



CERN-THESIS-2011-299

UNIVERSITY OF CALGARY

Detection of Trapped Antihydrogen

by

Richard Allan Hydomako

A DISSERTATION

SUBMITTED TO THE FACULTY OF GRADUATE STUDIES
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE
DEGREE OF DOCTORATE OF PHILOSOPHY

DEPARTMENT OF PHYSICS AND ASTRONOMY

CALGARY, ALBERTA

September, 2011

© Richard Allan Hydomako 2011

Abstract

The ALPHA experiment is an international effort to produce, trap, and perform precision spectroscopic measurements on antihydrogen (the bound state of a positron and an antiproton). Based at the Antiproton Decelerator (AD) facility at CERN, the ALPHA experiment has recently magnetically confined antihydrogen atoms for the first time. A crucial element in the observation of trapped antihydrogen is ALPHA's silicon vertexing detector. This detector contains sixty silicon modules arranged in three concentric layers, and is able to determine the three-dimensional location of the annihilation of an antihydrogen atom by reconstructing the trajectories of the produced annihilation products.

This dissertation focuses mainly on the methods used to reconstruct the annihilation location. Specifically, the software algorithms used to identify and extrapolate charged particle tracks are presented along with the routines used to estimate the annihilation location from the convergence of the identified tracks. It is shown that these methods can determine the annihilation location with a spatial resolution between about 0.6 to 0.8 cm (depending on the coordinate being measured). Furthermore, a robust analysis to identify and reduce cosmic ray background events is described. The cosmic ray background can obscure the trapped antihydrogen signal, and its suppression leads to a significant increase in the annihilation detection sensitivity. The background suppression analysis involves examining the reconstructed detector event based on several selection criteria, including: the number of charged particle tracks, the radial vertex position, and a fit of a straight line to the event hit positions. By carefully optimizing these criteria, $(99.54 \pm 0.02)\%$ of cosmic ray events are rejected, while $(64.4 \pm 0.1)\%$ of antihydrogen annihilation events are retained. Finally, the experimental results demonstrating the first-ever magnetic confinement of antihydrogen atoms are presented. These results rely heavily on the silicon

detector, and as such, the role of the annihilation vertex reconstruction is emphasized.

Acknowledgments

First off, I'd like to acknowledge my colleagues on the ALPHA experiment. Working on this experiment has shaped me as a scientist, and I've benefited enormously from the knowledge and support of everyone involved. To the senior scientists, including: Dr. Paul Bowe, Prof. Joel Fajans, Dr. Dave Gill, Prof. Jeff Hangst, Prof. Niels Madsen, Prof. Scott Menary, and Prof. Art Olin, I've learned a tremendous amount from all of you. To all the students and post-docs: Dr. Gorm Andresen, Mohammad Ashkezari, Marcelo Baquero-Ruiz, Dr. Eoin Butler, Dr. Will Bertsche, Steve Chapman, Tim Friesen, Andrea Gutierrez, Andrew Humphries, Alex Povilus, Dr. Petteri Pusa, Dr. Matt Jenkins, Prof. Daniel de Miranda Silveira, Dr. James Storey, Sarah Seif El Nasr-Storey, and Chukman So – I could not have asked for a better group. I have fond memories with all of you.

Particular thanks goes to my supervisors: Prof. Robert Thompson and Prof. Makoto Fujiwara. Both of you have inspired and pushed me with your drive, enthusiasm, and knowledge. I've received phenomenal opportunities under your tutelage, all of which have helped me grow as an academic and as a person.

I've enjoyed my time in the Department of Physics and Astronomy at the University of Calgary. The students, staff, and faculty have all been wonderful. I have made many friends and we have shared some great times in Calgary. I thank you all.

This work was supported by the Alberta Ingenuity Fund, the Natural Sciences and Engineering Research Council of Canada (NSERC), and the Killam Trust. I am grateful to all of these funding bodies and scholarship organizations, whose support allowed me to focus fully on this research.

I must thank my family. Everyone has provided support and encouragement, which has meant a great deal. It saddens me that my father, William Hydomako, and grandfather, Lloyd McDonald, could not be here to share this with me. Their memory and

influence remains, and has helped me through this whole endeavor.

Last, but not least, to my wife Kelly, without whom this could not have happened. Nobody believed in me more, or helped me as much as you have. I dedicate this to you.

Publications and presentations

The following list contains the ALPHA papers published (and upcoming) during the author's Ph. D. program. These works are the result of a collaborative effort, and following ALPHA collaboration policy, the authors lists are in alphabetic order.

1. **“Antihydrogen Annihilation Reconstruction with the ALPHA Silicon Detector”**¹

G. B. Andresen, M. D. Ashkezari, W. Bertsche, P. D. Bowe, E. Butler, C. L. Cesar, S. Chapman, M. Charlton, A. Deller, S. Eriksson, J. Fajans, T. Friesen, M. C. Fujiwara, D. R. Gill, A. Gutierrez, J. S. Hangst, W. N. Hardy, M. E. Hayden, R. S. Hayano, A. J. Humphries, R. Hydomako, S. Jonsell, L. V. Jørgensen, L. Kurchaninov, N. Madsen, S. Menary, P. Nolan, K. Olchanski, A. Olin, A. Povilus, P. Pusa, E. Sarid, S. Seif el Nasr, D.M. Silveira, C. So, J. W. Storey, R. I. Thompson, D. P. van der Werf, and Y. Yamazaki, *Submitted to Nucl. Instrum. Meth. A* (2011).

2. **“Confinement of antihydrogen for 1,000 seconds”**

G. B. Andresen, M. D. Ashkezari, M. Baquero-Ruiz, W. Bertsche, P. D. Bowe, E. Butler, C. L. Cesar, M. Charlton, A. Deller, S. Eriksson, J. Fajans, T. Friesen, M. C. Fujiwara, D. R. Gill, A. Gutierrez, J. S. Hangst, W. N. Hardy, R. S. Hayano¹, M. E. Hayden, A. J. Humphries, R. Hydomako, S. Jonsell, S. L. Kemp, L. Kurchaninov, N. Madsen, S. Menary, P. Nolan, K. Olchanski, A. Olin, P. Pusa, C. Ø. Rasmussen, F. Robicheaux, E. Sarid, D. M. Silveira, C. So, J. W. Storey, R. I. Thompson, D. P. van der Werf, J. S. Wurtele, and Y. Yamazaki, *Nat. Phys.* **7** 558 (2011). doi:10.1038/nphys2025

3. **“Centrifugal Separation and Equilibration Dynamics in an Electron-Antiproton Plasma”**

G. B. Andresen, M. D. Ashkezari, M. Baquero-Ruiz, W. Bertsche, P. D. Bowe, E. Butler, C. L. Cesar, S. Chapman, M. Charlton, A. Deller, S. Eriksson, J. Fajans, T. Friesen, M. C. Fujiwara, D. R. Gill, A. Gutierrez, J. S. Hangst, W. N. Hardy, M. E. Hayden, A. J. Humphries, R. Hydomako, S. Jonsell, N. Madsen, S. Menary, P. Nolan, A. Olin, A. Povilus, P. Pusa, F. Robicheaux, E. Sarid, D. M. Silveira, C. So, J. W. Storey, R. I. Thompson, D. P. van der Werf, J. S. Wurtele, and Y. Yamazaki, *Phys. Rev. Lett.* **106**, 145001 (2011). doi:10.1103/PhysRevLett.106.145001

4. **“Autoresonant Excitation of Antiproton Plasmas”**

G. B. Andresen, M. D. Ashkezari, M. Baquero-Ruiz, W. Bertsche, P. D. Bowe, E. Butler, P. T. Carpenter, C. L. Cesar, S. Chapman, M. Charlton, J. Fajans, T. Friesen, M. C. Fujiwara, D. R. Gill, J. S. Hangst, W. N. Hardy, M. E. Hayden, A. J. Humphries, J. L. Hurt, R. Hydomako, S. Jonsell, N. Madsen, S. Menary, P. Nolan, K. Olchanski, A. Olin, A. Povilus, P. Pusa, F. Robicheaux, E. Sarid, D. M. Silveira,

¹The author took the lead role in the preparation and writing of this article.

C. So, J. W. Storey, R. I. Thompson, D. P. van der Werf, J. S. Wurtele, and Y. Yamazaki, *Phys. Rev. Lett.* **106**, 025002 (2011). doi:10.1103/PhysRevLett.106.025002

5. **“Search for trapped antihydrogen”**

G. B. Andresen, M. D. Ashkezari, M. Baquero-Ruiz, W. Bertsche, P. D. Bowe, C. C. Bray, E. Butler, C. L. Cesar, S. Chapman, M. Charlton, J. Fajans, T. Friesen, M. C. Fujiwara, D. R. Gill, J. S. Hangst, W. N. Hardy, R. S. Hayano, M. E. Hayden, A. J. Humphries, R. Hydomako, S. Jonsell, L. V. Jørgensen, L. Kurchaninov, R. Lambo, N. Madsen, S. Menary, P. Nolan, K. Olchanski, A. Olin, A. Povilus, P. Pusa, F. Robicheaux, E. Sarid, S. Seif El Nasr, D. M. Silveira, C. So, J. W. Storey, R. I. Thompson, D. P. van der Werf, D. Wilding, J. S. Wurtele, and Y. Yamazaki, *Phys. Lett. B* **695**, 95 (2010). doi:10.1016/j.physletb.2010.11.004

6. **“Trapped antihydrogen”**

G. B. Andresen, M. D. Ashkezari, M. Baquero-Ruiz, W. Bertsche, P. D. Bowe, E. Butler, C. L. Cesar, S. Chapman, M. Charlton, A. Deller, S. Eriksson, J. Fajans, T. Friesen, M. C. Fujiwara, D. R. Gill, A. Gutierrez, J. S. Hangst, W. N. Hardy, M. E. Hayden, A. J. Humphries, R. Hydomako, M. J. Jenkins, S. Jonsell, L. V. Jørgensen, L. Kurchaninov, N. Madsen, S. Menary, P. Nolan, K. Olchanski, A. Olin, A. Povilus, P. Pusa, F. Robicheaux, E. Sarid, S. Seif el Nasr, D. M. Silveira, C. So, J. W. Storey, R. I. Thompson, D. P. van der Werf, J. S. Wurtele and Y. Yamazaki, *Nature* **468**, 673 (2010). doi:10.1038/nature09610

7. **“Evaporative Cooling of Antiprotons to Cryogenic Temperatures”**

G. B. Andresen, M. D. Ashkezari, M. Baquero-Ruiz, W. Bertsche, P. D. Bowe, E. Butler, C. L. Cesar, S. Chapman, M. Charlton, J. Fajans, T. Friesen, M. C. Fujiwara, D. R. Gill, J. S. Hangst, W. N. Hardy, R. S. Hayano, M. E. Hayden, A. Humphries, R. Hydomako, S. Jonsell, L. Kurchaninov, R. Lambo, N. Madsen, S. Menary, P. Nolan, K. Olchanski, A. Olin, A. Povilus, P. Pusa, F. Robicheaux, E. Sarid, D. M. Silveira, C. So, J. W. Storey, R. I. Thompson, D. P. van der Werf, D. Wilding, J. S. Wurtele, and Y. Yamazaki, *Phys. Rev. Lett.* **105**, 013003 (2010). doi:10.1103/PhysRevLett.105.013003

8. **“Antihydrogen formation dynamics in a multipolar neutral anti-atom trap”**

G. B. Andresen, W. Bertsche, P. D. Bowe, C. Bray, E. Butler, C. L. Cesar, S. Chapman, M. Charlton, J. Fajans, M. C. Fujiwara, D. R. Gill, J. S. Hangst, W. N. Hardy, R. S. Hayano, M. E. Hayden, A. J. Humphries, R. Hydomako, L. V. Jørgensen, S. J. Kerrigan, L. Kurchaninov, R. Lambo, N. Madsen, P. Nolan, K. Olchanski, A. Olin, A. Povilus, P. Pusa, F. Robicheaux, E. Sarid, S. Seif El Nasr, D. M. Silveira, J. W. Storey, R. I. Thompson, D. P. van der Werf, J. S. Wurtele and Y. Yamazaki, *Phys. Lett. B* **685**, 141 (2010). doi:10.1016/j.physletb.2010.01.066

9. **“Antiproton, positron, and electron imaging with a microchannel plate / phosphor detector”**
G. B. Andresen, W. Bertsche, P. D. Bowe, C. C. Bray, E. Butler, C. L. Cesar, S. Chapman, M. Charlton, J. Fajans, M. C. Fujiwara, D. R. Gill, J. S. Hangst, W. N. Hardy, R. S. Hayano, M. E. Hayden, A. J. Humphries, R. Hydomako, L. V. Jørgensen, S. J. Kerrigan, L. Kurchaninov, R. Lambo, N. Madsen, P. Nolan, K. Olchanski, A. Olin, A. P. Povilus, P. Pusa, E. Sarid, S. Seif El Nasr, D. M. Silveira, J. W. Storey, R. I. Thompson, D. P. van der Werf, and Y. Yamazaki (ALPHA Collaboration), *Rev. Sci. Inst.* **80**, 123701 (2009). doi:10.1063/1.3266967
10. **“Magnetic multipole induced zero-rotation frequency bounce-resonant loss in a Penning-Malmberg trap used for antihydrogen trapping”**
G. B. Andresen, W. Bertsche, C. C. Bray, E. Butler, C. L. Cesar, S. Chapman, M. Charlton, J. Fajans, M. C. Fujiwara, D. R. Gill, W. N. Hardy, R. S. Hayano, M. E. Hayden, A. J. Humphries, R. Hydomako, L. V. Jørgensen, S. J. Kerrigan, J. Keller, L. Kurchaninov, R. Lambo, N. Madsen, P. Nolan, K. Olchanski, A. Olin, A. Povilus, P. Pusa, F. Robicheaux, E. Sarid, S. Seif El Nasr, D. M. Silveira, J. W. Storey, R. I. Thompson, D. P. van der Werf, J. S. Wurtele, and Y. Yamazaki, *Phys. Plasmas* **16**, 100702 (2009). doi:10.1063/1.3258840
11. **“Compression of Antiproton Clouds for Antihydrogen Trapping”**
G. B. Andresen, W. Bertsche, P. D. Bowe, C. C. Bray, E. Butler, C. L. Cesar, S. Chapman, M. Charlton, J. Fajans, M. C. Fujiwara, R. Funakoshi, D. R. Gill, J. S. Hangst, W. N. Hardy, R. S. Hayano, M. E. Hayden, R. Hydomako, M. J. Jenkins, L. V. Jørgensen, L. Kurchaninov, R. Lambo, N. Madsen, P. Nolan, K. Olchanski, A. Olin, A. Povilus, P. Pusa, F. Robicheaux, E. Sarid, S. Seif El Nasr, D. M. Silveira, J. W. Storey, R. I. Thompson, D. P. van der Werf, J. S. Wurtele, and Y. Yamazaki, *Phys. Rev. Lett.* **100**, 203401 (2008). doi:10.1103/PhysRevLett.100.203401
12. **“A Novel Antiproton Radial Diagnostic Based on Octupole Induced Ballistic Loss”**
G. B. Andresen, W. Bertsche, P. D. Bowe, C. C. Bray, E. Butler, C. L. Cesar, S. Chapman, M. Charlton, J. Fajans, M. C. Fujiwara, R. Funakoshi, D. R. Gill, J. S. Hangst, W. N. Hardy, R. S. Hayano, M. E. Hayden, A. J. Humphries, R. Hydomako, M. J. Jenkins, L. V. Jørgensen, L. Kurchaninov, R. Lambo, N. Madsen, P. Nolan, K. Olchanski, A. Olin, R. D. Page, A. Povilus, P. Pusa, F. Robicheaux, E. Sarid, S. Seif El Nasr, D. M. Silveira, J. W. Storey, R. I. Thompson, D. P. van der Werf, J. S. Wurtele, and Y. Yamazaki, *Phys. Plasmas* **15**, 032107 (2008). doi:10.1063/1.2899305
13. **“Production of Antihydrogen at Reduced Magnetic Field for Anti-atom Trapping”**
G. B. Andresen, W. Bertsche, A. Boston, P. D. Bowe, C. L. Cesar, S. Chapman, M. Charlton, M. Chartier, A. Deutsch, J. Fajans, M. C. Fujiwara, R. Funakoshi, D. R. Gill, K. Gomberoff, J. S. Hangst, R. S. Hayano, R. Hydomako, M. J. Jenkins,

L.V. Jørgensen, L. Kurchaninov, N. Madsen, P. Nolan, K. Olchanski, A. Olin, A. Povilus, F. Robicheaux, E. Sarid, D. M. Silveira, J.W. Storey, R. I. Thompson, D. P. van der Werf, J. S. Wurtele, and Y. Yamazaki, *J. Phys. B: At. Mol. Opt. Phys.* **41**, 011001 (2008). doi:10.1088/0953-4075/41/1/011001

The following conference proceedings were prepared and submitted by the author during his Ph. D. program:

1. **“Antihydrogen Detection in ALPHA”**

R. Hydomako, G. B. Andresen, M. D. Ashkezari, M. Baquero-Ruiz, W. Bertsche, E. Butler, P. D. Bowe, C. L. Cesar, S. Chapman, M. Charlton, J. Fajans, T. Friesen, M. C. Fujiwara, D. R. Gill, J. S. Hangst, W. N. Hardy, R. S. Hayano, M. E. Hayden, A. J. Humphries, S. Jonsell, L. Kurchaninov, N. Madsen, S. Menary, P. Nolan, K. Olchanski, A. Olin, A. Povilus, P. Pusa, F. Robicheaux, E. Sarid, D. M. Silveira, C. So, J. W. Storey, R. I. Thompson, D. P. van der Werf, J. S. Wurtele, and Y. Yamazaki, *Submitted to the LEAP2011 conference proceedings* (2011).

2. **“Search for Trapped Antihydrogen: First Candidate Events”**

R. Hydomako, G. B. Andresen, M. D. Ashkezari, M. Baquero-Ruiz, W. Bertsche, P. D. Bowe, C. C. Bray, E. Butler, C.L. Cesar, S. Chapman, M. Charlton, J. Fajans, T. Friesen, M. C. Fujiwara, D. R. Gill, J. S. Hangst, W. N. Hardy, R. S. Hayano, M. E. Hayden, A. J. Humphries, S. Jonsell, L. V Jørgensen, L. Kurchaninov, R. Lambo, N. Madsen, S. Menary, P. Nolan, K. Olchanski, A. Olin, A. Povilus, P. Pusa, F. Robicheaux, E. Sarid, S. Seif El Nasr-Storey, D. M. Silveira, C. So, J. W. Storey, R. I. Thompson, D. P. van der Werf, D. Wilding, J. S. Wurtele, Y. Yamazaki, *PoS(ICHEP 2010)* 479 (2010).

The author delivered the following presentations throughout his Ph. D. program:

1. **Antihydrogen Detection in ALPHA**

R. Hydomako, (invited talk), 10th International Conference on Low Energy Antiproton Physics, TRIUMF, Vancouver, Canada (2011).

2. **The ALPHA Experiment and the Trapping of Antihydrogen**

R. Hydomako, (invited talk), TRIUMF Users Group Annual General Meeting, TRIUMF, Vancouver, Canada (2010).

3. **Search for Trapped Antihydrogen: First Candidate Events**

R. Hydomako, (poster), International Conference for High Energy Physics, Paris, France (2010).

4. **“Status of the ALPHA Antihydrogen Trapping Analysis”**

R. Hydomako, (contributed talk), Winter Nuclear and Particle Physics Conference 2010, Banff, Canada (2010).

Table of Contents

Abstract	i
Acknowledgments	iii
Publications and presentations	v
Table of Contents	ix
List of Tables	xii
List of Figures	xiii
List of Symbols, Acronyms and Abbreviations	xxv
1 Introduction	1
1.1 Motivation for the trapping and study of antihydrogen	2
1.1.1 Precision comparison of the properties of hydrogen and antihydrogen	2
1.1.2 Gravitational interaction between matter and antimatter	3
1.2 Short review of antimatter studies and antihydrogen experiments	4
1.3 Dissertation overview	8
1.4 Contributions of the author	9
2 Theory	11
2.1 Formation of antihydrogen with trapped plasmas	11
2.1.1 Three-body recombination	15
2.2 Charged (anti)particle and neutral-atom trapping	16
2.2.1 Penning-Malmberg trap for charged particles	17
2.2.1.1 Single particle confinement	18
2.2.1.2 Confinement of charged plasmas	21
2.2.2 Neutral-atom trapping	24
2.3 Particle-antiparticle annihilation	28
2.3.1 Electron-positron annihilation	28
2.3.2 Antiproton annihilation	30
2.4 Semiconductor detectors	32
2.4.1 Double-sided silicon microstrip detectors	32
2.5 Summary	33
3 The ALPHA apparatus & procedures	35
3.1 Infrastructure overview	37
3.2 Diagnostic and particle detection devices	39
3.2.1 Charged particle temperature measurement	42
3.3 Production and accumulation of antiparticles	44
3.3.1 Antiproton Decelerator (AD)	44
3.3.2 Positron accumulator	47
3.4 Charged particle manipulation and trapping	50
3.4.1 Penning-Malmberg trap	50
3.4.2 Catching and cooling of antiprotons	51
3.4.3 Rotating wall compression and electron kick-out	57
3.4.4 Evaporative cooling	64
3.4.5 Autoresonant mixing	68

3.5	Magnetic neutral-atom trap	71
3.5.1	Axial magnetic field	71
3.5.2	Radial magnetic field	73
3.5.3	Fast shutdown	75
3.6	Trapping experiment overview	76
3.7	Summary	78
4	The ALPHA silicon detector	81
4.1	Detector hardware and software for control and data collection	83
4.1.1	Silicon detector modules	84
4.1.2	VA1TA ASIC readout chips	86
4.1.3	Detector configuration	87
4.1.4	Readout chain	88
4.1.5	TA and readout triggers	90
4.1.6	Analog readout	95
4.1.6.1	Separating signal strips from noise	95
4.1.6.2	Hit positioning and strip clustering	100
4.1.6.3	Cosmic rays and hit efficiency	103
4.1.7	Summary of detector hardware	106
4.2	Monte Carlo simulation	107
4.2.1	Overview	107
4.2.2	Materials and geometry	108
4.2.3	Event generation	109
4.2.4	Digitization	113
4.3	Summary	119
5	Event reconstruction in the ALPHA detector	120
5.1	Challenges for the ALPHA event reconstruction	120
5.2	Overview of the event reconstruction	124
5.3	Enumeration and filtering of track candidates	125
5.4	Determining the helix parameters	133
5.4.1	Helix parametrization	134
5.4.2	Radial helix parameters	134
5.4.3	Axial helix parameters	138
5.5	Track pruning	138
5.6	Vertexing	143
5.6.1	Closest approach between two helices	143
5.6.2	Closest approach between $N_{\text{tracks}} > 2$ helices	144
5.6.3	Mean distance of closest approach to the vertex	146
5.6.4	Track exclusion	147
5.7	Vertex distributions	149
5.7.1	Annihilation on background gas	149
5.7.2	Octupole-induced antiproton annihilation	150
5.7.3	Antihydrogen formation in the neutral-atom trap field	153
5.7.4	Reconstructed vertex position resolution	155
5.7.4.1	Monte Carlo resolution estimate	156

5.7.4.2	Experimental resolution estimation	156
5.8	Summary	161
6	Rejection of background events	162
6.1	Backgrounds to the antihydrogen annihilation signal	162
6.1.1	Environment and hardware background	163
6.1.2	Mirror-trapped antiprotons	166
6.1.3	Cosmic Rays	167
6.2	Importance of background rejection	169
6.2.1	Blind analysis	169
6.3	Cosmic background rejection	170
6.3.1	Discriminating variables	170
6.3.1.1	Number of charged tracks, N_{tracks}	170
6.3.1.2	Combined linear fit residual, δ	171
6.3.1.3	Vertex radius, R	172
6.3.2	‘Training’ datasets	173
6.3.3	Cut placement and optimization	173
6.3.4	Results of the background rejection	176
6.4	Summary	178
7	The trapping of antihydrogen	179
7.1	Mirror-trapped antiprotons	180
7.2	Simulation of particle trajectories during magnetic trap shutdown	183
7.3	Control measurements	185
7.3.1	Neutral trap magnets disengaged	186
7.3.2	Without antiprotons and/or positrons	186
7.3.3	Electric potential biasing during fast magnetic shutdown	187
7.3.4	Heated positrons	189
7.4	Initial search for trapped antihydrogen (2009)	190
7.5	Observation of trapped antihydrogen (2010)	195
7.6	Summary	198
8	Future prospects and conclusion	199
8.1	Recent measurement: long-time confinement of antihydrogen	199
8.2	Upcoming measurement: microwave spectroscopy	201
8.3	Future measurement: new apparatus and laser spectroscopy	203
8.4	Conclusion	204
	Bibliography	205
A	Charged particles passing through matter	216
B	Particle detection using silicon p-n junctions	218
C	Particle propagation in GEANT3	223
D	Simulated particle momenta through the silicon detector	225

List of Tables

0.1	List of Acronyms and Abbreviations	xxv
0.2	List of Symbols	xxix
2.1	Particle oscillation frequencies for $B_0 = 1$ T, $V_0 = 10$ V, and $d = 1$ cm (similar to the typical operating conditions in the ALPHA trap).	21
2.2	Branching ratios of pion final states from antiproton-proton [106] and antiproton-neutron [109] annihilation at rest (the k denotes the grouping together of multiple π^0 channels). Note that there is a contribution of $\sim 2\%$ due to kaons, which are ignored here.	31
4.1	Descriptions of TTC NIM trigger outputs as defined for the ALPHA detector readout system.	91
4.2	Geometry element characteristics included in the Monte Carlo simulation.	109
5.1	List of all possible hit combinations in the example event of Figure 5.2. The combinations are referenced by the hit labels given in the above figure. If each hit is found in a different detector layer, a check mark is given under the ‘Different Layers?’ column. Likewise, if the hits clustered together in position, a check mark is given under the ‘Small Hit Separation?’ column. The lone combination that passes both tests (EFG) is highlighted in bold-face font.	127
5.2	Summary of the simulated and experimental reconstruction vertex position resolutions estimates.	158
6.1	Final parameter cut conditions. Events satisfying these conditions are classified as annihilation signal.	176
7.1	Summary table of the relevant experimental parameters for the 2009 trapping experiment events.	194
7.2	Summary table of the relevant experimental parameters for the 2010 trapping experiment events.	194
D.1	The number of tracks through the silicon detector, probable momenta (from Figs. D.2-D.8) for the various particle types.	227

List of Figures

2.1	Schematic illustration of the a) Penning, and b) Penning-Malmberg trap configurations. The Penning trap (a) consists of two hyperbolic endcaps and a hyperbolic ring electrode. A quadrupole electric field is established by biasing the endcaps to the ring electrode. A strong axial magnetic field provides radial confinement. The Penning-Malmberg trap (b) consists of a series of cylindrical electrodes. A strong axial magnetic field provides radial confinement, while axial confinement is provided by biasing the cylindrical electrodes.	19
2.2	X, Y, Z spatial illustration of single particle motion in a Penning-Malmberg trap. The motion is a superposition of three oscillatory modes (Eq. 2.22): the slow magnetron oscillation, the axial vibration, and the fast cyclotron orbit.	22
2.3	Radial variation in the magnetic magnitude for various multipole magnets (reproduced with permission from [99]). The curves are normalized to the field magnitude at the wall radius.	27
2.4	Lowest order Feynman diagrams for electron-positron annihilation for the (a) two- γ , and (b) three- γ channels.	28
2.5	A simplified diagram of a double-sided silicon microstrip detector. The p-side strips on the top side of the silicon wafer run orthogonally to the n-side strips on the opposite side of the silicon.	33
3.1	Schematic diagram of the ALPHA apparatus (Image from the ALPHA collaboration). Antiprotons delivered from the AD are depicted as the \bar{p} on the left. Several of the important components are identified and labeled. The length scale of the figure is shown at the bottom right. . . .	36
3.2	Flowchart of the interrelation between the major components of the ALPHA apparatus.	38
3.3	Cartoon schematic of the MCP/phosphor/CCD camera imaging system. The MCP is shown as the grey chevron, the phosphor screen is represented as the red rectangle, the mirror is the blue rectangle, and the CCD camera is the yellow rectangle. The red line is illustrative of electron multiplication in an MCP channel.	39
3.4	The number of antiprotons released from the confining potential as a function of the on-axis well depth (reproduced with permission from [64]). The antiprotons losses are integrated and corrected for the 25% scintillator detection efficiency. The black lines show the fitted region, and the temperatures found are: A: 1040, B: 325, C: 57, D: 23, E: 19, and F: 9 K. The inset shows a magnification for the low-energy distributions.	43
3.5	Antiproton momentum during over the course of a deceleration cycle in the AD, including the various stages of particle cooling. Figure adapted from [121].	45

3.6	Cartoon diagram of the AD system. The (simplified) beamline is shown as the black curves, the stochastic cooling transmission line is shown in grey, the electron cooler in blue, and the ALPHA experimental zone in red.	46
3.7	Overview of the positron accumulator (figure adapted from [103,127]). The accumulator apparatus is shown as a). The inset, b), shows a diagram of the positron buffer gas cooling within the long electrode stack, where the red curve shows the electric potential as a function of position, and the blue axis describes the pressure gradient.	48
3.8	The ALPHA superconducting solenoid magnet (Image from the ALPHA collaboration). A large magnetic field is produced along the bore of the magnet. The cryostat (Fig. 3.1) is located within the solenoid.	52
3.9	The electrode stack used to provide the electric potential as part of the principal ALPHA Penning-Malmberg trap (Image from the ALPHA collaboration). Using this series of electrodes, a variety of bias potentials can be established along the trap axis. The two narrow electrodes at either ends, between the white electrical insulating material, are used to provide the high-voltage (3-5 kV) potentials during antiproton catching (Sec. 3.4.2).	53
3.10	Cartoon illustration of the antiproton catching sequence. a) Antiprotons are delivered from the AD (left) into the catching trap, where electrons are already pre-loaded into a cooling well and the far HV-electrode is already engaged; b) antiprotons with kinetic energies < 3 keV rebound off the HV potential (higher energy antiprotons are lost); c) the second HV barrier is quickly engaged, confining the antiprotons between the two barriers; d) the antiprotons cool, via Coulomb collisions with the electrons, into the low-voltage electron well and the HV barriers are disengaged.	55
3.11	Antiproton catching efficiency as a function of solenoid magnetic field (reproduced with permission from [49]). The values have been normalized to the 3 T value.	58
3.12	The on-axis z component of the magnetic field with (red) only the external solenoid energized, and (blue) both the external and inner solenoids energized.	58
3.13	Typical setup for the Rotating Wall technique. An sinusoidal voltage is provided by a waveform generator. This voltage is then split, and the offshoots phase-shifted, before finally being applied to a special segmented electrode.	59
3.14	Electron and antiproton compression using the Rotating Wall technique (reproduced with permission from [66]). The images of antiproton and electron radial distributions are generated with the MCP system and are provided along with integrated radial particle density profiles (the location of the MCP is shown in Fig. 3.1). Various Rotating Wall application times, as indicated by the figure, are also shown.	60

3.15	Cartoon illustration of the electron kick-out sequence. a) A mixture of electrons and antiprotons are confined in a shallow well (see Sec. 3.4.2); b) the confining potentials are raised above ground; c) one of the confining potentials is lowered, allowing the electrons to quickly escape (being more massive, the antiprotons do not escape as quickly); d) the confining potential is reestablished after the electrons are gone, but before the antiprotons can leave.	62
3.16	The electric potential used to confine antiprotons during evaporative cooling (reproduced with permission from [64]). The antiprotons are shown at the bottom of the potential well (Red). The well potentials shown are labeled by their on-axis depths.	65
3.17	Evaporative cooling of antiprotons (reproduced with permission from [64]). (a) shows the final measured antiproton temperature as a function of on-axis well depth, while (b) gives the corresponding survival fraction of antiprotons. The measured datapoints are shown as black squares, while the evaporative cooling model, Eqs. 3.16, is shown as the solid black curve.	66
3.18	Nested potential well used to confine both positrons and antiprotons. In a), the on-axis vacuum electric potential is shown, with the positrons confined in the central well, and the antiprotons in a side-well; b) shows a schematic of the electrode stack, with c) showing the applied voltages and the chirp signal generator for the autoresonant drive.	69
3.19	An energy-frequency plot of autoresonantly-driven antiprotons (reproduced with permission from [65]). The left figure a) shows the longitudinal energy distributions, $f(U)dU$, of $\sim 15\,000$ antiprotons driven to several different final frequencies. b) shows the calculated longitudinal energy as a function of drive frequency for the vacuum potential (solid blue), $15\,000$ antiprotons (dashed green), and $50\,000$ antiprotons (dot-dashed red). The open squares show the mean of the measured longitudinal energy distributions. Frequencies are normalized to $\omega_0/2\pi = 410$ kHz.	70
3.20	A schematic diagram of the magnetic neutral-atom trap and the surrounding apparatus (reproduced with permission from [54]). The mirror coils (cut away at the centre to allow visualization of the interior of the apparatus) are shown in green, and the octupole in red. The Penning-Malmberg trap electrodes are shown in yellow, and the surrounding silicon detector in light blue. (Not to scale.)	72
3.21	The on-axis magnetic field magnitude for the ALPHA neutral-atom trap. The total field magnitude is shown as the blue curve, while the 1 T contribution from the external Penning-Malmberg trap solenoid is shown as the dashed grey line. The fields of the mirror coils produce the large peaks at ± 14 cm.	72
3.22	Total magnetic field magnitude in the ALPHA apparatus as a function of radius at the axial center ($z = 0$ cm). The dotted line indicates the electrode wall radius.	74

3.23	Magnet current decay following the fast shutdown of the neutral trap (note the log scale). The current in the octupole magnetic is shown as in red, the left mirror in black, and right mirror in blue. The shaded region represents the expected time period where trapped antihydrogen atoms will be released.	75
3.24	Timeline of the entire trapping experiment. Important events and operations are labeled. (a) shows the entire experiment, while (b) gives a magnification of 660 - 662 s, showing the details of the particle mixing and final magnetic trap operations.	80
4.1	Cut-away rendering of the silicon detector in relation to the cryostat region of the ALPHA apparatus. Each outlined blue rectangle represents the location of a silicon module. Not shown are the readout chips or support structure. The structure of the apparatus has also been simplified for clarity.	82
4.2	Diagram of the silicon modules (not to scale). The silicon wafers and readout chips are highlighted. Not shown are the auxiliary readout electronics, and details of the bond wires and signal strips.	84
4.3	Photograph of two detector modules mounted on the detector support structure at the University of Liverpool clean room. (Image credit: University of Liverpool)	85
4.4	Overview of one of the 128 readout channels in the VA1TA chip, following [137]. Detailed circuitry has been suppressed to emphasize the VA/TA division.	86
4.5	Diagram illustrating the radial and azimuthal placement of the silicon modules. Modules 1-30 are located on the ‘AD’ half of the detector, while modules 31-60 are located on the ‘Positron’ half.	88
4.6	Schematic diagram of the detector readout chain. The connections have been simplified for clarity.	89
4.7	Logic diagram of $Si > 1$ TA trigger combination and detector readout trigger decision. Note that the VF48 busy signal is high when the VF48 are engaged in the detector readout, and the NOT logic inverts the busy signal.	92
4.8	Simplified diagram of the clock timing during the readout of the analog strip levels. The 60 MHz sample clock is shown in black, the 3.33 MHz multiplexer clock is shown in red, and the analog strip voltage is shown in blue. One full strip readout is shown, along with the shifting to the next strip.	93
4.9	Raw strip sample output from the VF48 modules for an (a) n-side ASIC, and (b) p-side ASIC. Strip samples are shown as the dark trace, while the mean ADC value for the entire ASIC is shown as the dashed red line.	96
4.10	$\langle \text{ADC} \rangle_{\text{strip}}$ values (from the same readout as Fig. 4.9), relative to the ASIC sample mean, for an (a) n-side ASIC and (b) p-side ASIC. Note that the vertical axis has been magnified to highlight the structure.	97

