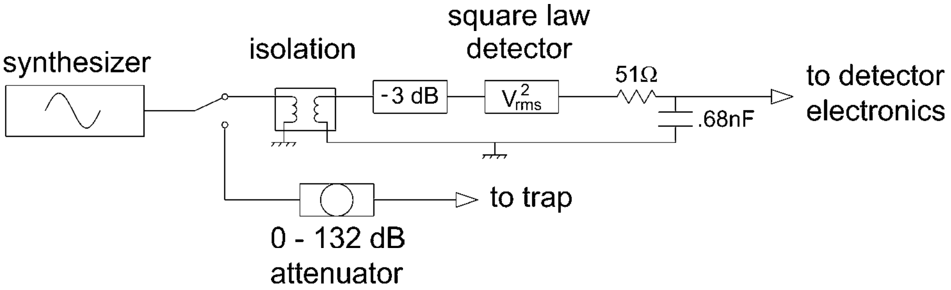

Figure 5.3: Hardware setup to communicate the drive state to the Jülich detector system.

5.2.3 Single Frequency Drive Capabilities

To illustrate the various processes involved when driving antiprotons across the nested well, I summarize the results of two experiments. In the first, 1.1×10^5 antiprotons are placed in a nested side well and 2×10^5 positrons are contained in the central well, so that the experiment will follow as closely as possible a real recombination experiment. A total of 2.5 million electrons are confined with the antiprotons. Figure 5.4(a) and (b) shows the arrangement. The antiprotons are driven with a power of 7 dBm at 700 kHz for 4 minutes, after which the T8 end of the nested well is ramped open to energy analyze the antiprotons driven out of the T6 electrode at his time (Fig. 5.4(c)). With the drive still applied, the T6 side of the nested well is then ramped open to determine the energy distribution of the antiprotons which remain on T6. This distribution is shown in Fig. 5.4(d) along with the distribution found in a similar experiment, but without the radio-frequency drive applied.

From Fig. 5.4(c), one can see that the energy distribution of antiprotons excited out of the T6 well is heavily peaked at low energies, i.e., at the bottom of the T8 side well, with a long tail extending beyond the potential at the center of the nested well. No antiprotons or electrons started on the T8 side, and no electrons will be driven over to T8. The fact that the bulk of the antiprotons are at low energies is a result of antiproton - antiproton collisions which decrease the energy of some antiprotons and increase the energy of others. The antiproton which loses energy remains in the T8 side well, but at this reduced energy. The heated antiproton oscillates the length of the nested well and can be cooled by the positrons on electrode T7, and electrons on T6. After many antiproton - antiproton collisions the net effect is that antiprotons will collect at the bottom of the T8 well, as seen.

The energy distribution of antiprotons on the T6 side of the nested well (Fig. 5.4(d)) indicates that the drive excites all the antiprotons from the bottom of the well, leaving no particles with low energies. This is exactly as one would hope, since, as can be seen from Fig. 5.2(a), a drive frequency of 700 kHz is resonant with low energy antiprotons in the side of the nested well, whether on or off axis. The distribution stops at the potential of the center of the nested well since all antiprotons with higher

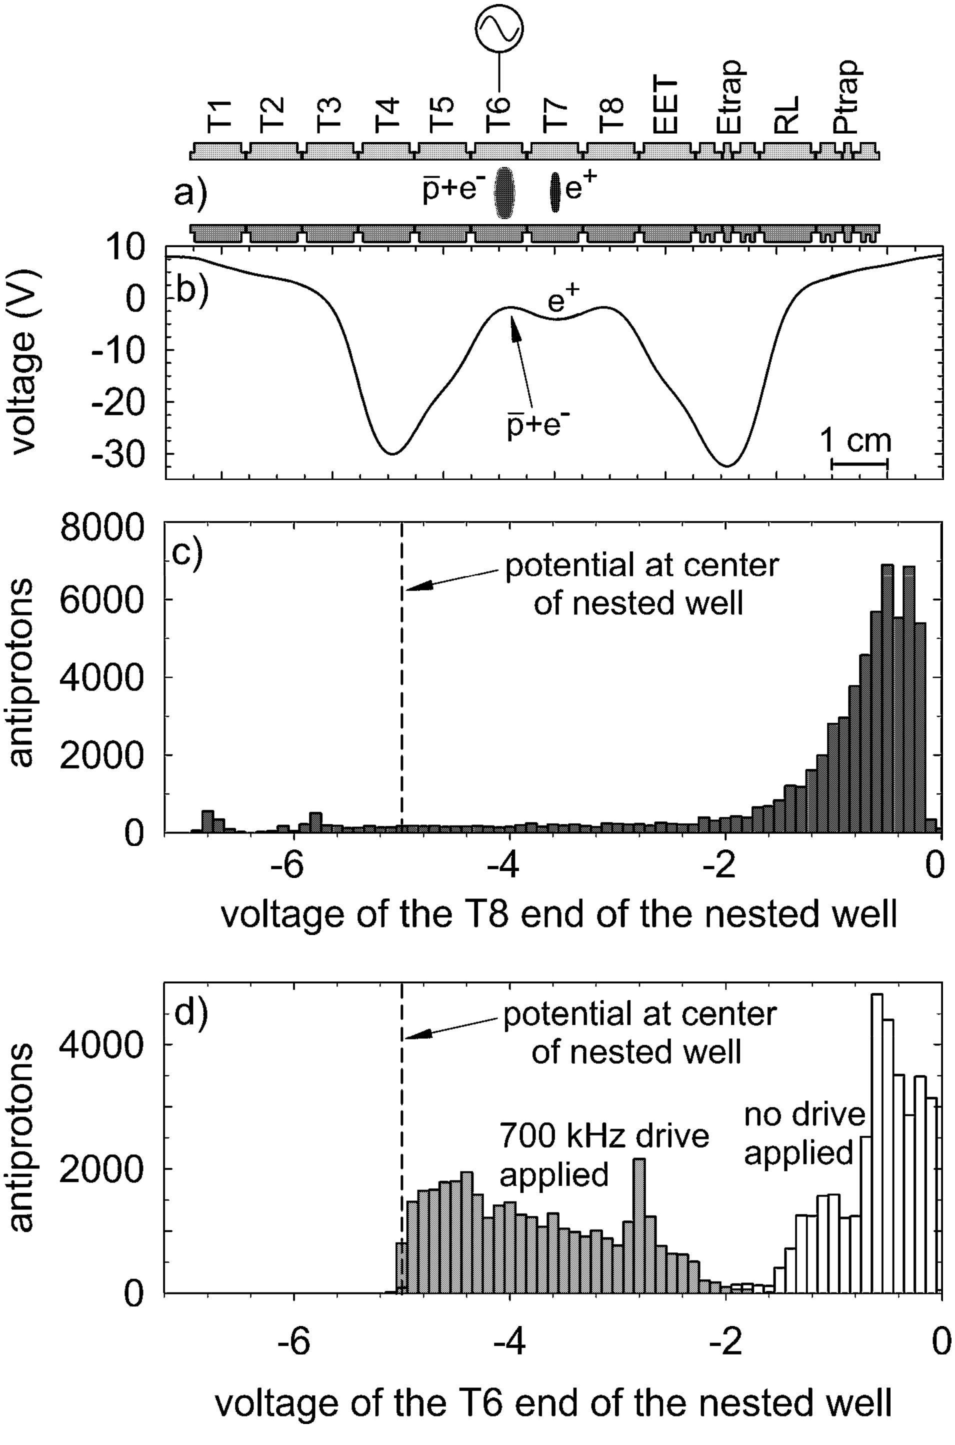

Figure 5.4: a) Electrodes, b) nested well potentials during the driving, and c) the resulting energy distribution of antiprotons excited out of the T6 well after being driven with 7 dBm of power at 700 kHz for 4 minutes. d) The energy distribution of antiprotons on the T6 side of the nested well after driving (grey bars), and the energy distribution from a separate experiment without the radio-frequency drive (unfilled bars).

energies were allowed to escape to the degrader when the T8 side of the nested well was investigated earlier.

Since the drive is only resonant with antiprotons at low energies, collisions between antiprotons play an important role in transferring particles over the center of the nested well. The picture to have in mind is of the drive energy percolating up through the antiproton energy distribution via collisions. In this way the highest energy antiprotons are jostled over the top of the central well.

In an experiment to determine an appropriate drive strength, antiprotons are contained on the T6 side of the nested well (as in Fig. 5.4(a) and (b)) and the drive strength is periodically increased while the T8 end of the well is ramped open and closed every 10 s. The drive is strong enough to transfer particles when a 10 s periodic variation in the loss from the trap appears: high loss when the T8 end of the nested well is open and antiprotons driven out of T6 can escape to the degrader, and low loss when the nested well is closed. This is shown in Fig. 5.5. A drive strength of 8 dBm is seen to stimulate substantial transfer. After nearly 700 s of driving at 8 dBm the transfer has reduced dramatically. This is because the total number of antiprotons remaining on T6 is significantly reduced, and because the drive has specifically depopulated the low energy part of the antiproton distribution. This second point is confirmed by turning off the drive for 5 minutes. When the drive is restarted the transfer also restarts, but at a higher rate than it left off, indicating a repopulation of low energies.

A closer look at the loss at early times in Fig. 5.5 reveals that, before the periodic structure appears, the loss rate has increased substantially above background. This

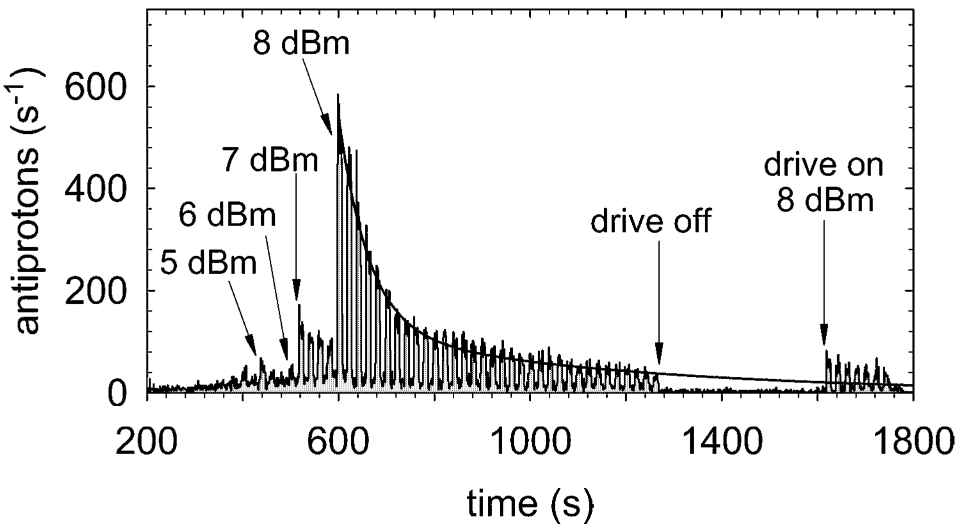

Figure 5.5: The observed antiproton loss rate as the drive strength is increased from -17 dBm at t = 0 s to 8 dBm at t = 600 s, while the end of the nested well is cycled open and closed. The re-population of resonant antiproton energies is shown by the increased loss at t = 1620 s

loss must be in the radial, rather than the axial, direction. If the loss were axially directed then more would certainly occur when the nested well is open than when it is closed, and the periodic loss pattern would be evident. High loss such as this occurs in all experiments when the drive strength is increased over many minutes. A cleaner technique is to immediately apply a drive which is strong enough to stimulate transfer.

These experiments demonstrate the effective transfer of radio-frequency drive energy to antiprotons forcing them to pass over the center of the nested well. The antiproton energy distributions are measured and understood in terms of the applied drive frequency and the collisional processes involved.

5.2.4 Alternative Driving Techniques

The above driving technique and a number of variations on it have been tried with no positrons present in the nested well. This enables a comparison between experiments of the antiproton transfer and antiproton loss (the latter without the possibility of loss due to antihydrogen). The results are summarized in Table 5.1. There was insufficient time to investigate in detail all of the different possibilities and to optimize each for drive strength, drive frequency, and drive duration. Most time was invested in the single frequency drive and the noise broadened drive, and, as noted in Table 5.1, the noise drive tends to suffer from a higher loss rate. As the experiments with a single drive frequency show, collisions redistribute drive energy, making it unnecessary to have a drive resonant with every antiproton energy. The many frequencies present in the noise drive may just serve to trigger instabilities in the cooling electrons (see Sec. 5.3) resulting in additional antiproton loss.

Noise broadened drive (500 - 1000 kHz).	Effective transfer, but, in general, higher antiproton loss observed than with a single frequency drive.
Single frequency drive applied to the electrode (T5) adjacent to the antiprotons.	Ineffective transfer even at maximum synthesizer power (13 dBm).
Electrons removed before driving.	Higher loss rate than when electrons are present.
Drive at twice the axial frequency (parametric drive).	Large, sudden, and uncontrollable losses can occur. Transfer is very sensitive to drive strength.

Table 5.1: Summary of alternative driving techniques.

A single frequency drive applied to the electrode upon which antiprotons and electrons are held is concluded to be the best driving technique, and experiments progressed swiftly to testing the efficiency of this method, as described in the next section.

5.2.5 Efficient Transfer of Antiprotons Across a Nested Well

The number of antihydrogen atoms that will be produced is proportional to the number of antiprotons that can be made to interact with the positrons. However, adding energy to the antiprotons to induce this interaction can cause antiproton loss that is not associated with antihydrogen and should, in general, be kept as low as possible. The aim, therefore, is to maximize the number of antiprotons transferred over the center of the nested well, while keeping to a minimum antiproton losses which are not due to antihydrogen. To this end two figures of merit are considered: the antiproton loss rate above the cosmic ray background (L_{net}), and the ratio, R_{net}, of the number of antiprotons transferred over the center of the nested well, to the loss incurred while transferring. The loss rate, L_{net}, is normalized to the number of antiprotons when comparing experiments using different particle numbers.

Cycling the nested well open and closed, as was done in the previous experiment, rapidly reduces the number of antiprotons, making it impossible to gauge the transfer and loss rates for many cycles. Also, to produce as much antihydrogen as possible, it is necessary to apply drives to both sides of the nested well so that the antiprotons are driven back and forth many times through the positrons. To address both of these issues a set of experiments is performed with particles and drives present on both of the nested side wells. During the driving, antiproton loss is continuously monitored and the nested well is only opened after the driving has ceased to discover the effectiveness of the transfer. In this way one can obtain values for L_{net} and R_{net}.

As a first step of this set of experiments a total of 1.8×10^5 antiprotons are loaded and cooled into two wells by two clouds of $\sim$ 2 million electrons. The antiprotons are then simply ejected and counted to determine the ratio of the number in each of the two wells. This ratio is confirmed by repeating the experiment. Antiprotons are then loaded as before and set up in the nested well, on electrodes T6 and T8 (Fig. 5.6(a) and (b)). T6 is driven with a power of 7 dBm at 700 kHz for 10 s, after which both ends of the nested well are consecutively ramped open to discover the number of antiprotons on each side. The experiment is repeated, but driving T6 for 10 s, waiting for 10 s with no drive on, and then applying the drive for 10 s to T8. The number of antiprotons on each side of the nested well is then measured. This procedure is continued, increasing the number of drive iterations each time.

Initially, 42% of the antiprotons start on T6, while 58% start on T8. After one iteration, driving T6 for 10 s, only 4% remain on T6. The fraction of antiprotons in each side well after 1, 2, 3, 4, 9, 20, and 49 iterations is shown in Fig. 5.6, indicating that repeatable transfer of roughly 1.7×10^5 antiprotons is possible. The 49 iteration experiment is therefore equivalent to a total of nearly 10 million antiprotons excited over the center of the nested well. As can be seen in Fig. 5.6(c) for each experiment the majority of antiprotons finish on the opposite side of the nested well from the previous experiment. A measure of whether or not the transfer is becoming less efficient as the number of iterations increases is therefore the difference in the fraction of antiprotons on say, the T8 side, between one experiment and the next. This is shown to decrease linearly in Fig. 5.7(a), although only reducing to 80% after a total of 49 iterations.