4.11	The strip samples after being mean corrected (from the same readout as Fig. 4.9) for an (a) n-side ASIC and a (b) p-side ASIC. The red line shows the mean ADC value.	98
4.12	Final pedestal-subtracted strip samples (from the same readout as Fig. 4.9) for an (a) n-side ASIC and a (b) p-side ASIC. The strip sample are show as the black trace, the 3.75σ threshold is shown as the grey trace, and the strips that pass the threshold are highlighted as the red trace.	99
4.13	Histogram of the ADC values for an (a) n-side ASIC and (b) p-side ASICs over the course of an experiment. The light red shows the samples that passed the 3.75σ threshold, while the grey shows the values that did not pass the threshold. Note the log scale on the vertical axis.	100
4.14	An example strip readout from part of the n-side of a detector module demonstrating charge sharing and strip clustering. Here a passing particle has resulted in signal in three consecutive strips. The red axis gives the position of the strips in the Z coordinate.	101
4.15	Spatial representation of an example event showing the combination of p- and n-side strips that produces three-dimensional set of position data for the hits. The active silicon is shown as the grey blocks, the strips as the red lines, and the hits as the blue markers.	102
4.16	Histogram of the angle (from vertical) of cosmic tracks passing through the detector. The smooth curve shows the $\cos^2\theta_v$ fit.	104
4.17	Plots of the p-side (black) and n-side (red) strip efficiencies per module, calculated by looking for where a comic track passed through, but did not result in a strip cluster about threshold. Note that the vertical axis starts at 96.5%. Data for Module 10 is missing, as it was disconnected from the data collection and analysis system (see footnote in Sec. 4.1.4).	106
4.18	Quadrant of the Monte Carlo radial geometry (to scale). Uncoloured volumes are treated as vacuum.	110
4.19	Length-wise cross-section of the Monte Carlo simulation geometry. Uncoloured volumes are treated as vacuum.	111
4.20	An event frequency histogram showing a comparison of the pion multiplicity resulting from the Monte Carlo event generator (black histogram) to the literature values (blue shaded squares) [106].	112
4.21	Generation process vs. particle type illustration of the various processes in GEANT3 [140], which generates particles during the Monte Carlo simulation. This data was collected by simulating 10000 proton-antiproton annihilations and recording the generating process for every simulated particle. Then mean number of particles per event per process is shown as colour and bin content.	114

4.22	Mean number of particles per event involved in the Monte Carlo simulation. This data was collected by simulating 10000 proton-antiproton annihilations, tracking every produced particle through its simulated trajectory, and determining the mean number of simulated particles which passed through each region. The yellow histogram shows all the generated and tracked particles. Also plotted are the mean number of particles per event which deposit energy in the 1st layer (blue), 1st and 2nd layers (green), and all three layers (brown). Bin contents are only shown for the yellow and brown histograms for clarity.	115
4.23	Cartoon diagram illustrating how the digital signal is produced by particles passing through a silicon detector module during the Monte Carlo simulation. The simple distribution of deposited energy to approximate charge sharing is shown by the grey shaded area.	116
4.24	Histograms of number of strips per cluster for the (a) n-side ASICs, and (b) a p-side ASICs. The same data is also shown on a log scale, and an expanded x-axis, for the (c) n-side ASICs, and (d) p-side ASICs.	117
4.25	Comparison of the number of hits measured during antiproton annihilation (red) to the number of hits found in the Monte Carlo simulation (black). On the left (a), the distribution is shown for 0-20 hits per event, while (b) extends the distribution to 120 hits per event.	118
5.1	Illustration of the effect of multiple scattering on the track extrapolation. The actual annihilation position is shown as the red star, whereas the extrapolated vertex is given as the blue star.	121
5.2	Spatial illustration of an example event, shown as an unsorted collection of hits. The hits are labelled sequentially (A-G), and in the Front view (left), the detector layers are shown in different line styles (dotted lines for layer 1, dashed lines for layer 2, and solid lines for layer 3). For brevity and clarity, only the modules with hits are shown in the side view (right).	127
5.3	Logarithmic histogram of the hit separation distribution along the z coordinate, $ z_{\text{layer } 1} - z_{\text{layer } 2} $ and $ z_{\text{layer } 2} - z_{\text{layer } 3} $. The unfiltered track candidates are shown along with tracks generated by the Monte Carlo simulation, as labelled in the figure. The range of accepted hit separations limits at the grey dashed line.	129
5.4	Logarithmic histogram of the total hit separation in the ϕ coordinate, $ \phi_{\text{layer } 1} - \phi_{\text{layer } 2} + \phi_{\text{layer } 2} - \phi_{\text{layer } 3} $. The unfiltered track candidates are shown along with tracks generated by the Monte Carlo simulation, as labelled in the figure. The range of accepted hit separations limits at the grey dashed line.	130
5.5	Logarithmic histogram of the hit correlation coefficient distribution (as described in Sec. 5.3). The unfiltered track candidates are shown along with tracks generated by the Monte Carlo simulation, as labelled in the figure. The range of accepted correlation coefficients limits at the grey dashed line.	131

5.6 An example of the ‘ghost hit’ ambiguity. Two particles passing through the same module (represented by the grey rectangle) leave signal in two n-side and two p-side strips (represented by the red lines in the modules). The intersections of n- and p-side strips gives four hit possibilities, where the red markers show the two spurious hits. 133

5.7 Illustration of the helix parameters in the radial projection. ρ_c gives the radius of curvature, D is the distance of closest approach to the origin, and ϕ_0 is the azimuthal angle at the point of closest approach. 135

5.8 Illustration of the helix parameters in the axial projection. z_0 is the z coordinate at the point of closest approach to the origin, and $\lambda = \cot \theta$ relates the azimuthal to the axial motion. 137

5.9 Logarithmic histogram of the $\tilde{\chi}^2$ distribution for: all unfiltered track candidates (blue), all simulated particle tracks (red), and only the simulated charged pion tracks (black). 139

5.10 Logarithmic histogram of the distances of closest approaches to the electrode wall radius distribution for: all unfiltered track candidates (blue), all simulated particle tracks (red), and only the simulated charged pion tracks (black). 141

5.11 Position space representation of the result of the track decision algorithm on an event data set. Grey curves represent track candidates which failed the distance of closest approach cut, green curves (obscured behind the other curves) represent track candidates that failed the χ^2 cut, and red curves are tracks candidates that share a hit with another track (and are subsequently rejected in favour of the other track). Finally, the blue curves pass all of the track selection criteria, and are accepted as charged-particle tracks. 142

5.12 Spatial illustration of the point of closest approach between two helical tracks. The blue curves show the track extrapolation, while the green line shows where the helices pass closest to each other. The green box shows the midpoint along the green line, which for a two-track event is taken as the vertex position. 144

5.13 Spatial illustration of the determination of the vertex position for an event with more than two tracks. The blue curves show the track extrapolations, and the green lines show where pairs of helices pass closest to each other (following Fig. 5.12). The black box shows the mean of points of closest approach, which is used as a seed position for the vertex minimization. 145

5.14 Spatial illustration of the determination of the mean distance of closest approach. The blue curves show the track extrapolations, while the red lines show the closest approach of the track extrapolations to the vertex position (shown as the black box). The mean length of the red lines, d_{vertex} , is used as a figure of merit in the final vertex determination. 146

5.15	Example reconstruction of an annihilation event, where (a) shows the $x-y$ (front) view, and (b) shows the $x-z$ (side) view. Here, the grey track has been considered, but ultimately excluded according to the algorithm described in Sec. 5.6.4. The blue cross shows the reconstructed vertex position.	149
5.16	Vertex distribution for annihilations on background gas; (a) shows the $x-y$ projection of the vertex distribution, with the colour palette representing the number of vertices in each bin. The dashed white circle represents the position of the electrode walls. (b) shows the z distribution of vertices, where the unhatched region shows the extent of the electrodes providing the confining electric potential.	151
5.17	Vertex distribution for annihilations induced by the impinging octupole field lines on the confinement region; (a) shows the $x-y$ projection of the vertex distribution, with the colour palette representing the number of vertices in each bin. The dashed white circle represents the position of the electrode walls. (b) shows the z distribution of vertices, where the unhatched region shows the extent of the electrodes providing the confining electric potential.	152
5.18	$R-\phi$ projection of the octupole-induced vertex distribution (same distribution as shown in Fig. 5.17). The eight peaks correspond to the turning points in the octupole winding.	153
5.19	Vertex distribution during antihydrogen formation (with the neutral trap magnets engaged); (a) shows the $x-y$ projection of the vertex distribution, with the colour palette representing the number of vertices in each bin. The dashed white circle represents the position of the electrode walls. (b) shows the z distribution of vertices, where the unhatched region shows the extent of the electrodes providing the confining electric potential.	154
5.20	Linear histograms illustrating the distributions of differences between simulated and reconstructed vertex distributions for the (a) axial, (b) radial, and (c) azimuthal coordinates. The dashed line shows a double-Gaussian fit to the distributions.	157
5.21	Plots of vertex frequency versus (a) vertex axial position and (b) vertex radial position. This data enables an experimental determination of the reconstruction vertex resolutions for the axial and radial coordinates. The distributions are taken from annihilations on background gas (Sec. 5.7.1), where the antiproton cloud is assumed to be well localized. The dashed red curves show the fitting of a double-Gaussian function to the distributions.	159

6.1	Detector read-out data representing a small amount of detector activity triggering the full readout. Recalling the description of the readout trigger in Sec. 4.1.5, two or more p-side TA signals are required to trigger the analog readout. The blue crosses indicate hit positions, while the red lines show the signal strips. The left view (a) shows the $x - y$ projection of the event, while the right view (b) shows the $z - y$ projection. The black ring shows the electrode radius.	163
6.2	Detector read-out data representing a readout event with an overwhelming number of hits (possibly due to an electromagnetic particle shower within the apparatus or detector). The blue crosses indicate hit positions, while the red lines show the signal strips. The left view (a) shows the $x - y$ projection of the event, while the right view (b) shows the $z - y$ projection. Accurate reconstruction is unlikely, and events of this type are discarded. The black ring shows the electrode radius.	164
6.3	Detector read-out data representing an example of a background event generated by electrical noise associated with the fast shut-down of the neutral atom trap magnets. The blue crosses indicate hit positions, while the red lines show the signal strips. The left view (a) shows the $x - y$ projection of the event, while the right view (b) shows the $z - y$ projection. The black ring shows the electrode radius.	165
6.4	Detector read-out data representing an example of a cosmic event, shown in the a) $x - y$ projection, and b) $y - z$ projection. Two tracks are found (which are segments of a single cosmic track), and shown as the blue curves. The black diamond gives the reconstructed vertex position, following the method described in Chapter 5. The blue crosses indicate hit positions, while the red lines show the signal strips. The black ring shows the electrode radius.	168
6.5	Example detector events for an (a) annihilation (accepted signal), and a (b) cosmic ray (rejected background). The blue crosses show the hit positions, the blue curves show the reconstructed tracks, and the red diamonds show the reconstructed vertex position.	171
6.6	Contour plots for the expected signal significance as a function of the cuts on the vertex radius, R_{cut} , and combined linear residual, δ_{cut} , for a) events with $N_{\text{tracks}} = 2$ ($N_{\text{tracks}} > 2$ events held constant), and b) events with $N_{\text{tracks}} > 2$ ($N_{\text{tracks}} = 2$ events held constant). The final cut decisions are shown as the white crosses.	175

6.7	Linear histograms of the measured signal and background distributions for the discriminating variables. Shown are (a) the distribution of the number of charged tracks, N_{tracks} per event, (b) the radial component of the reconstructed vertex, R , and the combined linear residuals for the cases where (c) $N_{\text{tracks}} = 2$, and (d) $N_{\text{tracks}} > 2$. The annihilation signal sample is shown as the solid black trace, while the signal sample after the application of the cuts is shown as the solid grey filled trace. The cosmic background sample is shown as the solid red trace, and the red dashed trace shows the background sample after the application of the cuts. All distributions are normalized to the number of events in each sample set. The inset in (a) shows a magnification of the y-axis, highlighting the effect of the cuts on the background sample. Likewise, both background traces have been multiplied by a factor of 20 in (d) to appear for comparison. The green shadings represent the regions rejected by the cuts.	177
7.1	Electromagnetic potentials and fields versus position during the antiproton clearing pulses (reproduced with permission from [53]). Fig. (a) gives the on-axis electric potential (solid black line) and magnetic field magnitude (dashed red line) across the trap. The electric potential and magnetic field magnitude are combined via Eq. 7.1 to give a pseudopotential (b), for various transverse kinetic energies, E_{\perp} . Charged particles with large transverse kinetic energies can be confined by the magnetic field of the neutral trap.	182
7.2	A flow chart outlining the various control and cross-check measurements and how they correspond to apparatus operations during the trapping experiments.	185
7.3	Scatter plot of the reconstructed axial vertex coordinate, z , and time after the start of the fast magnet shutdown, t , for intentionally mirror-confined antiprotons. Simulated antiproton distributions are shown as the small dots (the simulated distributions have been convolved with the reconstruction resolution function to reflect the uncertainty in z). The measured no-bias events are shown as the upside-down triangles, the left-bias events as the green rightside-up triangle, and the right-bias events as the blue circles. The simulated no-bias events are shown as the red dots, the simulated left-bias events as the green dots, and the simulated right-bias events as the blue dots.	188
7.4	Diagram illustrating the intended detector $\text{Si} > 1$ trigger module radial and azimuthal locations for use during the 2009 experimental run. The red dashed modules were intended to be included in the primary detector trigger algorithm.	191
7.5	Diagram illustrating the actual detector $\text{Si} > 1$ trigger module radial and azimuthal locations for use during the 2009 experimental run. The red dashed modules were the actual modules included in the primary detector trigger algorithm.	191

7.6	The $z-t$ distributions for a) simulated antihydrogen atoms (grey dots) and b) simulated mirror trapped antiprotons (red dots). The simulated distributions have been convoluted with the reconstruction resolution function to reflect the uncertainty in z . The observed events (satisfying all cuts) are shown as the black triangles. The region enclosed by the grey rectangle is excluded by the time-cut.	193
7.7	The $z-t$ distributions for a) simulated antihydrogen atoms (grey dots), and b) simulated mirror trapped antiprotons (coloured dots). The simulated distributions have been convoluted with the reconstruction resolution function to reflect the uncertainty in z . The measured no-bias events are shown as the upside-down triangles, the left-bias events as the green right-side-up triangle, and the right-bias events as the blue circles. Events during heated-positron control experiments are shown as cyan(purple) for the no-bias(right-bias) cases. The simulated no-bias events are shown as the red dots, the simulated left-bias events as the green dots, and the simulated right-bias events as the blue dots. The region enclosed by the grey rectangle is excluded by the time-cut.	196
8.1	The long-time confinement of antihydrogen (reproduced with permission from [55]). Figure (a) shows the trapping rate as a function of confinement time. The errors bar give the counting uncertainty only. Figure (b) shows the statistical significance of the observed events with respect to the cosmic ray background. The point for 0.4 s is off the scale ($> 20\sigma$), and not shown.	200
8.2	The $z-t$ scatter-plot, showing the results of trapping experiments with various confinement times, as indicated in the upper right scale in the figure (reproduced with permission from [55]). The simulated antihydrogen distribution is shown as the grey dots.	201
8.3	Breit-Rabi diagram for ground state hyperfine levels for (anti)hydrogen. The states shown in red correspond to the low-field seeking spin states, which are anti-parallel to the magnetic field, while the blue states correspond to the high-field seeking spin states, which are parallel to the magnetic field.	202
B.1	Diagram of a p-n junction: (a) before thermal diffusion of mobile charge carriers (unbiased); (b) after thermal diffusion, with a net positive charge on the n-side and a net negative charge on the p-side, and a depletion region in between; (c) in the case of a reverse-biased junction with an enhanced depletion region and a passing charged particle generating an electron-hole pair.	220
D.1	Generation proces vs. particle type illustration for the simulated processes that resulted in tracks through all three layers of the detector for 10 000 simulated annihilations.	226
D.2	The (a) transverse momentum, and (b) absolute axial momentum for simulated π^- particles that passed through all three detector layers.	227

D.3	The (a) transverse momentum, and (b) absolute axial momentum for simulated π^+ particles that passed through all three detector layers.	228
D.4	The (a) transverse momentum, and (b) absolute axial momentum for simulated e^- particles that passed through all three detector layers.	228
D.5	The (a) transverse momentum, and (b) absolute axial momentum for simulated e^+ particles that passed through all three detector layers.	228
D.6	The (a) transverse momentum, and (b) absolute axial momentum for simulated μ^- particles that passed through all three detector layers.	229
D.7	The (a) transverse momentum, and (b) absolute axial momentum for simulated μ^+ particles that passed through all three detector layers.	229
D.8	The (a) transverse momentum, and (b) absolute axial momentum for simulated protons that passed through all three detector layers.	229

List of Symbols, Acronyms and Abbreviations

AEgIS	Antimatter Experiment: Gravity, Interferometry, Spectroscopy	Sec. 1.2	pg. 8
AD	Antiproton Decelerator	Sec. 3.3.1	pg. 44
ADC	Analog-to-Digital Converter	Sec. 4.1.4	pg. 90
ALPHA	Antihydrogen Laser PHysics Apparatus	Sec. 1	pg. 1
ASACUSA	Atomic Spectroscopy And Collisions Using Slow Antiprotons	Sec. 1.2	pg. 7
ASIC	Application Specific Integrated Circuit	Sec. 4.1.1	pg. 84
ATRAP	Antihydrogen TRAP	Sec. 1.2	pg. 6
CCD	Charge-Coupled Device	Sec. 3.2	pg. 40
C. L.	Confidence Limit	Sec. 2.1	pg. 13
CERN	European Organization for Nuclear Research	Sec. 3.3.1	pg. 44
FPGA	Field Programmable Gate Array	Sec. 3.6	pg. 76
FRC	Front-end Repeater Card	Sec. 4.1.4	pg. 88
LEAR	Low Energy Antiproton Ring	Sec. 1.2	pg. 5
LIFO	Last In, First Out	Appx. C	pg. 223
MC	Monte Carlo Simulation	Sec. 4.2.1	pg. 107
MCP	Multi-Channel Plate	Sec. 3.2	pg. 39
MIPs	Minimum Ionizing Particles	Sec. A	pg. 216
NIM	Nuclear Instrumation Module standard	Sec. 4.1.5	pg. 91
n-side	Negative-charge measuring silicon strips	Sec. 2.4.1	pg. 33
p-side	Positive-charge measuring silicon strips	Sec. 2.4.1	pg. 33
PCA	Principal Component Analysis	Sec. 5.3	pg. 132
PCB	Printed Circuit Board	Sec. 4.1.1	pg. 84
PMT	Photo-Multiplier Tube	Sec. 3.2	pg. 41
PS	Proton Synchrotron	Sec. 3.3.1	pg. 44
SEU	Single Event Upset	Sec. 4.1.2	pg. 87
TTC	Timing, Trigger & Control module	Sec. 4.1.4	pg. 88
VF48	48-Channel VME-based digitizer	Sec. 4.1.4	pg. 90
VMC	Virtual Monte Carlo	Sec. 4.2.1	pg. 108
VME	VMEbus computer bus standard	Sec. 4.1.4	pg. 88
WEP	Weak Equivalence Principle	Sec. 1.1.2	pg. 4

Table 0.1: List of Acronyms and Abbreviations

(x, y, z)	Coordinates in a Cartesian reference system		
(R, ϕ, z)	Coordinates in a cylindrical reference system		
H/\bar{H}	Hydrogen/Antihydrogen	Sec. 1	pg. 1
p/\bar{p}	Proton/Antiproton	Sec. 1	pg. 1
$e^-/(e^+, \beta^+)$	Electron/Positron	Sec. 1	pg. 1
\mathcal{C}	Charge-conjugation operator	Sec. 1.1.1	pg. 2
\mathcal{P}	Parity-inversion operator	Sec. 1.1.1	pg. 2
\mathcal{T}	Time-reversal operator	Sec. 1.1.1	pg. 2
E	Particle total energy	Sec. 1.2	pg. 4
P	Particle momentum	Sec. 1.2	pg. 4
m	Particle rest mass	Sec. 1.2	pg. 4
c	Speed of light	Sec. 1.2	pg. 4
k_B	Boltzmann constant	Sec. 2.1.1	pg. 15
T	Temperature	Sec. 2.1.1	pg. 15
e	Fundamental electric charge	Sec. 2.1.1	pg. 15
ϵ_0	Vacuum permittivity	Sec. 2.1.1	pg. 15
Γ_{SRR}	Spontaneous radiative recombination formation rate	Sec. 2.1	pg. 11
Γ_{TBR}	Three-body recombination formation rate	Sec. 2.1.1	pg. 15
$\Gamma_{\text{Ps}-\bar{p}}$	Charge exchange formation rate	Sec. 2.1.1	pg. 15
G	Laser-simulated recombination enhancement factor	Sec. 2.1	pg. 11
I	Laser intensity	Sec. 2.1	pg. 11
F	Laser cross-sectional area	Sec. 2.1	pg. 11
ν	Laser frequency	Sec. 2.1	pg. 11
$\Delta\nu$	Spread in the laser frequency	Sec. 2.1	pg. 11
a_0	Bohr radius	Sec. 2.1	pg. 11
N_{Ps}	Number of positronium atoms	Sec. 2.1	pg. 11
Ps	Positronium	Sec. 2.1	pg. 11
$\sigma_{\text{Ps}-\bar{p}}$	Positronium-antiproton collision cross-section	Sec. 2.1	pg. 11
Φ	Electric potential	Sec. 2.2.1.1	pg. 18
V_0	Applied electric potential	Sec. 2.2.1.1	pg. 18
d	Characteristic well length for the single particle Penning trap description	Sec. 2.2.1.1	pg. 18
\vec{E}	Electric field	Sec. 2.2.1.1	pg. 18
\vec{B}	Magnetic field	Sec. 2.2.1.1	pg. 18
\vec{F}	Force vector	Sec. 2.2.1.1	pg. 18
ω_c	Cyclotron oscillation	Sec. 2.2.1.1	pg. 18
ω_z	Axial particle oscillation	Sec. 2.2.1.1	pg. 18
ω_k	Magnetron particle oscillation	Sec. 2.2.1.1	pg. 18
λ_D	Debye length	Sec. 2.2.1.2	pg. 21
P_θ	Canonical angular momentum	Sec. 2.2.1.2	pg. 21

$\pi^+/\pi^-/\pi^0$	Charged and neutral pi mesons	Sec. 2.2	pg. 31
n_γ	Number of photons	Sec. 2.3.1	pg. 28
$\vec{\mu}_{\bar{H}}$	Magnetic dipole moment of antihydrogen	Sec. 2.2.2	pg. 24
$\vec{\mu}_B$	Bohr magneton	Sec. 2.2.2	pg. 24
$\vec{\mu}_S$	Spin magnetic moment	Sec. 2.2.2	pg. 24
g_J	Landé's g-factor	Sec. 2.2.2	pg. 24
\vec{L}	Orbital angular momentum	Sec. 2.2.2	pg. 24
\vec{S}	Spin angular momentum	Sec. 2.2.2	pg. 24
\vec{J}	Total angular momentum	Sec. 2.2.2	pg. 24
U_{mag}	Magnetic potential energy	Sec. 2.2.2	pg. 24
U_{max}	Magnetic trap depth	Sec. 2.2.2	pg. 24
E_{\parallel}	Parallel (to the magnetic field) kinetic energy	Sec. 3.3.1	pg. 44
E_{kin}	Total kinetic energy	Sec. 3.2.1	pg. 42
m_p	Mass of the proton	Sec. 3.3.1	pg. 44
γ	Gamma-ray photon	Sec. 3.3.2	pg. 47
ν	Electron neutrino	Sec. 3.3.2	pg. 47
\vec{a}	Particle acceleration	Sec. 3.4.2	pg. 51
γ_c	Cyclotron cooling rate	Sec. 3.4.2	pg. 51
τ_c	Cyclotron cooling constant	Sec. 3.4.2	pg. 51
τ_e	Electron-antiproton collisional cooling constant	Sec. 3.4.2	pg. 51
n_e	Electron density	Sec. 3.4.2	pg. 51
$n_{\bar{p}}$	Antiproton density	Sec. 3.4.2	pg. 51
Λ	Coulomb logarithm	Sec. 3.4.2	pg. 51
N	Number of antiprotons	Sec. 3.4.4	pg. 64
τ_{ev}	Evaporation time constant	Sec. 3.4.4	pg. 64
γ_{ann}	Antiproton-background gas annihilation rate	Sec. 3.4.4	pg. 64
$\mathcal{P}_{\text{Joule}}$	Evaporation time constant	Sec. 3.4.4	pg. 64
α_{esc}	Energy removed by an escaping particle	Sec. 3.4.4	pg. 64
η	Evaporative cooling well depth	Sec. 3.4.4	pg. 64
κ	Evaporative cooling excess particle energy	Sec. 3.4.4	pg. 64
δ_{avg}	Average kinetic energy	Sec. 3.4.4	pg. 64
$ADC_{\text{strip},i}$	The ADC value for a particular <i>strip</i> on the <i>i</i> -th readout event	Sec. 4.1.6.1	pg. 96
$\langle ADC \rangle_{\text{strip}}$	The mean ADC value over an entire run for a particular <i>strip</i>	Sec. 4.1.6.1	pg. 96
σ_{strip}	Standard deviation of ADC values for a particular <i>strip</i> over an entire run	Sec. 4.1.6.1	pg. 96
$\langle ADC \rangle_{\text{ASIC}}$	Mean ADC value over the 128 strips of a particular <i>ASIC</i>	Sec. 4.1.6.1	pg. 97

x_{cluster}	Centroid position of a group of signal strips in the local frame of a detector module	Sec. 4.1.6.2	pg. 101
h_i	Digitizer pulse height	Sec. 4.1.6.2	pg. 101
σ_x	Standard deviation of the position of the hit position in the local frame of a detector module	Sec. 4.1.6.2	pg. 101
θ_v	Track angle from vertical	Sec. 4.1.6.3	pg. 103
p_{Si}	Silicon strip pitch width	Sec. 4.1.6.2	pg. 101
S/N	Ratio of signal to noise	Sec. 4.1.6.2	pg. 101
Z_{eff}	Effective atomic number	Sec. 4.2.2	pg. 109
A_{eff}	Effective atomic weight	Sec. 4.2.2	pg. 109
ρ	Material density	Sec. 4.2.2	pg. 109
X_0	Radiation length	Sec. 4.2.2	pg. 109
W_i	Weight of i -th event in the Monte Carlo simulation	Sec. 4.2.3	pg. 112
W_{max}	Maximum weight of an event in the Monte Carlo simulation	Sec. 4.2.3	pg. 112
\mathcal{D}	Diffusion constant in silicon	Sec. 4.2.4	pg. 113
μ_e	Electron charge mobility in silicon	Sec. 4.2.4	pg. 113
h	Silicon wafer thickness	Sec. 4.2.4	pg. 113
V_{dep}	Depletion voltage across the silicon wafer	Sec. 4.2.4	pg. 113
$\langle E \rangle$	Mean electric field magnitude across the silicon wafer	Sec. 4.2.4	pg. 113
B	Magnetic field magnitude	Sec. 5.1	pg. 120
ρ_c	Radius of curvature for a particle trajectory	Sec. 5.1	pg. 120
P_T	Tangential (to the magnetic field) particle momentum	Sec. 5.1	pg. 120
σ_{P_T}	Resolution of the particle's tangential momentum	Sec. 5.1	pg. 120
ϕ_0	Azimuthal angle for a helix at the point of closest approach to the origin	Sec. 5.4.1	pg. 134
D	Radial distance of closest approach for a helix	Sec. 5.4.1	pg. 134
λ	Helix dip angle	Sec. 5.4.1	pg. 134
z_0	Axial coordinate of the point of closest approach from a helix to the origin	Sec. 5.4.1	pg. 134
s	Arc length parameter	Sec. 5.4.1	pg. 134
χ^2	Goodness-of-fit measure	Sec. 5.5	pg. 138
D_w	Radial distance to the electrode wall	Sec. 5.5	pg. 142
N_{tracks}	Number of charged tracks in an event	Sec. 6.1.1	pg. 163
N_{hits}	Number of hits in an event	Sec. 5.6.2	pg. 144
d_i	Smallest distance between the i th track and the vertex	Sec. 5.6.3	pg. 146

d_{vertex}	Mean distance between the closest approaches of the tracks to the vertex	Sec. 5.6.3	pg. 146
δ	Fractional improvement in the vertex measure through the removal of one of the tracks	Sec. 5.6.4	pg. 148
α	P-value for a Poisson distribution	Sec. 6.3.3	pg. 173
$\langle dE/dx \rangle$	Stopping power	Appx. A	pg. 216
I_{avg}	Mean excitation energy	Appx. A	pg. 216
$\delta(\beta\gamma)$	Density effect correction	Appx. A	pg. 216
T_{max}	Maximum kinetic energy that can be transferred to an electron in a single collision	Appx. A	pg. 216
θ_0	Mean multiple scattering angle	Appx. A	pg. 216
β	Particle velocity	Appx. A	pg. 217
z_c	Particle charge number	Appx. A	pg. 217
E_f	Fermi energy	Appx. B	pg. 218
E_g	Band gap energy	Appx. B	pg. 218
w	Depletion width	Appx. B	pg. 218
V_b	Bias voltage	Appx. B	pg. 218
N_a	Concentration of acceptor atoms	Appx. B	pg. 218
N_d	Concentration of donor atoms	Appx. B	pg. 218
N_{Av}	Avogadro's number	Appx. C	pg. 223
N_λ	Number of interaction lengths for a process in the Monte Carlo simulation	Appx. C	pg. 223
λ_{process}	Mean free path for a process in the Monte Carlo simulation Monte Carlo	Appx. C	pg. 223
Z	Atomic number	Appx. C	pg. 223
A	Atomic weight	Appx. C	pg. 223
$\sigma(E, Z, A)_{\text{process}}$	Cross-section for a process in the Monte Carlo simulation	Appx. C	pg. 223

Table 0.2: List of Symbols

Chapter 1

Introduction

Antihydrogen ($\bar{\text{H}}$) is the bound state of an antiproton (\bar{p}) and a positron (e^+). It is the simplest anti-atomic system, the antimatter analog of hydrogen, and the object of study for the ALPHA experiment. The ALPHA (Antihydrogen Laser PHysics Apparatus)¹ collaboration (of which the author is a member) is an international experimental effort to produce, confine, and perform precision measurements on antihydrogen. In total, the ALPHA collaboration involves 15 institutions from all over the globe (including 5 Canadian institutions²), with the ALPHA apparatus being located at the CERN facility, just outside of Geneva, Switzerland. This dissertation focuses on the demonstration of magnetic confinement of antihydrogen atoms in the ALPHA apparatus. Specifically, aspects of the experiment are presented that allow for the detection of the annihilation products and determination of the annihilation location. These methods are an important part of identifying antihydrogen atoms that have been magnetically confined in the ALPHA apparatus (and subsequently released).

This short introduction will first give a brief summary of the history of the study of antimatter, as it pertains to antihydrogen. Next, several interesting avenues of study are presented, each of which motivates the desire to trap and hold antihydrogen. The current state of the antihydrogen field is then outlined, with particular attention paid to experiments attempting to magnetically confine antihydrogen. Finally, the overall structure of this dissertation is given, and the author's contribution is described.

¹ALPHA website: alpha.web.cern.ch/alpha

²ALPHA-Canada website: www.triumf.ca/alpha

1.1 Motivation for the trapping and study of antihydrogen

Because of its neutrality and correspondence with the well-known hydrogen atom, antihydrogen is a desirable system for experimental study. Moreover, the (presumed) stability of antihydrogen, along with its small magnetic dipole moment, mean that it should be possible to magnetically confine the atoms (Sec. 2.2.2) for extended periods of time. Several interesting experimental possibilities become realizable with confined antihydrogen³, with two of current interest: precision comparisons of spectroscopic properties of hydrogen and antihydrogen, and the measurement of the gravitational interaction between matter and antimatter.

1.1.1 Precision comparison of the properties of hydrogen and antihydrogen

A large amount of antimatter research has been focused on making comparisons between the physical properties of corresponding matter and antimatter systems. Several example comparisons include: electron/positron mass [3], charge [4], and gyromagnetic ratio [5]; neutral-kaon/neutral-antikaon mass and lifetime [6]; proton/antiproton magnetic moment [7], mass and charge [8]; etc. (see Ref. [9] for an extensive list). None of these measurements have revealed statistically significant differences between any of the properties studied. The lack of observed differences between particles and their antiparticles serves as a validation of the \mathcal{CPT} theorem, which states that any local Lorentz-invariant quantum field theory (in a flat space-time) will be invariant under the combined operations of charge-conjugation (\mathcal{C}), parity-inversion (\mathcal{P}), and time-reversal (\mathcal{T}) [10]. As a corollary, the \mathcal{CPT} theorem predicts that for every particle species, there is a corresponding antiparticle species, with identical mass and total lifetime, but equal and opposite charge and magnetic moment. For this reason, comparisons of the properties of matter

³It should be noted that some proposed measurement schemes of the properties of antihydrogen do not require confined antihydrogen atoms. Rather, the measurement can be performed in-flight [1], or with a dedicated spectrometer [2].

and antimatter systems test, with varying precision, this CPT symmetry.

Similarly, the antihydrogen system can be examined and compared with the corresponding hydrogen system as a test of CPT symmetry. Indeed, the antihydrogen $1S - 2S$ two-photon transition line provides an ideal candidate for precision comparison. This transition line has a very sharp (1.3 Hz) natural line width, which has been measured in hydrogen to parts in 10^{14} using Doppler-free two photon spectroscopy [11]. However, high-precision measurements require that the antihydrogen atoms are (1) isolated and (2) at (or nearly at) rest. These two conditions help to minimize disturbances to the system and enable the lengthy measurements needed for precision studies. The magnetic confinement of cold antihydrogen suitably addresses the two measurement conditions. Moreover, to assist spectroscopic efforts, there are a number of schemes to further cool and de-excite the antihydrogen atoms to the ground state [12–14].

Although the $1S - 2S$ transition line is the ultimate experimental goal, CPT -violating effects may also appear in other hydrogen energy splittings [15]. Other attractive comparative measurements include, the $2S_{1/2} - 2P_{1/2}$ Lamb shift [1] and the ground-state hyperfine structure [16–18]. All of these quantities have been very precisely measured for hydrogen (the $1S$ - $2S$ transition is measured as $2\,466\,061\,413\,187.103(46)$ kHz [1.9×10^{-14} rel. uncert.] [11], the $2S_{1/2} - 2P_{1/2}$ Lamb shift as $1\,057\,844.2(34)$ kHz [3.2×10^{-6} rel. uncert.] [19], and the ground-state zero field hyperfine splitting as $1\,420\,405\,751.7662(30)$ Hz [2.1×10^{-12} rel. uncert.] [20]), and similar measurements could conceivably be made on antihydrogen.

1.1.2 Gravitational interaction between matter and antimatter

Although the CPT theorem makes rigid predictions about the correspondence of quantum properties of particles and antiparticles, it does not say anything about the gravitational interaction between matter and antimatter. Indeed, there has not yet been any direct

experimental observations⁴ of the effect of gravity, for example, the gravitational acceleration in the field of the Earth, on antimatter [24, 25]. This is perhaps surprising, given that charged antiparticles have been available since the 1930s. However, the experimental difficulty lies in the overwhelming relative strength of the electromagnetic interaction over the gravitational interaction. For example, even small electric patch fields (variations in surface electric potential due to the crystal structure or absorbed contaminants) are found to significantly perturb the charged particles under study [26].

Antihydrogen, however, is electrically neutral, therefore avoiding many of the experimental difficulties associated with charged particles. This makes antihydrogen an ideal candidate for antimatter gravitational studies. A measurement of the gravitational interaction of antihydrogen with the Earth’s gravitational field would constitute a test of the Weak Equivalence Principle (WEP), which states that freely falling bodies of the same mass and mass distribution should follow the same trajectory [27]. This principle is a cornerstone of general relativity, and any test of the gravitational interaction of antimatter is of general interest.

1.2 Short review of antimatter studies and antihydrogen experiments

Antimatter is the conjugate to ‘normal’ matter (that is, it has identical mass but opposite quantum numbers), and was first predicted by Dirac in 1931 [28]. Dirac was initially attempting to generalize Schrödinger’s equation for relativistic electrons – specifically, by developing an equation that is first-order in both the space and time derivatives. This approach led to the successful Dirac Equation, which, however, provides negative-energy solutions to the relativistic energy equation, $E^2 = P^2c^2 + m^2c^4$, where E is the

⁴It is argued in [21, 22] that equivalence-principle experiments on normal matter place stringent limits on the gravivector acceleration of antimatter (although this interpretation has been contested by some authors [23]). Ultimately, these measurements do not provide a direct observation of the gravitational acceleration of antimatter [24].

particle energy, P is the particle momentum, m is the particle's rest mass, and c is the speed of light. Dirac proposed that these negative energy solutions might correspond to a new particle, the anti-electron, which would have the same mass as the electron, but with opposite electric charge. The anti-electron was first observed by Anderson in 1933, who coined it the 'positron' [29]. This observation was made by examining cloud chamber images of cosmic ray particles⁵ as they passed through a lead plate in a magnetic field. The curvature of the cloud chamber tracks matched the charge-to-mass ratio of the electron, but the direction of curvature was opposite to that of the electron, indicating that the observed particle had opposite electric charge. Dirac shared the Nobel Prize in physics in 1933 for "*the discovery of new productive forms of atomic theory*", while Anderson received a Nobel Prize in physics in 1936 for "*his discovery of the positron*".

Because of its larger mass, and therefore higher production threshold (see Sec. 3.3.1), the antiproton required more sophisticated particle accelerators to produce. It wasn't until 1955 that Chamberlain and Segrè produced and observed the antiproton at Lawrence Berkeley National Laboratory's Bevatron accelerator⁶ [30]. For this experiment, high-energy protons were directed onto a copper target, and a dedicated detection apparatus, consisting of several plastic scintillators, Čerenkov counters, and deflection magnets, was used to observe the produced antiprotons. Chamberlain and Segrè were awarded the 1959 Nobel Prize in physics for "*their discovery of the antiproton*".

The first antihydrogen atoms were observed in-flight at the Low Energy Antiproton Ring (LEAR) at CERN in 1995 [31] and at Fermilab's Antiproton Accumulator in 1998 [32]. For both of these experiments, antiprotons were directed through a gas-jet target. An antiproton passing close to a target nucleus can (with a small probability) produce an electron-positron pair and capture the produced positron. The newly formed antihydrogen atoms are then separated from the antiproton beam and detected by dedi-

⁵See Sec. 6.1.3 for more about cosmic rays.

⁶The Bevatron operated as a proton synchrotron from 1954 to 1993.

cated antihydrogen spectrometers. In this way, the CERN experiment (PS210) observed 9 antihydrogen atoms, while the Fermilab experiment (E862) observed 99. However, the antihydrogen atoms are formed with roughly the momentum of the antiproton beam and they only exist for a short amount of time (on the order of a hundred nanoseconds) before they reach the antihydrogen spectrometer and annihilate. This short survival time makes the investigation of in-flight antihydrogen atoms very difficult. To circumvent this difficulty, an alternate avenue of experimentation was started by the TRAP collaboration, which involved the merging of cold, trapped antiparticle plasmas [33]. The initial success of the TRAP experiment involved the dynamic capture [34] and cooling [35] of antiprotons in a cylindrical Penning Trap (more about these techniques in Sec. 2.2.1-3.4.2).

In 2002, both the ATHENA and ATRAP experiments, located in the Antiproton Decelerator (AD) facility at CERN, succeeded in producing low-energy antihydrogen through the merging of trapped positron and antiproton plasmas [36, 37]. This production method resulted in antihydrogen atoms which survived for much longer (on the order of hundreds of microseconds) than in-flight antihydrogen. Moreover, both ATHENA and ATRAP were able to obtain high antihydrogen production rates (with peak rates on the order of hundreds per second) [38, 39], which allowed for the detailed study of the formation temperature [40, 41] and cooling dynamics [42], as well as for the investigation of modulated [43] and stimulated [44, 45] formation. The ATHENA apparatus also contained a versatile antihydrogen detector, allowing for the three-dimensional imaging of antiproton annihilations [46], which was, in turn, used to study the spatial and temperature profile of the produced antihydrogen distribution [47].

Following their success with antihydrogen production, the experiments then began to focus on the goal of magnetically trapping antihydrogen. Many of the members of the ATHENA collaboration formed a core for the new ALPHA experiment, which directed its effort primarily on its octupole-based magnetic neutral-atom trap (discussed

in detail in Sec. 2.2.2). After demonstrating antiparticle storage [48] and antihydrogen production [49] in the inhomogeneous atom-trapping field, ALPHA also studied the antihydrogen formation dynamics in the multipole trap [50]. The competing ATRAP collaboration pursued a similar approach, but with a quadrupole-based Penning-Ioffe trap, and succeeded in demonstrating antiproton confinement [51] and antihydrogen production [52] in their magnetic trapping fields. More recently, ALPHA has been occupied with the systematic search for trapped antihydrogen in their apparatus [53]. This was followed shortly thereafter by the achievement of a long-awaited milestone: the observation of magnetically confined antihydrogen atoms in the ALPHA apparatus [54] (and of long survival times of those trapped atoms [55]).

Additionally, to complement the progress of the mainline antihydrogen experiments, the AD collaborations have also spent considerable effort developing diagnostic methods and techniques to manipulate positrons and antiprotons. For example, ATHENA focused heavily on their positron plasmas, with work including: the development of non-destructive temperature diagnostics [56], efficient positron accumulation [57], and spatial control [58]. Likewise, ATRAP studied the stacking [59] and adiabatic cooling [60] of antiprotons, along with the density and geometry of positron and antiproton plasmas [61]. ALPHA has also studied an antiproton diagnostic based on octupole-induced ballistic loss [62], zero-rotation frequency bounce-loss resonance [63], evaporative cooling [64] and autoresonant manipulation [65] of antiprotons. Both ALPHA and ASACUSA (another AD experiment) investigated the radial compression of antiproton clouds [66, 67], while ATRAP and ALPHA observed centrifugal separation of electron-antiproton plasmas [68, 69]. All of these techniques contributed, in various ways, to the overall progress of the antihydrogen experiments.

While the antihydrogen field has recently been largely dominated by ATHENA, ATRAP, and ALPHA experiments, there are several collaborations that are either just

starting, or branching into the antihydrogen field. The ASACUSA experiment, for example, had previously focused on the spectroscopy of antiprotonic helium, but has recently demonstrated antihydrogen production in the fields of a magnetic cusp trap [2], and intends to pursue spectroscopy of in-flight antihydrogen. The new AEGIS experiment, on the other hand, is planning on measuring the gravitational interaction between matter and antimatter by passing an antihydrogen beam through an atomic deflectometer [70]. This experiment is currently in the construction phase, and will not be operational until 2012 at the earliest.

1.3 Dissertation overview

This dissertation will describe the first-ever trapping of antihydrogen atoms in the ALPHA apparatus. The emphasis is given to the detection of trapped antihydrogen with the silicon vertex detector (the analysis for which was the responsibility of the author). First, Chapter 2 will give a short overview of the theory behind several selected topics pertaining to the production, trapping, and detection of antihydrogen. Next, Chapter 3 describes the ALPHA apparatus and the experimental components needed to produce and trap antihydrogen, including some novel techniques developed specifically for this initiative. The topic then turns to antihydrogen detection: Chapter 4 describes the ALPHA silicon antihydrogen detector, Chapter 5 outlines the annihilation event reconstruction algorithms, and Chapter 6 describes the cosmic ray background suppression analysis and signal optimization. Chapter 7 then presents the experimental results, describing the demonstration of the magnetic trapping of antihydrogen. Finally, Chapter 8 outlines several possible measurements involving trapped antihydrogen, and ends with some short concluding remarks.

Since this dissertation deals with an interdisciplinary subject, some of the discussions,

both in the main body and appendixes, are more rudimentary than usually expected for theses with more narrowly focused topics. This introductory approach was taken for the benefit of readers from different disciplines.

1.4 Contributions of the author

As a member of the ALPHA collaboration, the author worked closely with many individuals to produce the results described in this dissertation. It is then important to separate the work of author from the overall experimental effort of ALPHA as a whole. The major contributions of the author to the ALPHA project include:

- Design, implementation, and optimization of the annihilation vertexing algorithms for the silicon detector, Chapter 5.
- Implementation and configuration of the Monte Carlo-based detector simulation, Sec. 4.2.
- The detector analysis for observation of trapped antihydrogen, in particular, the analysis and optimization of the cosmic-ray background rejection routines, Chapter 6.

In addition, several technical contributions by the author are presented in this dissertation, including:

- Design and implementation of the dynamic analog strip thresholding routines, Sec. 4.1.6.1 (in collaboration with Dr. James Storey and Dr. Petteri Pusa).
- Implementation of the three-dimensional hit position determination, Sec. 4.1.6.2 (in collaboration with Dr. Petteri Pusa).
- Analysis of the silicon module efficiency and rotational alignment, Sec. 4.1.6.3 (in collaboration with Tim Friesen).

The author spent five summers (although two summers were part of his M. Sc. program [71]) on-site at the ALPHA experiment at CERN. This participation included: assisting in the construction, commissioning, and maintaining of the ALPHA device; day-to-day operations, such as cryogenic handling and shift preparation; routine shift-work, i. e., operation of the device when conducting experiments, online data analysis, development of new experiments; and offline development and analysis (primarily related to the data acquisition system and silicon detector).

As described above, the author contributed in a number of ways to the construction and operation of the apparatus described in Chapter 3 and the data collection, analysis, and presentation of the experimental results of Chapter 7. However, both of these chapters describe the efforts of the ALPHA collaboration as a whole.

Chapter 2

Theory

Experimental work on antihydrogen bridges a number of fields, including: atomic, particle, plasma, and trapping physics. With such a diverse range, it is important to establish the basics of the concepts to be discussed.

This chapter will mirror the typical lifecycle of antihydrogen in the ALPHA experiment: first, formation mechanisms for the synthesis of antihydrogen will be described, with a focus on three-body recombination; next, traps for charged and neutral antimatter are outlined; finally, the positron and antiproton annihilation processes are presented, and the principles of semiconductor detectors to detect the annihilation products are described.

2.1 Formation of antihydrogen with trapped plasmas

Antihydrogen is the bound state of a positron and an antiproton. There are a number of antihydrogen formation mechanisms involving the merging of positron and antiproton plasmas. However, these are low-energy processes, in contrast to the first observed antihydrogen atoms (which were produced relativistically by passing antiprotons through gas jet targets, Sec. 1.2). The intention is to use cold (low-energy) plasmas to produce low-energy antihydrogen that is amenable to neutral-atom trapping (Sec. 2.2.2).

A commonality shared by all of the mechanisms presented here is that, in addition to the requisite positrons and antiprotons, the presence of a third party is also required. This third party is needed to carry away the binding energy associated with placing the positron and antiproton in a bound state. Several important mechanisms are listed here

(all of these processes are reviewed at length in [72, 73]):

Radiative recombination: in the case of radiative recombination, an emitted photon acts as the third body and carries away the excess binding energy (the inverse process to photo-ionization). The simplest process is spontaneous radiative recombination (SRR),

$$e^+ + \bar{p} \rightarrow \bar{\text{H}} + \hbar\omega, \quad (2.1)$$

where \hbar is Planck's constant and ω is the angular frequency of the photon. Here, the photon is emitted spontaneously, which results in a relatively low production rate (compared to a mechanism not relying on spontaneous photon emission). After integrating over the allowed photon states and positron velocity distribution, the formation rate can be give as [33]

$$\Gamma_{\text{SRR}} = (3 \times 10^{-11}) \left(\frac{4.2 \text{ K}}{T_{e^+}} \right)^{1/2} \left(\frac{n_{e^+}}{\text{cm}^{-3}} \right) \text{ s}^{-1}, \quad (2.2)$$

where T_{e^+} and n_{e^+} are the positron temperature and density. For a hypothetical 4.2 K plasma with a density of 10^7 cm^{-3} , this gives a rate of $\Gamma_{\text{SRR}} = 3 \times 10^{-4} \text{ s}^{-1}$ per antiproton.

Radiative recombination can be stimulated through the addition of a laser source providing k photons of energy $\hbar\omega$,

$$e^+ + \bar{p} + k\hbar\omega \rightarrow \bar{\text{H}} + (k + 1)\hbar\omega. \quad (2.3)$$

This is an attractive process, since the laser frequency might be tuned to produce antihydrogen atoms in specific energy states [74]. The enhancement, or gain G , through the addition of laser-stimulated recombination can be expressed as [75]

$$G = \frac{Ic^2}{F\Delta\nu 8\pi h\nu^3}, \quad (2.4)$$

where I is the laser power (in watts), F is the cross-sectional area of the laser beam, ν is the laser frequency, and $\Delta\nu$ represents the 'frequency spread of the

beam' (that is, $\Delta\nu$ takes into account that only photons and positrons with energies within the width of the recombination energy level will engage in this process, and since the energy spread of the positrons usually dominates, it can be expressed as $\Delta\nu = (m_{e^+}v_{e^+}/h)\Delta v_{e^+}$, where Δv_{e^+} is the positron velocity spread [76]). Laser-stimulated radiative recombination was attempted by the ATHENA experiment [45]. ATHENA used a $100 \text{ W}\cdot\text{cm}^{-2}$, $\lambda = 10.96 \text{ }\mu\text{m}$ CO₂ laser to attempt to stimulate the $n = 11$ state and predicted an enhancement, given the positron plasma density used, from 24 s^{-1} to 60 s^{-1} (although according to Eq. 2.4, the expected enhancement would be from 24 s^{-1} to $\sim 224 \text{ s}^{-1}$ [73]). A 90% C. L. upper limit for a laser enhancement of 0.8% (or 0.3 s^{-1}) was found. This result, along with measurements of the temperature-dependent antihydrogen production in ATHENA [40], indicate that radiative recombination processes are of little importance when merging cold charged plasmas to produce antihydrogen (at least for the ATHENA parameter range).

Three-body recombination (TBR): with a dense positron plasma, an additional positron can efficiently carry away the excess binding energy,



This mechanism is suspected as being the dominant formation process for experiments using dense positron plasmas, given the observed high production rates [37, 40, 54]. Because of the importance of this process, it is discussed in detail in Sec. 2.1.1.

Charge transfer: here, a positronium atom Ps (the bound state between a positron and an electron), collides with an antiproton. The antiproton can then exchange with the electron to form antihydrogen:



Moreover, since the binding energy for the resultant antihydrogen atom is expected to be very similar to the positronium binding energy [77], the final energy state of the antihydrogen atom could be controlled by manipulating the state of the positronium atom. This manipulation can be accomplished in a two-step process, where the first step involves exciting cesium atoms using an infrared laser system, such that the excited atoms, Cs^* , are directed through the electrode aperture into the positron plasma to form excited positronium, Ps^* [78]:



Being electrically neutral, the excited positronium atoms are not significantly affected by the electric fields in the trap, and a small fraction will reach the antiprotons. Antihydrogen can then be formed through the exchange of a positronium electron for an antiproton:



where $\bar{\text{H}}^*$ is an excited antihydrogen atom. The exchange process is efficient for high- n state charge transfers, where the cross-section, $\sigma_{\text{Ps}-\bar{p}}$, is calculated (using classical trajectory Monte Carlo simulations with 4 K antiprotons and positronium atoms [78]) as

$$\sigma_{\text{Ps}-\bar{p}} = 58 n_{\text{Ps}}^4 \pi a_0^2, \quad (2.9)$$

where n_{Ps} is the positronium principal quantum number, and a_0 is the Bohr radius. Thus, for $N_{\text{Ps}} = 10^6$ Ps atoms with quantum numbers $n_{\text{Ps}} = 34$, the expected number of antihydrogen atoms would be $N_{\text{Ps}} \sigma_{\text{Ps}-\bar{p}} / (4\pi d^2) = 10^{-2}$ antihydrogen atoms per antiproton (where $d = 0.2$ cm is the separation between the two trapped species). However, the majority of the production is expected to occur in the first few milliseconds, corresponding to an instantaneous production rate of $\Gamma_{\text{Ps}-\bar{p}} \simeq 10 \text{ s}^{-1}$ per antiproton.

Following this procedure, the ATRAP experiment demonstrated the formation of antihydrogen atoms by using selectively laser-excited cesium atoms to produce excited positronium in a region adjacent to an antiproton cloud [44].

2.1.1 Three-body recombination

As mentioned above, the three-body recombination mechanism (Eq. 2.5) is suspected of being the dominant antihydrogen formation process when merging cold, dense plasmas. This process can occur when two positrons collide close to an antiproton. If one of the positrons loses the majority of its kinetic energy in the collision, it can become bound to the antiproton, while the other positron carries away the excess energy. This collision-driven processes is efficient for dense positron plasmas, but mostly forms high- n , weakly-bound antihydrogen atoms [33].

Following the discussion in [73], the steady-state equilibrium three-body recombination rate, Γ_{TBR} , can be written as:

$$\Gamma_{\text{TBR}} = Cn^2vb^5, \quad (2.10)$$

$$= C(8 \times 10^{-12}) \left(\frac{4.2 \text{ K}}{T} \right)^{9/2} \left(\frac{n_{e^+}}{\text{cm}^{-3}} \right)^2 \text{ s}^{-1}, \quad (2.11)$$

where n_{e^+} is the positron density, $v = \sqrt{k_B T / m_{e^+}}$ is the positron thermal velocity (where k_B is Boltzmann's constant, T is the temperature, and m_{e^+} is the positron mass), $b = e^2 / (4\pi\epsilon_0 k_B T)$ (where e is the fundamental electric charge, and ϵ_0 is the permittivity of vacuum), and C is a constant of proportionality. This expression can be viewed as a flux of positrons, nvb^2 (where b^2 functions as the cross section), times the probability of another positron being in the same volume, nb^3 (in this analysis, all numerical constants have been folded into C). While this rate scales with density and temperatures as $\Gamma_{\text{TBR}} \propto n^2 T^{-9/2}$, the constant of proportionality in Eqn. 2.10 requires careful evaluation.

Glinsky and O'Neil [79] calculated this rate for strong magnetic fields ($B \rightarrow \infty$, such

that the charged particles are pinned to the field lines) and found $C = 0.07 \pm 0.01$ (or, for a 4.2 K positron plasma with density of 10^7 cm^{-3} , $\Gamma_{\text{TBR}} \sim 56 \text{ s}^{-1}$ per antiproton). Later, Robicheaux and Hanson [80] found $C = 0.11 \pm 0.01$ ($\Gamma_{\text{TBR}} \sim 80 \text{ s}^{-1}$ per antiproton) using the next-order guiding-center approximation (that is, the inclusion of the $\vec{E} \times \vec{B}$ drift). Robicheaux [81] proceeded to perform full motion classical Monte Carlo simulations of the positron trajectories, finding (for a 1 T magnetic field) $C = 0.11$ ($\Gamma_{\text{TBR}} \sim 88 \text{ s}^{-1}$ per antiproton) at 4 K, $C = 0.15$ ($\Gamma_{\text{TBR}} \sim 120 \text{ s}^{-1}$ per antiproton) at 8 K, and $C = 0.19$ ($\Gamma_{\text{TBR}} \sim 150 \text{ s}^{-1}$ per antiproton) at 16 K.

However, Robicheaux [81] also found that the formation of low-field seeking states (i. e. trappable states, Sec. 2.2.2) is suppressed for three-body recombination in strong magnetic fields. Using a classical Monte Carlo method, less than 8% of the simulated antihydrogen atoms were shown to have large (greater than 10 K T^{-1}) low-field seeking magnetic moments [73]. Similarly, Robicheaux also found that the finite size of the plasma is important for three-body recombination in nested potentials (Sec. 3.4.5), which is likely the result of the capture process being arrested as the antiprotons enter, then exit, the positron plasma [82].

2.2 Charged (anti)particle and neutral-atom trapping

Having described the mechanisms of low-energy antihydrogen production, the methods to hold the constituent positrons and antiprotons and trap the resulting antihydrogen atoms are now discussed.

Ion and atom traps are invaluable tools for precision measurements on rare particles and exotic atoms. These traps provide precise control, isolation, and long-time confinement for charged particles, ions, and neutral atoms, affording the experimenter the opportunity to conduct complicated manipulations or time-intensive measurements.

There are many types of trapping systems, including (non-extensively): Paul traps [83], magneto-optical traps (MOT) [84], electron beam ion traps (EBIT) [85], etc. However, the specific traps that are relevant to the ALPHA system and thus will be described here are the Penning-Malmberg trap for charged particles and plasmas, and the minimum-B neutral atom trap.

In the following descriptions, the term ‘antiparticle’ is dropped, and no distinction is made between matter and antimatter. This follows from conservation of charge conjugation symmetry in the electromagnetic interaction. That is, antiparticles are expected to follow the same equations of motion (with the appropriate sign for charge) as their matter counterparts.

2.2.1 Penning-Malmberg trap for charged particles

The widely successful Penning trap (Fig. 2.1(a)) consists of hyperbolic electrodes, which provide a quadrupole electric field, and a strong homogeneous axial magnetic field [86]. The axial magnetic field prevents charged particles from escaping radially, while the electric field restricts the axial motion of the particles. Penning traps have been used very successfully for precision measurement on confined particles and ions [5, 87–89]

The Penning-Malmberg variant of this trap (Fig. 2.1(b)) replaces the hyperbolic electrodes with a series of hollow cylindrical electrodes [90]. Similar to above, the Penning-Malmberg trap includes a strong axial magnetic field for radial confinement, while the electrodes provide an axially-confining electric field. This configuration has two experimental advantages: ease of access to the trap volume, as instruments can be placed at the ends of the electrode cylinder; and with a segmented electrode stack, the axial electric fields can be varied to precisely manipulate the trapped species, or even simultaneously confine multiple bunches or differently charged particles/ions. The remainder of this subsection will focus on the theory of charged particle and plasma confinement in a

Penning-Malmberg trap.

2.2.1.1 Single particle confinement

Alone, an electrostatic potential is inadequate to provide three-dimensional confinement for a charged particle. This is a consequence of Laplace's equation, which forbids local electrostatic extrema for two- and three-dimensional systems. However, an electric potential Φ that satisfies Laplace's equation, and provides a restorative force in the \hat{z} direction, is¹:

$$\Phi(x, y, z) = \frac{V_0}{2d^2} \left(z^2 - \frac{1}{2}x^2 - \frac{1}{2}y^2 \right), \quad (2.12)$$

where V_0 is the applied potential, and d is the characteristic well length. The restorative, and therefore confining, force in \hat{z} for a massive particle with charge q can be seen by examining the gradient of the potential, which acts as the force on the particle,

$$\vec{F} = q\vec{E} = -q\vec{\nabla}\Phi = \frac{V_0}{2d^2} (2z\hat{z} - x\hat{x} - y\hat{y}). \quad (2.13)$$

Then isolating the z component of $\vec{F} = m\frac{d^2\vec{r}}{dt^2}$, an oscillatory equation of motion is found:

$$\frac{d^2z}{dt^2} + \omega_z^2 z = 0, \quad (2.14)$$

where $\omega_z = \sqrt{\frac{qV_0}{md^2}}$. In the \hat{x} and \hat{y} directions, however, the equations of motion have only exponential, and therefore unconfined, solutions. For an appropriate choice of applied potential either positively and negatively charged particles will oscillate in \hat{z} with characteristic frequency, ω_z (usually referred to as the axial frequency).

For this system, the \hat{x} and \hat{y} motion of the particle can be coupled, via the Lorentz force, by adding a magnetic field, $\vec{B} = B_0\hat{z}$. The force equation is then:

$$m\frac{d^2\vec{r}}{dt^2} = q \left(-\vec{\nabla}\Phi + B_0\frac{d\vec{r}}{dt} \times \hat{z} \right), \quad (2.15)$$

¹Note that the discussion here uses a harmonic potential (which is normally associated with the potential used in Penning traps). However, the electric potentials used in Penning-Malmberg traps are often more complicated and include anharmonic contributions. For this reason, the results presented here can be considered a first-order approximation to the full Penning-Malmberg equations of motion.




Figure 2.1: Schematic illustration of the a) Penning, and b) Penning-Malmberg trap configurations. The Penning trap (a) consists of two hyperbolic endcaps and a hyperbolic ring electrode. A quadrupole electric field is established by biasing the endcaps to the ring electrode. A strong axial magnetic field provides radial confinement. The Penning-Malmberg trap (b) consists of a series of cylindrical electrodes. A strong axial magnetic field provides radial confinement, while axial confinement is provided by biasing the cylindrical electrodes.

and the equations of motion can be expanded as,

$$\frac{d^2x}{dt^2} - \omega_c \frac{dy}{dt} - \frac{1}{2}\omega_z^2 x = 0, \quad (2.16)$$

$$\frac{d^2y}{dt^2} + \omega_c \frac{dx}{dt} - \frac{1}{2}\omega_z^2 y = 0, \quad (2.17)$$

$$\frac{d^2z}{dt^2} + \omega_z^2 z = 0, \quad (2.18)$$

where $\omega_c = \frac{qB_0}{m}$ is the cyclotron frequency, and the force equation in z is unchanged.

Following Kretzschmar [91], a complex substitution $u = x + iy$ can be made, allowing the \hat{x} and \hat{y} equations to be written together as:

$$\frac{d^2u}{dt^2} + i\omega_c \frac{du}{dt} - \frac{1}{2}\omega_z^2 u = 0. \quad (2.19)$$

Eq. 2.19 is a complex, ordinary differential equation, that can be solved using the ansatz $u = e^{-i\omega t}$, which reduces the differential equation to a quadratic,

$$\omega^2 - \omega_c \omega + \frac{1}{2}\omega_z^2 = 0, \quad (2.20)$$

which has roots at

$$\omega_{\pm} = \frac{1}{2} \left(\omega_c \pm \sqrt{\omega_c^2 - 2\omega_z^2} \right). \quad (2.21)$$

Thus, when $\omega_c^2 > 2\omega_z^2$ the solutions are oscillatory, and therefore confined, in the \hat{x} and \hat{y} directions. Then, for $\omega_c \gg \omega_z$, the positive root is identified as the cyclotron motion, $\omega_+ \simeq \omega_c$. As well, noting that $\omega_+ \omega_- = \omega_z^2/2$, the negative root can be expressed as $\omega_- \simeq \omega_z^2/(2\omega_c)$. This root is often called the magnetron frequency ($\omega_m = \omega_-$), and can be identified as the $\vec{E} \times \vec{B}$ drift motion. The motion of a charged particle in a Penning-Malmberg trap can then be seen to be the superposition of the three oscillatory modes (Fig. 2.2), with frequencies given by:

$$\omega_c = \frac{qB_0}{m}, \quad (\text{Cyclotron}) \quad (2.22)$$

$$\omega_z = \sqrt{\frac{qV_0}{md^2}}, \quad (\text{Axial}) \quad (2.23)$$

$$\omega_m = \frac{\omega_z^2}{2\omega_c}. \quad (\text{Magnetron}) \quad (2.24)$$

Particle Type	Cyclotron frequency, ω_c	Axial frequency, ω_z	Magnetron frequency, ω_m
Electron	28 GHz	21 MHz	8 kHz
Positron	28 GHz	21 MHz	8 kHz
Antiproton	15 MHz	500 kHz	8 kHz

Table 2.1: Particle oscillation frequencies for $B_0 = 1$ T, $V_0 = 10$ V, and $d = 1$ cm (similar to the typical operating conditions in the ALPHA trap).

These modes define the confined single particle motion in the electric and magnetic fields of the Penning-Malmberg trap. Moreover, Eqs. 2.22-2.24 set the timescales for the three motional degrees of freedom (some example frequency values are tabulated in Table 2.1). Of note are the several orders of magnitude differences between the electron/positron versus antiproton cyclotron and axial frequencies. However, the mass-independent magnetron motion is identical for both species (just as the $\vec{E} \times \vec{B}$ drift is mass and charge independent).

2.2.1.2 Confinement of charged plasmas

The ALPHA experiment deals with large numbers of particles; electrons, positrons, and antiprotons are often confined simultaneously. The temperatures and densities of the collections of particles are often such that their Debye lengths, $\lambda_D = \sqrt{\epsilon_0 k_B T / n_0 e^2}$ (where ϵ_0 is the vacuum permittivity), are small compared to the dimensions of the collection of particles used. As such, the many-body collections of charged particles used in ALPHA often need to be treated as plasmas. Moreover, these plasmas are considered to be nonneutral [92], as they are composed of charged particles of the same sign, and there is no overall charge neutrality. The single particle description of Sec. 2.2.1.1 still applies for antiproton clouds with small densities, such as before radial compression (Sec. 3.4.3).

Fortunately, excellent radial confinement of nonneutral plasmas in a strong axial magnetic field follows from angular momentum conservation. Summarizing Dubin and

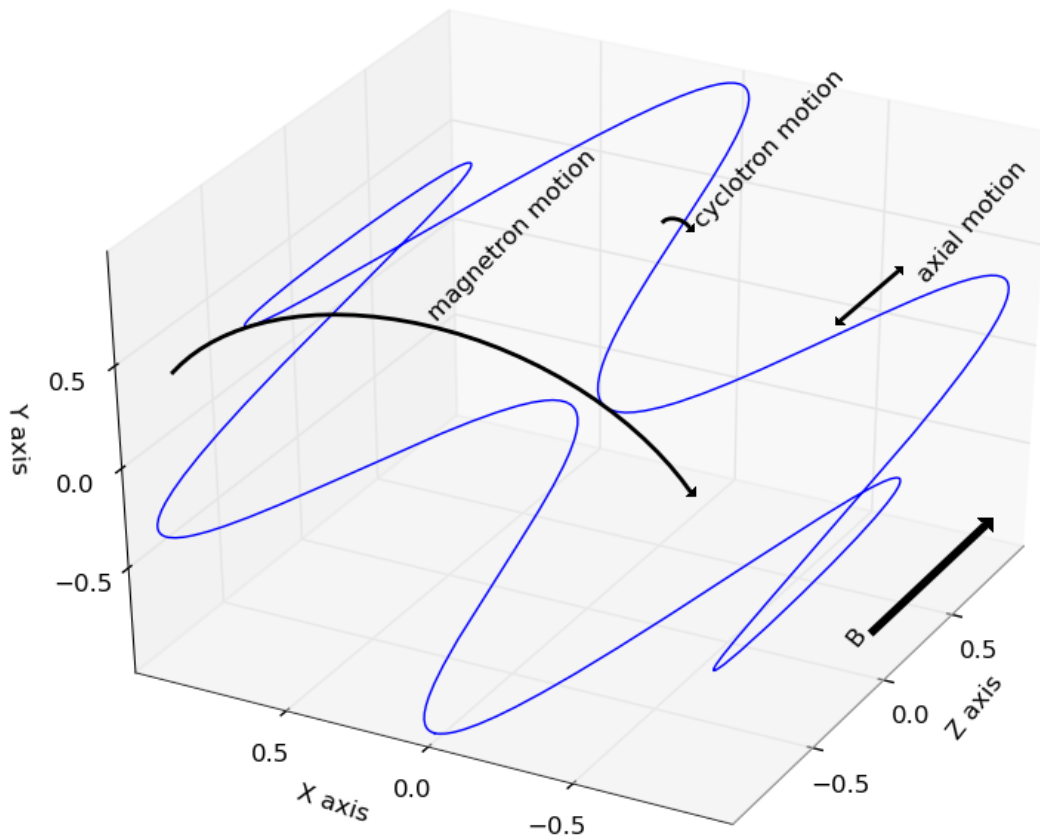


Figure 2.2: X, Y, Z spatial illustration of single particle motion in a Penning-Malmberg trap. The motion is a superposition of three oscillatory modes (Eq. 2.22): the slow magnetron oscillation, the axial vibration, and the fast cyclotron orbit.

O'Neil [93, 94], the total canonical angular momentum, P_θ , for N charged particles can be written (in cylindrical coordinates) as,

$$P_\theta = \sum_{j=1}^N \left(m v_{\theta j} r_j + \frac{q}{c} A_\theta(r_j) r_j \right), \quad (2.25)$$

where $v_{\theta j}$ is the angular velocity, r_j is the distance from the axis of symmetry (or radius) of the j th particle, and $A_\theta(r)$ is the azimuthal component of the vector potential (i. e. $\vec{A} = \hat{\theta} A_\theta(r)$). For a uniform, axial magnetic field $A_\theta(r) = Br/2$, and noting that for a strong magnetic field the second term in Eq. 2.25 will dominate, the canonical angular momentum reduces to

$$P_\theta \simeq \sum_{j=1}^N \frac{qB}{2c} r_j^2. \quad (2.26)$$

Moreover, since the non-neutral plasmas under consideration have only one species of charge, all of the non-subscripted coefficients are constants and thus can be taken out of the sum to produce

$$P_\theta \simeq \frac{qB}{2c} \sum_{j=1}^N r_j^2. \quad (2.27)$$

Since angular momentum is conserved for systems with rotational symmetry (such as Penning-Malmberg traps), a single species plasma in a strong magnetic field is subject to the constraint,

$$\sum_{j=1}^N r_j^2 \simeq \text{const.} \quad (2.28)$$

Eq. 2.28 places a constraint on the overall radius of the plasma, prohibiting expansion or contraction of the mean squared plasma radius without a corresponding change to the angular momentum, which cannot happen without an external force (outside the trapping fields) driving the system, collisions with neutral particles, or field errors. In practice this leads to excellent confinement properties for nonneutral plasmas, with possible confinement times being on the order of 10^5 s [95]. This is in contrast to neutral or

quasineutral plasmas, for which the conservation of angular momentum implies

$$\sum_{j=1}^N q_j r_j^2 \simeq \text{const}, \quad (2.29)$$

where q_j is the charge of the j th particle. This means that for neutral or quasi-neutral plasmas, it is possible for an electron and an ion (for example) to travel together outward and leave the confinement region, all while conserving the overall angular momentum. In this way, we see that the excellent confinement characteristics for nonneutral plasmas in Penning-Malmberg traps are not shared by neutral or quasineutral plasmas.

It should be stressed that this argument holds only when there is azimuthal symmetry and in the absence of collisions with neutral particles. Although collisions with similarly charged particles conserve an adiabatic invariant and do not affect the angular momentum conservation [96, 97], collisions with neutral particles (or atoms) can change the angular momentum of the system. Plasma confinement is also questionable if the trap system contains some small field or construction errors. These errors negate the symmetry condition and angular momentum will not be exactly conserved. The effects of azimuthal symmetry breaking will be discussed further in the description of ALPHA's magnetic neutral trap in Sec. 2.2.2.

2.2.2 Neutral-atom trapping

Antihydrogen is electrically neutral and is therefore not confined by Penning-Malmberg traps. However, antihydrogen atoms have a small permanent magnetic dipole moment, $\vec{\mu}_{\bar{\text{H}}}$, due to their angular momentum,

$$\vec{\mu}_{\bar{\text{H}}} = -g_J \mu_B \frac{\vec{J}}{\hbar}, \quad (2.30)$$

where g_J is the Landé g-factor, $\mu_B = e\hbar/2m_e$ is the Bohr magneton, and $\vec{J} = \vec{L} + \vec{S}$ is the total angular momentum, with the terms in the sum representing orbital, \vec{L} , and spin, \vec{S} , angular momentum. Quantum mechanics restricts the atoms to having magnetic dipole

moments $\mu_{\bar{H}} = -m_J g_J \mu_B$, where m_J is the total angular momentum projection quantum number, which can assume values: $-J, -(J-1), \dots, J-1, J$.

For ground state antihydrogen, the orbital angular momentum vanishes ($L = 0$). The spin contribution to the magnetic moment is

$$\vec{\mu}_S = -g \frac{q}{2m} \vec{S}, \quad (2.31)$$

where g depends on the particle ($g_e = 2.0023$ for the electron, $g_p = 5.586$ for the proton). Thus, the antihydrogen magnetic moment is dominated by the positron spin, due to the small positron-antiproton mass ratio, $m_{e^+}/m_{\bar{p}} \simeq (1800)^{-1}$. The magnetic dipole moment of antihydrogen can then be expressed as,

$$\vec{\mu}_{\bar{H}} \simeq \vec{\mu}_{S,e^+} = -g_e \mu_B \frac{\vec{S}}{\hbar}. \quad (2.32)$$

In a magnetic field, \vec{B} , a torque $\vec{\tau} = \vec{\mu}_{\bar{H}} \times \vec{B}$ will align the magnetic moment with the field. Thus, with $J = S = \pm \hbar/2$ and $g_e \simeq 2$, a ground-state antihydrogen atom can have two states in a magnetic field: $\mu_{\bar{H}} = \pm \mu_B$.

The magnetic potential energy, U_{mag} , between the external magnetic field and the atom's magnetic dipole moment is given as,

$$U_{\text{mag}} = -\vec{\mu}_{\bar{H}} \cdot \vec{B}, \quad (2.33)$$

$$= \mp \mu_B B, \quad (2.34)$$

since the magnetic dipole moment will either be aligned or anti-aligned with the field. For a positive dipole moment, the interaction energy decreases with increasing field. These states are known as 'high-field seekers', as they are attracted to regions with high magnetic fields. On the other hand, atoms with negative dipole moments will be repulsed from regions with high magnetic fields, and are called 'low-field seekers'. It is not possible to construct a 3-dimensional static local magnetic maximum in vacuum, so the high-field

seeking atoms are non-trappable [98]. However, it is possible to construct a local magnetic field minimum. This is the minimum-B (or ‘magnetic bottle’) trap, where the magnetic field magnitude is at a minimum in the trapping region and increases outward in every direction. Thus, if the kinetic energy of a low-field seeking antihydrogen atom is less than the magnetic potential difference, the atom will be magnetically confined within the minimum-B region. The maximum kinetic energy, or trap-depth, U_{\max} , is given by

$$U_{\max} = \mu_B(B_{\max} - B_{\min}), \quad (2.35)$$

where B_{\max} and B_{\min} are the maximum and minimum trap magnetic field magnitudes, respectively. Eq. 2.35 can also be re-expressed (by dividing by Boltzmann’s constant) to be expressed in units of temperature,

$$U_{\max} = 0.67\Delta B \text{ [K]}, \quad (2.36)$$

where ΔB is the difference between the maximum and minimum field magnitudes in units of Teslas.