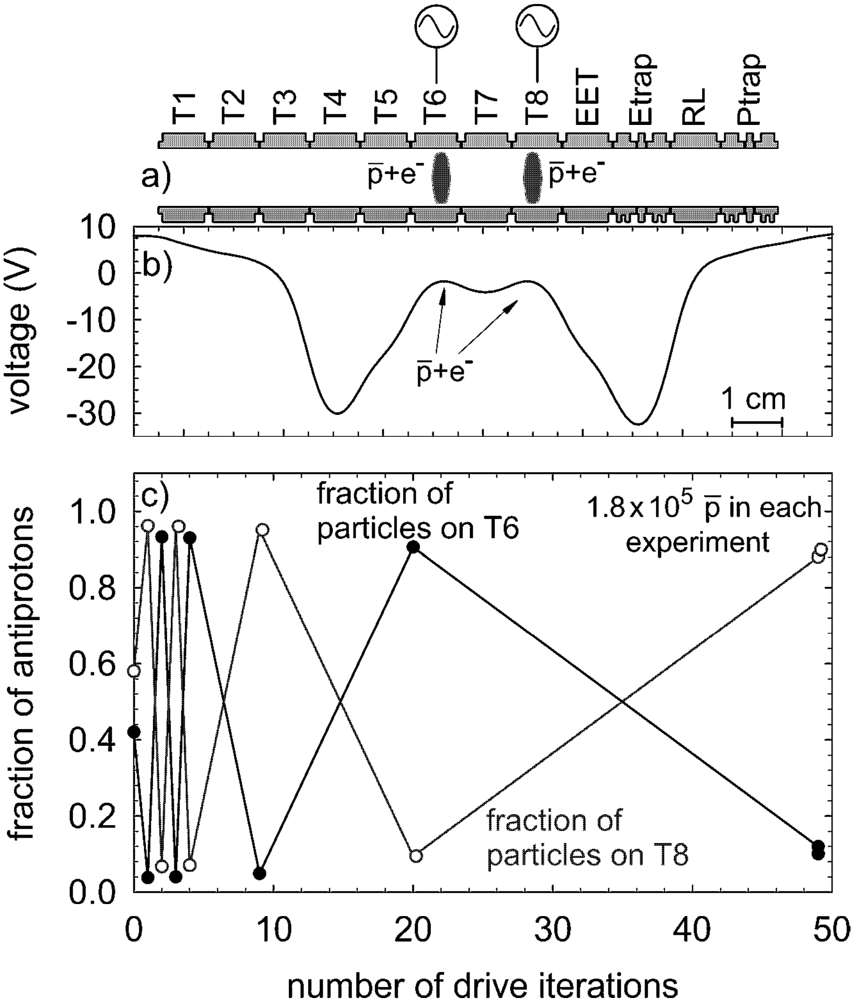

Figure 5.6: a) Electrodes, b) nested well potentials, and c) the fraction of antiprotons on each side of the nested well after the given number of drive iterations.

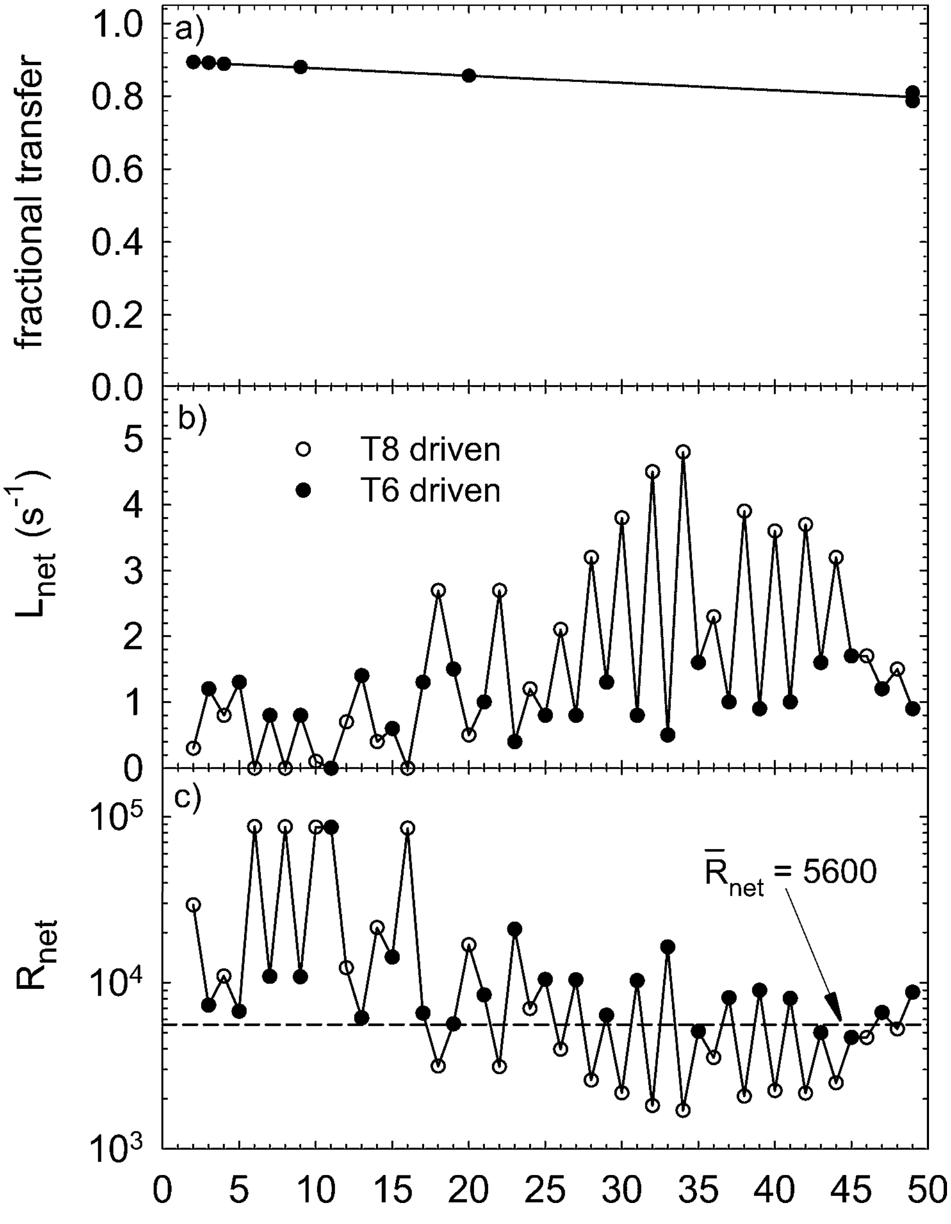

Figure 5.7: a) Fractional transfer of antiprotons over the center of the nested well, b) net loss incurred while driving, and c) the ratio of the antiproton transfer to the net loss.

The loss rate above background (L_{net}) during the periods when either drive is on is seen in Fig. 5.7(b) from an experiment when the antiprotons are transferred back and forth 49 times. For the first 15 iterations the loss rate above background is as low as 0.5 counts s^{-1} (1 antiproton s^{-1}). Later iterations show an increased loss each time T8 is driven, indicating that the T8 drive is a little too strong.

The number of antiprotons transferred at each iteration can be deduced from the linear fit in Fig. 5.7(a) and is used to find the ratio of the antiproton transfer to the loss above background, as seen in Fig. 5.7(c) (when the net loss is zero R_{net} is taken to be the number of antiprotons transferred). This ratio of transfer to net loss, when averaged over the whole experiment ($\overline{R}_{net}$), is equal to 5600, indicating very efficient transfer of antiprotons.

R_{net} is measured for an identical experiment with the same number of antiprotons, but with a drive frequency of 650 kHz, instead of 700 kHz. The results are shown in Fig. 5.8. Only 3% of the antiprotons remain on the side of the nested well driven last, indicating that effective transfer has taken place. The measured values of R_{net} are comparable to those obtained using the higher frequency drive. The transfer is not, therefore, highly sensitive to the frequency of the drive used.

In summary, highly efficient transfer of antiprotons across a nested well is stimulated by a single frequency drive. Approximately 95% of antiprotons are transferred when the drive is first applied, and the transfer is repeated up to 50 times, with only a small reduction in efficiency. The loss of antiprotons during the transfer is very low: as little as 1 antiproton s^{-1}. This efficient and repeatable transfer makes the driven nested well a very promising environment for antihydrogen production.

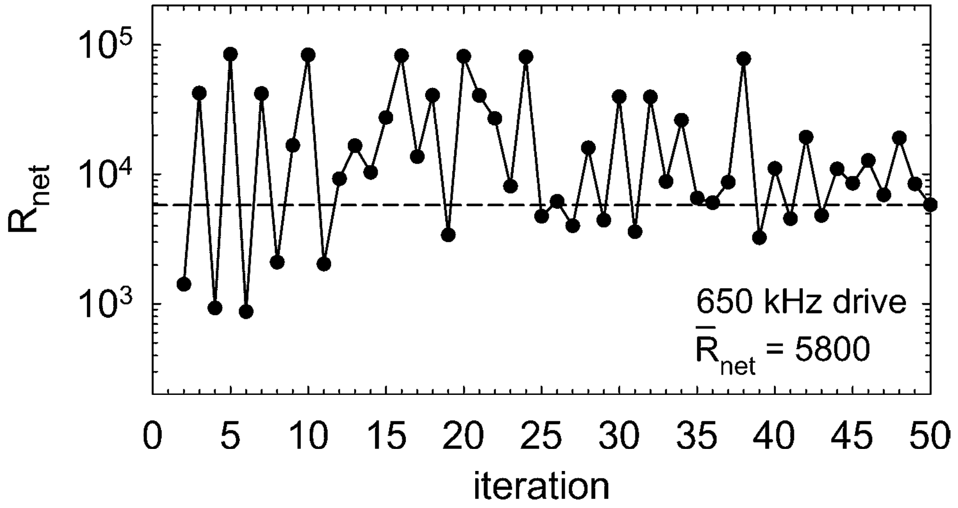

Figure 5.8: The ratio of the antiproton transfer to the net loss for a 650 kHz drive.

5.3 Anomalous Driving Loss

Occasionally large losses occur while driving antiprotons across the nested well. Preliminary studies suggest that this is related to the number of electrons mixed in with the antiprotons. Figure 5.9 shows the antiproton loss incurred during two consecutive experiments where the antiprotons are again driven back and forth 49 times between the sides of the nested well. The number of electrons is halved in the second experiment which is seen to have a very beneficial effect.

Large electron clouds can be unstable and stimulate instabilities in the antiproton cloud, as noted in Sec. 3.1.1. If excessive antiproton losses are observed during driving, the first thing to check is the number of electrons and whether or not the electron cloud is anomalously hot. This can be done using the radio-frequency detection techniques described in Ch. 2.

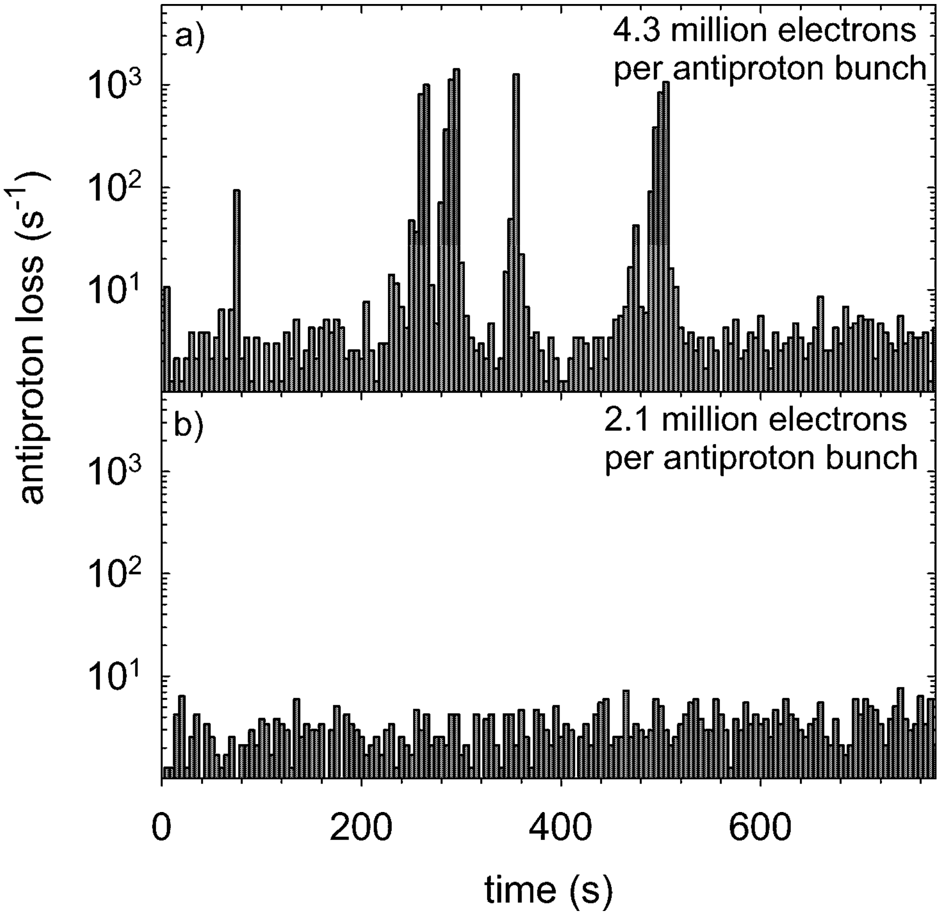

Figure 5.9: The antiproton loss observed from the trap during two nested well experiments. Antiprotons are cooled into two wells by two clouds of, a) 4.3 million electrons, and b) 2.1 million electrons.

5.4 Loss Analysis

One of the experiments described above indicates that antiproton loss from the nested well occurs in the radial, not the axial, direction. This can be confirmed by a more general technique which relies on using the BGO detector to detect antiproton annihilations (rather than positron annihilations for which it was designed).

Antiprotons exiting the trap annihilate on the gold plated copper electrodes, creating on average 3.5 charged pions and many high energy gamma rays. When the pions and gamma rays pass through a BGO crystal they can deposit a large amount of energy. Unfortunately, such events blind the detector to the small amount of energy (from a 511 keV gamma ray) produced by a positron - electron annihilation. Fortunately, however, if one measures only events where a large amount of energy is absorbed, then one has a measure of the number of antiproton annihilations.

Every time a trigger count (a temporal coincidence between an antiproton event in the fiber detector and one in the outer scintillators) occurs, the Jülich data acquisition system is triggered. The energy deposited in each of the 12 BGO crystals is then measured, along with the state of all other detector elements. A quantity of interest for antiproton detection is the number of crystals in which a large amount of energy is deposited at the time of a trigger. This is called the multiplicity (μ) and can be any number from 0 (no crystals show a count) to 12 (all crystals show counts).

The main background for this measurement is due to cosmic rays, since these also deposit a large amount of energy in BGO. However, since cosmic rays come from above, and given that the BGO detector is physically much smaller than the fiber detector and outer scintillators, the chances of a cosmic ray striking all three detectors is reasonably small. This means that if a cosmic ray causes a trigger then it is unlikely to also hit the BGO and will therefore have a multiplicity of zero. In contrast, antiprotons leaving the trap annihilate much closer to, and often within, the BGO detector and so the majority of these annihilations produce $\mu > 0$ events. Therefore, an effective technique to eliminate cosmic rays is to reject (or cut) all trigger counts associated with $\mu = 0$ events. If this cut is made, 92% of the cosmic rays are rejected at the expense of $\sim$ 20% of real antiproton annihilations.

The lower trap where all driving experiments take place is 19 cm long, while the BGO crystals are 12 cm long and centered on the Etrap electrodes. This can be seen in Fig. 2.8(b). The BGO detector therefore has different solid angle efficiencies (and therefore different multiplicities) for antiproton annihilations which occur at different positions in the trap. For example, antiproton annihilations on the ball valve are detected with a much smaller solid angle than, say, those on the inner surface of electrode T7. The multiplicity of ball valve annihilations is therefore peaked at a lower value than that of annihilations on T7.

In order to determine the direction of loss in a nested well driving experiment, it is necessary to have as a reference the multiplicity spectrum for antiprotons annihilating on the ball valve and on the inner surface of electrodes T6, T7, and T8. The spectrum for annihilations on the ball valve is found by simply ejecting a cloud of antiprotons to this end of the trap with the ball valve closed. The spectrum for radial loss from T6 and T7 is found by placing antiprotons in a well formed by applying the same potential to T6 and T7. Such a well is radially unstable as described in Sec. 4.1. The procedure is repeated for antiprotons on T7 and T8, and the two multiplicities are averaged to find the spectrum for loss on T6, T7, and T8. These spectra are shown in Fig. 5.10(a) and (b) after the cosmic background has been reduced by the cut $\mu > 0$. As expected the spectrum for annihilations on the ball valve has more low multiplicity counts than that of annihilations on T6, T7, and T8.

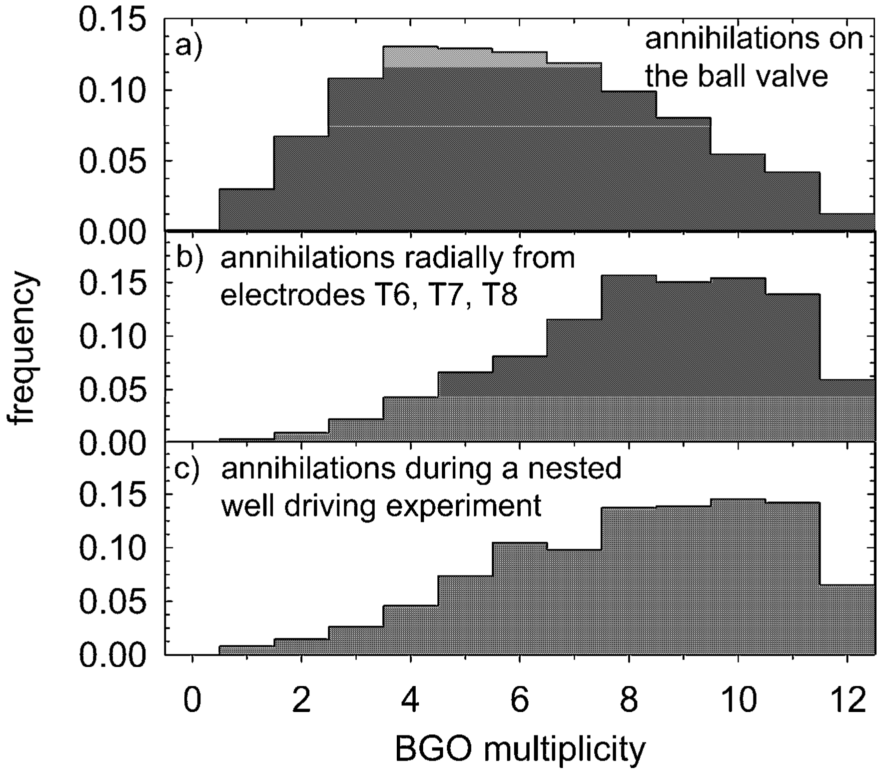

Figure 5.10: The multiplicity spectrum for antiprotons annihilating, a) on the ball valve, b) radially from T6, T7, and T8, and c) during a nested well driving experiment. A $\mu > 0$ cut has been applied.