In a cylindrical geometry, a minimum-B trap requires a radially increasing magnetic field. Such a radially increasing field can be accomplished by adding a multipolar magnet around the exterior of the trap. In this configuration, the magnetic field magnitude as a function of radius $|\vec{B}(r)|$ scales as [99],

$$|\vec{B}(r)| \propto r^{s-1}, \quad (2.37)$$

where s is the order of the multipole ($s = 2$ for a quadrupole field, $s = 3$ for sextapole field, $s = 4$ for an octupole, etc.). Figure 2.3 shows a comparison of the magnetic field magnitudes as a function of radius for several multipole orders. In this comparison, the field magnitudes have been normalized to the maximum field magnitude and radius, B_w and r_w , which correspond to the field magnitude and radius at the electrode surface in the apparatus.

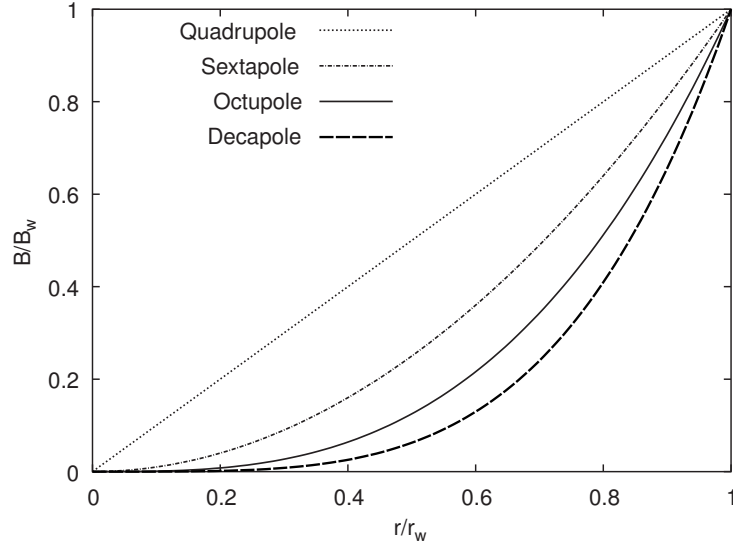


Figure 2.3: Radial variation in the magnetic magnitude for various multipole magnets (reproduced with permission from [99]). The curves are normalized to the field magnitude at the wall radius.

In order to simultaneously utilize both the Penning-Malmberg trap to synthesize antihydrogen (described further in Sec. 3.4.5) and the minimum-B trap to confine the produced antihydrogen atoms, the minimum-B trap must be superimposed over the Penning-Malmberg trap. However, recalling Sec. 2.2.1.2, one of the conditions for stable plasma confinement is azimuthal symmetry in the confining fields, which is violated with the addition of a multipolar magnetic field. The extent of azimuthal symmetry breaking is an important consideration when choosing the multipole order. Higher order multipole fields are expected to have better confinement properties for on-axis charged plasmas, since their perturbations are much smaller near-axis. However, due to the necessary inclusion of the magnet support form, vacuum chamber wall, and trap electrodes, the radial extent of the neutral trap will be less than the maximum radius shown in Fig. 2.3. This means that higher-order multipole magnets will have smaller trap depths as compared to a quadrupole trap (the lowest multipole order configuration) [99].




Figure 2.4: Lowest order Feynman diagrams for electron-positron annihilation for the (a) two- γ , and (b) three- γ channels.

2.3 Particle-antiparticle annihilation

Having described how to make and store antihydrogen, the question of detection is now addressed. When an antiparticle comes into contact with a particle of its conjugate species (e. g., a positron encountering an electron), the particle and antiparticle will annihilate. That is, since the particle and antiparticle have opposite additive quantum numbers, their interaction will cause the particle and the antiparticle to destroy each other. However, both energy and momentum must be conserved, and these conditions are satisfied through the production of additional (anti)particles (whose quantum numbers must also sum to zero). Many of the resultant (anti)particles can be detected, and a significant portion of this dissertation is devoted to the detection and analysis of antiproton annihilation products (Chapter 5).

The annihilations in ALPHA predominately occur between particles and antiparticles with very low kinetic energies. To simplify this discussion, it is assumed that particle-antiparticle systems are at rest, and higher energy cases are not considered.

2.3.1 Electron-positron annihilation

A positron incident on normal matter will eventually encounter an electron and annihilate (unless, of course, it subsequently exits the matter sample). Within matter, the amount

of time that the positron survives before annihilation is inversely proportional to the local electron density. The time-to-annihilation is typically about 500 ps for condensed matter, but can be much longer within gases where the annihilation usually proceeds through the formation of positronium. The average lifetime para-positronium (1S_0) is 125 ps (in vacuum), while ortho-positronium (3S_1) has an average lifetime of 142 ns (in vacuum) [100].

Low-energy electron-positron annihilation generally results in the production of two or three γ -ray photons (Fig. 2.4). Following from \mathcal{C} -conservation [101], the number of photons produced is set by the angular momentum state of the electron-positron pair,

$$(-1)^{n_\gamma} = (-1)^{L+S}, \quad (2.38)$$

where n_γ is the number of photons, and L and S are the orbital and spin angular momentum of the electron-positron pair, respectively (note that the production of a single photon is forbidden by the conservation of momentum). Thus, for the common S-wave ($L = 0$) annihilation, the singlet 1S_0 state will usually produce two photons (Fig. 2.4(a)),

$$e^+ + e^- \rightarrow \gamma + \gamma, \quad (2.39)$$

while the triplet 3S_1 state will usually result in three photons (Fig. 2.4(a)),

$$e^+ + e^- \rightarrow \gamma + \gamma + \gamma. \quad (2.40)$$

The 1S_0 two- γ annihilation produces a characteristic signature of two back-to-back 511 keV photons, which provides a signal that can be readily discriminated. Indeed, the ATHENA antihydrogen detector contained an array of 192 CsI crystals used to detect the photons produced in positron annihilation [36, 102]. However, this system had the disadvantage of a low photon detection efficiency ($\sim 25\%$ per CsI modules, or $\sim 5\%$ total efficiency for the detection the two simultaneous photons [103]).

2.3.2 Antiproton annihilation

Similar to the positron, an antiproton will eventually annihilate when brought into contact with normal matter. However, the antiproton is a composite system, made up of two anti-up quarks (\bar{u}), an anti-down quark (\bar{d}), and gluons. For this reason, the annihilation of the antiproton is decidedly more complicated than that of the positron.

Low-energy antiproton annihilation involves the rearrangement and/or annihilation of the constituent quarks and antiquarks. Moreover, in addition to proton-antiproton annihilation, the antiproton can also interact with the up quark and two down quarks of the neutron. The result of antiproton-nucleon annihilation is the transition from interacting baryons to a system comprised entirely of mesons [104]. In low-energy antiproton annihilation, charged and neutral pions are the dominant products (for example, Table 2.2 gives pion final-state branching ratios for antiproton-proton ($\bar{p}p$) and antiproton-neutron ($\bar{p}n$) annihilations at rest).

On average, an isolated antiproton-proton annihilation (at rest) produces 1.5 π^+ , 1.5 π^- , and 2 π^0 particles, while an antiproton-neutron annihilation produces 1 π^+ , 2 π^- , and 2 π^0 particles [105, 106]. However, in the ALPHA experiment it is much more likely that an antiproton will annihilate in the nucleus of a heavy atom (the material surrounding the antiproton trap is gold-plated aluminum). Annihilation in the nucleus introduces further complications, including: absorption of the produced pions [107], or the fragmentation of the nucleus itself [108]. The effect of the final state interaction (that is, the interaction of the pions with the nucleus and each other immediately after annihilation) is that $\bar{p}p$ and $\bar{p}n$ annihilations become hard to distinguish. This is the result of nuclear charge-exchange processes ($\pi^+n \leftrightarrow \pi^0p$ and $\pi^-p \leftrightarrow \pi^0n$) and pion absorption processes ($\pi^-pp \leftrightarrow np$, $\pi^-pn \leftrightarrow nn$, $\pi^+pn \leftrightarrow pp$, and $\pi^+nn \leftrightarrow np$) [105]. As a result, the low-energy antiproton annihilation is often modeled as the statistical release of, on average, three charged and two neutral pions [109].

Antiproton-proton, $\bar{p}p$ [106]		Antiproton-neutron, $\bar{p}n$ [109]	
Pion final state	Branching ratio	Pion final state	Branching ratio
$\pi^0\pi^0$	0.00028	$\pi^-\pi^0$	0.0075
$\pi^0\pi^0\pi^0$	0.0076	$\pi^-k\pi^0$ ($k > 1$)	0.169
$\pi^0\pi^0\pi^0\pi^0$	0.03	$\pi^-\pi^-\pi^+$	0.023
$\pi^+\pi^-$	0.0032	$\pi^-\pi^-\pi^+\pi^0$	0.17
$\pi^+\pi^-\pi^0$	0.069	$\pi^-\pi^-\pi^+k\pi^0$ ($k > 1$)	0.397
$\pi^+\pi^-\pi^0\pi^0$	0.093	$\pi^-\pi^-\pi^-\pi^+\pi^+$	0.042
$\pi^+\pi^-\pi^0\pi^0\pi^0$	0.233	$\pi^-\pi^-\pi^-\pi^+\pi^+\pi^0$	0.12
$\pi^+\pi^-\pi^0\pi^0\pi^0\pi^0$	0.028	$\pi^-\pi^-\pi^-\pi^+\pi^+k\pi^0$ ($k > 1$)	0.066
$\pi^+\pi^-\pi^+\pi^-$	0.069	$\pi^-\pi^-\pi^-\pi^-\pi^+\pi^+\pi^+k\pi^0$ ($k \geq 0$)	0.0035
$\pi^+\pi^-\pi^+\pi^-\pi^0$	0.196		
$\pi^+\pi^-\pi^+\pi^-\pi^0\pi^0$	0.166		
$\pi^+\pi^-\pi^+\pi^-\pi^0\pi^0\pi^0$	0.042		
$\pi^+\pi^-\pi^+\pi^-\pi^+\pi^-$	0.021		
$\pi^+\pi^-\pi^+\pi^-\pi^+\pi^-\pi^0$	0.019		

Table 2.2: Branching ratios of pion final states from antiproton-proton [106] and antiproton-neutron [109] annihilation at rest (the k denotes the grouping together of multiple π^0 channels). Note that there is a contribution of $\sim 2\%$ due to kaons, which are ignored here.

Once produced, the pions (that escape the nucleus) travel away from the annihilation location. The charged pions have lifetimes ($\sim 10^{-8}$ s) that are long enough that they can pass through and exit the apparatus. The neutral pions, however, will quickly ($\sim 10^{-16}$ s) decay into γ -rays, which in turn, can often produce e^-e^+ pairs in the apparatus material (neutral pions can also undergo Dalitz decay, $\pi^0 \rightarrow e^- + e^+ + \gamma$, such that the electron, positron, and γ -ray all originate from the annihilation location). The charged pions and e^-e^+ pairs can be detected as they travel away from the annihilation location (to be discussed further in Sec. 2.4 and Chapter 4), and they provide a characteristic annihilation signal.

2.4 Semiconductor detectors

Having dealt with the processes around the annihilation of antihydrogen, and specifically detailed the range of high-energy charged and neutral particles produces therein, the efficient detection of these annihilation events is now discussed.

Charged particles lose energy as they pass through matter (see Appx. A). This energy loss comes from electronic (e. g. excitation and/or ionization of atomic electrons), as well as radiative (e. g. bremsstrahlung, Compton scattering, pair-production) interactions [9]. In some materials the energy deposited by the passing charged particle can be converted to an electronic signal and quantified. For example, semiconductor (such as silicon) detectors operate by reverse-biasing a p-n junction (see Appx. B), which creates a region of the silicon that is depleted of mobile charge carriers. Thus, charged particles passing through the depletion region generate electron-hole pairs. The liberated electrons and holes drift in the electric field (in opposite directions) to signal collection electrodes where they take the form of an electrical signal that is read out by the signal processing electronics. Minimum Ionizing Particles (Appx. A), for example, will on average deposit $388 \text{ eV}/\mu\text{m}$ while passing through through silicon, generating about a hundred electron-hole pairs per micron [110].

2.4.1 Double-sided silicon microstrip detectors

Segmenting the charge collection electrodes can make silicon detectors position-sensitive. That is, by using an array of strips or segments, the location of passing charged particles can be identified as the strips or segments that register electric signal. Popular choices of segmentation include long, thin strips [111] or pixels [112]. The utility of the electron-hole pair generation can be extended by collecting both types of charge carrier with collection electrodes on opposite sides of the applied bias [113]. Moreover, the passage of a charged particle can be localized in two dimensions by making the collection electrode




Figure 2.5: A simplified diagram of a double-sided silicon microstrip detector. The p-side strips on the top side of the silicon wafer run orthogonally to the n-side strips on the opposite side of the silicon.

strips run in orthogonal directions (the third dimension can be inferred with knowledge on the detector position and geometry). The strip spacing, or ‘pitch’, can range between 50-1000 μm depending on the application [114].

Figure 2.5 shows an example of a double-sided microstrip detector consisting of n-type bulk silicon with p-type strips on one side of the wafer (‘p-side’), and orthogonal AC-coupled strips on the opposite side (‘n-side’). In this case, the p-n junction is established at the point of contact between the p-side strips and the n-type bulk. The n-side strips, however, are externally AC-coupled in order to separate their signal from the DC bias on the silicon backplane. In this way, signal from the electrons and holes can be collected simultaneously. Overall, double-sided silicon detectors provide fast and stable particle detection with excellent spatial resolution (to be discussed further in Chapter 4).

2.5 Summary

In this chapter, the basic theoretical concepts relevant to antihydrogen formation and trapping, the goal of this project, were presented. After first discussing the mechanisms that can be used to form antihydrogen atoms, the theory of containing antimatter (both charged (Penning-Malmberg traps) and neutral (magnetic bottles)) was presented. The

second half of this chapter dealt with detection – both the annihilation processes used to produce the antihydrogen signature and the technology needed to detect this signature are discussed. In the next chapter, the theory presented here will be employed in the description of the ALPHA apparatus, which has been designed to produce and confine antihydrogen with these concepts in mind.

Chapter 3

The ALPHA apparatus & procedures

The ALPHA apparatus combines a variety of components and experimental techniques in order to produce, trap, and study antihydrogen. The physics involved in these components and techniques is interesting in its own right – but here it is important to give a general description of the ALPHA experiment and antihydrogen trapping scheme.

This chapter provides the information necessary to understand the operation of the ALPHA apparatus and the experiments conducted therein. First, an overview is given to familiarize the reader with the overarching functioning and interplay of the experimental components. Next, the various components and methods are described in detail. Topics include: the diagnostic and detection devices, production and accumulation of antiparticles, trapping and manipulation of charged particles and antiparticles, and the magnetic neutral-atom trap. Finally, the interrelation of these elements is tied together and an overview of the entire trapping experiment is given at the end of the chapter. Antihydrogen detection systems, core to this dissertation, are dealt with independently in Chapter 4.




Figure 3.1: Schematic diagram of the ALPHA apparatus (Image from the ALPHA collaboration). Antiprotons delivered from the AD are depicted as the \bar{p} on the left. Several of the important components are identified and labeled. The length scale of the figure is shown at the bottom right.

3.1 Infrastructure overview

The ALPHA apparatus (Fig. 3.1), involves the interplay between a number of devices, including, but not limited to:

Antiproton Decelerator (Sec. 3.3.1): the antiprotons used in the ALPHA experiment are produced, initially decelerated, and delivered by the Antiproton Decelerator (AD). The AD is a specialized machine for low-energy antiproton physics located at the CERN facility in Geneva, Switzerland. This machine provides experiments with about 10^7 5.3 MeV antiprotons every 100 s.

Positron accumulator (Sec. 3.3.2): the ALPHA apparatus includes a Surko-type device [115, 116] for the accumulation of positrons from a radioactive source. This device combines a frozen neon moderator and a nitrogen buffer-gas cooling chamber to provide continuous accumulation of positrons, which are confined in a Penning-Malmberg trap. When accumulation is complete, the positrons are transferred to the mixing region of the apparatus (see Fig. 3.1).

Penning-Malmberg trap for charged particles (Sec. 2.2.1): this trap combines electric and magnetic fields to confine and manipulate charged particles and plasmas. Positrons and antiprotons can be confined simultaneously and mixed together to form antihydrogen (Sec. 3.4.5).

Magnetic neutral-atom trap (Sec. 3.5): superimposed over the Penning-Malmberg trap is an array of superconducting magnets. These magnets create a minimum-B trap to confine antihydrogen atoms. This trap is composed of two mirror coils and an octupole magnet, and when combined with the axial magnetic field from the Penning-Malmberg trap, a local three-dimensional magnetic minimum is established along the trap axis.




Figure 3.2: Flowchart of the interrelation between the major components of the ALPHA apparatus.

Diagnostic and detection devices (Sec. 3.2): there are a number of particle detectors and diagnostic devices attached to the ALPHA apparatus. These devices provide important information about the conditions and parameters of the particles within the apparatus, for example, the temperature, dimensions and particle density for a plasma under manipulation. A vertexing silicon detector, used to detect antihydrogen annihilation products, is core to this thesis, and discussed in detail in Chapter 4.

A simplified description of the operation of this apparatus is as follows (and shown as Fig. 3.2): antiprotons are produced, decelerated, and delivered by the Antiproton Decelerator; in parallel, positrons are collected in the positron accumulator. Both antiparticle species are confined and manipulated within a long Penning-Malmberg trap.




Figure 3.3: Cartoon schematic of the MCP/phosphor/CCD camera imaging system. The MCP is shown as the grey chevron, the phosphor screen is represented as the red rectangle, the mirror is the blue rectangle, and the CCD camera is the yellow rectangle. The red line is illustrative of electron multiplication in an MCP channel.

The positrons and antiprotons are mixed together within the magnetic neutral-atom trap to generate antihydrogen atoms and atoms with energies below the trap-depth of the neutral trap are magnetically confined. The neutral-trap is quickly de-energized, and trapped antihydrogen atoms can then be detected by the silicon detector.

3.2 Diagnostic and particle detection devices

Knowledge of the particle and plasma distributions is an experimental necessity. Without accurate measurements of the conditions and parameters involved in the experiment it would be very difficult to make any progress. As such, the ALPHA apparatus includes a number of diagnostic and detection devices. Several of these devices are listed below:

MCP/phosphor/CCD camera imaging system. Figure 3.3 gives a diagram of the setup of the imaging system used to measure the profiles of plasmas and particle clouds directed onto the microchannel plate (MCP). The MCP consists of two plates of

highly resistive material, with an array of channel paths through the plates. A high-voltage bias is placed across the plates to stimulate the emission of secondary electrons when an incoming particle collides with the plate wall. The emitted electrons continue to collide with the channel walls, and the resulting electron cascade amplifies the original signal. Upon exiting the MCP, the electrons are directed onto a phosphor screen. The electrons excite the phosphor atoms, which in-turn emit visible light. A 45 degree mirror directs the light to a charge coupled device (CCD) camera, which records the image.

The active area of the MCP is a circle with a diameter of 41.5 mm. The channels are arranged in a hexagonal array with 15 μm spacing (each channel has a 12 μm diameter). With a maximum applied voltage of 1 kV, the MCP provides a gain of 8×10^5 . However, because of their mass difference electrons/positrons have a different gain behaviour than antiprotons. The gain and collected charge calibration for each species is given in [117].

Faraday cup. This detector consists simply of a piece of aluminum foil near the vacuum window connecting the ALPHA apparatus to the AD beamline. The foil is located within the trap vacuum, such that particles and plasmas can be directed onto the foil surface. The voltage induced by incident electrons or positrons can be amplified and recorded. However, more than about 10^6 particles are required for the signal to be discerned from the measurement noise. Antiprotons are not good candidates for this device, as they can fragment aluminum nuclei upon annihilation, resulting in several charged fragments and confusing the measurement.

Silicon vertex detector. The main system for the identification of antihydrogen is the ALPHA silicon vertex detector. This detector consists of 60 silicon modules arranged in three coaxial layers around the neutral-atom trap. The silicon modules

detect the passage of ionizing radiation through the collection of electron-hole pairs, liberated by passing high-velocity charged particles, within a charge-depleted region of the silicon wafer (Appx. B). The tiered geometry allows for the reconstruction of the trajectories of charged particles through the detector. In the case of annihilations within the trap system, the trajectories of the annihilation products can be combined to find their origin point. The silicon detector is discussed in detail in Chapter 4 and the reconstruction is described in Chapter 5.

Scintillating detectors. Like the silicon detector, the scintillation detectors are sensitive to the passage of ionizing radiation. The plastic scintillating material is excited by the deposition of energy by the passing ionizing radiation. Upon de-excitation, light is produced and collected into a Photo-Multiplier Tube (PMT). The PMT first uses a photocathode to convert the scintillation light to photo-electrons, which are then cascade-multiplied to produce a measurable electric signal. ALPHA has twelve large area scintillation detectors, arranged in sets of two, placed around the exterior of the apparatus. The pair arrangement allows for coincident detection, which reduces readout noise.

CsI detectors. ALPHA has several small CsI detectors that are sensitive to the 511 keV γ -rays resulting from positron annihilation. The CsI crystals operate as scintillating detectors and are coupled to photodiodes to measure the light yield. These detectors are placed near the positron accumulator, just outside the cryostat (next to the silicon detector), and at the transfer section between the positron accumulator and main apparatus.

3.2.1 Charged particle temperature measurement

The temperature of a particle cloud or plasma is an important parameter, as, for example, the antihydrogen formation processes scale with temperature. As such, a measure of the plasma's temperature is extremely valuable.

Following [118], ALPHA can measure the parallel temperature of the particle cloud or plasma by slowly lowering the confining potential and fitting the distribution of escaping particles. Since the particles are released on-axis, their energies should follow a one-dimensional Maxwell-Boltzmann distribution:

$$f(E_{\parallel}) \propto e^{\left(-\frac{E_{\parallel}}{k_B T}\right)}, \quad (3.1)$$

where E_{\parallel} is the axial or parallel energy, k_B is Boltzmann's constant, and T is the temperature. The temperature can then be found through the relation,

$$\frac{d \ln(f(E_{\parallel}))}{dE_{\parallel}} \simeq -\frac{1}{k_B T}. \quad (3.2)$$

This derivative can be measured directly, since the parallel energy is known from the confining voltage and the number of escaped particles detected reflects the distribution. Figure 3.4 shows several measurements of antiproton temperature, where the number of escaping antiprotons is plotted as a function of confining potential. The exponential sections of these distributions are fitted using Eq. 3.2, with temperatures between 9 K and 1040 K being observed to date. The non-exponential behaviour seen for the low-energy part of curves is due to the effect of the space charge of the remaining particles as they are released. Likewise, correction factors can be applied to the temperatures to account for adiabatic cooling of the particles as the confining well changes, and also corrections to the potential height due to the space charge. More details about these corrections, and the calculations of the confining potentials, are described in [119].




Figure 3.4: The number of antiprotons released from the confining potential as a function of the on-axis well depth (reproduced with permission from [64]). The antiprotons losses are integrated and corrected for the 25% scintillator detection efficiency. The black lines show the fitted region, and the temperatures found are: A: 1040, B: 325, C: 57, D: 23, E: 19, and F: 9 K. The inset shows a magnification for the low-energy distributions.

3.3 Production and accumulation of antiparticles

Positrons and antiprotons are both required to synthesize antihydrogen. These antiparticles are rare, and in the case of antiprotons require a large facility to produce. This section describes how positrons and antiprotons are produced, delivered to, and accumulated by the ALPHA apparatus.

3.3.1 Antiproton Decelerator (AD)

Antiprotons, \bar{p} , can be produced through the following process:

$$p + p \rightarrow p + p + p + \bar{p}, \quad (3.3)$$

where the colliding protons, p , have energy above the antiproton production threshold ($E > 7m_p c^2 \approx 6.6$ GeV in the lab frame, where m_p is the proton mass [120]). Conventionally, high-energy protons are accelerated and directed into a stationary target, where antiprotons (along with protons and other secondary products) are released. This method is used to supply CERN's Antiproton Decelerator¹ (AD) with antiprotons. First, a pulse of 10^{13} protons, with momentum of 26 GeV/c, is supplied by the Proton Synchrotron (PS). These protons are directed to a target near the AD, consisting of a thin iridium rod embedded in graphite. From this pulse, about 5×10^7 antiprotons are produced (along with many other unwanted particles, largely pions and muons). From the pool of produced particles, about 3×10^7 antiprotons are focused in a magnetic horn, and selected according to their charge-to-mass ratio and momentum (in this case, antiprotons with momentum of 3.57 GeV/c are selected with a momentum spread of $\Delta p/p \sim 1.5\%$, corresponding to the peak antiproton yield for this target and incident proton momentum [120]). Finally, the collected antiprotons are injected into the AD storage ring.

The AD is a specialized machine designed to decelerate and cool antiprotons from

¹AD website: <http://psdoc.web.cern.ch/PSdoc/acc/ad/>

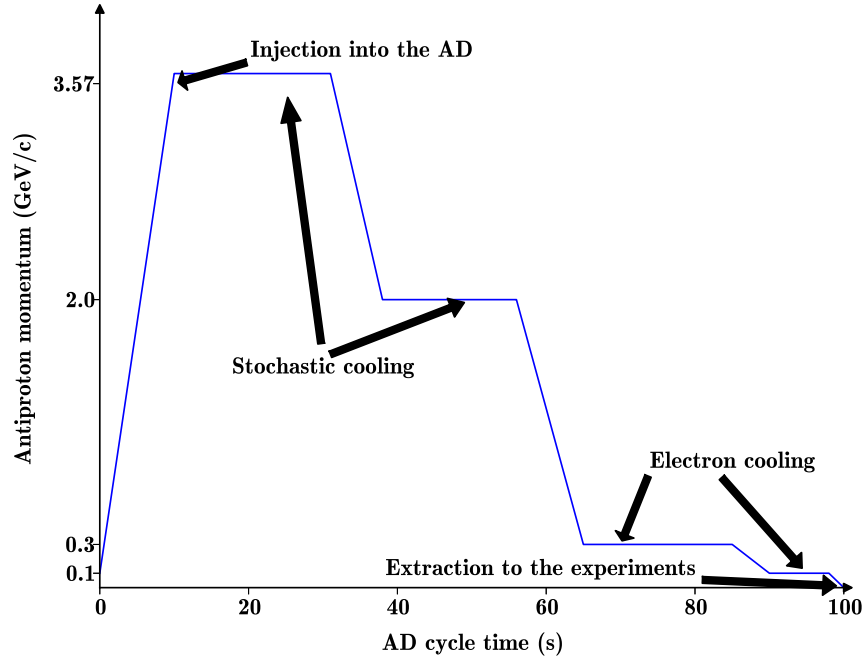


Figure 3.5: Antiproton momentum during over the course of a deceleration cycle in the AD, including the various stages of particle cooling. Figure adapted from [121].

the GeV to the MeV energy range. This is accomplished through several alternating stages of rf-deceleration and beam cooling (Fig. 3.5 and 3.6). For rf-deceleration, an oscillating electric field, which opposes the direction of motion, is applied as the circulating antiproton pulse passes through a special rf-driven cavity. Beam cooling is necessary to compensate for the adiabatic blow-up of the beam during deceleration [122]. Due to the conservation of phase-space density, the decrease in momentum during deceleration expands the conjugate position spread of the particles. With beam cooling, the antiproton beam can be reduced in size and energy spread. By reducing the beam size, stability is added to the circulating beam. Likewise, an increase in the phase-space density allows for more precise manipulation and control of the beam.

Two beam cooling techniques are used by the AD: stochastic [123,124] and electron cooling [125,126]. Stochastic cooling involves a series of corrections to the mean position




Figure 3.6: Cartoon diagram of the AD system. The (simplified) beamline is shown as the black curves, the stochastic cooling transmission line is shown in grey, the electron cooler in blue, and the ALPHA experimental zone in red.

and momentum of the circulating antiprotons. At each turn, a pickup measures the divergence of the beam from its ideal path, and a correction is applied, using an electric field, further down the beamline². Since the pickup measures the mean deviation of the antiproton ensemble, each antiproton experiences a coherent correction, due to its own deviation, along with an incoherent correction from the other antiprotons in the bunch. By carefully choosing the gain of the electronics and reducing the system noise, the effect of the incoherent correction (which will ‘heat’ the bunch) can be minimized and, over many turns, the coherent corrections will combine to reduce the beam size and energy spread. Conversely, electron beam cooling relies on the interaction between the antiproton beam and a cold electron beam to reduce the antiproton energy spread. Cold electrons are injected into the storage ring, matching the velocity of the antiprotons.

²As shown in Fig. 3.6, a transmission line diagonal to the storage ring is used for the stochastic cooling system. This shortcut allows the beam correction to be applied to the same particles that were measured at the pickup.

In a reference frame moving with the particles, the higher energy antiprotons transfer energy to the electrons via Columbic collisions. The electrons are then removed from the antiprotons, taking with them the excess energy transferred from the antiprotons. In this way, the electrons act as an external heat sink to the higher-energy antiprotons.

The methods of stochastic and electron cooling are complementary – especially in the case of the deceleration of antiprotons in the AD. Electron cooling is difficult for high-energy antiprotons, as the power needed to produce a velocity-matching electron current becomes difficult to maintain. Moreover, stochastic cooling is more efficient on higher-energy beams, as the electron cooling rate is strongly dependent on the relative velocity spread. For these reasons, the AD uses stochastic cooling for the 3.57 and 2 GeV/c antiproton momentum stages, and electron cooling for the 0.3 and 0.1 GeV/c stages (see Fig. 3.5). Using these methods, bunches of 10^7 antiprotons at energies of 5.3 MeV are delivered to the AD experiments every 100 s.

3.3.2 Positron accumulator

In stark contrast to the production and collection of antiprotons, positrons can be accumulated in much larger numbers at much lower energies. As well, particle accelerators are not required to produce positrons, which are emitted in the spontaneous beta-plus decay of many radioisotopes. For example, the beta-plus decay channel of sodium-22 proceeds as follows:



where a proton in the sodium nucleus converts to a neutron and emits a β^+ particle (positron) along with a electron-neutrino, ν_e . A 1.274 MeV gamma ray, γ , follows shortly after (3.7 ps) as the excited Ne^* nucleus undergoes an isomeric transition to the ground




Figure 3.7: Overview of the positron accumulator (figure adapted from [103, 127]). The accumulator apparatus is shown as a). The inset, b), shows a diagram of the positron buffer gas cooling within the long electrode stack, where the red curve shows the electric potential as a function of position, and the blue axis describes the pressure gradient.

state. The ^{22}Na radioisotope is a common choice of positron source (and used in the ALPHA experiment), as the β^+ branching ratio is $\sim 90\%$, and the half-life is about 2.6 years [100], allowing one source for several annual seasons of experimental operation.

The ALPHA apparatus includes a dedicated Surko-type positron accumulator [115, 116] (Fig. 3.7), a device which continuously collects, cools, and stores positrons emitted from a ^{22}Na source. The positron accumulator works by directing the emitted positrons through a solid moderator and into a neutral buffer gas, before finally confining the positrons in a Penning-Malmberg type trap (see Sec. 2.2.1).

The first stage in the accumulation process is the moderation of the positrons, which involves passing the positrons through a layer of solid neon deposited directly on the radioactive source. While passing through the neon, the positrons quickly thermalize and diffuse through the material. A small fraction ($\sim 0.4\%$) of the incoming positrons will reach the surface and be released into the vacuum, providing as much as 5×10^6 positrons per second [72]. The positrons that emerge from the moderator will do so in

the eV range and will follow external magnetic field lines towards the positron trapping region.

The trapping region contains a series of electrodes, which provide the axial electrostatic potential, along with a 0.15 T axial magnetic field which acts to radially confine the positrons (this configuration is known as a Penning-Malmberg trap, as discussed in Sec. 2.2.1). With such a low magnetic field magnitude, positrons have a long synchrotron cooling constant, necessitating a different mechanism for cooling the positrons into the trap. Without further cooling, positrons with enough energy to enter the electrostatic trap will have enough energy to leave. Here, cooling is provided by a nitrogen buffer gas, where the positrons lose kinetic energy by collisionally exciting the molecular degrees of freedom in the nitrogen gas. About 20% of the incoming positrons will cool into the trap, where they continue to cool via further collisions with the buffer gas.

The trap electrodes vary in diameter along the trap axis, such that a pressure gradient exists in the buffer gas across the electrodes. At the same time, a step-wise decreasing electrostatic potential is applied, where the electrodes are biased such that the low gas-pressure region corresponds to the bottom of the electrostatic well. As a consequence, the cooling positrons will eventually come to reside in the low-pressure region. Surrounding the low-pressure region is a special six-way segmented electrode, used to provide an azimuthally varying and oscillating electric field. This field applies a torque to the positron plasma, which acts to compresses the plasma radially. This “rotating wall” technique (to be discussed further in Sec. 3.4.3) increases the positron density, but also heats the plasma. Fortunately, the nitrogen buffer gas continues to provide cooling while the rotating wall compression is applied.

The accumulator is operated in a steady-state mode, where it continuously collects the positrons emitted by the ^{22}Na source. The positrons can then be ballistically transferred into the antiproton trap. The positron transfer involves pumping out the nitrogen gas

before opening the separation between the positron accumulator antiproton trap. A pulsed magnet is then fired in order to force the positrons across the low radial magnetic field region between the accumulator and the antiproton trap. This transfer has an efficiency of about 50%, with as many as about 10^8 positrons surviving the journey across the apparatus. In total, the accumulation cycle takes about 300 s and typically provides about 7×10^7 positrons/pulse to the experiment. This preaccumulation and ballistic-transfer method (along with the stacking of several accumulation cycles) can give the highest density positron plasmas currently available [57].

3.4 Charged particle manipulation and trapping

In order to synthesize and study antihydrogen, precise control over the constituent antiparticles is required. Critically, because any interaction with the surrounding matter is destructive, the positrons and antiprotons must be isolated at all times. Likewise, antiparticle containment must be sustained on timescales matching the duration of the experiments being performed (often, these experiments can exceed 800 s in duration).

This section will describe the elements of apparatus and methods used to confine and manipulate electrons, positrons, and antiprotons in the ALPHA experiment. First, the specific Penning-Malmberg trap for charged particles used in ALPHA is briefly described. Next, the catching, manipulation, and cooling of antiprotons (as well as positrons, where appropriate) is outlined. Finally, several techniques used to control the temperature and density of the antiparticle plasmas, such as rotating wall compression and evaporative cooling, are described.

3.4.1 Penning-Malmberg trap

ALPHA uses a Penning-Malmberg trap (Sec. 2.2.1) to confine and manipulate charged (anti)particles. These traps rely on a strong axial magnetic field for radial confinement of

the charged particles, with simultaneous axial confinement using applied electric potentials. In the ALPHA apparatus, the large axial magnetic field is provided by an external superconducting solenoid (Fig. 3.8) and the electric potentials are provided by a series of cylindrical electrodes (Fig. 3.9). Precise axial control of the charged particles is achieved by varying the bias voltages applied to the electrodes. Particles can be transported axially, all while maintaining radial confinement, by coordinating the potential changes in several electrodes.

The static external magnetic field is typically held at 1 T, while the electric potentials can vary between -140 V to +140 V (there are also special high-voltage electrodes that can sustain kV potentials, as will be discussed in Sec. 3.4.2).

3.4.2 Catching and cooling of antiprotons

For stable confinement of antiprotons in a Penning-Malmberg trap, the longitudinal kinetic energy of the antiprotons must be less than the axial electric potential. Moreover, cold (low-energy) antiprotons are needed to produce antihydrogen that can be confined by the magnetic neutral-atom trap. For these reason, the antiprotons received from the AD must be cooled well below the AD extraction energy before either the charged antiparticles or neutral atoms can be trapped.

The antiprotons delivered from the AD (Sec. 3.3.1) arrive at the ALPHA apparatus in a several-hundred-nanosecond pulse with a mean kinetic energy of 5.3 MeV. While the AD significantly reduces the momentum of the antiprotons (from momenta 3.57 GeV/c to 100 MeV/c), further deceleration is needed before they can be caught and confined. An initial reduction in kinetic energy is accomplished as the incoming antiprotons enter the ALPHA apparatus through a thin vacuum window. A fraction of antiparticles inelastically scatter in the vacuum window and exit with degraded kinetic energy. Antiprotons with sufficiently reduced kinetic energy can be dynamically ‘caught’ between




Figure 3.8: The ALPHA superconducting solenoid magnet (Image from the ALPHA collaboration). A large magnetic field is produced along the bore of the magnet. The cryostat (Fig. 3.1) is located within the solenoid.




Figure 3.9: The electrode stack used to provide the electric potential as part of the principal ALPHA Penning-Malmberg trap (Image from the ALPHA collaboration). Using this series of electrodes, a variety of bias potentials can be established along the trap axis. The two narrow electrodes at either ends, between the white electrical insulating material, are used to provide the high-voltage (3-5 kV) potentials during antiproton catching (Sec. 3.4.2).

two fast-switching, high-voltage electrodes (which can be seen in Fig. 3.9).