Figure 5.10(c) shows the spectrum observed during a nested well driving experiment using the potentials shown in Fig. 5.4(b). It can be seen that the potential of the end of the nested well closest to the degrader is higher than the end closest to the ball valve. A potential 5 V more negative is applied to electrodes Etrap than to T4 to achieve this, and is to ensure that any antiproton loss in the axial direction results in annihilations on the ball valve. It is easy to see from Fig. 5.10 that the loss in the driving experiment much more closely resembles radial loss from T6, T7, and T8, than it does axial loss at the ball valve. Quantitatively, the driving loss observed is $92\ +8, -10\%$ radial, confirming the previous conclusions.

This example illustrates the use of the BGO detector to measure antiproton annihilations, and shows a technique to determine their spatial distribution. In order to do this effectively the cosmic ray trigger count rate of 1.2 counts s^{-1} is reduced by an order of magnitude. It can be reduced further by extending the multiplicity cut to $\mu > 1$, but at the expense of cutting more real antiproton events.

Chapter 6

Antihydrogen Experiments in a Driven Nested Penning Trap

The previous chapter gives examples of efficiently driving antiprotons across the nested well with minimal loss. Transferring large numbers of antiprotons is essential to maximize the number of atoms created. Maintaining a low antiproton loss during the transfer is always desirable but will not be essential if the loss does not interfere with antihydrogen detection. This chapter describes such a detection technique, where antihydrogen is identified long after all driving and associated antiproton loss have ceased.

6.1 Background-Free Antihydrogen Detection

6.1.1 Nested Well With Antihydrogen Detection

A modification to the experiments detailed above is now made which provides for background-free antihydrogen detection. Again, antiprotons and electrons are confined in the sides of a nested well with positrons in the center, and the antiprotons are repeatedly driven through the positrons. Now, however, an "ionization well" (or "detection well") is added, outside the nested well region. An antihydrogen atom formed in the nested well will move in the initial direction of its antiproton. Those which pass through the detection well, and are weakly bound enough, will be field ionized, depositing their antiproton in the well. The ionization well can be investigated after the driving and all related particle loss have ceased. As shown below this is done with effectively zero background. With the ionization fields applied, however, this technique is only sensitive to a subset of possible antihydrogen atoms, those which are ionized by fields between ~ 20 Vcm^{-1} and up to ~ 150 Vcm^{-1}. Fields of ~ 20 Vcm^{-1} exist inside the nested well itself, ionizing the most weakly bound atoms before they can reach the detection well, and ~ 150 Vcm^{-1} is the maximum ionization field used in these studies.

Radiatively recombined atoms are preferentially formed in the ground and first excited state, since the cross-section for this process is inversely proportional to the principal quantum number of the recombined state [47]. The ground and first excited states require enormous fields to ionize – 1×10^9 and 3×10^8 Vcm^{-1}, and as a consequence any antiprotons in the ionization well cannot have come from radiatively recombined atoms. These atoms can only be detected during the driving process, as they leave the nested well and annihilate on the electrodes. In this case low antiproton loss is essential, and the experiments demonstrated in the previous chapter can be used to maximum effect.

An atom formed by three body recombination is, initially so loosely bound that it will be ionized by the fields which exist in the nested well. However, collisions between the atomic positron and positrons in the positron plasma will de-excite the atom as it continues through the plasma. If sufficient de-excitation occurs then the atom may escape the nested well intact, to be ionized and detected in the ionization well.

The potential structure of the nested well with the adjoining ionization well is shown in Fig. 6.1. The potentials are carefully designed to prevent any antiproton not bound in an antihydrogen atom from becoming trapped in the ionization well. The on-axis voltage of the left side of the nested well (T2) is lower by approximately 4 V than that on the right (T8). Any antiproton with enough energy to escape the nested well will therefore leave to the left, and annihilate on the ball valve. Even if an antiproton were able to escape to the ionization well it would have to lose energy while in the well in order to become trapped. The only mechanism for such energy loss is a strong collision between two antiprotons passing through the ionization well simultaneously, the probability of which is vanishingly small. Additionally, all experiments are performed twice, once with, and once without positrons. When positrons are present they will cool the antiprotons, reducing the likelihood that any antiproton will acquire enough energy to climb out of the nested well. Therefore, the

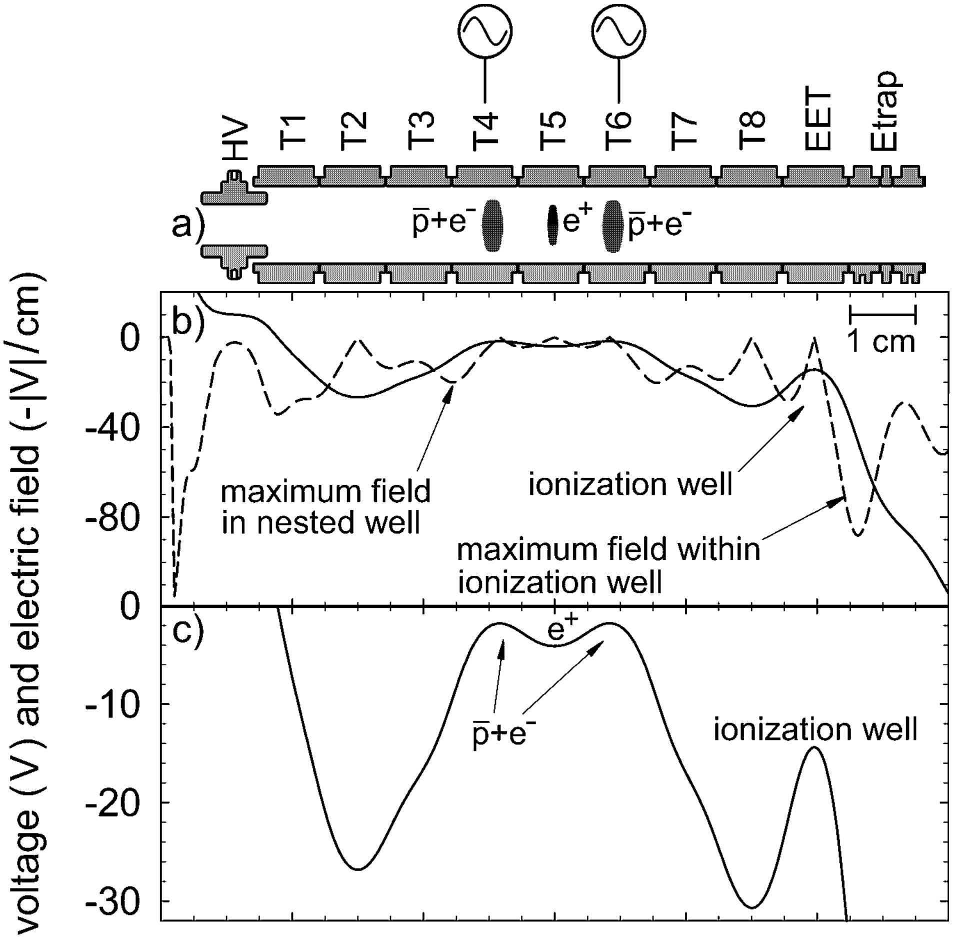

Figure 6.1: a) Electrodes, b) potential (solid line) and negative magnitude of the electric field (dashed line) of the nested well, and c) potentials in more detail.

chances of finding an antiproton not from antihydrogen in the ionization well will be lower when positrons are present. A careful study of Fig. 6.1 shows that the largest field an antihydrogen atom can experience and deposit its antiproton in the ionization well, is 73 Vcm^{-1}. While in the nested well the electric field on axis has a maximum value of 20 Vcm^{-1}.

6.1.2 Background-free Observations

A total of 3×10^5 positrons are contained in the center of the nested well and a mixture of 2×10^5 antiprotons with 4×10^6 electrons are divided between the side wells. A single frequency drive at 650 KHz and a power of 7 dBm is applied alternately to each side of the nested well for a total of 50, 10 s iterations. This is the driving procedure which achieved a high transfer to loss ratio in the experiments of Ch. 5. After 50 iterations, the ionization well is deepened to 15 V in order to safely store any captured antiprotons while a ramp is made to the degrader. In the final step the potential of the ionization well is swept from 15 V to -35 V in 10 ms, and any trapped antiprotons are expelled towards the degrader. There is essentially no background in this measurement since during the 10 ms ramp there will be on average only 0.01 cosmic ray background counts. The procedure is shown in Fig. 6.2.

Unfortunately, no antiprotons were detected in the ionization well, indicating that antihydrogen which could be ionized by fields between 20 Vcm^{-1} and 73 Vcm^{-1} was not produced.

The above scheme is repeated with two important, and several minor, changes. Firstly, the positron well is deepened from -5 V to -9 V. This increases the positron density from $9{\times}10^{12}\ \mathrm{m}^{-3}$ to $15{\times}10^{12}\ \mathrm{m}^{-3}$. Since the recombination rate for three body recombination scales quadratically with density the increase should have a significant effect. Secondly, the electrons are ejected prior to the experiment since without them as a cooling source the antiprotons can be driven more easily over the deeper positron well. Without electron cooling, the antiprotons will also be lost more easily from the nested well. However, this is not a great concern since the ionization well is

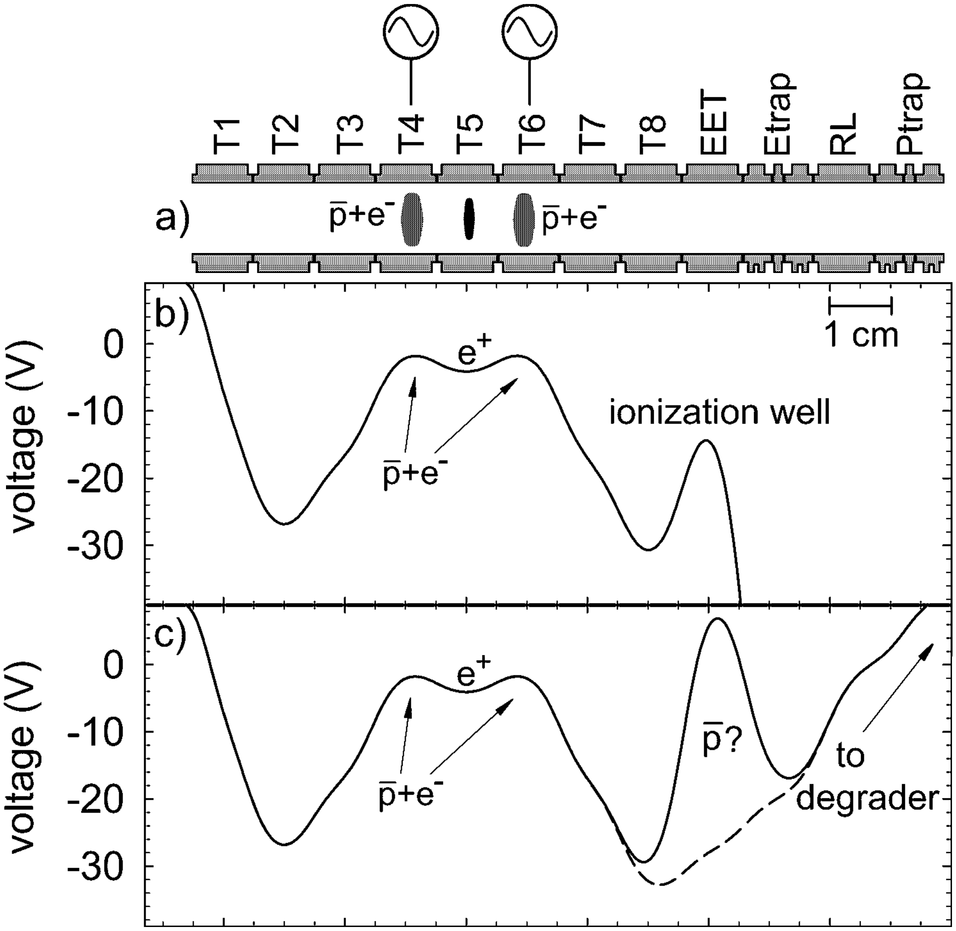

Figure 6.2: a) Electrodes, b) nested well potentials while driving, c) potentials just prior to ramping the ionization well (solid line) and potentials which ramp out any antiprotons caught in the ionization well (dashed line).

investigated after the radio-frequency drives are turned off and this loss has stopped.

The deeper positron well increases the resonant frequency of an antiproton in the side of the nested well. For this experiment a drive frequency of 825 kHz is appropriate. A lower drive strength (4 dBm) is used and all the antiprotons start on one side of the nested well. These minor changes should have little effect on the experiment. The drive is applied alternately to each side of the nested well for a total of 50, 10 s iterations, pausing for 5 s between each iteration. The ionization well is then investigated as described above, as are the sides of the nested well. The experiment is performed with 1.9×10^5 antiprotons and 3.3×10^5 positrons. It is repeated, but with no positrons present, and slightly fewer antiprotons (1.3×10^5) due to a stoppage in the AD.

Figure 6.3 shows the antiproton loss observed during the driving. It is immediately clear that the loss is far greater than that during the driving experiments of Ch. 5. This is because a strong drive is still used but no cooling electrons are present. The loss rate is high when the drive is on and falls abruptly when the drive is turned off. It maximizes on the second cycle, rather than the first, as not enough energy is introduced in the first 10 s to reach maximum transfer (and therefore maximum loss). After the second cycle the amplitude of the loss peaks falls approximately exponentially as the number of antiprotons is depleted. By the end of the 50^{th} cycle the loss and transfer have virtually stopped. The amplitude of alternate peaks is not the same. All the antiprotons started on one side of the nested well so the amplitude difference indicates that there is not 100% transfer of antiprotons. Other than the lower loss amplitude due to fewer antiprotons the spectrum without positrons shows the same characteristics. Figure 6.3(c) shows how the antiproton losses in the experiments with and without positrons compare for each driving iteration, normalized to the same number of antiprotons. During the first three drive cycles there is substantially

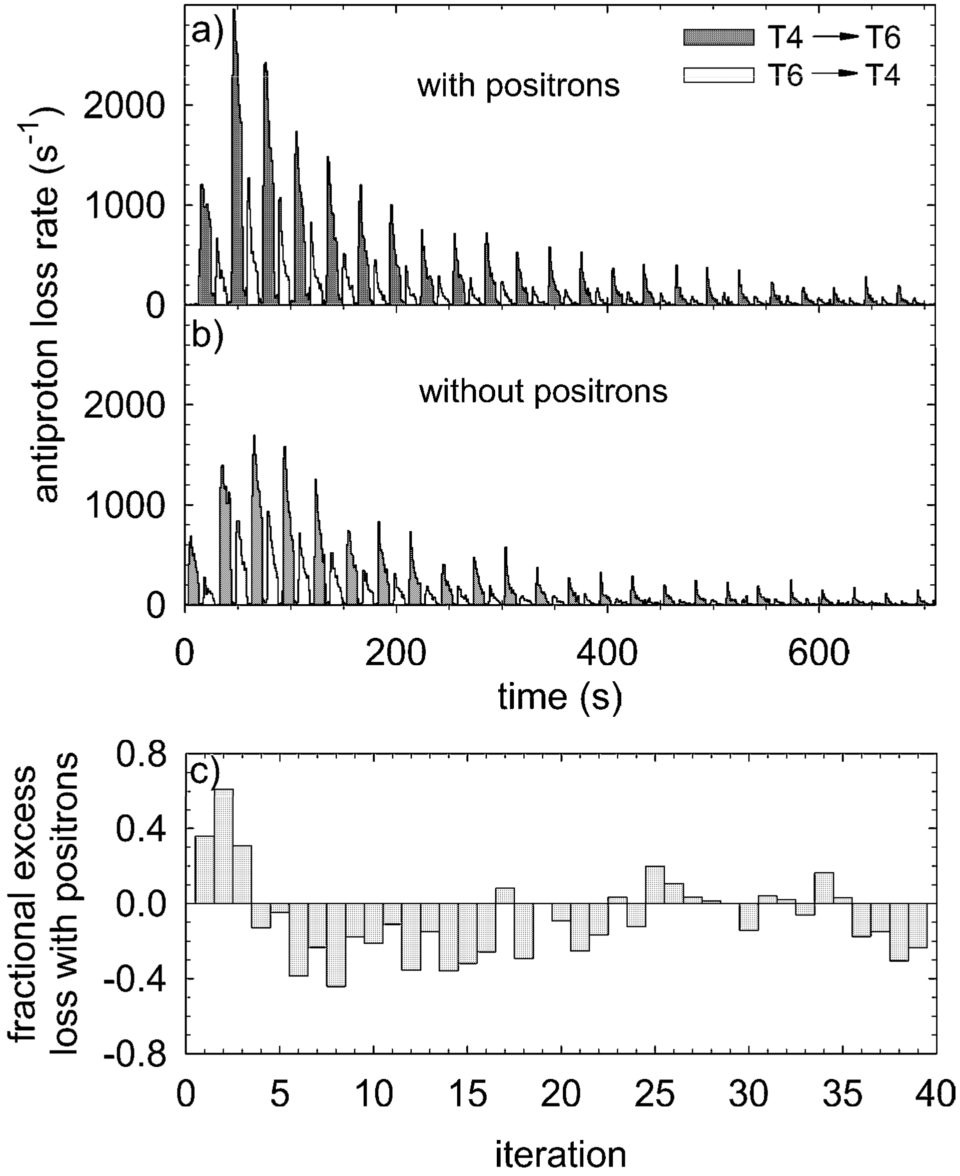

Figure 6.3: Antiproton loss from the nested well, a) with, and b) without positrons (fewer antiprotons are used in the experiment without positrons, hence the reduced amplitude of the loss). c) Excess loss when positrons are present in the center of the nested well, per driving iteration, and normalized to the same number of antiprotons.

more loss observed when positrons are present, than when they are not. The reverse is true for later cycles. The loss with positrons present could have two components - normal antiproton loss from the nested well and loss due to antihydrogen striking the trap electrodes. The loss without positrons can of course only have the first component. If antihydrogen is being produced it is presumably contributing more to the loss observed during the first few cycles that the later cycles. No stronger conclusion can be drawn, however, since positrons can modify the normal antiproton loss from the nested well making it impossible to separate this from any loss via antihydrogen.