The catching sequence is outlined in Fig. 3.10: the antiprotons enter the ALPHA apparatus and, with the downstream high-voltage barrier engaged, antiprotons with kinetic energy below ~ 3 keV will not be able to overcome the electric potential barrier and will rebound, travelling back towards the AD entrance. Another electrode with the same applied high-voltage is engaged before the antiprotons can leave the trap. The antiprotons will then be confined between the two potential barriers. However, another mechanism is required to cool the antiprotons; were the high-voltage potential barriers to be disengaged without such a cooling mechanism, the antiprotons would still have kinetic energies of about 3 keV, and would quickly leave the trap. Fortunately, charged particles self-cool in magnetic fields. Larmor's equation gives change in energy as radiated by a (non-relativistic) charged particle due to its acceleration, \vec{a} [128]:

$$\frac{dE}{dt} = -\frac{q^2}{6\pi\epsilon_0 c^3} |\vec{a}|^2. \quad (3.6)$$

For an (anti)particle in a Penning-Malmberg trap, the cyclotron motion dominates the cooling, such that the other modes can be ignored and the acceleration becomes $\vec{a} \simeq \vec{\omega}_c \times \vec{v}$. Following Brown and Gabrielse [86] (who also provide a calculation for the full particle motion), the change in energy in this system can be rewritten, using the fact that the system energy is entirely kinetic, $E = 1/2mv^2$, as

$$\frac{dE}{dt} = -\gamma_c E, \quad (3.7)$$

where γ_c gives the cyclotron cooling rate,

$$\gamma_c = \frac{q^2 \omega_c^2}{3\pi\epsilon_0 m c^3} = \frac{q^4 B_0^2}{3\pi\epsilon_0 m^3 c^3}. \quad (3.8)$$

In the last step, Eq. 2.22 was used to replace ω_c with the relevant particle and field parameters.




Figure 3.10: Cartoon illustration of the antiproton catching sequence. a) Antiprotons are delivered from the AD (left) into the catching trap, where electrons are already pre-loaded into a cooling well and the far HV-electrode is already engaged; b) antiprotons with kinetic energies < 3 keV rebound off the HV potential (higher energy antiprotons are lost); c) the second HV barrier is quickly engaged, confining the antiprotons between the two barriers; d) the antiprotons cool, via Coulomb collisions with the electrons, into the low-voltage electron well and the HV barriers are disengaged.

The solution to Eq. 3.7, $E(t) = E_0 e^{-\gamma t}$ where E_0 is the initial kinetic energy, shows that charged particles exponentially self-cool in a magnetic field. In the 3 T solenoid field used during capture, the antiproton cyclotron cooling constant is $1/\gamma_{c,\bar{p}} \sim 2 \times 10^9$ s, which is much too long to be experimentally useful. However, since $\gamma_{c,\bar{p}}/\gamma_{c,e} = (m_e/m_{\bar{p}})^3$, the electron self-cooling constant in the same field is $1/\gamma_{c,e} = 1/\gamma_{c,\bar{p}}(m_e/m_{\bar{p}})^3 \sim 0.3$ s, which motivates the use of electrons as a cooling medium for the antiprotons. Indeed, as shown in Fig. 3.10, electrons are pre-loaded within the catching region, where they can interact with the caught antiprotons. The temperature evolution for combined electrons and antiprotons can then be modelled as [129],

$$\frac{dT_{\bar{p}}}{dt} = -\frac{T_{\bar{p}} - T_e}{\tau_c}, \quad (3.9)$$

$$\frac{dT_e}{dt} = \frac{n_{\bar{p}}}{n_e} \frac{(T_{\bar{p}} - T_e)}{\tau_c} - \frac{(T_e - T_t)}{\tau_e}, \quad (3.10)$$

where T_e , $T_{\bar{p}}$, and T_t are the temperatures of the electrons, antiprotons, and the surrounding environment, respectively; n_e and $n_{\bar{p}}$ are the electron and antiproton densities; τ_e is the electron cooling constant ($1/\gamma_{c,e}$), and τ_c is the cooling constant from the overlap between electrons and antiprotons, given by [130]

$$\tau_c = \frac{3m_e m_{\bar{p}} c^3}{8\sqrt{2\pi} n_e e^4 \ln \Lambda} \left(\frac{kT_{\bar{p}}}{m_{\bar{p}} c^2} + \frac{kT_e}{m_e c^2} \right)^{3/2}, \quad (3.11)$$

where Λ is the Coulomb logarithm, which represents the cut-off in the integration over the impact parameter. This model indicates that 10^4 antiprotons, with initial energies in the keV range, can be cooled to the eV range in several hundred milliseconds using an electron cloud with a density of $\sim 10^8$ cm⁻³ [103]. In the absence of any external heating, the electrons would cool to the ambient temperature (4.2 K liquid helium cryogenic bath), and the antiprotons would follow. However, the ALPHA traps are coupled to the outside environment at both ends of the electrode stack, as well as through the electrode cables. Although an extensive effort has been made to heat-sink the cables to

the cryogenic bath, and minimize the heating sources at the trap ends, the thermalized antiproton temperature is measured to be 358 ± 55 K [53], which is well above the cryogenic temperature. The complementary technique of evaporative cooling of antiprotons can be used to reduce the antiproton temperature, and will be discussed in Sec. 3.4.4.

During normal operation, the external solenoidal magnet (Fig. 3.8) provides a static 1 T axial magnetic field. However, as shown in Fig. 3.11, the number of antiprotons caught increases proportional to the magnetic field. The increase in magnetic field allows for antiprotons with higher transverse momentum to be captured, and increases the electron cyclotron cooling rate quadratically. However, as will be discussed in Sec. 3.5, a higher solenoidal field will reduce the overall neutral-atom trap depth. As a compromise, ALPHA has included a 2 T internal solenoid that can be easily ramped up and down – providing the additional catching efficiency when energized, and disengaged otherwise. Figure 3.12 shows the axial magnetic field for the external Penning-Malmberg trap (red), and with the addition of the inner solenoid magnet (blue). The limit of 3 T is set such that the magnetic forces do not cause undue stress on the apparatus and magnet windings when all of the superconducting magnets are energized.

3.4.3 Rotating wall compression and electron kick-out

There are several advantages to having precise control over the radial size of confined particle clouds and plasmas: by keeping the particle number constant, the plasma density can be manipulated; also, by controlling the plasma radius, the overlap during particle mixing can be optimized for maximum antihydrogen formation efficiency.

Following from Eq. 2.26, the mean squared plasma radius is proportional to the canonical angular momentum, $P_\theta \simeq (-e)B/2c \langle r^2 \rangle$ (for electrons and antiprotons). It follows that

$$\frac{dP_\theta}{dt} > 0 \Rightarrow \frac{d\langle r^2 \rangle}{dt} < 0, \quad (3.12)$$




Figure 3.11: Antiproton catching efficiency as a function of solenoid magnetic field (reproduced with permission from [49]). The values have been normalized to the 3 T value.

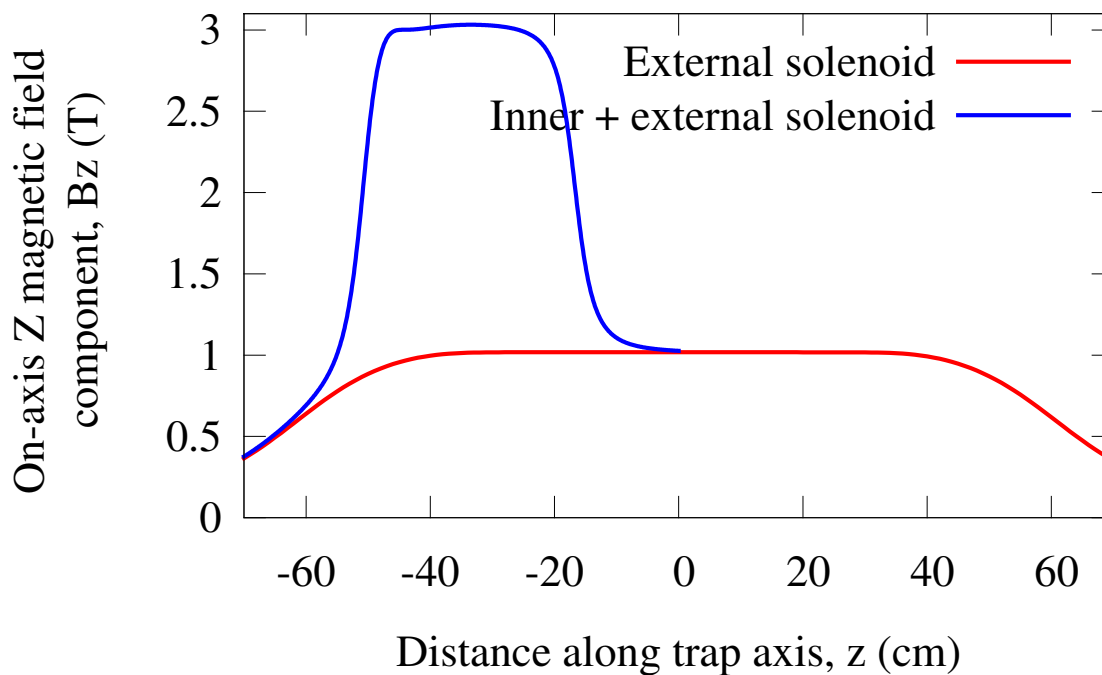


Figure 3.12: The on-axis z component of the magnetic field with (red) only the external solenoid energized, and (blue) both the external and inner solenoids energized.

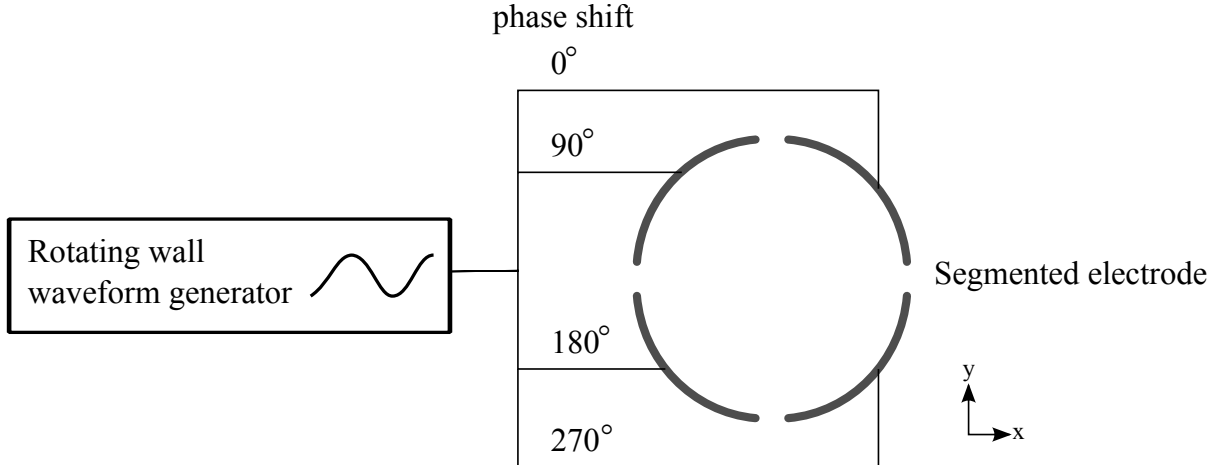


Figure 3.13: Typical setup for the Rotating Wall technique. An sinusoidal voltage is provided by a waveform generator. This voltage is then split, and the offshoots phase-shifted, before finally being applied to a special segmented electrode.

that is, the application of a positive torque (relative to the sign of the charge), $\mathcal{T} = dP_\theta/dt$, will have the effect of decreasing the plasma radius. It also follows that the application of a negative torque (as with drag due to collisions with neutrals and trap asymmetries) will cause the plasma to expand.

Figure 3.13 illustrates a system for generating such torques. An electrode is split into several sections and an oscillating potential (typically dipole in spatial distribution) is applied, with each of the sections receiving a different phase. In this way, a rotating radial electric field is generated and couples to the radial motion of the plasma. This is known as the ‘Rotating Wall’ method and is a popular nonneutral plasma technique for controlling the radial plasma parameters [94].

In ALPHA, antiproton clouds are radially compressed sympathetically using the Rotating Wall technique on the combined electron-antiproton plasma. The electrons and antiprotons are both brought into the sectored electrode and the rotating potential is applied to both. The collisions between the antiprotons and the electrons provides a mechanism for the cooling of antiprotons needed for such a compression (as the radial compression amounts to a compression of phase space density, and Liouville’s theorem




Figure 3.14: Electron and antiproton compression using the Rotating Wall technique (reproduced with permission from [66]). The images of antiproton and electron radial distributions are generated with the MCP system and are provided along with integrated radial particle density profiles (the location of the MCP is shown in Fig. 3.1). Various Rotating Wall application times, as indicated by the figure, are also shown.

must be circumvented [122]). Antiprotons transfer energy to the electrons through collisions, with the electrons then efficiently radiating away the additional energy in the strong magnetic field (Sec. 3.4.2). Figure 3.14 shows MCP images and their corresponding radial profiles for various Rotating Wall application times, for both electrons and antiprotons. Using Rotating Wall compression, the antiproton density can be increased by as much as a factor of ten, such that antiproton clouds with radii of 0.29 mm are produced [66].

The electron-antiproton mixture is ideal for the manipulation of antiprotons within the traps, as the electrons will keep the antiprotons cold. However, a plasma containing electrons is non-ideal when it comes to mixing with positrons: the electrons will quickly cool into the side-wells of the nested potential (Sec. 3.4.5) and drag the antiprotons into the side-wells with them (in the same way that the antiprotons cool into the electron well in Sec. 3.4.2). If the antiprotons cool into the side-wells the production of antihydrogen will be prohibited, as the antiprotons will be axially separated from, and therefore will not interact with, the positron plasma³ (this is only relevant for the mixing scheme, described in Sec. 3.4.5, where the electron-antiproton mixture is injected directly into the positron plasma). Finally, the presence of electrons in the plasma could result in the depletion of positrons due to the electron-positron annihilations. It is therefore necessary to remove the electrons before the antiprotons can be mixed with the positrons.

The large mass difference between the electrons and antiprotons can be exploited as an effective means to remove the electrons from the electron-antiproton mixture. The thermal velocity, v , of a massive particle in the longitudinal direction (as the electrons

³Some mixing schemes start with the antiprotons at the bottom of a side-well (to be discussed in Sec. 3.4.5). The antiprotons are then rf-driven into the positron plasma to produce antihydrogen. Here again, the presence of electrons would be counter-productive, as the antiprotons would be cooled back into the side-well and would need to be driven back into the positron plasma.




Figure 3.15: Cartoon illustration of the electron kick-out sequence. a) A mixture of electrons and antiprotons are confined in a shallow well (see Sec. 3.4.2); b) the confining potentials are raised above ground; c) one of the confining potentials is lowered, allowing the electrons to quickly escape (being more massive, the antiprotons do not escape as quickly); d) the confining potential is reestablished after the electrons are gone, but before the antiprotons can leave.

will escape along the axial magnetic field lines) is given by

$$v(T) = \sqrt{\frac{k_B T}{m}}, \quad (3.13)$$

where T is the temperature, k_B is Boltzmann's constant, and m is the particle mass. With a measured antiproton temperature of 358 K (Sec. 3.4.2), the mean thermal velocity of the antiprotons is $\sim 1.7 \times 10^3$ m/s. Since the thermal velocity scales as $(\text{mass})^{-1/2}$, the electrons have a mean thermal velocity $42\times$ larger than the antiprotons, or $\sim 7.4 \times 10^4$ m/s. With such discrepant velocities, the antiprotons and electrons will separate almost immediately when not confined by a trapping field. For example, it would take the electrons $\sim 2 \mu\text{s}$ to cross a typical electrode length of 2 cm, while the antiprotons would take 0.1 ms to traverse the same distance.

A simple procedure can be then employed to remove the electrons:

1. The potential well containing both antiprotons and electrons is changed to a potential where one side of the well is separated from the ground potential by a single electrode (Fig. 3.15a-b)).
2. The potential applied to the end electrode is quickly lowered to ground (Fig. 3.15c)).
3. With the mixed plasma axially unconfined, most of the electrons leave while the vast majority of the antiprotons remain in the same region.
4. The confining potential is reapplied before the antiprotons have time to exit (Fig. 3.15d)).

This procedure expels virtually all of electrons, and an all-but-pure antiproton cloud remains when the confining potential is reinstated. In practice, not all of the electrons will be expelled in a single cycle, as some very slow electrons may remain mixed with the antiprotons. Several iterations are often necessary to separate out the remaining electrons.

3.4.4 Evaporative cooling

For antihydrogen atoms to be trapped, their kinetic energies must be less than < 0.5 K (this is discussed in detail in Sec. 3.5). As such, the particle clouds and plasmas used in antihydrogen production need to be at low enough temperatures that the formation of trappable atoms is possible. One technique used to cool these particle distributions is evaporative cooling. Evaporative cooling proceeds as high-energy particles elastically scatter out of the confining potential, leaving the remaining distribution colder as a result. A common analogy is the cooling of a cup of coffee: the most energetic molecules are released as steam, cooling the coffee remaining in the cup. This technique has previously been used on neutral atoms, and is an essential part of the preparation of Bose-Einstein condensates [131].

Fig. 3.16 shows the confining potentials used for evaporative cooling of antiprotons. Here, the antiprotons are initially confined in a 1500 mV potential well. The potential on one side of the well is then slowly lowered, which stimulates evaporative boil-off as the well depth becomes comparable with the longitudinal kinetic energy of the most energetic particles. The remaining particles rethermalize to a lower temperature, which is a function of how shallow the potential well can be made.

Unfortunately, evaporative cooling relies intrinsically on particle loss – which is potentially devastating, as antiprotons are already rare particles. Moreover, due to the conservation of canonical angular momentum (similar to that described in Sec. 2.2.1.2), the loss of particles is accompanied by an expansion of the radius of the particle cloud. Escaping particles leave on-axis, which is where the confining potential well is shallowest, and therefore do not carry away angular momentum (Eq. 2.25). As some of the remaining particles diffuse inward into the ‘hollow’ on-axis region left by the escaping particles, other particles will move outwards to conserve the total canonical angular momentum

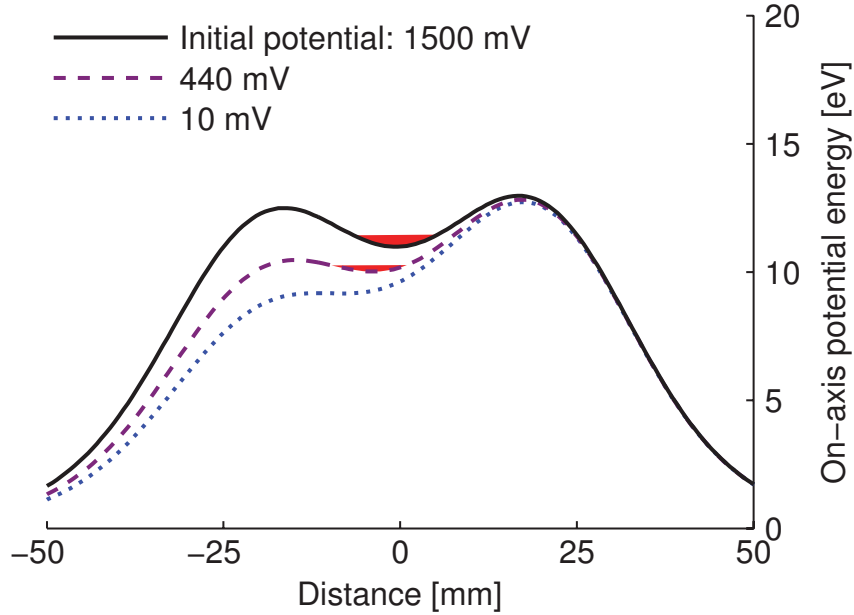



Figure 3.16: The electric potential used to confine antiprotons during evaporative cooling (reproduced with permission from [64]). The antiprotons are shown at the bottom of the potential well (Red). The well potentials shown are labeled by their on-axis depths.

(Eq. 2.28). This expansion can be expressed through the relation:


$$\frac{\langle r^2 \rangle}{\langle r_0^2 \rangle} = \frac{N_0}{N}, \quad (3.14)$$

where $\langle r_0^2 \rangle$ and N_0 are the initial mean squared radius and number of particles, respectively, and $\langle r^2 \rangle$ and N are the final mean squared radius and number of particles.

Figure 3.17 shows the antiproton temperature and fraction of remaining particles as a function of the final well depth. Clearly, both the distribution temperature and the number of remaining particles steadily decrease as the well depth is lowered (however, note that the plot of the surviving fraction of antiprotons is on a linear scale, while the plot of the antiproton temperature is log-scale, indicating that the antiproton temperatures drops significantly faster than the particle loss). These parameters can be modeled with



(a) Measured temperature (log scale).



(b) Surviving fraction of antiprotons.

Figure 3.17: Evaporative cooling of antiprotons (reproduced with permission from [64]). (a) shows the final measured antiproton temperature as a function of on-axis well depth, while (b) gives the corresponding survival fraction of antiprotons. The measured datapoints are shown as black squares, while the evaporative cooling model, Eqs. 3.16, is shown as the solid black curve.

the following system of equations [64]:

$$\frac{dN}{dt} = -\frac{N}{\tau_{\text{ev}}} - \gamma_{\text{ann}}N, \quad (3.15)$$

$$\frac{dT}{dt} = -\alpha_{\text{esc}}\frac{T}{\tau_{\text{ev}}} + \mathcal{P}_{\text{Joule}}, \quad (3.16)$$

where τ_{ev} is the evaporation timescale, γ_{ann} represents the rate of antiproton annihilation on background gas (10^{-4} s^{-1}), and α_{esc} gives the energy removed by the escaping antiprotons. Finally, the $\mathcal{P}_{\text{Joule}}$ term represents Joule heating during the radial expansion. Following Ketterle and Van Druten [132], the energy carried away by escaping antiprotons is calculated as,

$$\alpha_{\text{esc}} = \frac{\eta + \kappa}{\delta_{\text{avg}} + 3/2} - 1, \quad (3.17)$$

where η is the depth of the potential well, κ is the excess kinetic energy, and $\delta_{\text{avg}} + 3/2$ is the average energy of the distribution. Here, ν , κ , and $\delta_{\text{avg}} + 3/2$ represent the various energies present in the system, divided by $k_B T$. The results of this model are shown alongside the data in Fig. 3.17, and compare well with the measured results. An in-depth description and analysis of this model is given in [119].

Using the evaporative cooling technique on antiprotons, the lowest temperature measured was $(9 \pm 4) \text{ K}$, with $(6 \pm 1)\%$ of the initial 45 000 antiprotons remaining in the final 10 mV well [64]. Although the technique was developed and described here for antiprotons, ALPHA has also applied the evaporative cooling method to positrons [54]. Evaporative cooling of positrons is important for the production of very cold antihydrogen, as the positron plasma sets the temperature scale during mixing. That is, antiprotons will cool to the positron temperature during mixing, so it is important to ensure that the positron temperature is as low as possible.

3.4.5 Autoresonant mixing

In ALPHA, as with other AD experiments [2,36,37], antihydrogen is synthesized through the merging of an antiproton cloud with a positron plasma. This is normally done within a so-called ‘nested’ potential (Fig. 3.18a)), where both charged species can be confined simultaneously. In most particle mixing schemes, the antiprotons are initially held off to the side of the positron plasma, then slowly (or quickly) brought together. Two procedures that have previously been implemented by ALPHA are:

1. *Antiproton injection*: in order to overcome the central potential hill and interact with the positrons, the antiprotons are initially held at a much higher potential than the positron plasma [42,133]. The antiprotons are then launched into the nested potential where they interact with the positron plasma. Although this method will effectively mix the antiprotons with the positron plasma, it has the serious disadvantage that the antiprotons experience a potential drop on the order of tens of volts. This kick imparts a significant amount of kinetic energy in both the transverse and longitudinal degrees of freedom.
2. *‘Incremental mixing’*: this method involves gradually raising the voltage confining the positrons, such that the two species are slowly brought into contact (this would correspond to slowly adding a positive voltage on electrode E18 in Fig. 3.18a) and b)) [50,52]. The goal here is to minimize the axial energy of the antiprotons by keeping them as stationary as possible. However, the changing electric fields will invariably increase both the transverse and longitudinal antiproton energy.

Because the antiproton is much heavier than the positron, the velocity distribution of antihydrogen after formation will be dominated by the antiproton velocity. Thus, maintaining a cold antiproton distribution is important when designing mixing schemes. To this end, ALPHA developed a new method known as ‘autoresonant mixing’, which

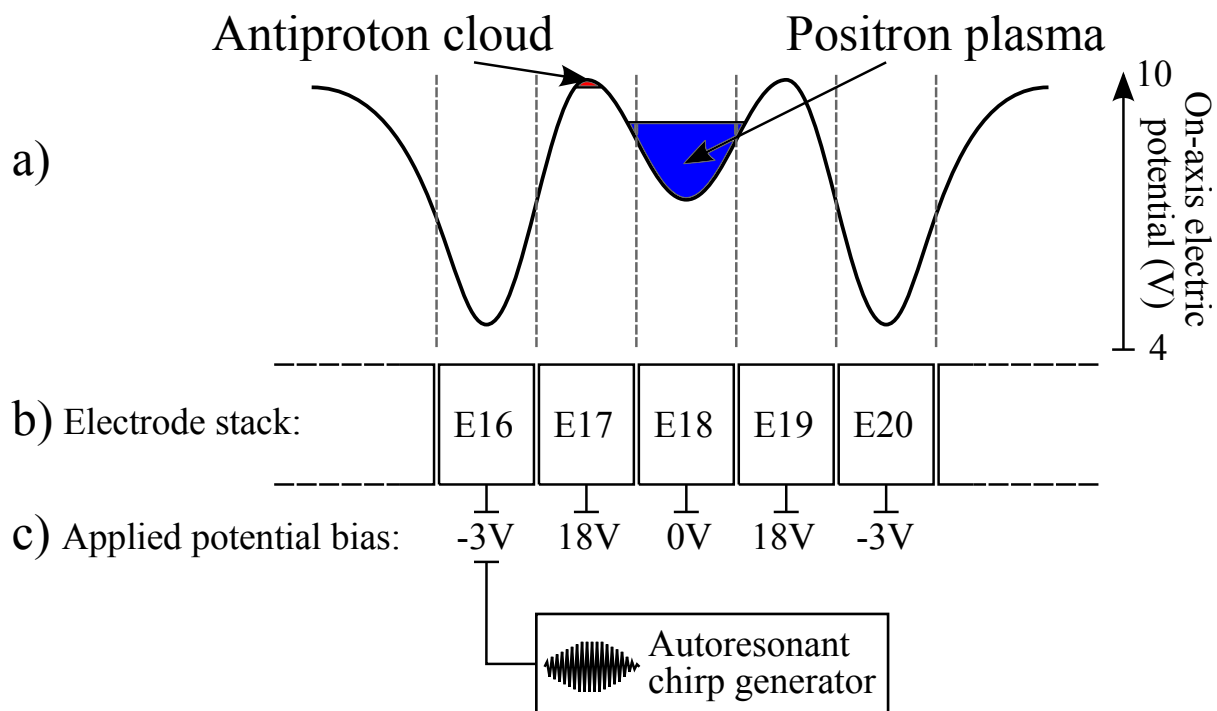


Figure 3.18: Nested potential well used to confine both positrons and antiprotons. In a), the on-axis vacuum electric potential is shown, with the positrons confined in the central well, and the antiprotons in a side-well; b) shows a schematic of the electrode stack, with c) showing the applied voltages and the chirp signal generator for the autoresonant drive.




Figure 3.19: An energy-frequency plot of autoresonantly-driven antiprotons (reproduced with permission from [65]). The left figure a) shows the longitudinal energy distributions, $f(U)dU$, of $\sim 15\,000$ antiprotons driven to several different final frequencies. b) shows the calculated longitudinal energy as a function of drive frequency for the vacuum potential (solid blue), 15 000 antiprotons (dashed green), and 50 000 antiprotons (dot-dashed red). The open squares show the mean of the measured longitudinal energy distributions. Frequencies are normalized to $\omega_0/2\pi = 410$ kHz.

aims to excite the longitudinal motion of the antiproton cloud, driving them into the positron plasma, while minimizing the perturbation to the transverse antiproton energy. This scheme takes advantage of the fact that the confining potential for the antiprotons is anharmonic, and the axial oscillation frequency is therefore related to the oscillation amplitude. Figure 3.18b) and c) show the setup of the autoresonant mixing scheme, highlighting the chirped sinusoidal drive which is used to excite the antiproton cloud. The sinusoidal drive frequency is swept down through the linear (small-amplitude) oscillation frequency, ω_0 . With a strong enough drive, the nonlinear antiproton oscillation frequency, ω , will at every instant match the drive frequency (this phase-locking is why this method is known as ‘autoresonant’ [134]).

Figure 3.19 shows the energy-frequency response for autoresonantly-driven antiprotons [65]. A very inviting feature of this scheme is that, as shown in Fig. 3.19a), the

longitudinal antiproton energy can be selected with an appropriate choice of drive frequency. When mixing antiprotons and positrons, the autoresonant chirp-drive is engaged for 200 μs at an amplitude of ~ 55 mV, and swept from 350 \rightarrow 200 kHz, which injects about 70% of the antiprotons into the position plasma. As stated earlier, the autoresonant mixing method should have very little effect on the transverse antiproton energy. Moreover, despite raising the antiproton potential energy, this technique leaves the temperature (i. e. the distribution width) largely unchanged, allowing for coherent and cold injection.

3.5 Magnetic neutral-atom trap

As antihydrogen atoms are electrically neutral, the Penning-Malmberg trap configuration (Sec. 2.2.1) is inadequate to confine them once they are formed. The ALPHA neutral-atom trap (as shown in Fig. 3.20) consists of several superconducting magnets [99]. The combination of the magnetic fields of these magnets creates a three-dimensional minimum along the trap axis where the antihydrogen atoms are formed. The axially confining magnetic field is established by two coils (Sec. 3.5.1), while the radially confining field is the result of an octupolar magnet (Sec. 3.5.2). The entire neutral-atom trapping field is superimposed on the 1 T solenoidal field of the Penning-Malmberg trap.

Following from Eq. 2.36, the trap depth ($U_{\text{max}} = 0.67\Delta B$) is set by the difference between the maximum and minimum magnetic field magnitudes. As will be seen in Sec. 3.5.1 and 3.5.2, the radial field gives the smallest difference, and therefore sets the overall trap-depth of the ALPHA neutral-atom trap.

3.5.1 Axial magnetic field

The axially confining magnetic field in ALPHA is produced via two ‘mirror’ or ‘pinch’ coils. These coils are essentially short-section solenoids, the main effect of which is to




Figure 3.20: A schematic diagram of the magnetic neutral-atom trap and the surrounding apparatus (reproduced with permission from [54]). The mirror coils (cut away at the centre to allow visualization of the interior of the apparatus) are shown in green, and the octupole in red. The Penning-Malmberg trap electrodes are shown in yellow, and the surrounding silicon detector in light blue. (Not to scale.)

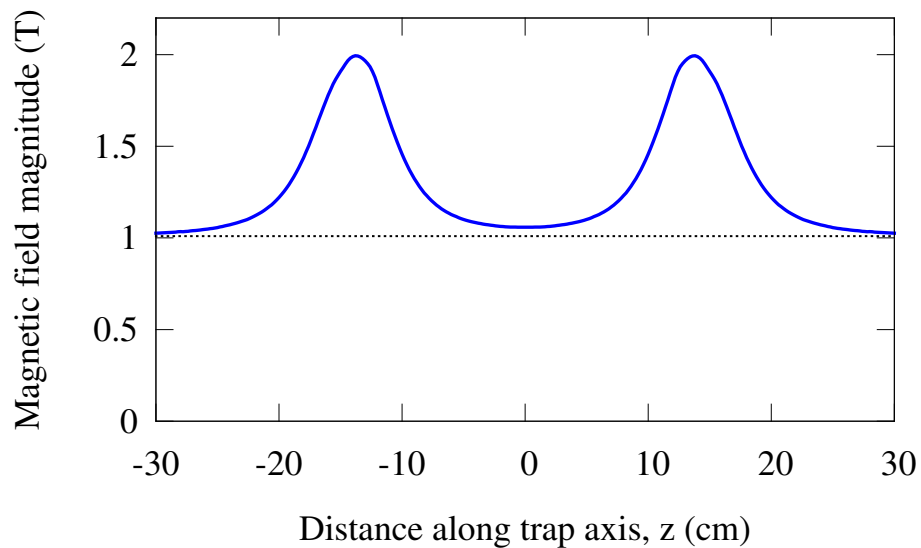


Figure 3.21: The on-axis magnetic field magnitude for the ALPHA neutral-atom trap. The total field magnitude is shown as the blue curve, while the 1 T contribution from the external Penning-Malmberg trap solenoid is shown as the dashed grey line. The fields of the mirror coils produce the large peaks at ± 14 cm.

increase the axial field component directly within the coil extent. While the magnetic field of the mirror coils also has a small radial component (and fringe fields), these coils effectively create an increasing magnetic field magnitude for atoms travelling in either axial direction from the middle of the neutral trap. Figure 3.21 shows the on-axis field magnitude. The large bumps are the contributions of the mirror coils, which are superimposed on the field of the external Penning-Malmberg solenoid. With the mirror coils fully energized (with a current of about 600 A), the maximum axial magnetic field is about $B_{\text{max,axial}} = 1.99$ T (which is the combination of the Penning-Malmberg trap field and the mirror coil field at the axial center of the mirror coils, $z = \pm 14$ cm). When compared to the minimum trap field magnitude of $B_{\text{min,axial}} = 1.06$ T (the combination of the Penning-Malmberg trap field and the fringe of the mirror coil field at the axial center of the neutral trap, $z = 0$), this gives an axial well depth of $54 \mu\text{eV}$, or 0.6 K.

3.5.2 Radial magnetic field

The radial trap-depth is set by the difference between the minimum and maximum magnetic field magnitudes. This difference can be written as [99],

$$\Delta B = B_{\text{max,radial}} - B_{\text{min,radial}} \quad (3.18)$$

$$= \sqrt{B_w^2 + B_z^2} - B_z, \quad (3.19)$$

where B_w is the magnetic field magnitude of the octupole magnet at the electrode wall, and B_z is the axial magnetic magnitude (these fields are spatially orthogonal, so the magnitudes is added in quadrature for $B_{\text{max,radial}}$). Equation 3.19 provides the motivation for setting the Penning-Malmberg solenoidal field at 1 T and using an additional 2 T internal solenoid for antiproton catching (Fig. 3.12), since reducing B_z amounts to increasing the neutral-atom trap depth.

Figure 3.22 shows the radial profile (at the axial center, $z = 0$) for the full neutral trap




Figure 3.22: Total magnetic field magnitude in the ALPHA apparatus as a function of radius at the axial center ($z = 0$ cm). The dotted line indicates the electrode wall radius.

magnetic field (the electrode wall radius is shown as the dotted line). With the octupole magnet fully energized at a current of about 880 A, the minimum field magnitude is on-axis with a value of $B_{\min,\text{radial}} = 1.06$ T (the combination of the Penning-Malmberg trap field and the fringe field of the mirror coils at the axial center of the neutral trap, $z = 0$ cm). The shallowest magnitude at the wall radius (there is azimuthal variation due to the octupole configuration) is $B_{\max,\text{radial}} = 1.82$ T (the quadrature sum of Penning-Malmberg trap and octupole field magnitudes at the wall radius), which gives a radial trap-depth of $44 \mu\text{eV}$, or 0.5 K. This value defines the overall trap-depth of the neutral-atom trap, as antihydrogen atoms with kinetic energy greater than the trap-depth will eventually leave the magnetic trap.

Despite the inevitable azimuthal symmetry breaking by the octupole magnetic field (Sec. 2.2.2), stable charged plasma confinement in the octupolar field was quickly demonstrated in the ALPHA neutral-atom trap [48].

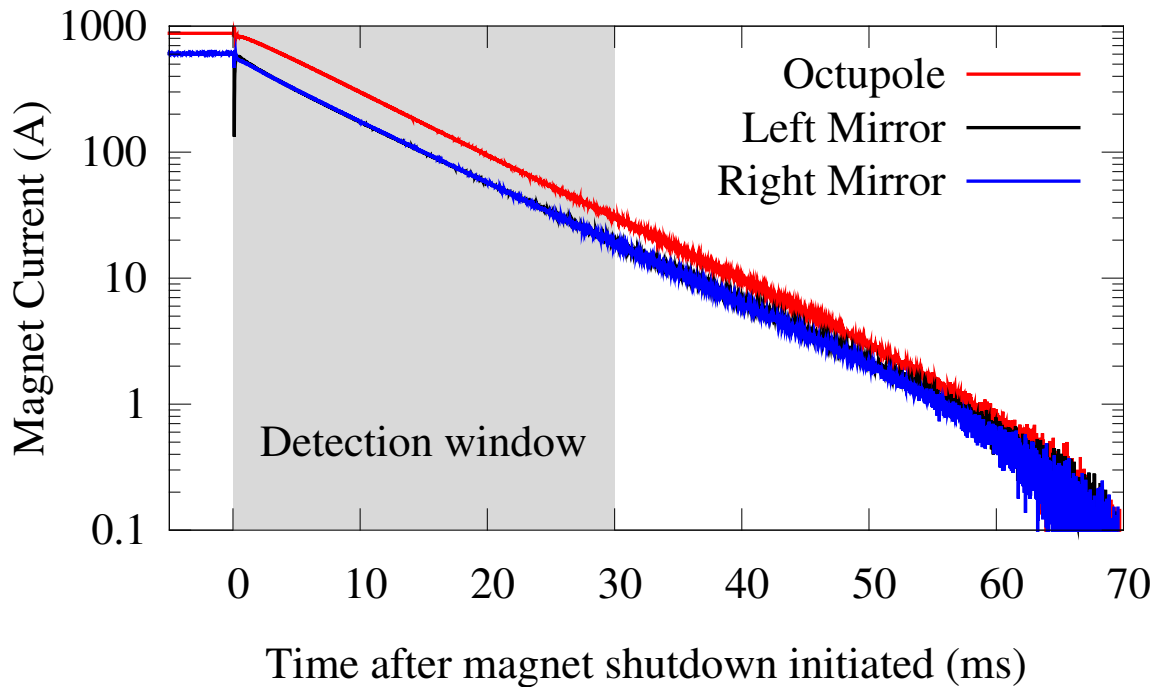


Figure 3.23: Magnet current decay following the fast shutdown of the neutral trap (note the log scale). The current in the octupole magnetic is shown as in red, the left mirror in black, and right mirror in blue. The shaded region represents the expected time period where trapped antihydrogen atoms will be released.

3.5.3 Fast shutdown

To protect the neutral trap magnets from damage during a ‘quench’ event (where the magnet winding loses superconductivity), the neutral trap control system is designed with the option to quickly, and safely, de-energize the magnets. This is achieved by constantly monitoring the voltage drop across the superconducting wire. If a voltage is registered (indicating that the wire has lost superconductivity), a fast insulated-gate bipolar transistor (IGBT) switch is engaged and the current is diverted to a resistor network, where it is quickly dissipated as heat. Figure 3.23 show the the response of the magnet currents during a fast shutdown. The magnet currents, and therefore the fields, decay with a time constant of ~ 10 ms.

In addition to protecting the magnet integrity, the fast shutoff option provides a

means to expel any confined antihydrogen atoms in a well-defined time window. The shaded region in Fig. 3.23 shows the time period where the vast majority of trapped atoms would exit the trap to be detected. The detection window is 30 ms long, and the analysis of the antihydrogen detector data (to be described in the following chapters) is focused on this time interval.

3.6 Trapping experiment overview

To complete an antihydrogen trapping experiment, every component and technique described in this chapter must function correctly and in the right order. Likewise, many operations proceed on the different timescales, and it is necessary to coordinate between various experimental components and measurement devices. To facilitate experimental operation, ALPHA uses a sequencing system consisting of a number of digital and analog inputs and outputs with a field-programmable gate array (FPGA) controller. This FPGA controller can be programmed to reproduce a sequence of digital and analog states as defined by the user with a graphical user interface (GUI). These states can be held indefinitely, or programmed to be as short as 37 ns.

Using this control sequencer, a trapping experiment can be programmed and executed repeatedly. The order and description of the major operations in a typical trapping experiment are given as follows:

1. Positrons are collected in the positron accumulator (Sec. 3.3.2). This operation usually takes about 250 s, or until the desired number of positrons (usually $\sim 7 \times 10^7$) are accumulated.
2. The inner solenoid is energized to 2 T, giving a total field (when combined with the external solenoid) of 3 T in the antiproton catching region. An electron plasma containing about 1.5×10^7 particles is pre-loaded into the catching region (Sec. 3.4.2).

3. About 3×10^7 antiprotons are delivered from the AD into the ALPHA apparatus and dynamically captured between two high-voltage barriers and cooled by the pre-loaded electrons (Sec. 3.4.2). The resulting antiproton-electron mixture contains about 4.5×10^4 antiprotons, which take about 80 s to cool to the electron temperature.
4. The positron plasma is ballistically transferred from the accumulator into the mixing trap, where the positrons remain until the antiprotons are ready for mixing.
5. The rotating wall potential is applied to the electron-antiprotons mixture for 20 s to radially compress the system (Sec. 3.4.3). The internal solenoid is then ramped down, and the mixture is transferred to mixing region.
6. The electrons are removed (Sec. 3.4.3), and the antiprotons are placed in a side-well of a nested potential, immediately adjacent to the positron plasma (Sec. 3.4.5). The antiproton cloud at this point has a temperature of about 200 K.
7. The superconducting neutral-atom trap magnets are energized over 25 s, resulting in a minimum-B configuration with a trap-depth of 0.5 K (Sec. 3.5).
8. The positron plasma is quickly evaporatively cooled to about 40 K (Sec. 3.4.4), and the antiprotons are autoresonantly driven (Sec. 3.4.5) into the positrons and allowed to mix for 1 s. As a control measurement, the positron plasma can be rf-heated to suppress antihydrogen production (to be discussed further in Chapter 7). To attempt to remove any charged particles, an electric field (2.5 V/cm) is pulsed across the trap four times.
9. The neutral-atom trap is then quickly de-energized over tens of milliseconds (Sec. 3.5), releasing anything remaining within. As another control measurement, an electric field can be placed across the trap to deflect any mirror-trapped antiprotons that


might have been confined by the neutral trap magnetic field. This field acts to spatially separate the bare, charged antiprotons from the neutral antihydrogen atoms (more details about this method are given in Chapter 7). The first 30 ms after the start of the neutral trap magnet rampdown is the signal window for any escaping antihydrogen atoms (Sec. 3.5.3). The now-unconfined antihydrogen atoms will transit to the electrode wall, taking on the order of hundreds of microseconds, and will most likely come into contact with a gold atom. Charged and neutral pions are produced in the low-energy antiproton annihilation on the heavy nucleus, and these produced particles will then travel outwards from the annihilation location. A large fraction of the produced particles will travel through the silicon detector, triggering its readout (Sec. 4.1.4). Events registered by the silicon detector are examined carefully, looking for antihydrogen annihilation signatures.

Figure 3.24(a) gives the timeline for the example antihydrogen trapping experiment. Each experiment is about 660 s in duration, starting from the beginning of positron accumulation to the fast shutoff of the neutral trap. Since there are a number of important events happening in a short period near the end of the experiment, Fig. 3.24(b) shows a magnification of the multiple operations that occur between 660 and 662 s. To accumulate statistics, several hundred cycles of this experiment are typically performed.


3.7 Summary

The ALPHA apparatus is a complex and versatile instrument. This chapter described the various components of this apparatus, as well as several important techniques used during antihydrogen trapping experiments. The production, confinement, and manipulation of charge antiparticles within the ALPHA apparatus comprised a large section of this chapter and included descriptions of rotating wall compression, evaporative cooling of

charged antiparticles, as well as autoresonant mixing of positrons and antiprotons. The composition and operation of the magnetic neutral-atom trap was also described, along with an overview of a typical trapping experiment.



(a) Entire trapping experiment.



(b) Magnification of 660 - 662 s.

Figure 3.24: Timeline of the entire trapping experiment. Important events and operations are labeled. (a) shows the entire experiment, while (b) gives a magnification of 660 - 662 s, showing the details of the particle mixing and final magnetic trap operations.

Chapter 4

The ALPHA silicon detector

An important component in the ALPHA experiment, and the technical focus of this dissertation, is the silicon annihilation detector. This detector is sensitive to the passage of high-velocity charged particles, and is primarily used to detect and identify antiproton annihilations within the ALPHA apparatus. This detector was designed to accommodate the following considerations:

- The low expected antihydrogen trapping rate requires that the detector be highly efficient, with large solid angle coverage.
- The ability to identify and reject background events is needed to enable the conclusive identification of trapped antihydrogen atoms.
- A sensitive diagnostic tool for plasma and atomic processes is highly desirable.

However, the integration of the detector with the rest of the apparatus imposes some challenges:

- Because of the cryostat housing the Penning-Malmberg and neutral-atom traps, the first layer of silicon modules in the detector has to be far relatively far (at least 5 cm) from the annihilation points.
- The neutral-atom trap superconducting magnet windings add a large amount of material between the annihilation points and the detector. This contributes to the multiple scattering experienced by the charged particles before they reach the detector (Sec. 5.1).
- Antiproton annihilation involves the production of π^0 particles (Sec. 2.3.2), which can result in complex event topologies (Sec. 5.1).




Figure 4.1: Cut-away rendering of the silicon detector in relation to the cryostat region of the ALPHA apparatus. Each outlined blue rectangle represents the location of a silicon module. Not shown are the readout chips or support structure. The structure of the apparatus has also been simplified for clarity.

To accommodate these considerations and compromises, a three-layer silicon detector with high granularity (30,720 channels) and large solid angle coverage (an active area of 0.8 m^2 , covering about 90% of the full 4π steradians solid angle) was designed and constructed. Silicon detectors generally also have a high particle detection efficiency [110], due to the low threshold for electron-hole generation (Appx. B).

This chapter starts by describing the silicon detector hardware, including the detector modules and readout electronics. This both provides an overview to the detector configuration (which imposes a number of constraints on the event reconstruction, to be discussed in Chapter 5), but also motivates the operation and readout of the detector. Next, the analog readout is discussed in detail, focusing on the determination of particle hits. Finally, the Monte Carlo simulation of the silicon detector and surrounding apparatus is described and some initial comparisons between the simulation and data presented.

4.1 Detector hardware and software for control and data collection

The ALPHA detector consists of 60 double-sided silicon microstrip detector modules arranged in three concentric layers around the ALPHA cryostat and trap region (Fig. 4.1). In order to provide the largest solid angle coverage around the magnetic neutral atom trap, the detector is positioned coaxially around the trapping region, and as axially centred as physically possible (due to space constraints, the detector is axially displaced 2.4 cm on-axis from the trap centre).

This detector was designed at TRIUMF and University of Liverpool in 2005 and built at the University of Liverpool between 2006-2009 [135, 136]. A partial detector (38 of 60 modules) was installed in the ALPHA apparatus at CERN in 2008, and the final full detector was installed in 2009 and utilized throughout the 2009 and subsequent




Figure 4.2: Diagram of the silicon modules (not to scale). The silicon wafers and readout chips are highlighted. Not shown are the auxiliary readout electronics, and details of the bond wires and signal strips.

experimental runs.

4.1.1 Silicon detector modules

Each detector module contains the two silicon wafers, along with four VA1TA [137] ASIC (Application Specific Integrated Circuit) readout chips and associated readout electronics, all mounted on a Printed Circuit-Board (PCB) support structure (Fig. 4.2). Each silicon wafer is $6.1 \text{ cm} \times 11.5 \text{ cm}$, so with two wafers per module, the active silicon area on a single module is $6.1 \text{ cm} \times 23.0 \text{ cm} = 140.3 \text{ cm}^2$.

Two ASICs read out the p-side strips, while the other two ASICs read out the n-side strips. Each ASIC chip reads out 128 signal strips, for a total of 256 p-side strips, and 256 n-side strips, per module. The pitch width (the distance between the centers of two adjacent strips) is $227 \mu\text{m}$ for the p-side strips, and $875 \mu\text{m}$ for the n-side strips. The p- and n-side signal strips are arranged orthogonally, with the p-side strips extending




Figure 4.3: Photograph of two detector modules mounted on the detector support structure at the University of Liverpool clean room. (Image credit: University of Liverpool)

across the long (23 cm) length of the modules, while the n-side strips extend across the short (6.1 cm) side of the modules. By combining the p- and n-side signal information, the position where a charged particle passed through the silicon can be localized in the plane of the module. Moreover, by including information about detector geometry (that is, the 3-dimensional position and orientation of the silicon module), the hit positions can be found in the global reference frame of the apparatus (to be discussed in detail in Sec. 4.1.6.2).




Figure 4.4: Overview of one of the 128 readout channels in the VA1TA chip, following [137]. Detailed circuitry has been suppressed to emphasize the VA/TA division.

4.1.2 VA1TA ASIC readout chips

After the silicon itself, the most important component on the detector modules is the VA1TA readout chip. These chips handle both the trigger signalling and strip readout. Each ASIC is essentially two chips in one: the ‘VA’ part of the VA1TA handles the analog strip signals, while the ‘TA’ part handles the digital triggers. The VA1TA chip has 128 channels, with each channel containing a preamplifier, slow shaper, fast shaper, discriminator, and monostable circuit.

Figure 4.4 shows an overview of how the two parts of the VA1TA interplay. Both the VA and TA parts share the same preamplifier. After that, the VA circuit proceeds to a slow ($\sim 1 \mu\text{s}$) shaper, while the TA circuit contains a fast (75 ns) shaper. If the pulse height of the fast shaper exceeds the preset discriminator threshold, a fixed width (105 ns) trigger pulse is generated by a monostable multivibrator. All 128 TA trigger signals are logically ORed, such that each ASIC has only one TA trigger signal.

Similarly, there is only one analog out (for all 128 channels) from each ASIC. This requires the VA channels to be multiplexed and read out serially. To achieve this, each analog pulse height must be held until it is its turn to be multiplexed. This is the basis of the ‘sample and hold’ method, where the analog signals are held at the shaping peak

until it is time to be multiplexed.

Note that VA1TA allows for adjustment in the shaping times, threshold and signal widths through a 680 bit control register. This control register is particularly sensitive to Single Event Upset (SEU) effects, e.g., bits of this control register being flipped due to external influences (energetic particles passing through transistor silicon, for example). For this reason, the VA1TA has a redundant SEU-correcting circuit [137, 138].

4.1.3 Detector configuration

As shown in Figs. 4.1 and 4.3, the detector modules are arranged in three layers around the particle trapping region of the apparatus. The reason for the three separate concentric layers is so that the detector can sample the trajectory of a charged particle in a magnetic field at three points in its flight, and thus be able to extrapolate its path back its origin, typically an annihilation vertex (to be discussed in Chapter 5).

The detector is effectively divided into two half-detectors along the axial direction. Each half contains 30 modules arranged with the silicon wafers positioned towards the middle of the detector, and the two sets of readout electronics towards either end. Together, the two half-detectors have an axial extent of 46 cm. The modules are numbered between 1-60, with modules 1-30 making up the half closest to the AD, and 31-60 being closest to the positron accumulator (see Fig. 3.1).

The radial configuration is shown as Fig. 4.5. The inner layer contains 8 modules (modules 1-8 and 31-38) on a radius of 7.5 cm, the middle layer has 10 modules (modules 9-18 and 39-48) at on a radius of 9.55 cm, and the outer layer 12 modules (modules 19-30 and 49-60), with half of the outer layer modules around a radius of 10.9 cm and the other half around a radius of 11.4 cm. The staggering in the outer layer is to allow for airflow in order to cool the modules.




Figure 4.5: Diagram illustrating the radial and azimuthal placement of the silicon modules. Modules 1-30 are located on the ‘AD’ half of the detector, while modules 31-60 are located on the ‘Positron’ half.

4.1.4 Readout chain

With 60 modules¹, each having 4 ASICs, and each ASIC managing 128 strips, there is a total of 30720 channels to be read out. An overview of the readout chain is shown as Fig. 4.6.

The overall timing and control of the detector readout chain is controlled by the Timing, Trigger and Control (TTC) module. The TTC is a special-purpose FPGA-based VME module which coordinates the detector readout chain by providing the central clock signal, as well as computing the readout trigger multiplicity logic. Each readout ASIC has a 680 bit control register that needs to be loaded for proper operation. These control bits are also provided by the TTC module.

The TTC communicates with silicon modules via Front-end Repeater Cards (FRCs). The FRCs are responsible for powering and communicating with the ASIC readout chips

¹While there are 60 modules included in the detector, module 10 is not read out. It is suspected that module 10 is poorly connected at the (presently inaccessible) interface between the readout cable and connector on the PCB board. This disconnect hinders the communication between FRC and ASIC, and renders the readout information unusable.

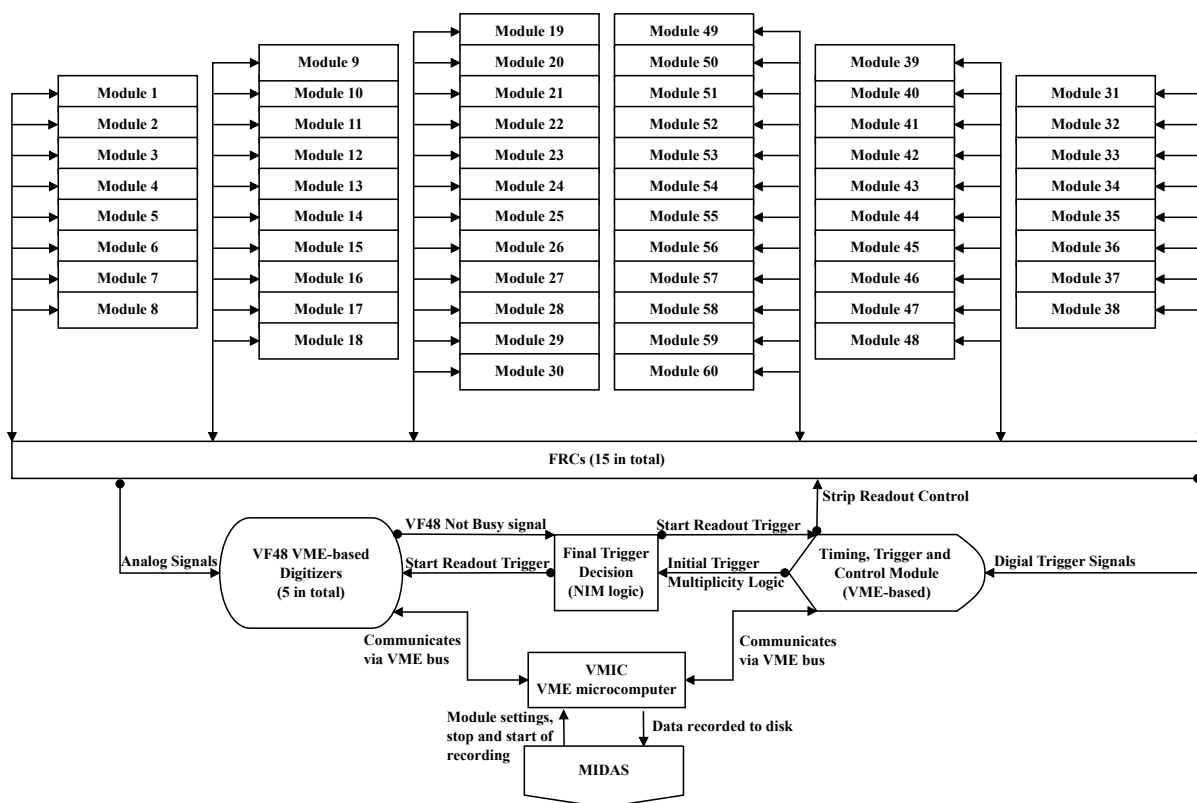


Figure 4.6: Schematic diagram of the detector readout chain. The connections have been simplified for clarity.

on the silicon modules. The FRCs are located just outside the apparatus, and are connected to the silicon modules via two ribbon cables (one cable for the digital trigger signal, and one cable for the analog strip readout signal) per module. Each FRC communicates with 4 silicon modules.

The analog strip signals of all 240 ASIC readout chips are read out in parallel into 5 VF48 [139] digitizers. Each VF48 has 48 10-bit analog-to-digital (ADC) channels. Moreover, each VF48 ADC is differential, allowing for both positive and negative signals relative to the reference pedestal voltage. Like the TTC, the VF48s are VME modules and the digitized detector strip signals are recorded to disk via a VMIC VME microcomputer networked with the ALPHA Linux computer cluster. The MIDAS² (Maximum Integrated Data Acquisition System) framework is used to coordinate and record the incoming data-streams in a computer-readable (and easily analyzable) format.

4.1.5 TA and readout triggers

There are several triggers used in the detector readout chain, with the base type being the TA triggers (Sec. 4.1.2). The TA triggers are digital signals that come directly from the VA1TA chips when the amount of charge collected in a silicon strip exceeds a preset threshold value (TA threshold, Fig. 4.4). Each VA1TA chip contributes only one TA trigger, which represents the logical OR of all of the 128 strips connected to that chip. The TA triggers are combined to form the $S_i > 1$ trigger (the logic of which is described below), which is intended to reduce the influence of spurious TA triggers. Finally, the readout trigger ('RO trigger') is what actually initiates the readout of the analog levels of all of the silicon strips.

All the TA triggers from the silicon modules are sent to the TTC. The TTC is programmed with the information about the module layering and groups the incoming

²MIDAS website: midas.psi.ch

Name	Description
$T1$	≥ 2 AD-side Layer 1 triggers
$T2$	≥ 1 AD-side Layer 1 triggers
$T3$	≥ 2 AD-side Layer 1 triggers AND ≥ 1 AD-side Layer 2 triggers AND ≥ 2 AD-side Layer 3 triggers
$T4$	≥ 2 Positron-side Layer 1 triggers
$T5$	≥ 1 Positron-side Layer 1 triggers
$T6$	≥ 2 Positron-side Layer 1 triggers AND ≥ 1 Positron-side Layer 2 triggers AND ≥ 2 Positron-side Layer 3 triggers

Table 4.1: Descriptions of TTC NIM trigger outputs as defined for the ALPHA detector readout system.

trigger signals according to which half-detector and which layer they originate from. Just as importantly, the TTC computes the signal multiplicity, that is, how many trigger signals arrive in coincidence. For technical reasons³, only p-side trigger signals are used for trigger decisions.

There are six programmable NIM outputs from the TTC which are summarized in Table 4.1. These outputs are directed to NIM logic modules, where the final trigger decision is made (Fig. 4.7). The most generally used trigger decision in ALPHA is the so-called ‘Si > 1’ trigger. The Si > 1 trigger is a logic combination of TTC outputs,

$$\text{Si} > 1 = T1 \text{ OR } T4 \text{ OR } (T2 \text{ AND } T5). \quad (4.1)$$

where the $T\#$ definitions (where $\#$ is between 1-6) are as described in Table 4.1.

Si > 1 is intended be a very loose trigger. That is, this trigger is intended to accept as many events as possible (while still rejecting most outright noise-induced triggers).

The idea being that useful events will contain two or more charged particles, as such the

³The n-side signal strips are connected to the readout ASICs via long signal wires (see Fig. 4.2). Because of the wire length and path, there is significant capacitive coupling between the signal wires. As a result, the n-side trigger signals are much noisier then their p-side counterparts.

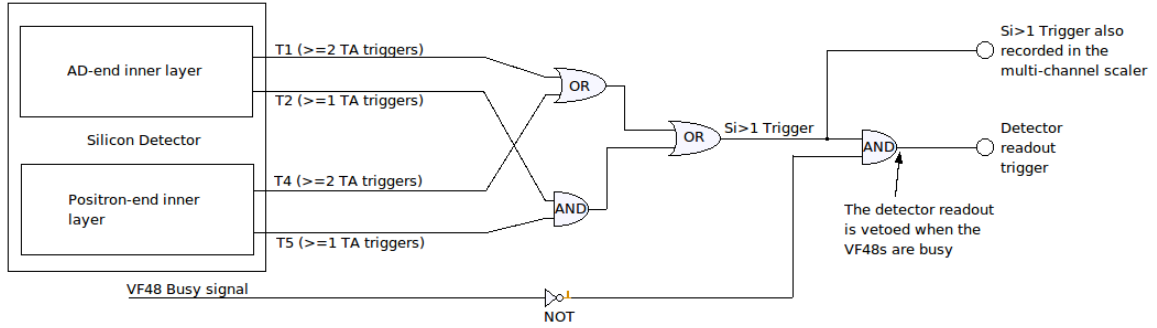


Figure 4.7: Logic diagram of $\text{Si} > 1$ TA trigger combination and detector readout trigger decision. Note that the VF48 busy signal is high when the VF48 are engaged in the detector readout, and the NOT logic inverts the busy signal.

trigger is based on having signal in two or more ASICs. $T1$ and $T4$ represent having two or more signals in either half of the detector, while $T2$ AND $T5$ deals with the case where there are two or more signals divided between the two half-detectors. During positron-antiproton mixing, the $\text{Si} > 1$ trigger will regularly reach a peak rate of ~ 30 kHz.

While the $\text{Si} > 1$ trigger might signal the presence of a desirable candidate for triggering the full detector readout (especially for low rates, where every $\text{Si} > 1$ could be read out and examined), for $\text{Si} > 1$ trigger rates higher than ~ 250 Hz, the digitization of the analog strip levels will not be able to keep up with the number of incoming triggerable events. So as not to interrupt a readout in progress, the $\text{Si} > 1$ trigger is combined with a ‘Not Busy’ signal from the VF48s to give the ReadOut trigger (‘RO trigger’). The Not Busy signal is engaged when the VF48s are not in the process of reading out the analog strip levels, and it acts to veto the $\text{Si} > 1$ trigger (Fig. 4.7).

When the full readout is triggered, the whole readout chain is engaged in order to record the analog signal of all the silicon strips. The TTC coordinates this operation by providing both the multiplexer clock to the ASICs (via the FRCs) and the sampling clock to the VF48s. That is, every falling edge of the multiplexer digital clock signal connects




Figure 4.8: Simplified diagram of the clock timing during the readout of the analog strip levels. The 60 MHz sample clock is shown in black, the 3.33 MHz multiplexer clock is shown in red, and the analog strip voltage is shown in blue. One full strip readout is shown, along with the shifting to the next strip.

the ASICs readout to the next strip and at the same time the sample clock starts the next ADC sample. In this way, all the ASICs are readout out in parallel (one VF48 channel per ASIC), while the strips are readout serially. This operation is limited by the ASIC multiplexer, which is limited to clock frequencies below 5 MHz. In operation, a multiplexer clock of 3.33 MHz and a VF48 sampling clock of 60 MHz is used. The VF48s are then set to only sample every 18th sample clock cycle (Fig. 4.8).

In principle, the 3.33 MHz multiplexer clock limits the minimum readout to $\sim 40\ \mu\text{s}$, or a peak readout rate of 25 kHz. However, because of shaping time and overhead in the data transfer and storage, the peak readout rate is limited to 500 Hz. This peak rate corresponds to the 30 Mbytes/s data-to-memory throughput (where the data goes from the VF48 internal buffer to the MIDAS shared-memory buffer). However, this peak rate can only be maintained until the shared-memory buffer is exhausted (the shared-memory buffer is typically allocated 200 MB, which can sustain the 500 Hz data-rate

for about 15 s – which is much longer than needed for the high-rate portion of a typical experiment), at which point, the data flow is written to disk, with a throughput of about 16 Mbytes/s, corresponding to a 250 Hz event read-out rate. Thus, the full detector read-out rate is effectively limited to 500 Hz, with that number falling to 250 Hz for sustained experiments.

For experimental purposes, accurate timing of the detector triggers is invaluable (the trapping results in Chapter 7, for example, rely heavily on accurate knowledge of the time distribution of detector readout events). All of the detector timing in ALPHA is coordinated by two VME-based multi-channel scaler modules (known as the ‘SIS’ modules due to their model number, SIS3840). Each SIS module has 32 input channels which can handle 250 MHz counting rates with 32-bit depth. Both the $S_i > 1$ and RO triggers are connected to SIS channels, along with a 10 MHz atomic clock line provided by CERN (the other SIS channels are occupied by various instruments for which accurate timing is also important). Accurate timing is accomplished by configuring the SIS module to run in first-in-first-out (FIFO) mode, with each incoming detector trigger causing the SIS to push all of the scaler channel values into the FIFO buffer by engaging the Load Next Event (LNE) SIS input line. As the FIFO buffer is read out into MIDAS (in time-order), each detector trigger is compared to the number of 10 MHz clock cycles to establish the time between LNE pushes, and therefore the time between all of the triggers. Since all of the VME modules are controlled by the MIDAS interface, the start of the data recording for each experiment is taken as the $t = 0$ reference time. However, to safeguard the SIS FIFO buffer from overflowing, a 250 μs veto is forced between each LNE push, which defines the time resolution of SIS readout.

Since the VF48 data-stream is separate from the SIS data-stream, each set of VF48 readouts (containing the silicon strip analog levels) must be time-aligned with the RO trigger recorded by the SIS module. This is done post-experiment by comparing the

number of RO triggers counted by the SIS module to an internal VF48 timestamp associated with each analog readout. All of the time-sensitive instruments in the ALPHA apparatus have their own SIS channel (or use a MIDAS data-stream) and can therefore be compared throughout the course of an experiment or during data analysis long after the completion of the experiment.

4.1.6 Analog readout

The analog readout of the silicon modules is particularly important as it is the analog readout that decodes the position information required for tracking and annihilation vertexing. This section will describe how particle positions are determined from the strip information, as well as describe some measurements that can be completed using only this particle hit information.

4.1.6.1 Separating signal strips from noise

The output of all 30720 strips is digitized each time the detector readout is triggered. From this data, the signal due to particles passing through the detector needs to be isolated. The detector readout comes in the form of 240 (60 modules \times 4 ASICs per module) sets of 128 ADC samples. Each ADC sample represents the integrated charge accumulated on an individual strip. Because this is an analog digitization involving long signal cables, leakage current from the silicon, and an overall noisy environment, any signal is superimposed on an approximately Gaussian-distributed noise background. Therefore the determination of which strips fired comes down to removing the background noise in order to reveal the signal.

Figure 4.9 shows raw strip samples from a p-side and an n-side ASIC for a single detector readout. Each experimental cycle will contain on the order of 2000 detector readouts. Since every ASIC is read out during each detector readout trigger, there will be about 2000 samples for each strip over the course of the cycle. Currently there is no




Figure 4.9: Raw strip sample output from the VF48 modules for an (a) n-side ASIC, and (b) p-side ASIC. Strip samples are shown as the dark trace, while the mean ADC value for the entire ASIC is shown as the dashed red line.

zero suppression, so the vast majority of the strips do not contain any signal (typically, only a couple dozen of the 30720 strips will be occupied by signal). These samples can be used to determine the noise width, correct for systematic dc offsets, and calculate the signal threshold.

To understand this procedure, assume that there are N readouts during an experiment. The signal determination procedure proceeds as follows:

1. A first run-through of all N readouts is made. The mean ADC value, $\langle \text{ADC} \rangle_{\text{strip}}$, and standard deviation, σ_{strip} , for each strip is determined,

$$\langle \text{ADC} \rangle_{\text{strip}} = \frac{1}{N} \sum_{i=1}^N \text{ADC}_{\text{strip},i}, \quad (4.2)$$

$$\sigma_{\text{strip}}^2 = \frac{1}{N-1} \sum_{i=1}^N \left(\text{ADC}_{\text{strip},i} - \langle \text{ADC} \rangle_{\text{strip}} \right)^2, \quad (4.3)$$

where $\text{ADC}_{\text{strip},i}$ is the ADC value for that strip for the i th readout.

Because the vast majority of the strip samples will not contain any signal (due to the absence of zero suppression), Eqs. 4.2 and 4.3 estimate the sample mean and standard deviation of the strip background, or pedestal.




Figure 4.10: $\langle \text{ADC} \rangle_{\text{strip}}$ values (from the same readout as Fig. 4.9), relative to the ASIC sample mean, for an (a) n-side ASIC and (b) p-side ASIC. Note that the vertical axis has been magnified to highlight the structure.

2. To remove signal contamination in the background estimates, a second pass through all the readouts is made, and samples with ADC values larger than $\langle \text{ADC} \rangle_{\text{strip}} \pm 3\sigma_{\text{strip}}$ are not included in this averaging. This filtering is intended to remove any ADC values due to signal from particles passing through the silicon. $\langle \text{ADC} \rangle_{\text{strip}}$ and σ_{strip} are then recalculated, and are then a good determination of the baseline strip behaviour. Two filtering passes are usually sufficient, as only a small fraction of samples will contain any signal.

Figure 4.10 shows an example of $\langle \text{ADC} \rangle_{\text{strip}}$ values for an entire ASIC (one plot shows an n-side ASIC, the other shows a p-side ASIC). These values are calculated relative to the overall ASIC sample mean, $\langle \text{ADC} \rangle_{\text{ASIC}}$, for the readout:

$$\langle \text{ADC} \rangle_{\text{ASIC}} = \frac{1}{128} \sum_{\text{strip}=1}^{128} \langle \text{ADC} \rangle_{\text{strip}}. \quad (4.4)$$

The advantage of expressing the strip means, $\langle \text{ADC} \rangle_{\text{strip}}$, relative to the ASIC sample mean, $\langle \text{ADC} \rangle_{\text{ASIC}}$, is that a change in the reference voltage somewhere in the electronics chain (for example, noise on the common ground) can cause a dc offset in all the ADC samples. By determining the strip means with respect to the

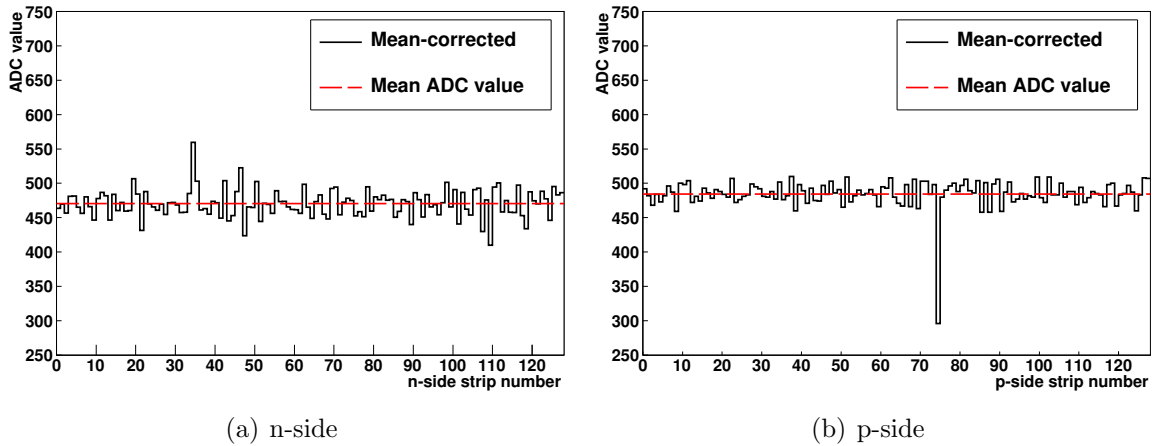


Figure 4.11: The strip samples after being mean corrected (from the same readout as Fig. 4.9) for an (a) n-side ASIC and a (b) p-side ASIC. The red line shows the mean ADC value.

ASIC sample mean, this dc offset can be corrected for on an event-by-event basis. Likewise, Fig. 4.10 shows that there is a microstructure present between the strips for any given ASIC. This microstructure is stable, and shows the variation of strip capacitance ASIC channel gain. This variation can be corrected for on a strip-by-strip basis by subtracting the strip mean from the raw readout. Figure 4.11 shows the application of these strip mean corrections to the raw strip samples.

3. Finally, the overall mean can be subtracted and the threshold can be calculated and applied. The ASIC sample mean for the readout, i , is calculated ($\langle \text{ADC} \rangle_{\text{ASIC},i} = 1/128 \sum_{\text{strip}=1}^{128} \text{ADC}_{\text{strip},i}$), and each sample ADC value is subtracted by this value. This is done to compensate for any shifting of the ADC baseline, and also to normalized the sample mean to zero. Figure 4.12 shows the result of performing this pedestal subtraction.

The signal threshold is then set to be $+3.75 \sigma_{\text{strip}}$ for n-side strips and $-3.75 \sigma_{\text{strip}}$ for p-side strips. Thus, any strip that has an ADC value that exceeds this threshold is considered to be signal. Figure 4.12 shows example readouts, including the signal

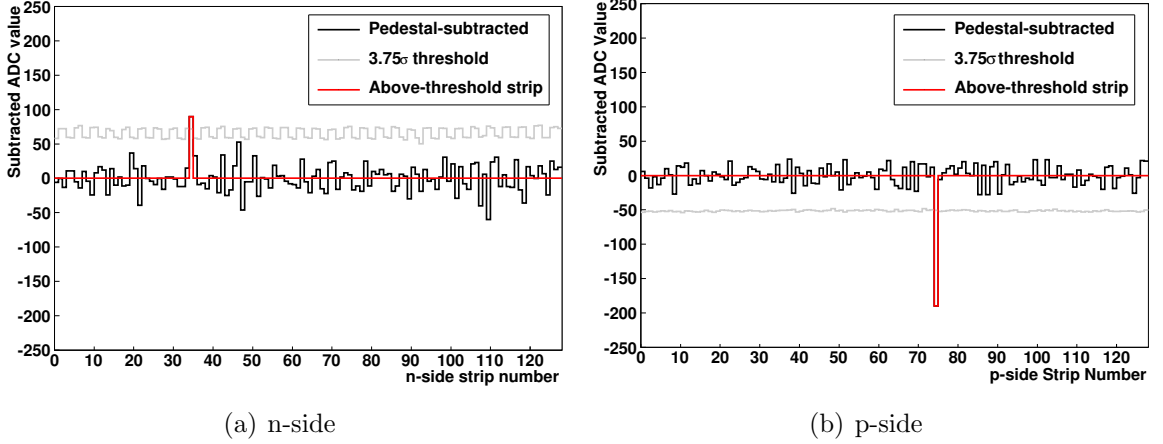


Figure 4.12: Final pedestal-subtracted strip samples (from the same readout as Fig. 4.9) for an (a) n-side ASIC and a (b) p-side ASIC. The strip samples are shown as the black trace, the 3.75σ threshold is shown as the grey trace, and the strips that pass the threshold are highlighted as the red trace.

threshold and highlighting strips crossing the threshold.