When the driving is finished the contents of the two sides of the nested well are analyzed separately. For the experiment with positrons contained in the T5 electrode, no antiprotons are found in the T6 side of the nested well but 3.5×10^4, or 19%, of the initial number are discovered on T4. Since the loss had virtually stopped after 50 iterations it is likely that these represent a fraction of the antiprotons which are never transferred by the drive. This is probably due to the fact that the drive frequency is 825 kHz, while antiprotons at the very bottom of the side well have a resonant frequency of 875 kHz and are therefore not strongly driven. Similarly, for the experiment without positrons in T5, no antiprotons are found on T6 but 15% of the initial number remained on T4.

The contents of the ionization well are then investigated by quickly ramping it towards the degrader as described above. The results of these fast ramps for the two experiments, along with the potentials used during the driving, are shown in Fig. 6.4. The results are clear: a total of 183 antiprotons from antihydrogen are caught in the ionization well when positrons are present, while the test with no positrons

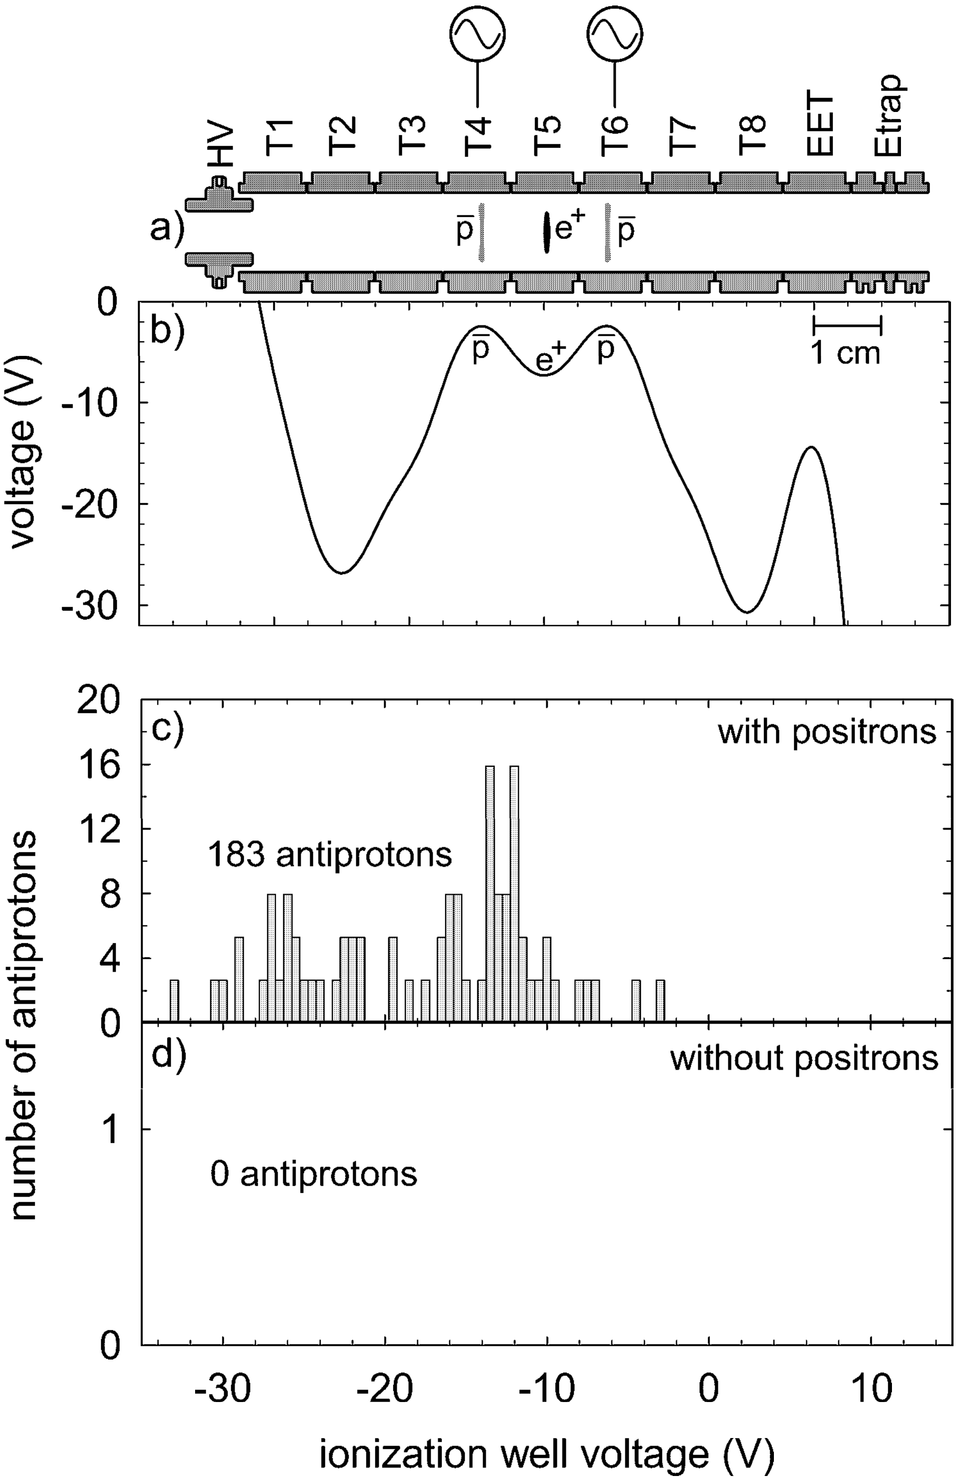

Figure 6.4: a) Electrodes, b) nested well potentials, c) ionization well contents with positrons, d) ionization well contents without positrons.

reassuringly showed no antiprotons captured. Further experiments are done, both with and without positrons. Antiprotons are always observed in the ionization well when positrons are present and none have ever been found without positrons. Even without these additional experiments the potential structure guarantees that any antiproton in the ionization well must be from antihydrogen. This one demonstration therefore already unambiguously demonstrates antihydrogen production.

6.1.3 Recombination Rate Dependencies

The number of atoms detected cannot simply be proportional to $n^2 * b$, as might be predicted for a three body recombination rate which scales as the square of the positron density (n), and an interaction time proportional to thickness of the plasma (b). Analysis similar to that in Sec. 4.3.4 shows that $n^2 * b$ increases in proportion to n ($n * b$ is approximately constant), meaning that in the initial experiment with a 5 V positron well, a total of 110 antiprotons would have been expected in the ionization well. It must be the case that in this initial experiment only antihydrogen atoms which ionize at fields of 20 Vcm^{-1} or less are created, and so all atoms quickly ionize in the nested well. The number of de-exciting collisions between the atomic positron and a plasma positron is, on average, proportional to both the positron plasma density and its thickness. Since this is approximately constant, the number of de-exciting collisions should not change as the well depth is increased. Despite this prediction a change in the collisional dynamics is still a likely candidate for the dramatic change in the number of atoms detected using the deeper positron well. Limited time prevented further investigation of the well depth dependency of recombination. Only one additional experiment was performed before our attention turned to investigating the atomic states of the antihydrogen produced.

This additional experiment retains the same potential structure and driving technique – the number of positrons is simply increased from 3.3×10^5 to 8.6×10^5. A total of 2.1×10^5 antiprotons are used. In this experiment as many as 718 atoms deposited an antiproton in the ionization well, or equivalently, 1 detected atom for every 290 antiprotons used. Again, no antiprotons were discovered when the experiment is repeated with no positrons present. The distribution of antiprotons in the ionization well is shown in Fig. 6.5.

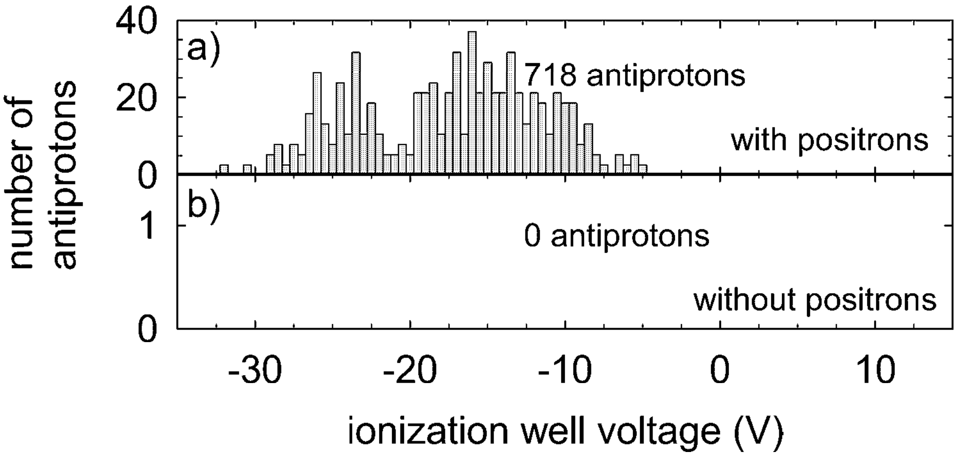

Figure 6.5: Ionization well contents with, a) 8.6×10^5, and b) 0 positrons in a 9 V well.

The numerical code for calculating equilibrium cloud shapes described in Ch. 4 is used to find the density, diameter, and axial extent of clouds of 3.3×10^5 and 8.6×10^5 positrons in the nested well. This is shown in Table 6.1. The higher density, axial extent and diameter of the 8.6×10^5 positron cloud will increase the antihydrogen production rate (recall that the antiproton cloud radius is approximately 4 mm –

$N \times 10^5$	n (10^{12}m^{-3})	a (mm)	b (mm)
3.3	15.4	3.2	0.45
8.6	17.0	3.7	0.79

Table 6.1: Cloud parameters for 3.3×10^5 and 8.6×10^5 positrons in a 9 V well.

larger than either of the positron clouds). Given the same energy input from the driven antiprotons a larger number of positrons will retain a lower temperature which will also translate into a higher recombination rate. All these effects will come into play, making for a complex dependence of the antihydrogen production rate on the physical parameters of the positron plasma.

6.1.4 Atom Spatial Distribution

In order to estimate the total number of atoms produced, not just the number which are caught in the ionization well, the spatial distribution of the antihydrogen must be known. If the antihydrogen were moving isotropically, then, since the ionization well covers roughly 1/260 of the solid angle, a total of 47 000 and 187 000 atoms were created in the two experiments. This second number is equivalent to creating 1 antihydrogen atom for every 1.14 antiprotons. Although technically possible, this figure strongly suggests that the antihydrogen distribution is peaked in the axial direction. This is perfectly reasonable since the radio-frequency drives increase the axial energy of the antiprotons which could preferentially direct the created atoms along the axis.

The spatial distribution of atoms is probed in an experiment described by Bowden [18] using a different method to induce antiproton - positron interaction. In these experiments the ionization well is moved progressively further from the nested well, resulting in a reduction in the number of antihydrogen atoms detected. This reduction is consistent with an isotropic distribution of atoms, and with an axial beam, which at the position of the ionization well is considerably larger than the electrode diameter.

Whatever the exact spatial distribution, the atoms detected are only those which are ionized by fields between 20 and 73 Vcm^{-1}. Other atoms, both more tightly, and more loosely bound, will have undoubtedly been created. The actual number of atoms produced is therefore rather uncertain. What is certain is that this was, at the time, by far the largest number of antihydrogen atoms directly detected in a 1 hour trial.

6.1.5 Estimated Limit on Atom Temperature

One can attempt to place a limit on the temperature of the recombined atom given that it must have been de-excited from an initial three body recombined state by collisions with plasma positrons. The time between these collisions (τ_{coll}) is given by,

$$\begin{aligned}
\tau_{coll} &= (n_{\overline{e}}\,\overline{v}\,r_T^2)^{-1} \qquad (6.1)\\
&= 0.5\ \mu s\ (T = 4.2\ K)\\
&= 1.5\ \mu s\ (T = 42\ K),
\end{aligned}$$

where $\overline{v}$ is the average velocity of positrons in the plasma and r_T is the initial capture radius of the atom, the Thompson radius. The collision time does not depend strongly on temperature and so a value $\tau_{coll} = 1\ \mu\mathrm{s}$ is a reasonable estimate. Numerical calculations of the positron cloud parameters (as described in Ch. 4) shows that it has a thickness of about 1 mm. In order to escape the nested well and be detected, the atom must have *at least* one de-exciting collision, and so on average remain within the positron plasma for 1 μs. The atomic velocity along the trap axis is therefore less than 1000 ms^{-1}, or equivalently, a temperature of 60 K. There is, however, the possibility that the atoms are predominantly created from antiprotons which, by chance, suffer a collision after much less than the average time between collisions. This would increase the limiting temperature above 60 K.

6.2 Electric Field Ionization Experiments

Attention is now turned to investigating the type of atoms that are being created. The high production rates observed in the previous experiments allow the ionization well to be moved further from the nested well and still collect an appreciable number of antiprotons from antihydrogen. This increased separation allows space for a variable "state-analysis" electric field. This field is used to pre-ionize atoms as they move from the nested well to the ionization well, thereby reducing the number of atoms which reach the ionization well and deposit their antiproton. Varying the state-analysis field and measuring the number of antiprotons in the ionization well determines the fraction of atoms that are pre-ionized by a given state-analysis field. The state-analysis field is directed along the z axis and is created by biasing several electrodes. In this way it varies by less than ±10% on and off the axis. In addition, a normalization well is included which acts very much like the ionization well in that it ionizes atoms and captures their antiprotons. The number of antiprotons from antihydrogen in this well provides a normalization for the number collected in the ionization well. Figure 6.6 shows the electrodes, potentials, and two examples of the state-analysis field.

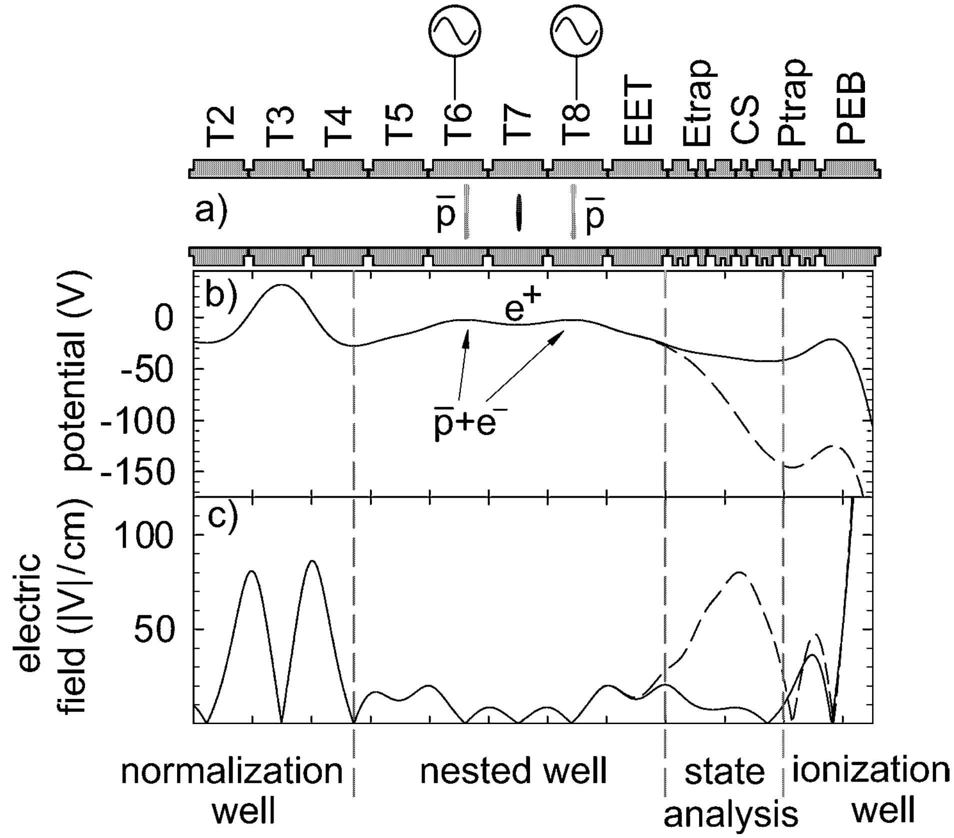

Figure 6.6: a) Electrodes, b) two values of the potential, and c) electric field magnitude on axis.

In this set of experiments the positrons are again contained in a 9 V well and 825 kHz drives are applied to the sides of the nested well just as in the experiments of the previous section. The only difference is that the antiprotons are divided evenly between the two sides of the nested well in order to be sure that as much antihydrogen moves towards the ionization well as towards the normalization well. In this case the loss observed from the nested well during driving does not have alternate loss peaks of very different amplitude, as was seen in Fig. 6.3.

The ratio, r, of the number of antiprotons in the ionization well (N_{ion}) to the number in the normalization well (N_{norm}) is plotted as a function of the state-analysis electric field (F) in Fig. 6.7. The measured values of r are consistent with a linear dependence on F up to a cut-off, and are zero at higher fields. This gives a constant number of atoms ionized per unit electric field (Fig. 6.7(b)). For example, the number of atoms which are ionized by fields between 30 and 40 Vcm^{-1} is equal to the number ionized by fields of 50 to 60 Vcm^{-1}. No atoms are observed which require fields greater than 62 Vcm^{-1} to ionize. This is seen in the distribution, dr/dF, which is zero above this field. The cut-off at 62 Vcm^{-1} is presumably not a discontinuous drop to zero as is shown in the figure, but the error bars on the data indicate that there is

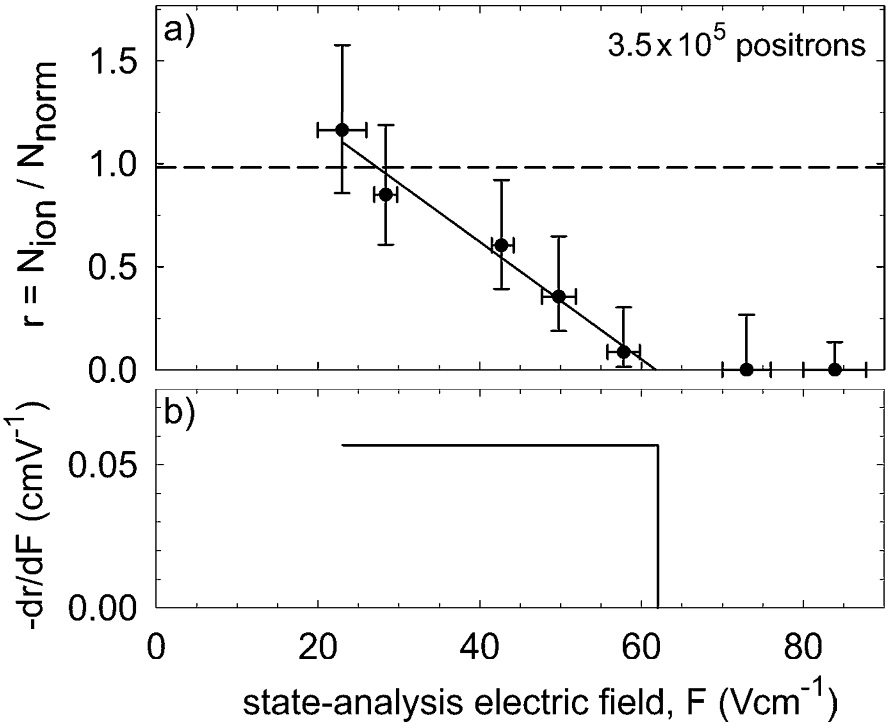

Figure 6.7: a) Ratio of the number of antiprotons in the ionization well to the number in the normalization well. The dashed line is the expected ratio for a state-analysis field of 25 Vcm^{-1}. b) dr/dF dependency on state-analysis field.

insufficient resolution to map out the edge.

Numerical modelling of antihydrogen trajectories from the nested well shows that antiprotons in the ionization well come from atoms which are stripped by fields of 25 to 150 Vcm^{-1}, while antiprotons in the normalization well are from atoms stripped by fields of 35 to 140 Vcm^{-1}. The ionization well and the normalization well also have different solid angle efficiencies for antihydrogen detection. When the state-analysis field is reduced to less than 25 Vcm^{-1} there is no pre-ionization of atoms. In this case the number of antiprotons in the ionization well should equal the number in the normalization well, once the different solid angles and ionization fields in the two wells have been taken into account. This is shown by the dashed line in Fig. 6.7(a). Indeed, the value of r for the lowest state-analysis field is consistent with this.

6.3 Ionization Theory for a Stationary Atom

Having measured the electric field which will ionize an antihydrogen atom one would like to parameterize the atom by its atomic state, the antiproton - positron separation, or the positron binding energy. This topic is addressed in a recent manuscript [6], and the remainder of the chapter is devoted to this end.

6.3.1 Guiding Center Atoms

The fact that a field as small as 60 Vcm^{-1} will ionize the atom indicates that it is rather loosely bound. Such highly excited atoms, when in a strong magnetic field, can be approximated as "guiding center" atoms [35], and described classically. Figure 6.8 shows the trajectory of the positron in such an atom with no external applied electric field. The positron performs fast cyclotron oscillations about the magnetic field,

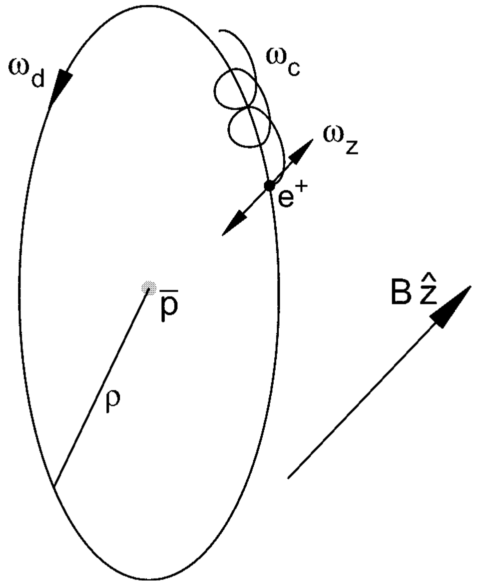

Figure 6.8: In a circular guiding center atom the positron oscillates axially in the Coulomb field of the antiproton, and the guided center of the cyclotron orbit drifts around the antiproton at a radius ρ.

oscillates in the z direction in the Coulomb potential well of the antiproton, and drifts around the antiproton with a velocity $\dot{\boldsymbol{\rho}} = (\mathbf{E} \times B\hat{z})/B^2$, where $\mathbf{E}$ is the Coulomb electric field. If the cyclotron motion is much faster than all other motions, and its radius much smaller than the drift radius, then the cyclotron motion can be averaged out. The positron trajectory then follows the guided center of the cyclotron orbit. The conditions for neglecting the cyclotron motion place limits on the minimum size of atom which is adequately described under the guiding center approximation. For a cyclotron frequency much greater than the drift frequency,

$$\omega_c = \frac{qB}{m} \gg \omega_d = \left(\frac{q}{4\pi\epsilon_0}\right)\frac{1}{\rho^3 B} \quad \Rightarrow \quad \rho \gg \left[\left(\frac{1}{4\pi\epsilon_0}\right)\frac{m}{B^2}\right]^{1/3} = 0.06\ \mu m, \tag{6.2}$$

for a magnetic field of 5.4 T. While for a drift radius much greater than the cyclotron radius,

$$\rho \gg r_c = \frac{\sqrt{2mkT}}{Bq} = 0.01\ \mu m, \tag{6.3}$$

for a positron temperature of 4.2 K, and field of 5.4 T. These limits will be found to (just) apply to the atoms created in the experiments detailed above. At smaller radii the guiding center approximation is not valid and, in addition, the orbits are chaotic [88].

6.3.2 Atomic Ionization of a Stationary Atom

A simple model of how a guiding center atom is ionized by an applied electric field, $F\hat{\boldsymbol{z}}$, is to assume a stationary antiproton (or at least one which does not have any momentum perpendicular to $B\hat{\boldsymbol{z}}$) and consider the axial potential well in which the positron resides. Figure 6.9 shows how the potential in the z direction changes with increasing atomic radius, ρ, for a fixed F. It can be seen that the axial well disappears as the radial separation of the positron and the antiproton increases. An alternative viewpoint is to consider an atom of fixed ρ and vary the applied field (Fig. 6.10). As the field is increased the potential minimum is displaced slightly from $z = 0$ as the atom is polarized by the field, and the depth of the Coulomb well is reduced. For axial displacements much smaller than ρ the polarizability of the atom, α_p, is easily found to be:

$$\alpha_p = \left(\frac{4\pi\epsilon_0}{q}\right)\rho^3, \tag{6.4}$$

which is reasonable in the sense that a larger atom is more weakly bound, and therefore easier to polarize.

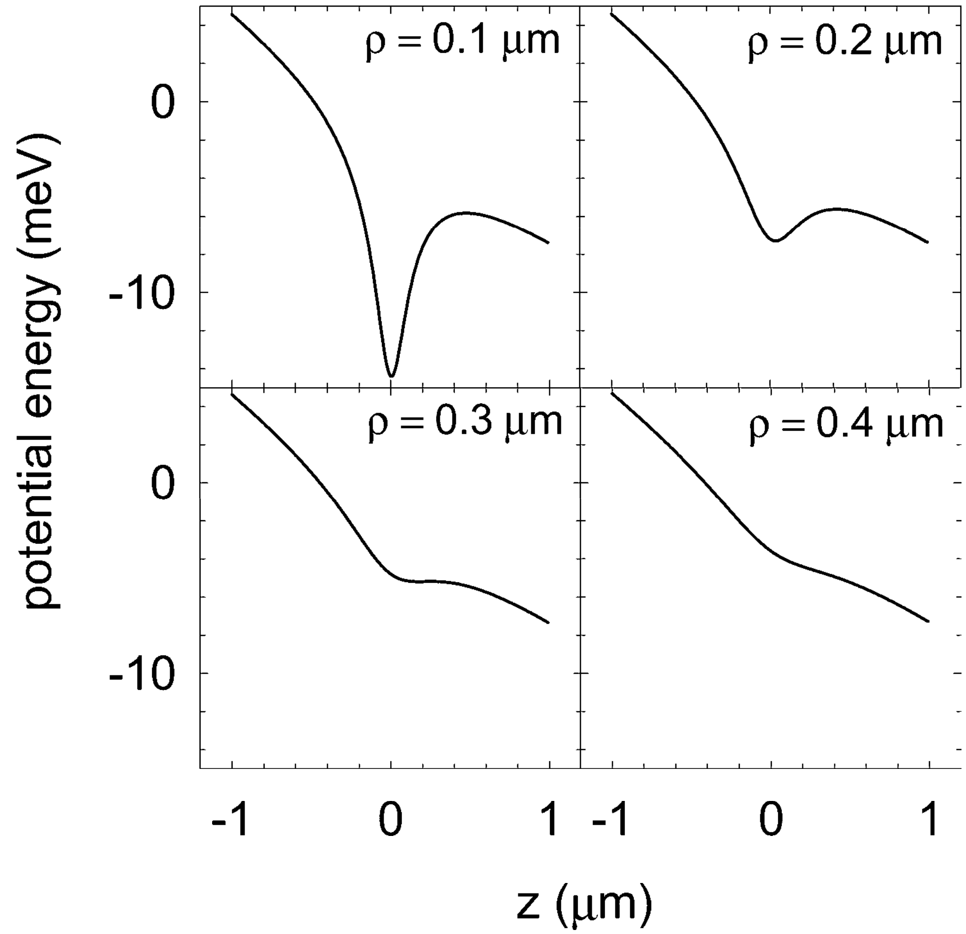

Figure 6.9: In an applied electric field the axial well depth decreases and moves towards positive z as the positron - antiproton radial separation increases. All plots are for an applied field of 60 Vcm^{-1}.

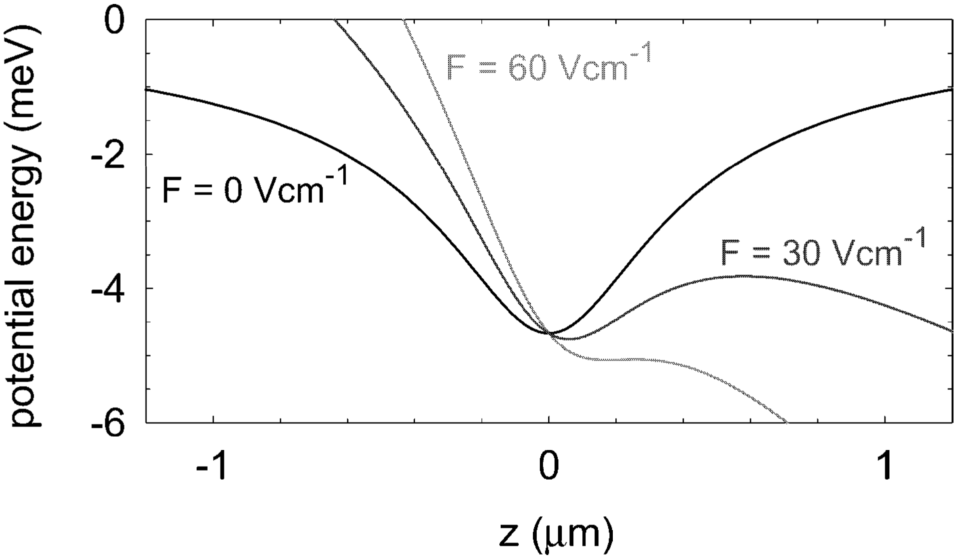

Figure 6.10: As the applied electric field increases the axial well depth decreases and moves towards positive z. All plots are for an atomic radius of 0.3 μm.

Clearly, once the axial well has disappeared the positron will escape and the atom is ionized. If the positron has some axial energy then it can leave the axial well before the well entirely disappears. In general this will be the case, since through collisions, the axial motion of the atomic positron is expected to come into thermal equilibrium with the positrons in the positron cloud [35]. Comparing the positron energy to the depth of the axial well allows the calculation of the field at which a positron at a given ρ will escape. This calculation is performed for the limiting positron temperature, 0 K, and for 4.2 K and 42 K, since the temperature of the positron plasma while antiprotons are being driven through is not known. Figure 6.11(a) shows the results. A field of 60 Vcm^{-1} is seen to ionize an atom with a positron - antiproton separation of 0.26 μm, if the atomic positron has an energy of 0.36 meV (4.2 K). At this temperature, this would be the smallest atom created experimentally, since none survive larger pre-ionizing fields. A higher positron temperature would indicate

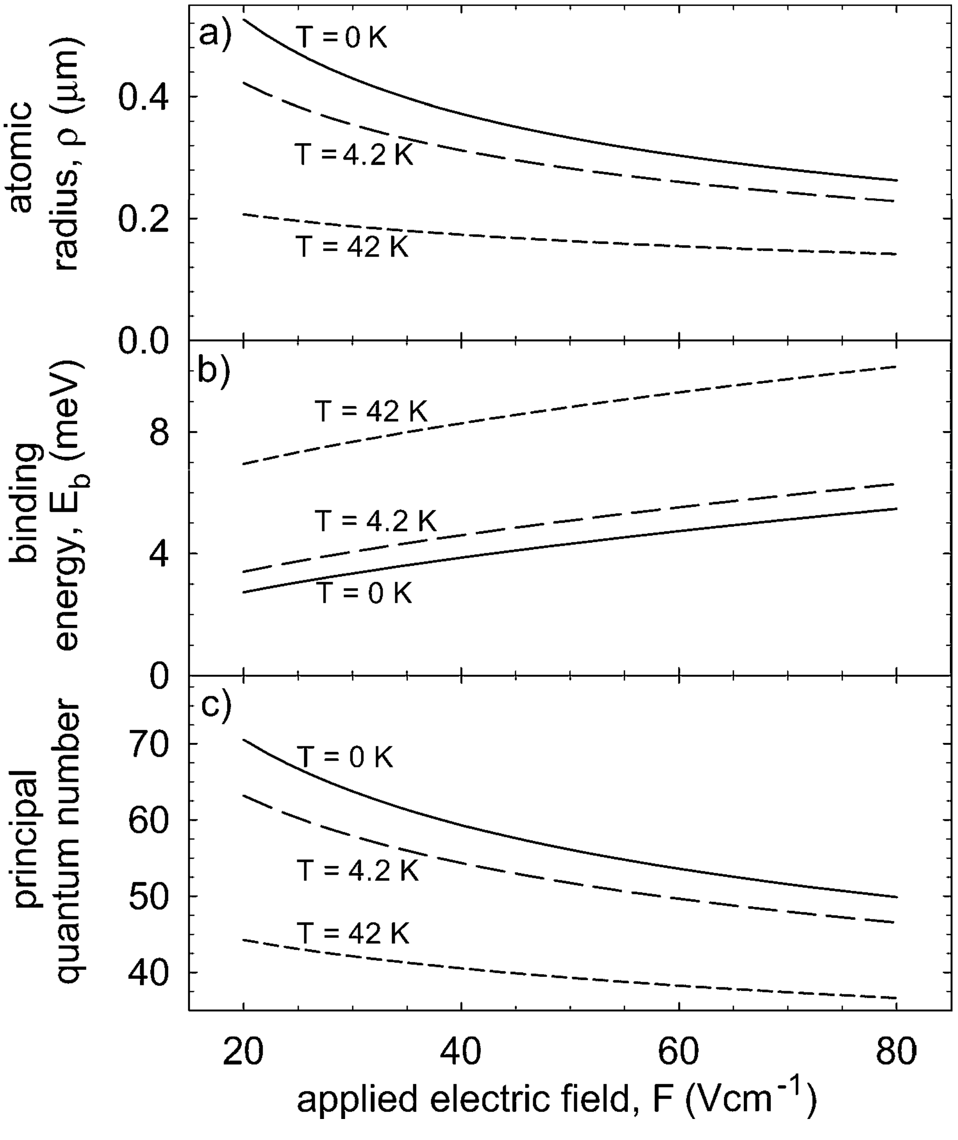

Figure 6.11: a) An applied electric field will ionize atoms with radii above the value shown. Equations 6.6 and 6.7 translate the atomic radius into, b) a binding energy, and c) a value for the principal quantum number of the atom.

a smaller atomic radius, while the limiting value of $T = 0$ indicates a slightly larger atom. In the $T = 0$ limit there is a simple relationship between the radius of the atom and the field which ionizes it,

$$\rho = \frac{2.35}{\sqrt{F[\mathrm{Vcm}^{-1}]}}\ \mu\mathrm{m}. \tag{6.5}$$

Also shown in Fig. 6.11 are the binding energy, E_b, and a value for the principal quantum number, n_q, where E_b is given by,

$$E_b = \frac{q^2}{4\pi\epsilon_0\rho}, \tag{6.6}$$

and n_q is given by,

$$n_q = \sqrt{\frac{13.6[\mathrm{eV}]}{E_b}}. \tag{6.7}$$

The relationship of Eq. 6.7 should be treated with caution since for such highly excited atoms in a large magnetic field n_q is not a good quantum number. Figure 6.11(c), which uses this relationship, is included only to give an estimate of this commonly used atomic parameter.