To visualize the result of this correction-subtraction-threshold method, Fig. 4.13 shows a histogram of the final corrected and subtracted ADC values for a large number of detector readouts (a full experimental cycle of ~ 2000 readouts). One feature to note are the peaks around $+550$ (for a)) and -500 (for b)), which result from the finite dynamic range of the VF48 digitizers. Pulse heights that exceed the maximum ADC value of 1024 (or are below the minimum of 0), are set to 1024 (or 0), and after corrections and subtractions, these end peaks are the result.

The distributions in Fig. 4.13 also show the division of signal from pedestal at the $3.75 \sigma_{\text{strip}}$ threshold. The pedestal background is well characterized by a Gaussian function, and as such, the expected number of strips exceeding the $3.75 \sigma_{\text{strip}}$ threshold (per ASIC), due to fluctuations in the background is,

$$\langle \text{strips} \rangle / \text{ASIC} = (128) \operatorname{Erfc} \left(\frac{3.75}{\sqrt{2}} \right) / 2, \quad (4.5)$$

$$= 1.1 \times 10^{-2} \text{ strips/ASIC}, \quad (4.6)$$




Figure 4.13: Histogram of the ADC values for an (a) n-side ASIC and (b) p-side ASICs over the course of an experiment. The light red shows the samples that passed the 3.75σ threshold, while the grey shows the values that did not pass the threshold. Note the log scale on the vertical axis.

where Erfc is the complementary error function. Although increasing the threshold will reduce the amount of background, this reduction is small compared to the amount of signal discarded, as seen in Fig. 4.13.

4.1.6.2 Hit positioning and strip clustering

The hit positions, that is, where the particle passed through the silicon modules, can be determined once the signal strips are identified. The location of each strip is well-known, so strips identified in Sec. 4.1.6.1 can be mapped directly into position space. However, particles traversing the silicon layer at a large incident angle can result in signal in more than one strip. In these cases, the charge cloud generated within the silicon depletion region is shared between adjacent strips. These groups of strips that share signal are referred to as strip clusters. An example is given as Fig. 4.14, where three adjacent strips all cross the hit threshold and form a cluster.

The hit position of a cluster can be better determined, as compared to the position of single strip, by weighting the strip positions by the ADC pulse heights. That is, the

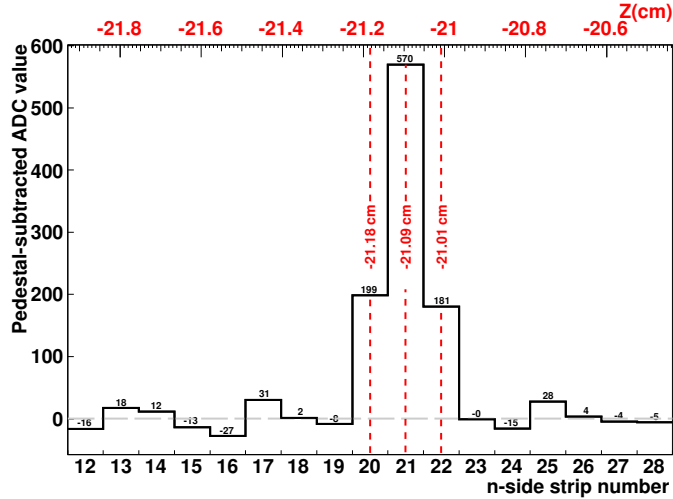


Figure 4.14: An example strip readout from part of the n-side of a detector module demonstrating charge sharing and strip clustering. Here a passing particle has resulted in signal in three consecutive strips. The red axis gives the position of the strips in the Z coordinate.

cluster position, x_{cluster} (be it either p- or n-side) is given as:

$$x_{\text{cluster}} = \frac{\sum_{i=1}^N (h_i)(x_i)}{\sum_{i=1}^N h_i}, \quad (4.7)$$

where h_i is the pulse height, and x_i is the position, of the i th strip. This allows for the hit position to be interpolated between the fixed strip locations. Moreover, the resolution of the cluster position can be also improved. When the entirety of the charge is collected by a single strip, the hit position is equally likely to have passed through the detector at any position in the range $[x - p_{\text{Si}}/2, x + p_{\text{Si}}/2]$, where x is the position of the center of the strip containing the charge, and p is the strip pitch. Thus, the resolution of this position is given by the variance in a uniform distribution with a finite range, or $\sigma_x^2 = p_{\text{Si}}^2/12$. On the other hand, when using the weighted mean interpolation the precision is set by how well the analog signal is determined. In this case, the resolution of the position is $\sigma_x^2 \approx p_{\text{Si}}^2/(S/N)^2$, where S/N is the ratio of signal to noise [110].

Once the strips are clustered, the hit can be determined in the local frame of the detector. In this frame, the p-side strips define the axis along the short side of the module,

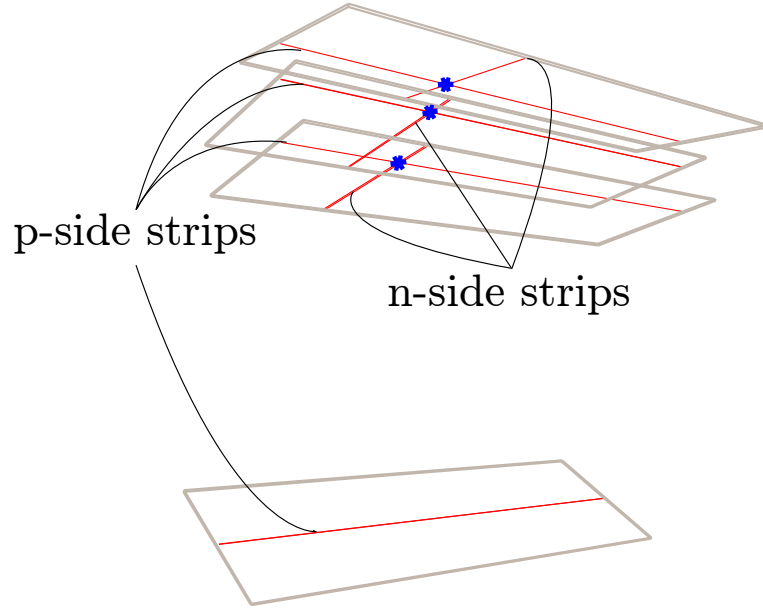


Figure 4.15: Spatial representation of an example event showing the combination of p- and n-side strips that produces three-dimensional set of position data for the hits. The active silicon is shown as the grey blocks, the strips as the red lines, and the hits as the blue markers.

while the n-side strips define the orthogonal axis along the long side. The position along the p-side axis can be combined with the position along the n-side axis to fully define the hit position in the plane of the detector. Likewise, since the placement of the module is known, the hit can be transformed from the local into the global coordinate frame with a rotation and a translation. Figure 4.15 shows an example event and highlights how both p- and n-side strips are required (along with good knowledge of the detector geometry) to fully localize the hit position.

Using Eq. 4.5, an estimate of the number of hits per event due to fluctuations in the ADC pedestal (exceeding the $3.75\sigma_{\text{strip}}$ threshold) is given by,

$$\langle \text{hits} \rangle / \text{event} = N_{\text{modules}} \langle \text{strips} \rangle_{\text{p-side}} \langle \text{strips} \rangle_{\text{n-side}} \quad (4.8)$$

$$= (60 \text{ modules}) \left((128)(2) \text{Erfc} \left(\frac{3.75}{\sqrt{2}} \right) / 2 \right)^2, \quad (4.9)$$

$$= 3.1 \times 10^{-2} \text{ hits/event}. \quad (4.10)$$

This suggests that the rate of pedestal-induced hits is very low. However, spurious hits can also be formed from signal strips combining with pedestal-induced strips. For this reason, it is important to have a robust placement of the signal threshold, which balances good signal acceptance with adequate noise rejection.

4.1.6.3 Cosmic rays and hit efficiency

Some initial measurements of the detector orientation and performance can be completed using the tracks of cosmic rays passing through the detector. Because of their ubiquity (see Sec. 6.1.3), collecting cosmic ray data is simply a matter of operating the detector without any antiparticles present. Cosmic ray particles that reach the detector (having passed through the ceiling of the AD hall) are dominated by Minimum-Ionizing Particles (Appx. A), typically muons.

MIPs that pass through the detector and satisfy the $S_i > 1$ readout trigger will generally result in 5 or 6 hits. Despite the large solenoidal magnetic field and scattering material present in the apparatus, the momentum of the MIPs is typically large enough that there is minimum track curvature, and the track is well represented by a straight line segment. The angular distribution can be determined by fitting a line to the event and extracting the angular information from the fit. Figure 4.16 shows the experimental angular distribution and a fit to the expected $\cos^2 \theta_v$ distribution, where θ_v is track angle from vertical. The angular displacement is consistent with a rotational offset in the placement of the detector of 7.2 degrees. This offset can be corrected for by adjusting the software geometry to reflect the actual positioning of the detector.

The analog hit efficiency can also be estimated using cosmic ray tracks. By selecting events with well-defined cosmic ray tracks, the trajectory of the cosmic ray can be extrapolated across the detector. The intersection points between the cosmic track and silicon modules are all examined, and any of the intersection points along the particle




Figure 4.16: Histogram of the angle (from vertical) of cosmic tracks passing through the detector. The smooth curve shows the $\cos^2 \theta_v$ fit.

trajectory that failed to register hits are counted as inefficiencies. This measures the efficiency of the analog readout and cluster determination – a silicon module that failed to collect enough charge to exceed the signal threshold, or mistiming in the sample and hold or multiplexing which lead to the absence of signal in the proper strip, for example, will all be counted as an analog hit inefficiency. The positions and active areas of the strips are modelled in our software, taking into consideration gaps between wafers and non-active areas of the silicon wafer. Only the active area of the strips is considered, which helps to ensure that any inefficiencies are the result of the strips themselves, and are not due to geometric effects resulting from tracks passing through non-active areas (edges of the modules, readout electronics, etc.).

To perform this measurement, cosmic events are first fitted with a line segment using the technique as outlined above. Only events with a well defined track containing 5 or 6 hits are considered. For each of these events, each silicon module is checked to see if it intersects with the cosmic trajectory and whether there are p-side or n-side clusters located at that intersection point. The efficiency can then be calculated as:

$$\text{Hit Efficiency} = \frac{\text{Events that contain a hit at the intersection point}}{\text{Events where a hit was expected}}. \quad (4.11)$$

Figure 4.17 shows this measurement as a function of module, and making the distinction between n-side and p-side. On average, the efficiency for the p-side is $\sim 99.5\%$, and $\sim 98.5\%$ for the n-side. The total hit efficiency then roughly the p-side efficiency times the n-side efficiency times the geometric efficiency, $0.995 \times 0.985 \times 0.9 = 0.88$, or $\sim 88\%$. Note that there is a trigger bias towards the p-sides for modules 1-8 and 31-38, which are included in the readout trigger. This means that events are only read out when there is a p-side trigger in more than one of those modules, and as such, the readouts are biased towards events that contain hits in those modules. Also, as shown in Fig. 4.5, several modules have a vertical, or near-vertical alignment, which drastically decreases




Figure 4.17: Plots of the p-side (black) and n-side (red) strip efficiencies per module, calculated by looking for where a cosmic track passed through, but did not result in a strip cluster about threshold. Note that the vertical axis starts at 96.5%. Data for Module 10 is missing, as it was disconnected from the data collection and analysis system (see footnote in Sec. 4.1.4).

the cosmic ray flux for these modules. This decreases the statistics for these modules, which is reflected in their large error bars. Another possible systematic error is due to the fact that this analysis only considers single-hit inefficiencies (that is, if a coincidence of inefficiencies caused two or more hits to be missed, that event would not be counted in either the numerator or denominator of Eq. 4.11). Although this would lower the measured efficiency, it is expected to be a small effect, as it would require the coincidence between two strip inefficiencies, each of which is individually a low probability event.

4.1.7 Summary of detector hardware

The ALPHA detector consists of an array of double-sided silicon microstrip modules arranged in three layers around the trapping region of the ALPHA apparatus. Each detector module contains 256 n-side and 256 p-side strips, which are managed by four VA1TA ASIC read out chips. The ASICs, in turn, are powered and controlled by an external FRC board, which communicates with the TTC control module and the VF48 signal digitizers. The digital triggers are processed and combined with the TTC module,

and the final readout trigger decision is made using NIM logic modules. The analog strip signals, representing the passage of charged particles, are extracted using an algorithm involving several filtering iterations and the determination of signal threshold levels for the individual strips. The final result of the analog detector readout is an readout event containing some number of hits due to charged particles passing through the detector modules.

4.2 Monte Carlo simulation

This section describes the Monte Carlo simulation used to emulate antiproton annihilation and the resulting particle propagation through the ALPHA apparatus and detector. Details of the simulation geometry and materials are provided along with details about the event generator and particle transporter. Finally, the digitization of the particle trajectories according to the physics of the silicon detectors is described.

4.2.1 Overview

To study the performance of our detector, we can simulate the passage of particles through the apparatus and detector using Monte Carlo (MC) methods. In this type of Monte Carlo simulation, particles are assigned initial conditions randomly, according to a pre-defined distribution. The system is then evolved; in our case, the particle trajectories are followed according to known physics, with the stochastic processes simulated by randomly sampling of known distributions. If a large number of simulations are performed, the range of possibilities stemming from the initial distribution can be explored.

In our case, detailed detector simulations were mainly used to evaluate the performance of the annihilation reconstruction algorithms (Chapter 5). Since it is not always possible to sufficiently control the experimental annihilation distribution, it can be difficult to measure the accuracy of the reconstructed annihilation position. However, by

using a simulation, the initial conditions are always known, along with the particle identities and trajectories. In this way, the reconstructed position can be compared with the actual simulated initial position, and the accuracy of the reconstruction evaluated.

The GEANT3 [140] software package along with the ROOT [141] Virtual Monte Carlo (VMC) interface [142] provide the framework to generate antiproton annihilations and simulate the passage of the resulting particles through the apparatus and detector. GEANT3 (GEometry ANd Tracking) is a suite of tools, written primarily at CERN in FORTRAN, to simulate the passage of elementary particles through matter, complete with methods for specifying complex detector geometries. ROOT is a data analysis framework and, with the VMC library, provides a C++ interface to the GEANT3 package.

Each annihilation is treated as a separate and independent ‘event’. The prescription for our Monte Carlo detector simulation is as follows:

1. Initialize the material definitions and geometry (Sec. 4.2.2).
2. Generate the initial position and momenta of the particles (Sec. 4.2.3).
3. Simulate the trajectories of all the particles and any secondary particles (Appx. C).
4. Record any energy deposition within the detector modules (Sec. 4.2.4).
5. Repeat steps 2-4 until a representative sample has been obtained.

Steps 1-4 will be detailed in the sections to follow, while the use of the Monte Carlo simulation to evaluate the primary annihilation position reconstruction will be discussed in Chapter 5.

4.2.2 Materials and geometry

Before particles can be transported within the simulation, the material properties and configuration must be specified. In GEANT3, the simulation geometry is defined as a collection of volumes. Each volume has two sets of parameters: the first being strictly

Description	Material	Z_{eff}	A_{eff}	Density ρ (g/cm^3)	Radiation Length X_0 (cm)
Electrode stack	Aluminium T6082	13.47	28.02	2.7	8.65
Magnet form, Vacuum walls	Stainless Steel 316LN	26.25	56.71	8.0	1.72
Left mirror, Right mirror, Octupole	Cu-NbTi	32.96	73.58	8.0	2.10
Liquid He chamber	Liquid He	2	4.00	0.125	684
Detector modules	Silicon	14	28.9	2.33	9.35

Table 4.2: Geometry element characteristics included in the Monte Carlo simulation.

material properties (effective atomic number, Z_{eff} ; effective atomic weight, A_{eff} ; density, ρ ; and radiation length, X_0), while the other set of parameters affect particle transport regardless of material (magnetic field, tracking precision, energy cut offs). Table 4.2 outlines the volume elements included in the simulation, and Fig. 4.18 and 4.19 show the placement of each element. For most of the simulations, the non-material parameters were fixed for all volumes, e.g., a 1 T solenoidal magnetic field, 5 degree per step maximum angular deflection due to magnetic field, 1 mm maximum and 10 μm minimum step size, 5% maximum fractional energy loss per step, and 5 μm boundary crossing precision.

4.2.3 Event generation

To start the Monte Carlo simulation, it is necessary to specify the number of initial particles, along with their positions and momenta. Since the ALPHA detector is primarily sensitive to charged particles resulting from antiproton annihilations, it is sufficient to focus on generating particles produced through antiproton annihilations (Sec. 2.3.2).

Although in the ALPHA apparatus the majority of antiprotons annihilate on the nucleus of a heavy atom (that is, on the gold-plated aluminum electrode surface), this




Figure 4.18: Quadrant of the Monte Carlo radial geometry (to scale). Uncoloured volumes are treated as vacuum.




Figure 4.19: Length-wise cross-section of the Monte Carlo simulation geometry. Un-coloured volumes are treated as vacuum.

process is not well characterized experimentally. Following [143], the well-studied pion distribution from $\bar{p}p$ annihilation at rest (Table 2.2) is used as the primary generator of high-energy charged and neutral particles with which the Monte Carlo simulation is initialized. This is an adequate approximation, in which nuclear interactions involving the antiproton with the atoms of the electrode stack are neglected.

Events are generated in proportion to the ratios given in Table 2.2, resulting in the pion multiplicity distribution as shown in Fig. 4.20. Each event must conserve the initial energy (~ 2 GeV) and momentum ($\sum_i \vec{P}_i = 0$) of an antiproton-proton system annihilating at rest. This is accomplished using the ROOT *TGenPhaseSpace* routine, based on [144], which returns a kinematically valid N-body configuration by iteratively splitting the system into two-body decays in their center of mass frame.

Although the *TGenPhaseSpace* method will return a kinematically valid event, the specific configuration generated may not be particularly likely. For example, an event might be generated with 7 pions, such that all the pions are produced with co-linear momentum, with 6 pions all travelling in one direction, and the 7th pion balancing the net momentum and travelling in the opposite direction. While this event is realizable,

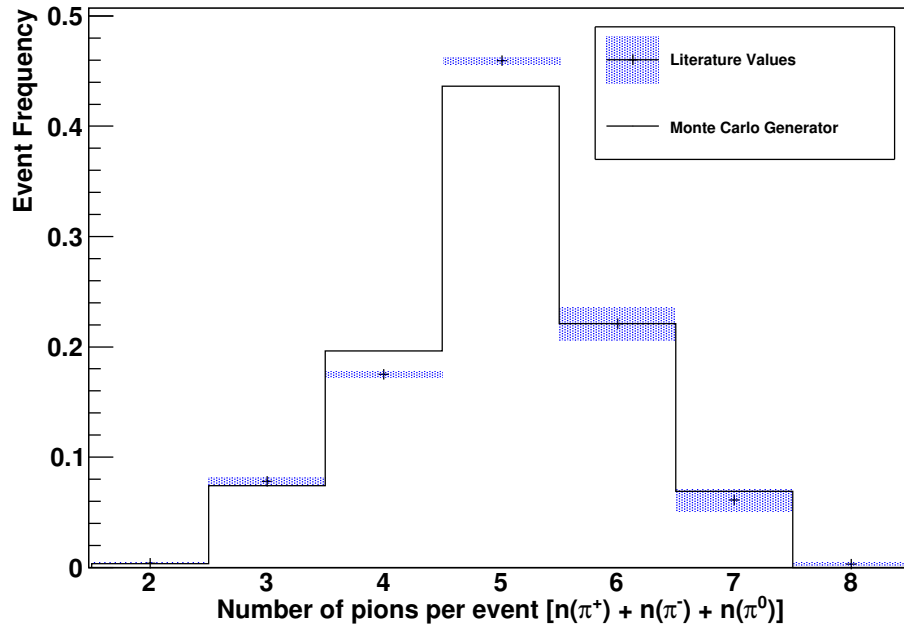


Figure 4.20: An event frequency histogram showing a comparison of the pion multiplicity resulting from the Monte Carlo event generator (black histogram) to the literature values (blue shaded squares) [106].

it is much less likely than, say, the 7 pions being distributed roughly isotropically. To manage for this, the *TGenPhaseSpace* routine also returns the weighting of the generated event. Configurations located in unlikely corners of the phase-space are assigned small weightings, while every $\pi^+\pi^-$ event, for example, will have a weight of 1, because the π^+ and π^- will always be emitted back-to-back, and every orientation is equally likely. To produce an unweighted sample, where every event is equally likely (though still generated at a rate proportional to its branching ratio), for each event a random number r is generated between zero and the maximum weight of the event, W_{\max} . If the event weight W_i is less than r , the event is rejected and another is generated (and its weight tested again). The result of this is that the probability of keeping an event is equal to its weight, and as such, the accepted events are all on the same footing.

Once the initial particles and their momenta have been specified, their passage through the apparatus geometry is simulated using the GEANT3 particle propagator (described

in Appendix C). Figure 4.21 shows the mean number of particles per event that are generated by various processes simulated by GEANT3. Shown in another way, Fig. 4.22 also shows the mean number of particles per event, but the particles that interact with the detector volumes are identified. Of particular interest are the π^+ and π^- , along with the e^- and e^+ , which all contribute heavily (1.06, 1.05, 0.888 and 0.879 particles per event, respectively) to the reconstructable detector signal (which will be discussed in Chapter 5).

4.2.4 Digitization

The connection between the GEANT3 simulation and what is actually measured by the detector is made in the digitization stage. This involves taking the position where the particle entered and left the detector volume, and distributing the deposited charge over the appropriate strips.

Figure 4.23 outlines the scheme used to simulate the detector response. The GEANT3 simulation provides the position of the particle as it enters and leaves the detector volume as well as the amount of energy deposited within the volume, which is proportional to the number of electron-hole pairs generated, and thus proportional to the signal recorded by the detector. In the case that the particle crosses several strips, the signal is distributed according to the fraction of the strip over which the particle traversed. This approximates the charge sharing that happens within the silicon depletion layer.

This scheme relies on a well defined length-wise extent of the charge cloud generated in the particle passage. However, due to diffusion, the electrons and holes will spread perpendicular to the direction of the strips. The length-wise spread, σ_x of the charge cloud can be estimated as [110, 145]:

$$\sigma_x = \sqrt{2\mathcal{D}t_{\text{drift}}} \approx \sqrt{2 \frac{k_B T}{e} \frac{h}{V_{\text{dep}}}}, \quad (4.12)$$




Figure 4.21: Generation process vs. particle type illustration of the various processes in GEANT3 [140], which generates particles during the Monte Carlo simulation. This data was collected by simulating 10000 proton-antiproton annihilations and recording the generating process for every simulated particle. Then mean number of particles per event per process is shown as colour and bin content.




Figure 4.22: Mean number of particles per event involved in the Monte Carlo simulation. This data was collected by simulating 10000 proton-antiproton annihilations, tracking every produced particle through its simulated trajectory, and determining the mean number of simulated particles which passed through each region. The yellow histogram shows all the generated and tracked particles. Also plotted are the mean number of particles per event which deposit energy in the 1st layer (blue), 1st and 2nd layers (green), and all three layers (brown). Bin contents are only shown for the yellow and brown histograms for clarity.




Figure 4.23: Cartoon diagram illustrating how the digital signal is produced by particles passing through a silicon detector module during the Monte Carlo simulation. The simple distribution of deposited energy to approximate charge sharing is shown by the grey shaded area.

where the diffusion constant in silicon, \mathcal{D} , is given as $\mathcal{D} = \frac{kT}{e} \mu_e$, with kT being the temperature (in eV), e is the charge of the electron, and μ_e is the electron charge mobility in silicon. The drift time, t_{drift} in the Eq. 4.12 can be expressed as $t_{\text{drift}} = \frac{h/2}{\mu_e V_{\text{dep}}/h}$, where h is the thickness of the silicon layer, and V is the depletion voltage. The approximation comes from using the mean electric field magnitude, $\langle E \rangle = V_{\text{dep}}/h$, and average drift time. Using the values $h = 300 \mu\text{m}$, $T = 300 \text{ K}$, and $V_{\text{dep}} = 65 \text{ V}$, the spread of charge due to diffusion is roughly $\sigma_x \approx 8.5 \mu\text{m}$. As such, the diffusion is minimal compared to the strip pitches of $227 \mu\text{m}$ and $875 \mu\text{m}$, and can safely be ignored.

Figure 4.24 shows the comparison between the simulated and experimental distributions of the number of strips included in a hit. The simulation result is tuned by rejecting strips that fall below a user defined fraction of the total energy deposited. This is to approximate the effect of the ADC threshold, where a hard 3.75σ cut is made. The effort was made mostly to reproduce the ratio of hits with one strip versus the number of hits

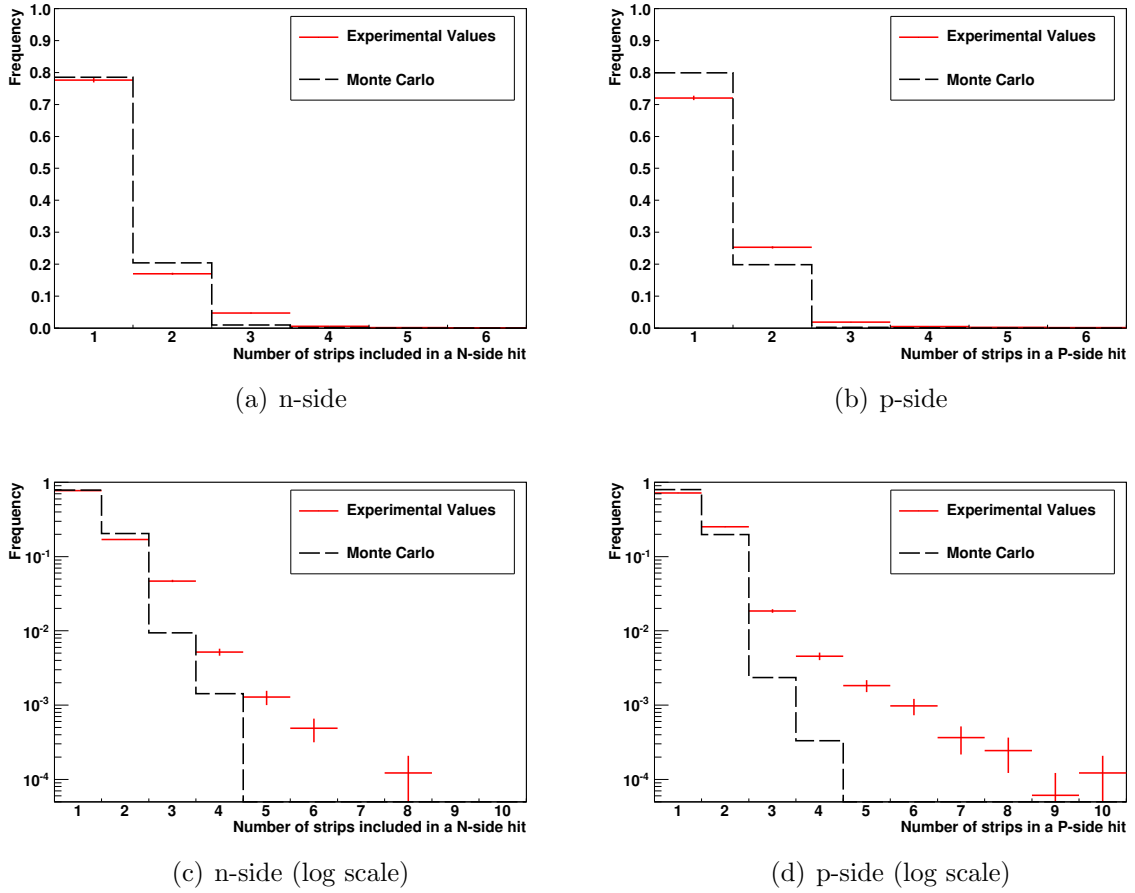


Figure 4.24: Histograms of number of strips per cluster for the (a) n-side ASICs, and (b) a p-side ASICs. The same data is also shown on a log scale, and an expanded x-axis, for the (c) n-side ASICs, and (d) p-side ASICs.

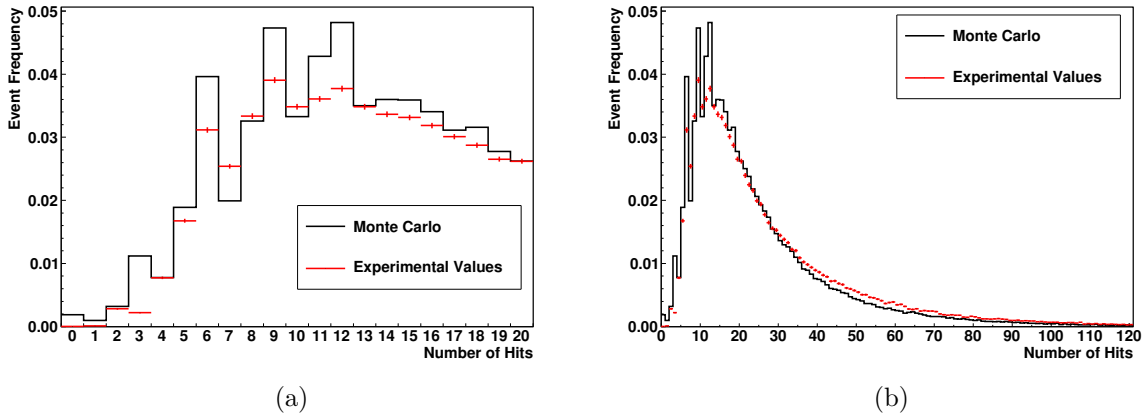


Figure 4.25: Comparison of the number of hits measured during antiproton annihilation (red) to the number of hits found in the Monte Carlo simulation (black). On the left (a), the distribution is shown for 0-20 hits per event, while (b) extends the distribution to 120 hits per event.

with two strips, as these bins represent the majority of hits. Figures 4.24 c) and d) both show that there are higher order contributions to the experimental values (that is, hits that contain five or more strips), which are rare and not accurately modelled with these simulations. It is possible that these events are the result of electromagnetic showers (Fig.6.2), where many electron-positron pairs are generated, and will pass through the same region of the detector. However, given the log scale, this higher-order contribution is small, and can be safely ignored.

Finally, Fig. 4.25 compares the distribution of number of hits generated in the Monte Carlo simulation to the experimental values measured for antiproton annihilation. Both figures show the same distribution, with Fig. 4.25 a) magnifying the 0-20 hit range, and Fig. 4.25 b) showing the extended distribution out to 120 hits per event. Note that the spiky structure in the Monte Carlo distribution for low numbers of hits comes from simulated events which only have hits from charged particles, resulting in events with 3, 6, 9, or 12 hits (that is, three hits per charged particle track). The simulation does not include: noise from the environment or electronics, which can result in spurious hits; or

contamination from cosmic ray particles, which can leave tracks with more than three hits. These contributions can wash away the spiky structure, which is not as prevalent in the measured distribution. Overall, the distribution for Monte Carlo events compares well enough for the intended purpose with the measured distribution.

4.3 Summary

This chapter discussed the ALPHA silicon detector, focusing on the analog readout of the signal microstrips. The detector hardware and readout chain is described in detail, including the readout chips, triggering control module and digitizers. The method used for extracting the particle signal from the analog readout, along with a description of the hit position determination, is then detailed. An example is given which estimates the hit efficiency of the silicon modules using cosmic ray data.

The second part of this chapter describes the Monte Carlo simulation of the silicon detector and surrounding apparatus. First, the materials and detector geometry used in the simulation are given. Then the generation of the antiproton annihilation events is detailed. Finally, the digitization of the particle trajectories into detector signal is described.

The detector design presented here represents a system optimized to resolve antiproton annihilation products. Likewise, the results of the Monte Carlo simulations indicate that the apparatus geometry and the physics behind the operation of the silicon detector is sufficiently well understood. With this in-hand, the task of determining the annihilation vertices could then be attempted.

Chapter 5

Event reconstruction in the ALPHA detector

The ALPHA silicon detector is a tracking detector. That is, using the position-sensitive capabilities of the silicon detector modules (Sec. 4.1.1), the particle trajectories can be determined, or reconstructed. Moreover, by examining the tracks from several particles, their common origin (called the vertex position) can be determined. Tracking detectors are widely used in high energy physics [146, 147] and astrophysics [148, 149].

This chapter will describe the reconstruction methods used to characterize an event and determine its annihilation position. The challenging aspects of the reconstruction, due to the specifics of the ALPHA detector and apparatus, are first outlined. The rest of the chapter covers, in detail, the various steps involved in reconstructing the charged particle tracks and annihilation vertices. Finally, to demonstrate the utility of the annihilation reconstruction methods, several example vertex distributions are presented and discussed.

5.1 Challenges for the ALPHA event reconstruction

There are two major challenges for the reconstruction, due to the specifics of the ALPHA apparatus and detector:

1. *There is a large amount of scattering material between the annihilation point and the detector.* In order to make an accurate determination of the vertex position, the particle trajectories need to be followed back to their origin. When there is dense material along the particle path, it can undergo scattering (both small angle multiple scattering and large angle hard scattering). Figure 5.1 illustrates the

possible effect of scattering on the vertex reconstruction. The largest contributor of scattering material is the superconducting octupole magnet (Figure 4.18 and Table 4.2). The resolution of the vertex position can be degraded, as tracks that cover a large distance before scattering (the lever arm) will not project accurately back to the original annihilation point.

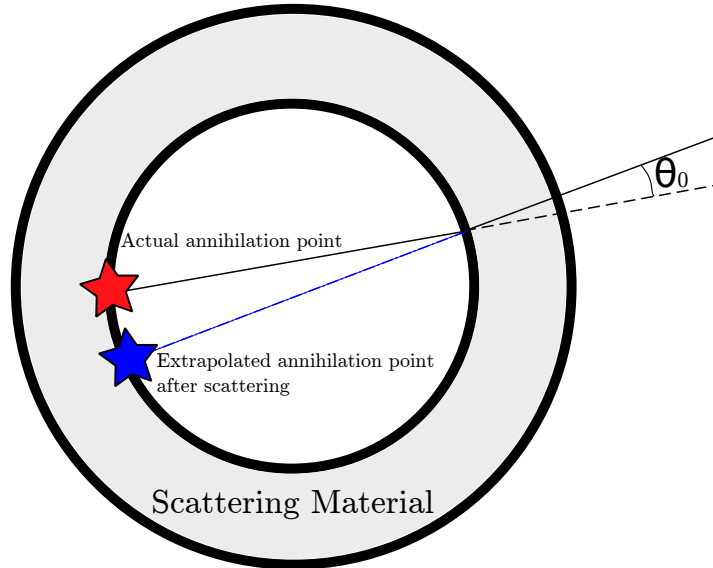


Figure 5.1: Illustration of the effect of multiple scattering on the track extrapolation. The actual annihilation position is shown as the red star, whereas the extrapolated vertex is given as the blue star.

Eq. A.3 gives the expected scattering angle for a charge particle undergoing small-angle multiple scattering. This formula can be applied to the ALPHA geometry, where the combined radiation length of the material radially between the trap volume and detector is $(x/X_0)_{\text{total}} \sim 0.7$ (this includes the octupole windings, although it is possible for a particle to travel radially and pass between the octupole windings). Therefore, for a particle with $z = 1$, $\beta \sim 1$, $P \sim 300 \text{ MeV}/c$, a deflection of $\theta \simeq 39 \text{ mrad}$ is found. If the particle travels diametrically across the trap volume before passing through the scattering material, the lever arm can span close to 4 cm, which corresponds to $\sim 2 \text{ mm}$ deviation from the proper annihilation point.

This effect can vary substantially depending on the amount of scattering material encountered by the particle. Specifically, a particle with a large axial component to its trajectory will pass through more scattering material than a particle traveling radially outward. The expected scattering then goes as $(\cos \theta_{\perp})^{-1/2}$, where θ_{\perp} is the incident angle from perpendicular. A particle traveling outwards with a 45 deg angle from perpendicular would encounter about 40% more scattering material, and would have a 20% greater expected scattering angle than a particle passing perpendicularly through the module.

The large amount of material also provides ample opportunity for gamma rays to convert to electron-positron pairs. These gamma rays result from the decay of neutral pions released during antiproton annihilation ($\pi^0 \rightarrow 2\gamma \rightarrow e^+e^-$). The electron and positron will often be energetic enough to travel outwards through the entire detector, resulting in a particle tracks similar to those resulting from charged pions. However, the electron and positron tracks do not necessarily follow back to the primary annihilation point; rather, these tracks will converge to where the pair-production occurred. These tracks can degrade the quality of the vertex determination, as the track trajectories do not follow back to the primary annihilation position.