The relationship between ρ and F displayed in Fig. 6.11(a) allows the results of Fig. 6.7(a) to be interpreted in terms of the atomic radius, rather than the ionizing electric field. This is shown in Fig. 6.12. The calculated atomic radii are seen to depend substantially on the energy of the atomic positron. For a positron temperature of $T = 42$ K, no atoms smaller than 0.15 μm are created, while none larger than 0.21 μm can escape the nested well intact to be detected. For $T = 0$ the corresponding limits on the atomic size are 0.30 μm and 0.52 μm.

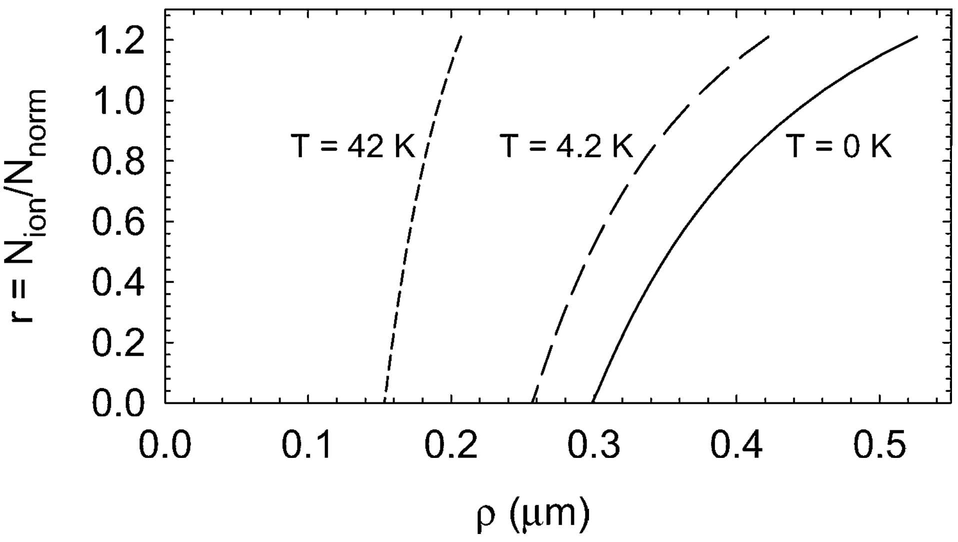

Figure 6.12: The range of atomic radii detected in the experiments of Sec. 6.2, as determined by the analysis of a stationary atom in an applied electric field. The ordinate is a measure of the number of atoms more tightly bound than the given ρ.

6.4 Ionization Theory for a Moving Atom

The analysis presented in the previous section requires that the antiproton is not moving in the magnetic field. If this restriction is lifted then the problem becomes significantly more difficult, since it is not possible to separate the center of mass motion of the atom from the internal motion. The motion of the center of mass will therefore affect the ionization process.

6.4.1 The Pseudomomentum

The problem is, however, simplified by the introduction of a pseudomomentum (**K**) which is a conserved quantity for the system, and leads to a pseudoseparation of the center of mass coordinate, **R**, and the relative coordinate, **r** [89, 90, 91, 6]. It is called a pseudoseparation since the positron then moves in an effective potential which depends only on the relative coordinate and the eigenvalue of **K**. The Hamiltionian for the atom in an applied magnetic and electric field is given by [90]:

$$H = \frac{1}{2\mu}(\mathbf{p} - qA(\mathbf{r}))^2 + \frac{(M\dot{\mathbf{R}})^2}{2M} + V(\mathbf{r}) - q\mathbf{F}.\mathbf{r}, \tag{6.8}$$

where μ is the reduced mass, **p** the relative momentum vector, $A(\mathbf{r})$ the vector potential, M is the sum of the positron and antiproton masses, $V(\mathbf{r})$ is the Coulomb potential energy, and F is the state-analysis field. The pseudomomentum is defined by,

$$\mathbf{K} \equiv M\dot{\mathbf{R}} + q\mathbf{B} \times \mathbf{r}, \tag{6.9}$$

allowing the replacement of the center of mass part of Eq. 6.8, leaving only relative coordinates:

$$H = \frac{1}{2\mu}(\mathbf{p} - qA(\mathbf{r}))^2 + \frac{1}{2M}(\mathbf{K} - q\mathbf{B} \times \mathbf{r})^2 + V(\mathbf{r}) - q\mathbf{F}.\mathbf{r}. \tag{6.10}$$

The last three terms on the right hand side are gauge independent (they do not contain $A(\mathbf{r})$) and can therefore represent an effective potential energy for the positron, given by,

$$\begin{aligned} V_{eff} &= \frac{1}{2M}(\mathbf{K} - q\mathbf{B} \times \mathbf{r})^2 + V(\mathbf{r}) - q\mathbf{F}.\mathbf{r} \\ &= \frac{\mathbf{K}^2}{2M} - \frac{q}{M}\mathbf{r}.(\mathbf{K} \times \ \mathbf{B}) + \frac{q^2}{2M}(\mathbf{B} \times \mathbf{r})^2 + V(\mathbf{r}) - q\mathbf{F}.\mathbf{r}. \end{aligned} \tag{6.11}$$

Without loss of generality, let $\mathbf{K} = K_\perp \hat{\boldsymbol{y}}$ and the effective potential energy becomes,

$$V_{eff} = \frac{K_\perp^2}{2M} - \frac{qK_\perp B}{M}x + \frac{B^2q^2}{2M}(x^2 + y^2) + \frac{q^2}{4\pi\epsilon_0}(x^2 + y^2 + z^2)^{-1/2} - qFz, \tag{6.12}$$

where all coordinates are those of the relative separation of the positron and the antiproton. The first term in the potential is constant, the second corresponds to a motional electric field which acts to ionize the atom in the x direction, and the third term is a confining potential which dominates at large distances.

An estimate of an upper limit for the value of $K_\perp$ of the atoms created can be calculated. The spatial distribution of atoms is peaked to some degree in the direction along the trap axis (Sec. 6.1.4). The center of mass velocity is therefore smaller perpendicular to the axis, than along it,

$$M\dot{Y} < M\dot{Z} < 1 \; a.u., \tag{6.13}$$

where the second inequality is from the requirement that the atom remain in the positron plasma for long enough to be de-excited (Sec. 6.1.5). One atomic unit (a.u.) of momentum is equal to 2×10^{-24} Kgms^{-1}, and only the Y component is considered, since $\mathbf{K}$ is chosen to be entirely in the Y direction. The value of $K_\perp$ is given by,

$$K_\perp = M\dot{Y} + qBx. \tag{6.14}$$

The relative co-ordinate, x, is an oscillatory function with an amplitude ≈ 0.25 μm, as estimated above. The contribution to $K_\perp$ from the second term is ≤ 0.1 a.u., and a limiting value of $K_\perp$ is therefore 1 a.u.. Of the atoms created, those which reach the ionization well have a tighter constraint on the value of $K_\perp$, as discussed below (Sec. 6.4.4).

6.4.2 Positron Trajectories in the Effective Potential

To demonstrate the full range of effects due to center of mass motion perpendicular to the magnetic field, values of $K_\perp$ from 0 a.u. to twice the expected maximum value, i.e., 2 a.u., are considered now. Figure 6.13 shows the effective potential in the x direction ($y = z = 0$) for values of $K_\perp$ in this range, with no external applied electric field. For small x the Coulomb well dominates, but as x increases the motional electric field acts to lower the potential. For large x, the quadratic confining potential dominates, and for intermediate values a potential well can develop. It is possible

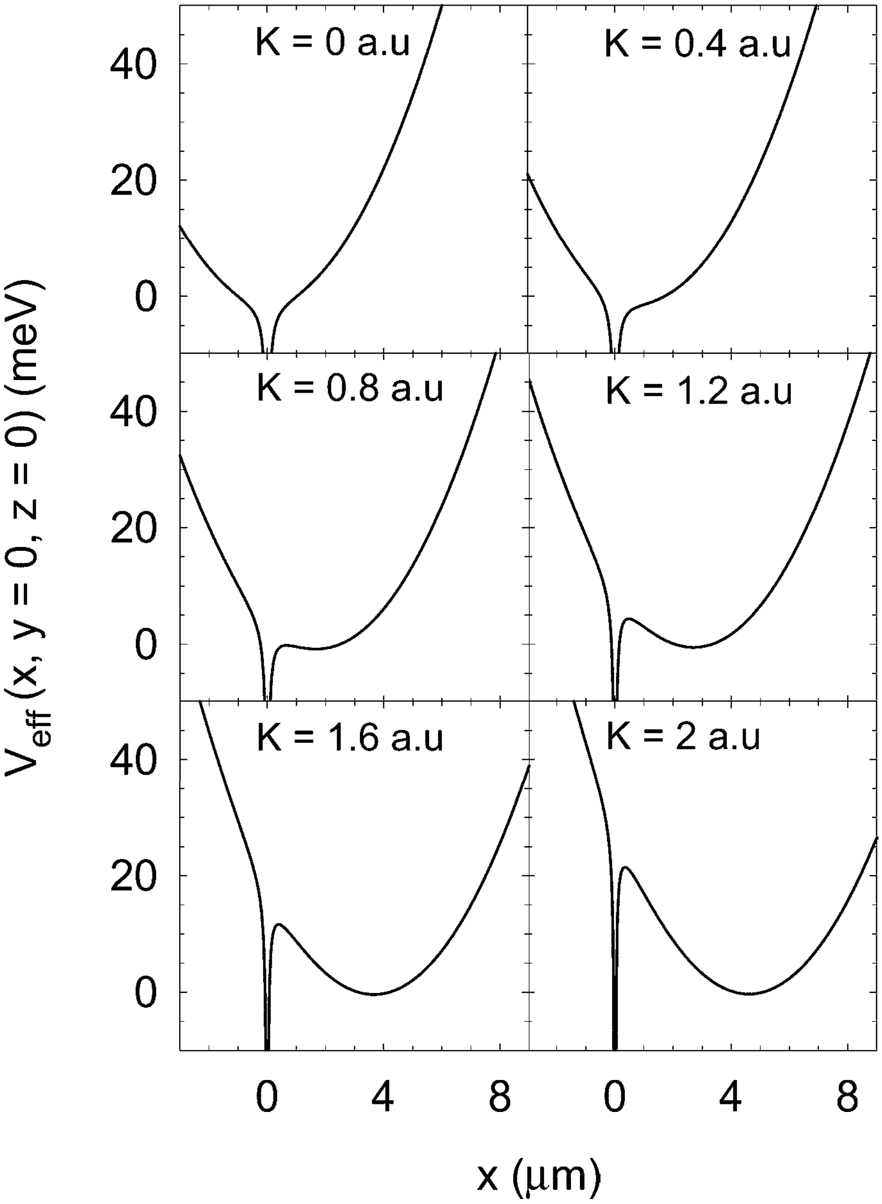

Figure 6.13: As the value of $K_\perp$ increases so a well, in addition to the Coulomb well, develops at large x.

for the atomic positron to become trapped in this outer well, and such an atom is referred to as a giant dipole. Depending upon the positron - antiproton separation and value of $K_\perp$ the outer well can affect the positron orbit even when it is confined in the Coulomb well at small x.

The classical equations of motion, which dictate the guiding center positron trajectory are,

$$\begin{aligned} \dot{\boldsymbol{\rho}} &= \frac{-1}{qB^2}(\nabla V_{eff} \times \mathbf{B}) \\ \mu \ddot{\boldsymbol{z}} &= \frac{\partial}{\partial z} V_{eff}. \end{aligned} \tag{6.15}$$

These are integrated numerically to solve for the trajectory of the positron with a starting position, $x = y = 0.5$ μm, $z = 0$, and an initial velocity in the z direction equivalent to an energy of 4.2 K. These initial x, y values are, by the analysis of Sec. 6.3, somewhat larger than the atom sizes found experimentally, but are chosen since they illustrate very well the range of possible trajectories. The equations are integrated for 0.2 μs (equivalent to 35 radial orbits and 1000 axial oscillations, when $K_\perp = 0$) and the trajectory is unchanged if this time is doubled. Values of $K_\perp$ ranging from 0 to 2 a.u. are used, and the trajectories are shown in Fig. 6.14. For $K_\perp = 0$ the positron orbit is simply circular (a "circular guiding center atom" of Ref. [6]), but as $K_\perp$ is increased the outer potential well develops which perturbs the motion, allowing the positron orbit to extend along the positive x axis. For $K_\perp = 2$ a.u. the positron no longer moves around the antiproton at all, but is entirely contained within the outer well - this is a giant dipole atom mentioned above. Such atoms are, however, ionized by an applied electric field of only 1 Vcm^{-1}, as confirmed by further numerical modelling. They will therefore always be ionized by the fields in the nested

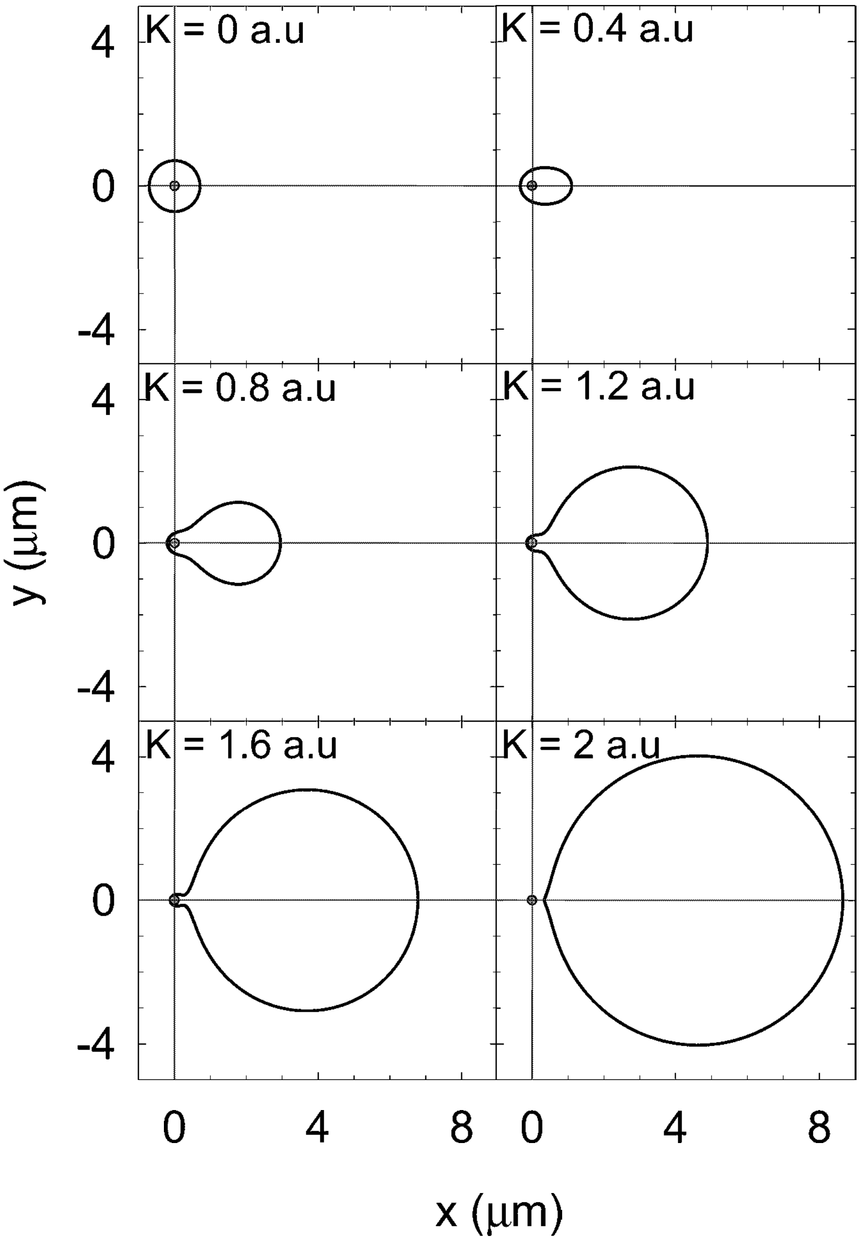

Figure 6.14: The positron radial trajectory starting at (0.5 μm, 0.5 μm) is influenced greatly by the value of $K_\perp$.

well.

The pseudomomentum also influences the motion in the z direction. For $K_\perp = 0$ the positron simply oscillates with a constant amplitude and frequency in the Coulomb well of the antiproton. As $K_\perp$ is increased the positron - antiproton separation varies in the $x - y$ plane as seen in Fig. 6.14. This variation modulates the strength of the Coulomb well, affecting the z motion: when the separation is small, the Coulomb well is stiff and the positron oscillates at a high frequency and a small amplitude, while for large separations, the Coulomb well is weak and the oscillation is much slower, and with an increased amplitude. The variation of axial frequency with radial separation is stated explicitly in Ref. [6], the relationship being $\omega_z \sim \rho^{-3/2}$. The effect of non-zero $K_\perp$ is shown in Fig. 6.15.

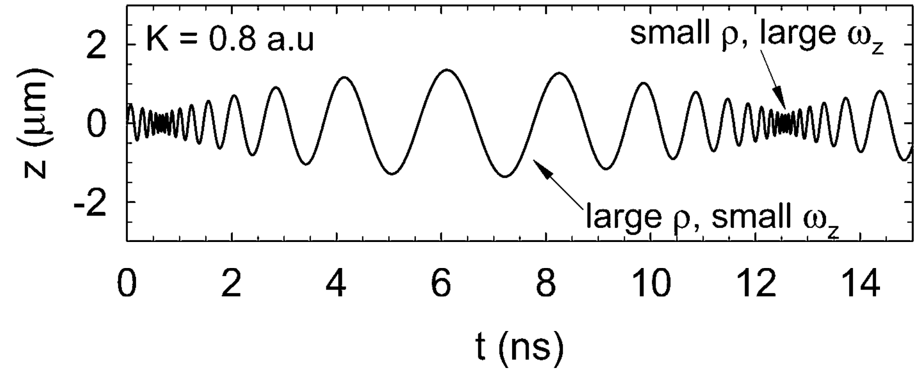

Figure 6.15: The positron axial motion is influenced greatly by the value of $K_\perp$.

6.4.3 Atomic Ionization of a Moving Atom

Having illustrated the influence of the center of mass motion on the positron trajectory the full experimental conditions can be modelled by including the applied electric field, F. Increasing F will polarize and eventually ionize the atom. This can be seen by integrating the equations of motion (Eq. 6.15) and examining the relative z coordinate to discover if ionization has occurred. Figure 6.16 shows the results of two such calculations. Both model a positron with 4 K axial energy, initially at (0, 0.2 μm, 0), with $K_\perp = 1.2$ a.u., but with two different fields: 30 Vcm^{-1} and 60 Vcm^{-1}. No initial velocities are assumed in the $x - y$ plane. Although the figure shows only the first nanosecond of the z trajectory, each calculation is run for at least 200 ns. The maximum integration time used is 3 μs and no trajectory has been discovered which was stable after 200 ns, but not after 3 μs. As can be seen a field of 60 Vcm^{-1} ionizes the atom after a fraction of a nanosecond (less than one radial orbit), while 30 Vcm^{-1} polarizes the atom (shifts the trajectory in the positive z direction) without ionizing it.

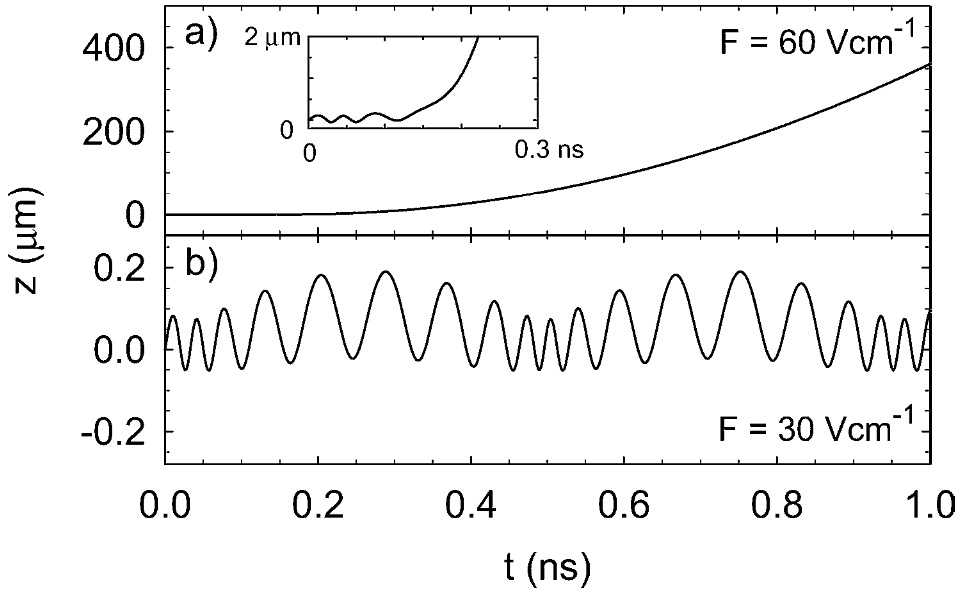

Figure 6.16: The axial motion of an atomic positron in a field of, a) 60 Vcm^{-1}, and b) 30 Vcm^{-1}. The latter field clearly ionizes the atom.

A series of similar calculations are made with fixed $K_\perp = 1$ a.u., $F = 60$ Vcm^{-1}, but varying the initial position (x_0, y_0) of the atomic positron. In this way a region of stability is found in the $x - y$ plane, as shown in Fig. 6.17(a). Also shown in the figure is the stability region in the case, $K_\perp = 0$. The calculations are repeated for a field of 30 Vcm^{-1} (Fig. 6.17(b)). One can see that for $K_\perp = 1$ the region of stability is diminished (relative to $K_\perp = 0$) in all parts of the $x-y$ plane except along the positive x axis. The effective potential is weakest along this half axis which means that if a trajectory starts at $(\rho_0, 0)$ then ρ_0 is the largest positron - antiproton separation at any time during the orbit. By contrast, for any other starting position (with the same initial separation, ρ_0) ρ will increase as the positron crosses the positive x axis. The extremity of the orbit will therefore be larger, the potential in the z direction weaker, and the likelihood of ionization greater. The co-incidence of the stability regions for $K_\perp = 0$ and $K_\perp = 1$ on the positive x axis is only approximate, and does not occur for much larger values of $K_\perp$, or x, where the term $eBK_\perp x/M$ in Eq. 6.12 has a greater effect.

Given that, experimentally, no atom survives a field greater than 60 Vcm^{-1} the boundary of the stable region in Fig. 6.17(a) represents the size of the most tightly bound atom produced in the experiments. For the case $K_\perp = 0$ this equates to an atomic radius of 0.25 μm, while for $K_\perp = 1$ the atomic radius is, on average, 0.21 μm.

6.4.4 Center of Mass Motion

Since one can solve for the positron trajectory $(x(t), y(t))$ then the relationship,

$$\dot{Y}(t) = \frac{K_\perp}{M} - \frac{qBx(t)}{M}, \tag{6.16}$$

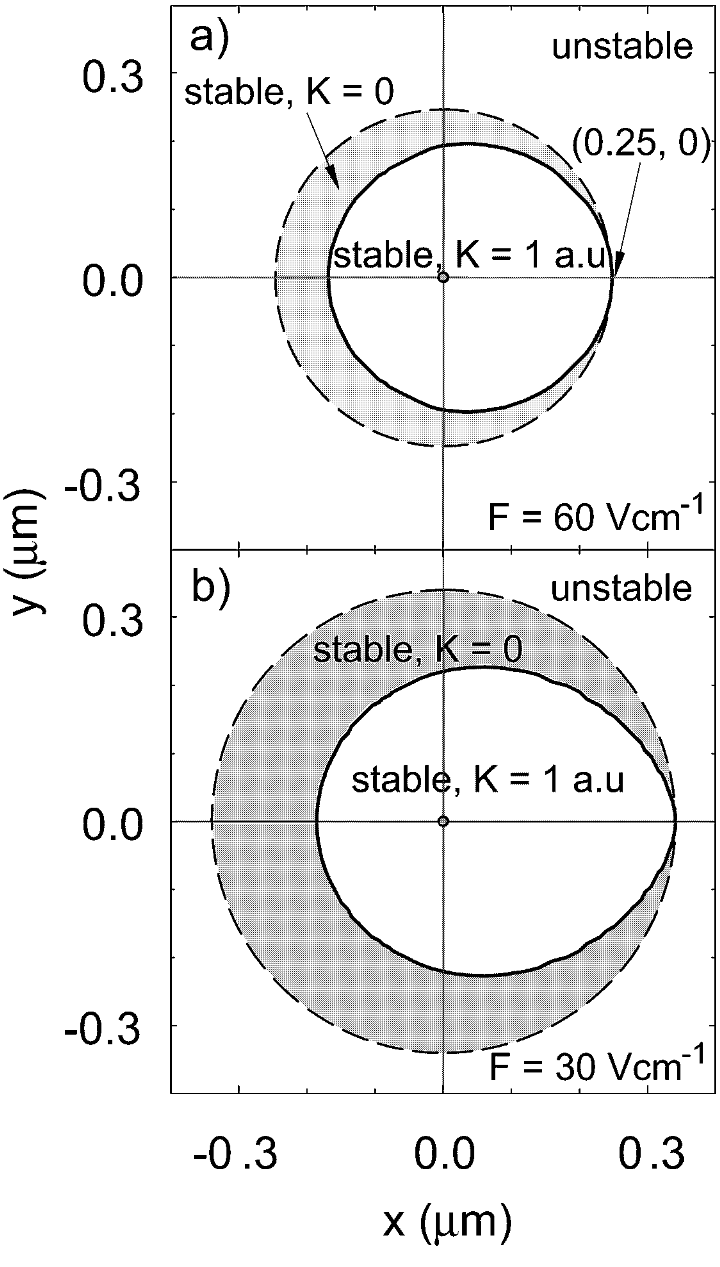

Figure 6.17: In an applied field of a) 60 Vcm^{-1}, and b) 30 Vcm^{-1}, the radial position of the atomic positron determines whether or not the atom is stable. Three regions are noted: for $K_\perp = 1$ all atoms within the solid boundary are stable, for $K_\perp = 0$ all atoms within the dashed boundary are stable, while outside the dashed boundary all atoms are unstable, irrespective of the value of $K_\perp$.

allows the calculation of the center of mass co-ordinate, Y. Since $K_\perp$ is a constant of the motion, $\dot{Y}(t)$ will mirror changes in $x(t)$, while $K_\perp$ remains constant. Equation 6.16 can be integrated numerically for a given $K_\perp$ to solve for the center of mass motion, once $x(t)$ is found from the equations of motion (Eqs. 6.15).

The most tightly bound atoms observed are those which will just survive a field of 60 Vcm^{-1}. To calculate the center of mass motion of these atoms Eqs. 6.15 are solved for $x(t)$ with the positron starting at (0.25 μm, 0). This position is the intersection of the boundary of the stability region with the positive x axis in Fig. 6.17(a). The resulting trajectory is shown in Fig. 6.18 for $K_\perp = 1$, $F = 0$.

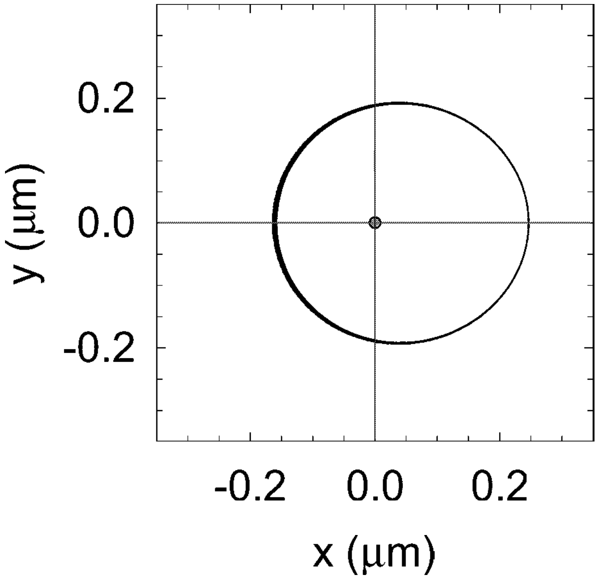

Figure 6.18: The trajectory of the positron in the most tightly bound atoms created experimentally, if $K_\perp = 1$.

The results of the numerical integration of Eq. 6.16, for $K_\perp = 1$, $F = 0$ are shown in Fig. 6.19. The center of mass motion deviates from that of a ballistic atom, the trajectory of which is found by neglecting the second term on the right hand side of Eq. 6.16. The internal motion is oscillatory, as seen in Fig. 6.18, and it is

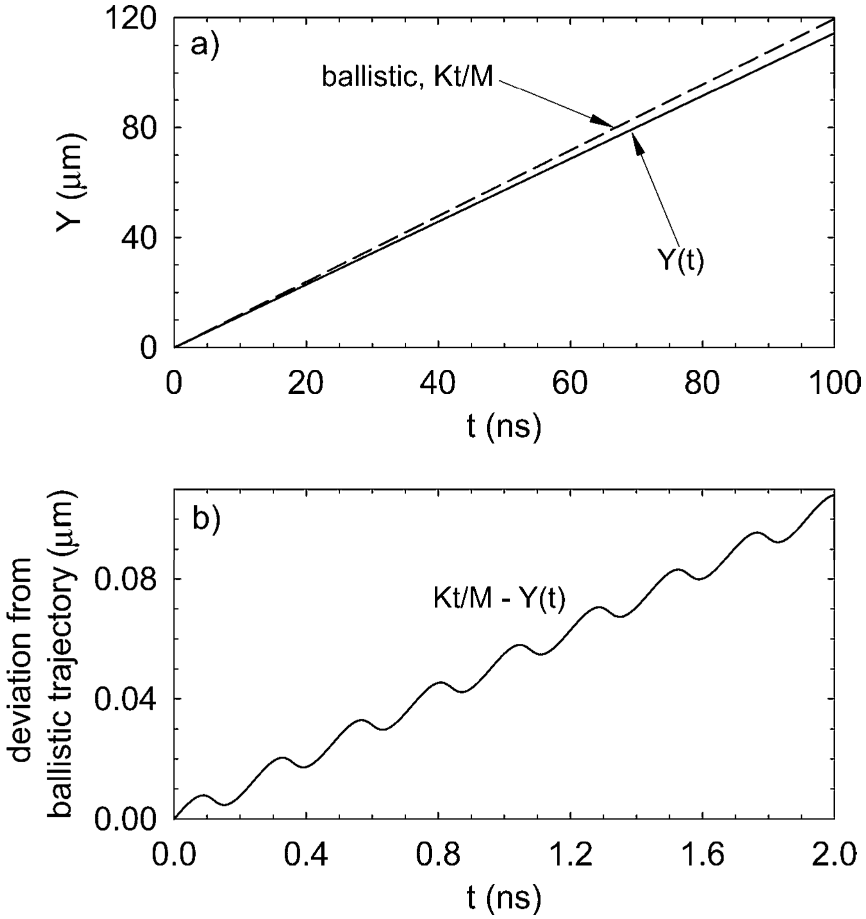

Figure 6.19: a) The center of mass (solid line) moves slightly slower than a ballistic atom (dashed line). b) The deviation from a ballistic trajectory exhibits oscillations due to the oscillatory internal motion.

this oscillation which is imprinted onto $Y(t)$ (Fig. 6.19(b)). The center of the $x(t)$ oscillation is offset from (0, 0) which reduces the average gradient of $Y(t)$ below that of a ballistic trajectory. The dynamics of the center of mass are only slightly altered when a field of 60 Vcm^{-1} is applied. The motion in the X direction is oscillatory with a small amplitude and no offset.

One can see from Fig. 6.19(a) that the average Y velocity is about 1000 ms^{-1}. Since an estimate of the maximum velocity along the trap axis is $\dot{Z} = 1000$ ms^{-1} (Sec. 6.1.5) the atom is not directed into the ionization well, which only accepts atoms travelling at less than 7° to the trap axis. If the value of $K_\perp$ is reduced and the calculation repeated, it is found that for $K_\perp \leq 0.1$ a.u. an atom is within the acceptance of the ionization well. Therefore, for the most tightly bound atoms created experimentally, $K_\perp = 0.1$ a.u. represents the maximum value of pseudomomentum that an atom can possess, and be detected in the ionization well. For a given $K_\perp$, an atom which ionizes in a field of 30 Vcm^{-1} has a $\sim 10\%$ smaller center of mass velocity than the most tightly bound atoms. $K_\perp = 0.1$ a.u. therefore represents an upper limit on the pseudomomentum for all atoms which reach the ionization well.

6.4.5 Summary

The analysis of the electric field ionization studies of antihydrogen presented here provides a framework for calculating the atomic dimensions given a value of an ionizing field. The inclusion of center of mass motion perpendicular to the magnetic field direction leads to a range of positron orbits which deviate from the circular orbits around a stationary antiproton. Experiments have placed limits on the value of pseudomomentum possible for all atoms created, and for those atoms which arrive at the detection region. The most tightly bound atoms produced experimentally have an atomic radius of 0.2 - 0.25 μm. This atomic size determines that the atoms are initially created by three body recombination, and then collisionally de-excited while passing through the positron plasma.

Chapter 7

Conclusions and Future Directions

ATRAP experiments at CERN have been extremely successful in the last two years. Slow antihydrogen is produced during the positron cooling of antiprotons in a nested Penning trap, a technique we first demonstrated during this time [5]. Large numbers of atoms are created and detected in a background-free measurement [1]. Field ionization allows the first investigation of the internal structure of the atoms [2]. Theoretical analysis of these experiments determines a positron - antiproton separation of $\sim 0.2\ \mu$m. This suggests that the atoms are initially formed in a three body recombination process, and then collisionally de-excited while they travel through the positron plasma.

The production and subsequent study of the antihydrogen internal structure represents the culmination of a number of advances in antiproton and positron trapping, manipulation, and diagnostics. In 1 hour it is now routine to simultaneously accumulate 3×10^5 antiprotons and 1 million positrons, both cooled to 4.2 K, for use in antihydrogen studies. The positron accumulation technique is the most efficient method of trapping large numbers of positrons directly into a cryogenic UHV environment [7]. The stacking technique used to accumulate antiprotons is the only way to capture more that 20 000 antiprotons for antihydrogen studies [10, 11]. Nearly half a million cold antiprotons have been accumulated in 1 hour using this technique. Given the linearity of the number of trapped, cold antiprotons with the number of bunches delivered by the CERN accelerator, this record number could be increased.

A new diagnostic establishes the physical parameters of the antiproton and positron plasmas confined in our Penning trap [17]. This allows the first determination of the antiproton plasma properties, and is the only applicable technique for measuring the positron plasmas confined in our trap. The knowledge of these parameters is essential for interpreting the results of antihydrogen experiments, for accurate calculations of recombination rates, and for the design of future experiments. Standard assumptions made when calculating plasma properties in a Penning trap do not accurately describe the antiproton clouds, as demonstrated by our numerical calculation.

The next step for antihydrogen experiments is to find a method to de-excite the highly excited atoms to the ground state. Two approaches are currently underway at CERN. The first is to use much larger numbers of positrons which will provide many more positron - atom collisions after the atom is first formed. With sufficient de-excitation spontaneous decay to the ground state will be possible. The second method is to resonantly transfer the excited atoms to a much lower state ($n = 3$) by applying a laser field. This lower state will then spontaneously decay to the ground state in a few nanoseconds.

Once ground state atoms have been produced, the hope is to trap them in the minimum of a magnetic gradient trap imposed upon the homogenous Penning trap field. The consequences for charged particle trapping in such a combined trap have been studied [92], and the results look favorable, at least for low particle densities. The trapping will be aided by laser cooling the atoms using a continuous Lyman alpha source [93]. The final step will be to measure the 1S-2S transition frequency and compare it to that of hydrogen for an extremely accurate test of CPT symmetry.

Many challenges lie ahead but great progress has been made.

Bibliography

[1] G. Gabrielse, N. S. Bowden, P. Oxley, A. Speck, C. H. Storry, J. N. Tan, M. Wessels, D. Grzonka, W. Oelert, G. Schepers, T. Sefzick, J. Walz, H. Pittner, T. W. Hansch, and E. A. Hessels, *Phys. Rev. Lett.*, 89:213401-1, 2002.

[2] G. Gabrielse, N. S. Bowden, P. Oxley, A. Speck, C. H. Storry, J. N. Tan, M. Wessels, D. Grzonka, W. Oelert, G. Schepers, T. Sefzick, J. Walz, H. Pittner, T. W. Hansch, and E. A. Hessels, *Phys. Rev. Lett.*, 89:233401-1, 2002.

[3] M. Amoretti, C. Amsler, G. Bonomi, A. Bouchta, P. Bowe, C. Carraro, C. L. Cesar, M. Charlton, M. J. T. Collier, M. Doser, V. Filippiniq, K. S. Fine, A. Fontanaq, M. C. Fujiwara, R. Funakoshi, P. Genova, J. S. Hangst, R. S. Hayano, M. H. Holzscheiter, L. V. Jrgensen, V. Lagomarsino, R. Landua, D. Lindelo, E. Lodi Rizzini, M. Macry, N. Madsen, G. Manuzio, M. Marchesottiq, P. Montagnaq, H. Pruys, C. Regenfus, P. Riedler, J. Rochet, A. Rotondiq, G. Rouleau, G. Testera, A. Variola, T. L. Watson, and D. P. van der Werf, *Nature*, 419:456, 2002.

[4] G. Gabrielse, S. L. Rolston, L. Haarsma, and W. Kells, *Phys. Lett.*, A129:38, 1988.

[5] G. Gabrielse, J. Estrada, J. N. Tan, P. Yesley, N. S. Bowden, P. Oxley, T. Roach, C. H. Storry, M. Wessels, J. Tan, D. Grzonka, W. Oelert, G. Schepers, T. Sefsick, W. Breunlich, M. Carngelli, H. Fuhrmann, R. King, R. Ursin, H. Zmeskal, H. Kalinowsky, C. Wesdorp, J. Walz, K. S. E. Eikema, and T. W. Hansch, *Phys. Lett. B*, 507:1, 2001, and see Ref. [7] in [1].

[6] D. Vrinceanu, D. Vinceanu, H. R. Sadeghpour, L. Cederbaum, R. Parrott, A. Mody, J. N. Tan, and G. Gabrielse, *to be published.*

[7] J. Estrada, T. Roach, J. N. Tan, P. Yesley, and G. Gabrielse, *Phys. Rev. Lett.*, 84:859, 2000.

[8] G. Gabrielse, X. Fei, K. Helmerson, S. L. Rolston, R. Tjoelker, T. A. Trainor, H. Kalinowsky, J. Haas, and W. Kells, *Phys. Rev. Lett.*, 57:2504, 1986.

[9] G. Gabrielse, X. Fei, L. A. Orozco, R. L. Tjoelker, J. Haas, H. Kalinowsky, T. Trainor, and W. Kells, *Phys. Rev. Lett.* 63:1360, 1989.

[10] G. Gabrielse, *Advances in Atomic, Mololecular and Optical Physics.*, Vol. 45, Edited by B. Bederson and H. Walther (Academic Press, New York, 2001), p. 1.

[11] G. Gabrielse, N. S. Bowden, P. Oxley, A. Speck, C. H. Storry, J. N. Tan, M. Wessels, D. Grzonka, W. Oelert, G. Schepers, T. Sefzick, J. Walz, H. Pittner, T. W. Hansch, and E. A. Hessels, *Phys. Lett. B*, 548:140, 2002.

[12] R. L. Tjoelker, Ph.D. thesis 1990.

[13] X. Fei, R. Davisson, and G. Gabrielse, *Rev. Sci. Instrum.*, 58:2197, 1987.

[14] G. Gabrielse, X. Fei, L. A. Orozco, R. L. Tjoelker, J. Haas, H. Kalinowsky, T. Trainor, and W. Kells, *Phys. Rev. Lett.*, 65:1317, 1990.

[15] G. Gabrielse, D. Phillips, W. Quint, H. Kalinowsky, G. Rouleau, and W. Jhe, *Phys. Rev. Lett.*, 74:3554, 1995.

[16] G. Gabrielse, A. Khabbaz, D.S. Hall, C. Heimann, H. Kalinowsky, and W. Jhe, *Phys. Rev. Lett.*, 82:3198, 1999.

[17] P. Oxley, N. S. Bowden, R Parrott, A. Speck, C. H. Storry, J. N. Tan, M. Wessels, G. Gabrielse, D. Grzonka, W. Oelert, G. Schepers, T. Sefzick, J. Walz, H. Pittner, T. W. Hansch, and E. A. Hessels, *to be published.*

[18] N.S. Bowden, Ph.D. thesis 2003.

[19] G. Lüders, *Ann. Phys.* 2:1, 1957.

[20] T. D. Lee and C. N. Yang, *Phys. Rev.* 104:254, 1956.

[21] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson, *Phys. Rev.* 105:1413, 1957.

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

[23] Don Colladay and V. Alan Kostelecky, *Phys. Rev. D*, 55:6760, 1997.

[24] Robert Bluhm, V. Alan Kostelecky, and Neil Russell, *Phys. Rev. Lett.*, 79:1432, 1997.

[25] J. Wess, *Hyperfine Interact.*, 44:3, 1988.

[26] G. Gabrielse, *Hyperfine Interact.*, 44:349, 1988.

[27] Robert S. Van Dyck, Jr., Paul B. Schwinberg, and Hans G. Dehmelt, *Phys. Rev. Lett.*, 59:26, 1987.

[28] R. Carosi *et al*, *Phys. Lett. B*, 237:303, 1990.

[29] L. K. Gibbons *et al*, *Phys. Rev. Lett.*, 70:1199, 1993.

[30] M. Niering, R. Holzwarth, J. Reichert, P. Pokasov, Th. Udem,, M. Weitz, T. W. Hansch, P. Lemonde, G. Santarelli, M. Abgrall, P. Laurent, C. Salomon, A. Clairon, *Phys. Rev. Lett.*, 84:5496, 2000.

[31] Robert Bluhm, V. Alan Kostelecky, and Neil Russell, *Phys. Rev. Lett.*, 82:2254, 1999.

[32] G. Baur, G. Boero, S. Brauksiepe, A. Buzzo, W. Eyrich, R. Geyer, D. Grzonka, J. Hauffe, K. Kilian, M. Lo Vetereb, M. Macrib, M. Moosburgerc, R. Nellen, W. Oelert, S. Passaggio, A. Pozzo, K. Rhrich, K. Sachs, G. Schepers, T. Sefzick, R. S. Simon, R. Stratmann, F. Stinzing, and M. Wolke, *Phys. Lett. B*, 368:251, 1996.

[33] G. Blanford, D. C. Christian, K. Gollwitzer, M. Mandelkern, C. T. Munger, J. Schultz, and G. Zioulas, *Phys. Rev. Lett.*, 80:3037, 1998.

[34] G. Gabrielse, S. L. Rolston, L. Haarsma, and W. Kells, *Phys. Lett. A*, 129:38, 1988.

[35] M. Glinsky and T. O'Neil, *Phys. Fluids B*, 3:1279, 1991.

[36] J. Stevefelt, J. Boulmer, and J. F. Delpech, *Phys. Rev. A*, 12:1246, 1975.

[37] P. O. Fedichev, *Phys. Lett. A*, 226:289, 1997.

[38] R. Neumann, H. Poth, A. Winnacker, and A. Wolf *Z. Phys. A*, 313:253, 1983.

[39] J. W. Humberston, M. Charlton, F. J. Jacobsen, and B. I. Deutch, *J. Phys. B*, 20:L25, 1987.

[40] E. A. Hessels, D. M. Homan, and M. J. Cavagnero, *Physical Review A*, 57:1668, 1998.

[41] C. Wesdorp, F. Robicheaux, and L. D. Noordam, *Phys. Rev. Lett.*, 84:3799, 2000.

[42] U. Schramm, J. Berger, M. Grieser, D. Habs, E. Jaeschke, G. Kilgus, D. Schwalm, and A. Wolf, *Phys. Rev. Lett.*, 67:22, 1991.

[43] F. B. Yousif, P. Van der Donk, Z. Kucherovsky, J. Reis, E. Brannen, and J. B. A. Mitchell, *Phys. Rev. Lett.*, 67:26, 1991.

[44] C. Heerlein, G. Zwicknagel, and C. Toepffer, *Phys. Rev. Lett.*, 89:083202-1, 2002.

[45] T. C. Killian, M. J. Lim, S. Kulin, R. Dumke, S. D. Bergeson, and S. L. Rolston, *Phys. Rev. Lett.*, 86:17, 2001.

[46] B. Zygelman, *J. Phys. B*, 36:L31, 2003.

[47] Hans A. Bethe and Edwin E. Salpeter. *Quantum Mechanics of One and Two-Electron Atoms.* Springer Verlag, 1957.

[48] J. J. Bollinger and D. J. Wineland, *Phys. Rev. Lett.*, 53:348, 1984.

[49] Joseph N. Tan, J. J. Bollinger, B. Jelenkovic, and D. J. Wineland, *Phys. Rev. Lett.*, 75:4198, 1995.

[50] J. M. Kriesel and C. F. Driscoll *Phys. Rev. Lett.*, 87:135003-1, 2001.

[51] C. F. Driscoll, F. Anderegg, D. H. E. Dubin, D. Z. Jin, J. M. Kriesel, E. M. Hollmann, and T. M. ONeil, *Phys. Plasmas*, 9:1905, 2002.

[52] E. M. Hollman, F. Anderegg, and C. F. Driscoll, *Phys. Plasmas*, 7:1767, 2000.

[53] Lowell S. Brown and Gerald Gabrielse, *Rev. Mod. Phys.*, 58:233, 1986.

[54] G. Gabrielse, L. Haarsma, and S. L. Rolston *Intl. J. of Mass Spec.*, 88:319, 1989.

[55] G. Gabrielse, D. S. Hall, T. Roach, P. Yesley, A. Khabbaz, J. Estrada, C. Heimann, and H. Kalinowsky, *Phys. Lett. B*, 455:311, 1999.

[56] P. Yesley, Ph.D. thesis 2002.

[57] J. K. Estrada, Ph.D. thesis 2002.

[58] Jim MacArthur, Harvard University Physics Dept. Electronics Engineer.

[59] D. S. Hall, Ph.D. thesis 1996.

[60] X. Fei, Ph.D. thesis 1990.

[61] P. B. Schwinberg, R. S. Van Dyck, Jr., and H. G. Dehmelt, *Phys. Lett.*, 81A:119, 1981.

[62] P. B. Schwinberg, R. S. Van Dyck, Jr., and H. G. Dehmelt, *Phys. Rev. Lett.*, 47:1679, 1981.

[63] R. S. Van Dyck, Jr., P. B. Schwinberg, and H. G. Dehmelt, *Phys. Rev. Lett.*, 59:26, 1987.

[64] F. J. Wysocki, M. Leventhal, A. Passner, and C. M. Surko, *Hyperfine Interact.* 44:185, 1988.

[65] T. J. Murphy and C. M. Surko, *Phys. Rev. A.* 46:5696, 1992.

[66] R. G. Greaves and C. M. Surko, *Nucl. Inst. and Methods B.* 192:90, 2002.

[67] L. Haarsma, K. Abdullah, and G. Gabrielse, *Hyperfine Interact.* 76:143, 1993.

[68] G. Gabrielse, L. Haarsma, and K. Abdullah, *Hyperfine Interact.* 89:371, 1994.

[69] L. Haarsma, K. Abdullah, and G. Gabrielse *Phys. Rev. Lett.* 75:806, 1995.

[70] K. Abdullah, L. Haarsma, and G. Gabrielse, *Physica Scripta.* T59:337, 1995.

[71] L. V. Jrgensen, M. J. T. Collier, K. S. Fine, T. L. Watson, D. P. van der Werf and M. Charlton, *Mat. Sci. Forum.* 634:363, 2001.

[72] Loren Haarsma, Ph.D. thesis 1994.

[73] K. G. Lynn, B. Nielsen, and J. H. Quateman, *Appl. Phys. Lett.*, 47:239, 1985.

[74] S. L. Rolston and G. Gabrielse, *Hyperfine Interact.*, 44:233, 1988.

[75] D. Mohl, G. Petrucci, L. Thorndahl, and S. van der Meer, *Phys. Rept,* 58:76:1980.

[76] Herbert Stelzer, *Nucl. Inst. and Methods*, 133:409, 1976.

[77] *www.srim.org.*

[78] R. L. Spencer, S. N. Rasband, and R. R. Vanfleet, *Physics of Fluids B*, 5:4267, 1993.

[79] H. E. Dubin and T. M. O'Neil, *Reviews of Modern Physics*, 71:87, 1999.

[80] L. R. Brewer, J. D. Prestage, J. J. Bollinger, W. M. Itano, D. J. Larson, and D. J. Wineland, *Phys. Rev. A*, 38:859, 1988.

[81] J. J. Bollinger, D. J. Heinzen, F. L. Moore, W. M. Itano, and D.J. Wineland, *Phys. Rev. A*, 48:525, 1993.

[82] F. Anderegg, X. P. Huang, C. F. Driscoll, E. M. Hollmann, T. M. ONeil, and D. H. E. Dubin, *Phys. Rev. Lett.*, 78:2128, 1997.

[83] R. Parrott, *Ph.D. thesis, to be published.*

[84] S. A Prasad and T. M. O'Neil, *Phys. Fluids*, 22:278, 1978.

[85] R. L. Spencer, *Private comminication.*

[86] Grant W. Mason, Ross L. Spencer, and Jonathan A. Bennett, *Phys. Plasmas*, 3:1502, 1996.

[87] Deborah L. Paulson and Ross L. Spencer, *Phys. Plasmas*, 5:346, 1998.

[88] J. B. Delos, S. K. Knudson, and D. W. Noid, *Phys. Rev. A*, 30:1208, 1984.

[89] P. Schmelcher and L. S. Cederbaum, *Chem. Phys. Lett.*, 208:548, 1993.

[90] O. Dippel, P. Schmelcher and L. S. Cederbaum, *Phys. Rev. A*, 49:4415, 1994.

[91] Peter Schmelcher and Lorenz S. Cederbaum, *Comm. At. Mol. Phys.*, 2, 1999.

[92] T. M. Squires, P. Yesley, and G. Gabrielse, *Phys. Rev. Lett.*, 86:5266, 2001.

[93] K. S. E. Eikema, J. Walz, and T. W. Hänsch, *Phys. Rev. Lett.*, 83:3828, 1999.