2. *The detector is limited to three layers.* The motivation for including three detector layers, as opposed to the two layers of silicon modules used in the ATHENA experiment [150], is to be able to sample the curvature of the particle bending in the magnetic field [151]. This allows for a much better measure of the particle trajectory, and consequently, a better determination of the annihilation point.

The radius of curvature of the track, ρ_c , which can also be used to determine the

transverse particle momentum with the relation [9],

$$B\rho_c [\text{T} \cdot \text{m}] = 3.3356 P_T [\text{GeV}/c], \quad (5.1)$$

where B is the magnetic field magnitude (in Teslas), ρ_c the radius of curvature of the particle (in meters), and P_T is the transverse momentum (in GeV/c). For example, with a field of $B = 1$ T, and a particle momentum of $P_T \sim 200$ MeV/c, the radius of curvature is $\rho_c \sim 0.7$ m.

A good determination of the track momentum would be extremely useful. For example, tracks originating from an antiproton annihilation are constrained to have momentum $P_T \lesssim 2$ GeV/c, or $\rho_c \lesssim 6$ m (this comes from the total energy available from antiproton annihilation at rest being $E = 2m_p c^2$). With such a constraint, tracks with much larger momentum could be dismissed, as they would likely be due to high energy cosmic rays. Momentum distributions of various simulated particles produced in antiproton annihilations are compiled in Appendix D. The distributions show a large spread of particle momenta which corresponds to a large range of track curvature. It is important to note that a two-layer detector would be unable to determine any track curvature for outgoing particles – as such, the three-layer detector configuration gives a substantial advantage in determining the particle trajectories.

However, the accurate determination of particle momentum is still limited by the number of samples of the curvature. For a particle exiting the apparatus radially, the radius of curvature will be sampled by the p-side detector strips once at each layer. This amounts to three measurements of the particle position in the radial projection, which is sufficient to determine the circle parameters, including the radius of curvature. The momentum resolution is then limited by the error in the sagitta (the depth of the measured arc), and the fractional momentum resolution

can be given as [152]:

$$\left(\frac{\sigma_{P_T}}{P_T}\right)_{\text{sagitta}} = \frac{\sqrt{3/2}\sigma_x 8P_T}{0.3BL^2}, \quad (5.2)$$

where L is length over which the trajectory is sampled, and σ_x is the position resolution of the samples. For a 300 MeV/c particle in the $B = 1$ T field, over the $L = 4$ cm sampling area, and with $\sigma_x = (227 \mu\text{m})/\sqrt{12}$ (Sec. 4.1.6.2), the fractional momentum resolution is $\sigma_{P_T}/P_T \sim 30\%$. This large uncertainty in the momentum makes it difficult to invoke any kinematic constraints when examining the tracks.

Despite these challenges the annihilation point can be located for a large fraction of events. The procedure for determining the annihilation location is described in the sections to follow.

5.2 Overview of the event reconstruction

Each event reconstruction follows the same general procedure, which proceeds roughly as follows (note that, as described earlier, a ‘hit’ is defined as the intersection of the n-side and p-side strips, and here is taken as the three-dimensional location in the global reference frame of the detector):

1. *Enumeration and filtering of track candidates* (Sec. 5.3). The total set of hits is assigned into groups of three hits, each in a different layer, and each group is called a track candidate. Each track candidate is evaluated based on the separation and correlation of its hits. Candidates that fail to meet some initial filtering conditions are rejected.
2. *Determining the helix parameters* (Sec. 5.4). The helix parameters are determined for all of the remaining track candidates.

3. *Helix pruning* (Sec. 5.5). Track candidates that do not form good helices are discarded. Likewise, helices that do not approach the trapping volume are discarded. Finally, when two or more helices share hits, the straightest track is kept, and the others discarded.
4. *Vertexing* (Sec. 5.6). Using the remaining helices, the vertex position is determined.

The overarching goal of the reconstruction is to return as many well-reconstructed events as possible. This requires robust algorithms which can deal with a wide range of events and with varying numbers of hits. These algorithms can be characterized and evaluated using the Monte Carlo simulation, then tuned and optimized using experimental data.

5.3 Enumeration and filtering of track candidates

This section will describe how the hits in an event are divided into track candidates. The initial criteria for filtering track candidates, based on their hit separation and correlation, will then be described.

In the ALPHA detector, a track candidate is defined as a group of three hits, where each hit is on a separate detector layer. This follows from the three layer detector configuration, wherein a particle travelling radially outwards from the inner trapping volume should pass through three detector modules before exiting the detector (detector inefficiencies and geometric acceptance can lead to trajectories with fewer hits on silicon modules). To this end, only three-hit tracks are considered, and combinations with fewer (or greater) hits are ignored.

There are many methods for recognizing particle tracks used by large-scale detectors [153, 154]. These methods deal mainly with the challenge of handling events with large numbers of tracks, and therefore, a large number of hits over many layers. This can

be combinatorially challenging, as an event with a large number of hits can have many combinations, and many possible track configurations. Fortunately, the number of tracks in an antiproton annihilation at rest in the ALPHA apparatus is small enough that the events are combinatorially tractable. That is, every combination of three hits can be evaluated individually, and combinations unlikely to have resulted from charged particles can be rejected.

Figure 5.2 and Table 5.1 give an example event, demonstrating a simple case of enumerating and filtering track candidates. In this example event, there is only one track, and therefore, only one correct combination of hits (*EFG*). By iterating over all the combinations of three hits, this track candidate can be identified, and the other combinations filtered out based on how closely their hits are clustered together.

The number of combinations of three hits, and therefore track candidates, for an event with N hits is given as ${}_N C_3 = \frac{N!}{3!(N-3)!}$. Likewise, the number of combinations where each layer has only one hit is given as $C_{\text{one per layer}} = M_1 M_2 M_3$, where M_1 , M_2 , and M_3 are the number of hits in the first, second, and third layers, respectively. The example event, with 7 hits, has ${}_7 C_3 = 35$ combinations (all of which are listed as Table 5.1), and there are $C_{\text{one per layer}} = (3)(1)(3) = 9$ combinations that have a hit in each of the three layers. However, several of the 9 combinations come from grouping together hits from opposite sides of the detector, and it is unlikely that they are due to a single particle. As such, combinations with large separation between hits should be rejected.

The optimal hit separation criteria can be investigated using the Monte Carlo simulation. Since the simulation keeps track of all the generated particle trajectories, and where they form hits in the simulated detector, the hit separation distribution can be generated and investigated. Figure 5.3 shows the distribution of the displacement between hits (that is, $|z_{\text{layer } 1} - z_{\text{layer } 2}|$ and $|z_{\text{layer } 2} - z_{\text{layer } 3}|$) along z for all the simulated particles passing through all three detector layers (red), only the charged pion tracks passing

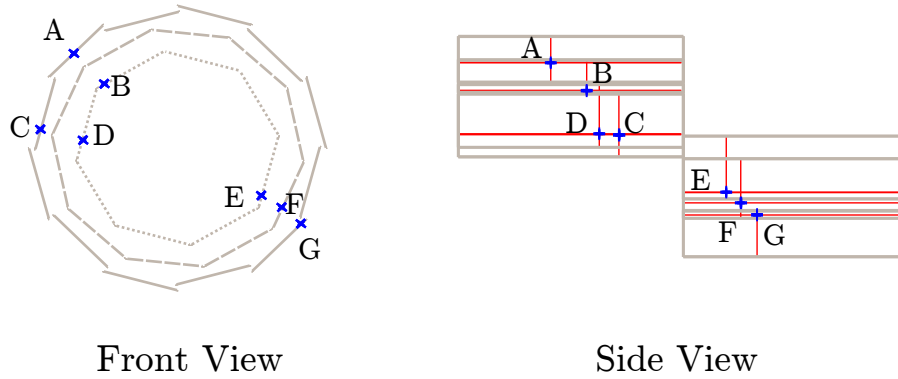


Figure 5.2: Spatial illustration of an example event, shown as an unsorted collection of hits. The hits are labelled sequentially (A-G), and in the Front view (left), the detector layers are shown in different line styles (dotted lines for layer 1, dashed lines for layer 2, and solid lines for layer 3). For brevity and clarity, only the modules with hits are shown in the side view (right).

Combination	Different Layers?	Small Hit Separation?	Combination	Different Layers?	Small Hit Separation?
ABC	-	-	BCG	-	-
ABD	-	-	BDE	-	-
ABE	-	-	BDF	-	-
ABF	✓	-	BDG	-	-
ABG	-	-	BEF	-	-
ACD	-	-	BEG	-	-
ACE	-	-	BFG	✓	-
ACF	-	-	CDE	-	-
ACG	-	-	CDF	✓	-
ADE	-	-	CDG	-	-
ADF	✓	-	CEF	✓	-
ADG	-	-	CEG	-	-
AEF	✓	-	CFG	-	-
AEG	-	-	DEF	-	-
AFG	-	-	DEG	-	-
BCD	-	-	DFG	✓	-
BCE	-	-	EFG	✓	✓
BCF	✓	-			

Table 5.1: List of all possible hit combinations in the example event of Figure 5.2. The combinations are referenced by the hit labels given in the above figure. If each hit is found in a different detector layer, a check mark is given under the ‘Different Layers?’ column. Likewise, if the hits clustered together in position, a check mark is given under the ‘Small Hit Separation?’ column. The lone combination that passes both tests (EFG) is highlighted in bold-face font.

through all three detector layers (black), and the hit displacement in z for all of the track candidates (blue). Since the charged pions from antiproton annihilation will extrapolate back to the annihilation point (disregarding multiple scattering and energy loss), the track candidate filter threshold should be set to reject candidates with hit displacements outside what is expected for charged pions. To that effect, any track candidate with a z displacement > 6 cm is rejected. This criterion amounts to removing candidates with large angles from radial ($\theta \gtrsim 72$ deg, with a corresponding reduction of 5% in the total solid angle). Although these candidates do sometimes correspond to particles with very shallow trajectories, these trajectories are difficult to detect and reconstruct as the particles pass through a large amount of scattering material over a long path-length.

Similarly, Figure 5.4 shows the distributions of azimuthal separation between hits (that is, $|\phi_{\text{layer 1}} - \phi_{\text{layer 2}}| + |\phi_{\text{layer 2}} - \phi_{\text{layer 3}}|$), for all the simulated particle tracks (red), only the charged pion tracks (black), and all the track candidates (blue). Candidates with azimuthal separations larger than expected for charged pion tracks (in this case > 0.35 radians) are rejected. This measure of azimuthal separation is small when the three hits are aligned radially. On the other hand, a large separation usually requires the hits to form a tangent to the radius, or be on opposite sides of the detector (both situations being undesirable). Another possibility for a large hit separation using this measure involves the hits on layer 1 and 3 to be radially aligned, and the layer 2 hit to be azimuthally shifted by more than 0.175 radians. However, this means the candidate has a very small radius of curvature, and is unlikely to be the result of a particle originating from within the apparatus.

Being closely separated in space does not ensure that three hits, combined, form a good track candidate. Particles with momentum ~ 100 MeV, and thus a radius of curvature ~ 1 m in the 1 T field, will produce tracks through the detector that have only slight curvature in the radial projection of the magnetic field, and are extremely straight.




Figure 5.3: Logarithmic histogram of the hit separation distribution along the z coordinate, $|z_{\text{layer } 1} - z_{\text{layer } 2}|$ and $|z_{\text{layer } 2} - z_{\text{layer } 3}|$. The unfiltered track candidates are shown along with tracks generated by the Monte Carlo simulation, as labelled in the figure. The range of accepted hit separations limits at the grey dashed line.

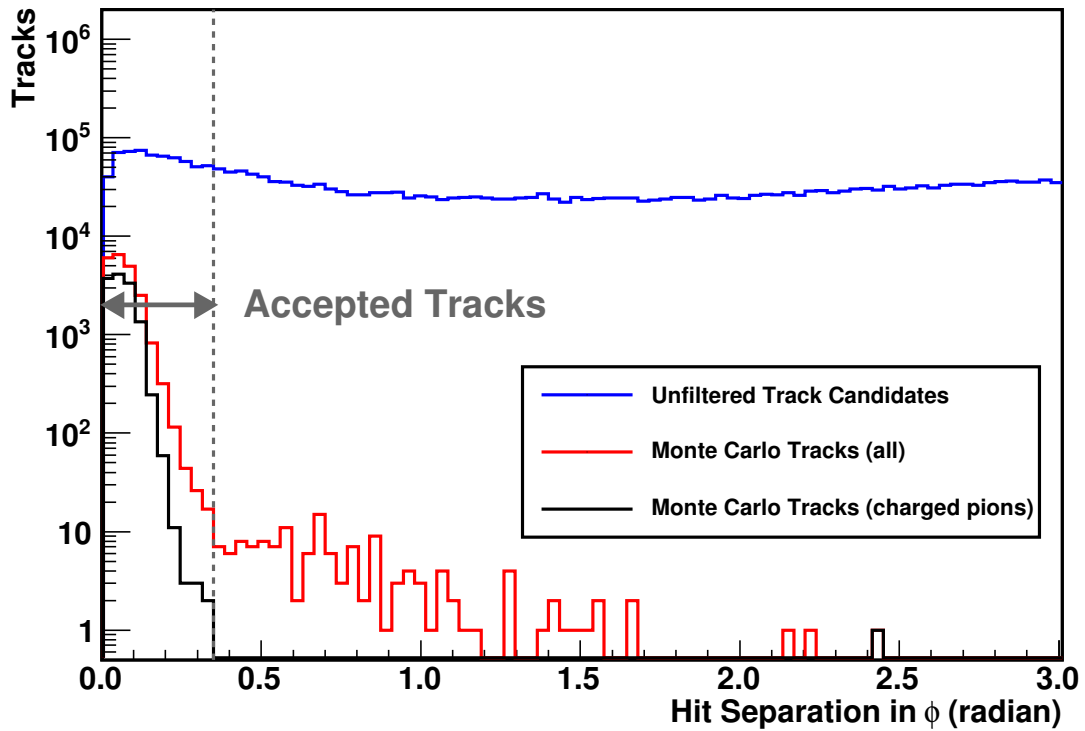


Figure 5.4: Logarithmic histogram of the total hit separation in the ϕ coordinate, $|\phi_{\text{layer 1}} - \phi_{\text{layer 2}}| + |\phi_{\text{layer 2}} - \phi_{\text{layer 3}}|$. The unfiltered track candidates are shown along with tracks generated by the Monte Carlo simulation, as labelled in the figure. The range of accepted hit separations limits at the grey dashed line.




Figure 5.5: Logarithmic histogram of the hit correlation coefficient distribution (as described in Sec. 5.3). The unfiltered track candidates are shown along with tracks generated by the Monte Carlo simulation, as labelled in the figure. The range of accepted correlation coefficients limits at the grey dashed line.

This allows for a strong filter on the ‘straightness’ of the track candidate. This is measured using a Principal Component Analysis (PCA) method, which minimizes the orthogonal distance between the hits and the regression line in all three dimensional components. If the three hits are highly correlated in position space, e. g. in a straight line, the PCA method returns a linear correlation coefficient close to one. Figure 5.5 shows the distribution of correlation coefficients for the unfiltered track candidates as well as for the true particle tracks, as recorded by the Monte Carlo simulation. This figure shows that the vast majority of real particle tracks return a correlation coefficient ≥ 0.9 , and as such the pattern recognition is set to reject hit combinations that return a correlation coefficient less than 0.9. This is a powerful filter, as it discriminates against spurious noise hits as well as ‘ghost hits’. Ghost hits are the result of a detector module having more than one strip cluster on either the p-side or the n-side. Figure 5.6 illustrates how ghost hits are present when there is an ambiguity caused by several intersecting strips. Likewise, from Figure 5.6 it can be seen how the PCA regression can help resolve the ghost hit ambiguity. For this example, there are four combinations that contain a single hit per layer and only two remain after the hit separation cuts. However, only the combination of the blue hits (labeled as Real Hits) satisfies the correlation cut, as every other combination has large perpendicular residuals to a fitted regression line. That is, only the blue hits generate a straight line to within our defined cut tolerance.

Using the Monte Carlo simulation to evaluate the overall efficiency of the track recognition showed that the algorithm found $(88 \pm 5)\%$ of the particle tracks that were reconstructable, that is, tracks where the particle passed through all three layers (regardless of the particle species). Tracks that were not found likely had a small radius of curvature (low momentum), or scattered in the middle layer, resulting in a crooked line and failing the correlation cut. Although not all particle tracks will be found by the pattern recognition algorithm, there are generally more track candidates for a given event than

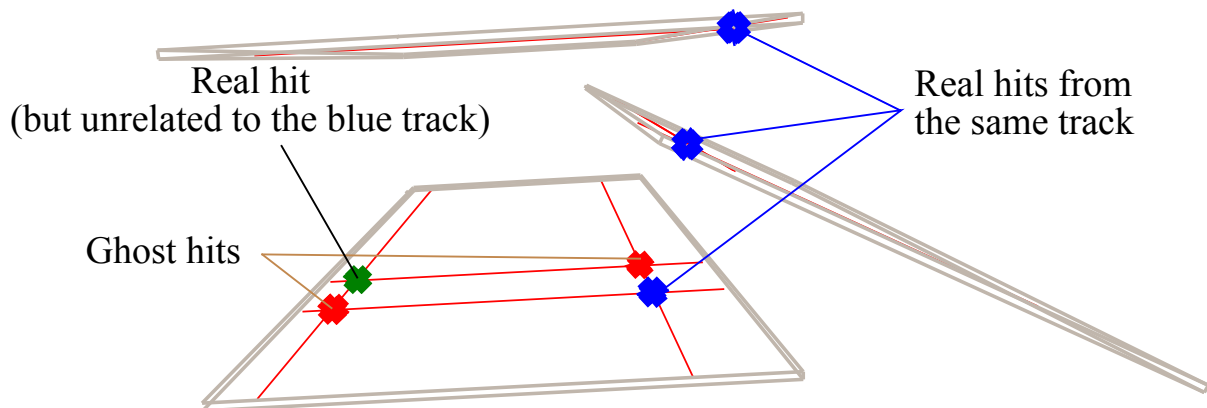


Figure 5.6: An example of the ‘ghost hit’ ambiguity. Two particles passing through the same module (represented by the grey rectangle) leave signal in two n-side and two p-side strips (represented by the red lines in the modules). The intersections of n- and p-side strips gives four hit possibilities, where the red markers show the two spurious hits.

real particle tracks. One reason for this is that the track candidates are not all forced to have unique hits. At the pattern recognition level, the track candidates can share hits, which artificially increases the number of candidates. However, at this stage it is sufficient to evaluate as many hit combinations as possible, and the method for differentiating between track candidates with shared hits will be described in Section 5.5.

5.4 Determining the helix parameters

In the absence of material interactions, i. e. multiple scattering and ionization/bremsstrahlung energy loss, charged particles will follow a helical path in a solenoidal magnetic field (the octupole and mirror fields are treated as perturbations). Thus, to first order, the particle tracks through the apparatus and detector can be approximated by the helix model. This allows for the extrapolation of particle trajectories, which is necessary to find the vertex position.

This section will describe the helix parametrization and how it can be used to supplement the pattern recognition in filtering track candidates. The final decision on which track candidates to consider in the vertex finding is also discussed.

5.4.1 Helix parametrization

The helical path followed by the charged particles (ignoring multiple scattering and energy loss through the material) can be described in terms of five parameters [155]: $(\rho_c, \phi_0, D, \lambda, z_0)$, where ρ_c is the radius of curvature, ϕ_0 is the azimuthal angle at the point of closest approach to the axis, D is radial distance of closest approach to the axis in the $x - y$ plane (D is a signed parameter, where the sign represents the particle charge), λ is related to the polar angle θ (where θ is measured from the positive z axis) by $\lambda = \cot \theta$, and z_0 is the z coordinate of the distance of closest approach of the track to the axis. With these parameters, the coupled equations for the helical path in the global reference frame are given as:

$$\begin{aligned} x(s) &= D \cos \phi_0 + \rho_c \left[\cos \left(\frac{1}{\rho_c} s + \phi_0 \right) - \cos \phi_0 \right], \\ y(s) &= D \sin \phi_0 + \rho_c \left[\sin \left(\frac{1}{\rho_c} s + \phi_0 \right) - \sin \phi_0 \right], \\ z(s) &= z_0 + \lambda s, \end{aligned} \tag{5.3}$$

where s parametrizes the system of equations by the arclength.

The helix equations (Eqns. 5.3) can be separated into two projections: radial and axial. The radial (or $x - y$ plane) projection contains the $\vec{v} \times \vec{B}$ motion of the particles in solenoidal magnetic field. The axial projection is along the z axis, parallel with, and unaffected by, the magnetic field.

5.4.2 Radial helix parameters

Figure 5.7 shows the radial helix parameters: ρ_c , ϕ_0 , and D . These parameters can be determined using only the p-side strip information, that is, the hit coordinates in the $x - y$ plane. In this projection, the charged particles travel in a circular trajectory, and as such, the determination of the radial parameters amounts to determining the circle




Figure 5.7: Illustration of the helix parameters in the radial projection. ρ_c gives the radius of curvature, D is the distance of closest approach to the origin, and ϕ_0 is the azimuthal angle at the point of closest approach.

which passes through the three hit coordinates. The equation for this system can be written as:

$$(x - a)^2 + (y - b)^2 = \rho_c^2, \quad (5.4)$$

where (a, b) is the center of circular motion in the $x - y$ plane. Eqn. 5.4 can always be reduced to the form:

$$x^2 + y^2 + Ax + By + C = 0, \quad (5.5)$$

where A , B , and C are real parameters. This equation can be further rearranged by completing the square such that,

$$\left(x + \frac{A}{2}\right)^2 + \left(y + \frac{B}{2}\right)^2 = \frac{A^2 + B^2 - 4C}{4}, \quad (5.6)$$

or, equating Eqns. 5.4 and 5.6, the circle parameters are given as $a = -A/2$, $b = -B/2$, and $\rho_c = 1/2\sqrt{A^2 + B^2 - 4C}$.

To determine A , B , and C , the three hit coordinates, (x_1, y_1) , (x_2, y_2) , (x_3, y_3) , can

be substituted into Eqn. 5.5 to form an exactly determined system of equations:

$$\begin{aligned}x_1^2 + y_1^2 + Ax_1 + By_1 + C &= 0, \\x_2^2 + y_2^2 + Ax_2 + By_2 + C &= 0, \\x_3^2 + y_3^2 + Ax_3 + By_3 + C &= 0.\end{aligned}\tag{5.7}$$

This system can then be rewritten in matrix notation:

$$\begin{pmatrix} x_1 & y_1 & 1 \\ x_2 & y_2 & 1 \\ x_3 & y_3 & 1 \end{pmatrix} \begin{pmatrix} A \\ B \\ C \end{pmatrix} = \begin{pmatrix} -(x_1^2 + y_1^2) \\ -(x_2^2 + y_2^2) \\ -(x_3^2 + y_3^2) \end{pmatrix}.$$

In this form, the parameters A , B and C can be solved for through matrix inversion.

With the circle parameters determined, the two remaining helix parameters can easily be determined. The distance of closest approach, D , is given as,

$$D = \sqrt{a^2 + b^2} - \rho_c.\tag{5.8}$$

Finally, the azimuthal angle for the point of closest approach, ϕ_0 , is given as

$$\phi_0 = \arctan(b/a),\tag{5.9}$$

and the reference point, (x_0, y_0) , that is, the point in the $x - y$ plane closest to the origin is then

$$x_0 = D \cos \phi_0,\tag{5.10}$$

$$y_0 = D \sin \phi_0.\tag{5.11}$$

These three parameters: ρ_c , D , and ϕ_0 , fully parametrize the helical motion in the radial projection.




Figure 5.8: Illustration of the helix parameters in the axial projection. z_0 is the z coordinate at the point of closest approach to the origin, and $\lambda = \cot \theta$ relates the azimuthal to the axial motion.

5.4.3 Axial helix parameters

Figure 5.8 shows the axial helix parameters. In this projection, the helical motion is parametrized by λ , which relates the azimuthal motion to the axial motion, and z_0 , which is the z coordinate of the point of closest approach of the helix to the origin.

With three hit coordinates (z_1, z_2, z_3) , the two axial helix parameters are over-constrained and the most appropriate solutions can be determined through a fitting procedure. This fitting is done by minimizing the χ^2 figure of merit,

$$\chi^2(\lambda, z_0) = \sum_{i=1}^3 \left(\frac{z(\lambda, z_0, s_i) - z_i}{\sigma_N} \right)^2, \quad (5.12)$$

where s_i are the helix arclength parameters found in Sec. 5.4.2. The χ^2 value is minimized using the ROOT *TMinuit* class, which implements a gradient minimization method [156]. The resulting χ^2 value provides a measure of the quality of fit, and likewise, a measure of how well the track candidate conforms to the helix model.

5.5 Track pruning

Many of the track candidates are not the result of a particle travelling through the detector. Similarly, even if a track candidate does match a particle track, it is not interesting in terms of the annihilation position reconstruction, unless the particle trajectory follows back to the vertex point as defined by other particle tracks generated by the same annihilation. For example, tracks due to e^-e^+ pairs are not desired, as they do not necessarily extrapolate back to the annihilation point. As such, it is advantageous to discard track candidates that do not fit well to the helix model, or candidates that are unlikely to make a good contribution to the vertex position determination.

The following three steps are repeated for each event to finalize which track candidates will be included in the vertex determination:




Figure 5.9: Logarithmic histogram of the $\tilde{\chi}^2$ distribution for: all unfiltered track candidates (blue), all simulated particle tracks (red), and only the simulated charged pion tracks (black).

1. *Reject track candidates that do not fit well to a helix.* Using Eqn. 5.12 as a measure of goodness-of-fit of the track candidate to a helix (strictly only for the axial helix parameters, as the radial parameters are exactly constrained), poor fitting candidates are rejected. Each track candidate fit has one degree of freedom, so the χ^2 measure is trivially normalized, $\tilde{\chi}^2 = \chi^2/(\text{degrees of freedom}) = \chi^2/1$. Figure 5.9 shows the $\tilde{\chi}^2$ distribution for unfiltered track candidates (blue), all simulated particle tracks (red), and the simulated charged pion tracks (black). Following from the charged pion distribution, a cut is set to reject any track candidates with $\tilde{\chi}^2 > 15.0$.
2. *Reject track candidates whose trajectories do not extrapolate back to the trap vac-*

uum volume or electrode walls. Only track candidates that are likely to have originated directly from the primary annihilation are of interest. Moreover, since the charged antiparticles and neutral antiatoms are captured and manipulated within the Penning-Malmberg trap volume, annihilation can only occur within the vacuum volume (on residual gases) or on the surface of the Penning-Malmberg electrodes. As such, the helix parameter of the distance of closest approach to the origin in the $x - y$ plane, D , can be used to reject tracks that do not approach the area of interest. Similarly, the radial distance of closest approach from the track to the electrode wall is given as $D_w = D - 2.2775$ (cm). Tracks with $D_w \leq 0$ survive this decision, as they pass through the trapping volume, and possibly follow back to the annihilation point. Figure 5.10 shows the distributions of $D_w > 0$, that is, for tracks candidates that do not pass directly through the trapping volume. Candidates where $D_w > 1.5$ cm are discarded as they are much less likely to lead back to the annihilation vertex. Since the vertex is physically constrained to be located within the trap vacuum volume or on the surface of the electrodes, track candidates that do not extrapolate to this region either: a) scattered in the apparatus material; b) correspond to a secondary particle that did not originate from the annihilation vertex; or c) do not correspond to a particle at all (mis-reconstruction).

3. *Resolve any conflicts between track candidates which share hits.* Because every combination of three hits is enumerated and included as a track candidate, occasionally track candidates will share hits. This is generally an artifact of the track candidate enumeration, as two (or more) particles crossing close enough together to register the same p- and n-side strips is highly unlikely (note that e^-e^+ pair production within a silicon wafer can create an undesirable situation where two tracks share a hit).

To resolve any conflicts, each track is compared with every other track, if there

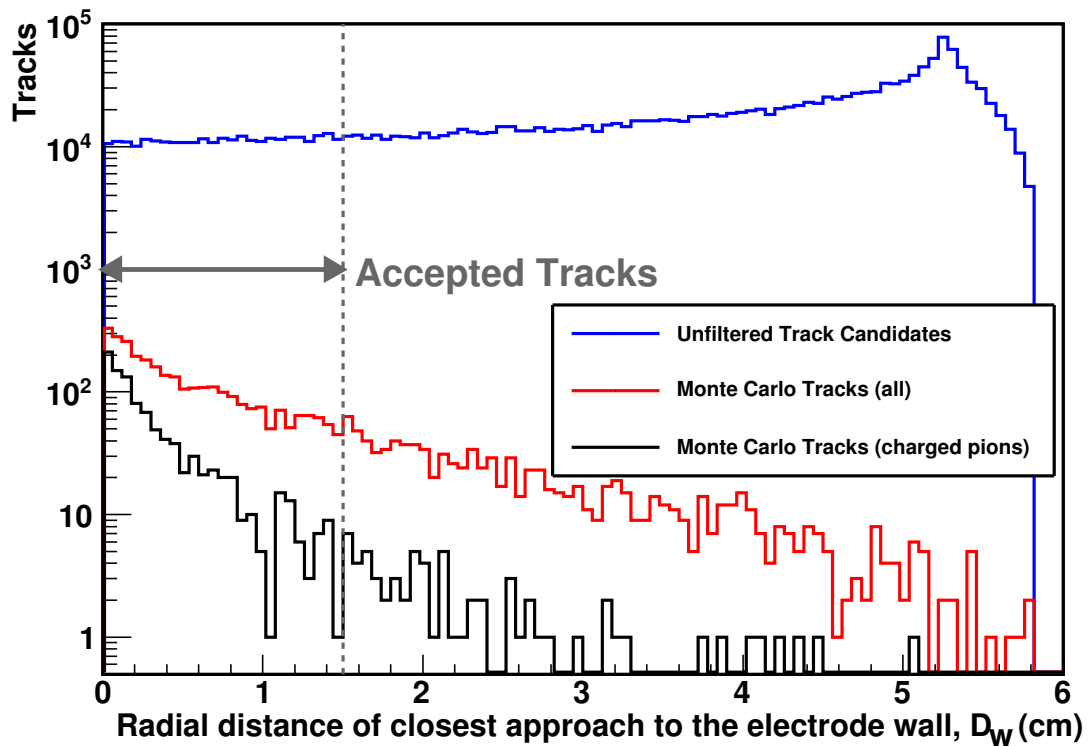


Figure 5.10: Logarithmic histogram of the distances of closest approaches to the electrode wall radius distribution for: all unfiltered track candidates (blue), all simulated particle tracks (red), and only the simulated charged pion tracks (black).




Figure 5.11: Position space representation of the result of the track decision algorithm on an event data set. Grey curves represent track candidates which failed the distance of closest approach cut, green curves (obscured behind the other curves) represent track candidates that failed the χ^2 cut, and red curves are tracks candidates that share a hit with another track (and are subsequently rejected in favour of the other track). Finally, the blue curves pass all of the track selection criteria, and are accepted as charged-particle tracks.

are any shared hits, the track with the smaller radius of curvature is discarded.

The choice to retain the track with the larger radius of curvature comes from the observation that, in the ALPHA detector, tracks with larger curvature are more likely to have resulted from charged pions (see Appendix D).

Figure 5.11 is an example of the full pruning decision. The grey track candidates shown have $D_w > 1.5$ cm, and clearly do not travel close to the trapping region. Track candidates shown in green (difficult to see behind the other candidates) failed the $\tilde{\chi}^2$ cut, while red track candidates shared hits with another track, but failed the radius of curvature comparison. Finally, the blue track candidates survived all the selection criteria, and are promoted to full track status. These promoted tracks will be passed to

the vertex determination algorithm.

It should also be noted that survival of all of these conditions does not guarantee that the track will ultimately be included in the final annihilation position determination. Many tracks from e^-e^+ production can survive all of these cuts, for example. It is left to the vertex determination algorithm to decide if any of these tracks should be excluded from the final vertex calculation. The goal of this section is to determine which tracks are appropriate to pass to the vertexing routines.

5.6 Vertexing

The foremost goal of the event reconstruction process is the determination of the position of the antiproton annihilation event. The method for determining the vertex position takes as input the tracks identified in Sections 5.3 - 5.5, and relies on the extrapolations of the track trajectories. The vertex position is determined by finding the point where the identified tracks pass closest to each other.

Because of the deviation of the track extrapolation from the true annihilation point due to multiple scattering (discussed in Sec 5.1), the reconstructed tracks will generally not have a single strict intersection point. Therefore, the common origin of two or more tracks will be localized to where these tracks pass closest to each other. As such, this method of vertex reconstruction attempts to determine the (x, y, z) coordinates of the annihilation position by finding the point to which all of the reconstructed particle tracks pass closest.

5.6.1 Closest approach between two helices

Figure 5.12 shows an example pair of helices, and the point of closest approach between them. The dark green line indicates the smallest distance between the two helices (within a gyro-orbit of the reference coordinates). Likewise, the green box gives the midpoint,




Figure 5.12: Spatial illustration of the point of closest approach between two helical tracks. The blue curves show the track extrapolation, while the green line shows where the helices pass closest to each other. The green box shows the midpoint along the green line, which for a two-track event is taken as the vertex position.

and represents the point of closest approach between the pair of helices. The point of closest approach is found by using the *TMinuit* local search routines to minimize the distance between the two helices in the 3-dimensional coordinate space of the detector.

In events with only two tracks, the point of closest approach is also interpreted as the annihilation vertex. That is, the point of closest approach gives an estimate of a common point of origin, and thus the 3-dimensional vertex of the two tracks.

5.6.2 Closest approach between $N_{\text{tracks}} > 2$ helices

For events with more than two tracks, the vertex can be estimated by combining the points of closest approaches for all pairs of helices. Figure 5.13 shows the an example reconstruction with three tracks and three green boxes representing the points of closest approach between the pairs of helices. The black box in Figure 5.13 shows the estimated vertex position, determined by taking the mean of the points of closest approach. This estimate is very sensitive to outliers due to spurious tracks, as every track is involved in $N_{\text{tracks}} - 1$ points of closest approach. As such, any poorly reconstructed track (or a




Figure 5.13: Spatial illustration of the determination of the vertex position for an event with more than two tracks. The blue curves show the track extrapolations, and the green lines show where pairs of helices pass closest to each other (following Fig. 5.12). The black box shows the mean of points of closest approach, which is used as a seed position for the vertex minimization.




Figure 5.14: Spatial illustration of the determination of the mean distance of closest approach. The blue curves show the track extrapolations, while the red lines show the closest approach of the track extrapolations to the vertex position (shown as the black box). The mean length of the red lines, d_{vertex} , is used as a figure of merit in the final vertex determination.

track that does not follow back to the annihilation point) can shift the vertex position significantly. It is therefore very important that all of the tracks included in the vertex determination be well chosen.

5.6.3 Mean distance of closest approach to the vertex

A useful vertex reconstruction quantity is the mean distance of closest approach from the tracks to the vertex. This measure is illustrated in Figure 5.14, where the dark red lines show the individual distances where the tracks pass closest to the vertex position.

The mean distance, d_{vertex} is then given as

$$d_{\text{vertex}} = \frac{1}{N_{\text{tracks}}} \sum_{i=1}^{N_{\text{tracks}}} d_i, \quad (5.13)$$

where N_{tracks} is the number of tracks included in the vertex determination, and d_i is the smallest distance between the i th track and the vertex. The d_{vertex} measure gives an estimate of the vertex quality, as tracks which do not extrapolate to a common position will return large d_{vertex} values, whereas d_{vertex} will be small if all the tracks converge well together.

The mean distance of closest approach measure can also be used to fine-tune the vertex position. By varying the vertex position so as to minimize d_{vertex} , the resulting position will represent the point of closest approach to all the tracks included in the vertex determination. Similarly, d_{vertex} can be used to identify and eliminate poorly reconstructed events. A large d_{vertex} value indicates that the reconstructed tracks do not extrapolate close to each other, and as such, there is no reason to believe that the reconstructed vertex position corresponds to the annihilation origin.

5.6.4 Track exclusion

Not all of the reconstructed tracks will extrapolate back to the primary annihilation vertex. As discussed above (Section 5.1), an example of this is when particles are scattered within the material of the apparatus (or detector). Because of this scattering, the reconstructed trajectory will often not pass close to the actual annihilation position. Another possibility is that the event contains secondary vertices, for example, when a γ -ray (resulting from the decay of a π^0 particle) passes through the apparatus (or detector) material, it can produce an electron-positron pair. The electron and positron can both leave tracks through the detector that point back to the secondary vertex, and can confuse the determination of the primary annihilation vertex.

Including spurious tracks in the vertex determination will bias the calculated position towards these tracks (as the vertex calculation is based on the finding the point of closest approach to all of the included tracks), and therefore, away from the proper tracks.

To this end, it is useful to identify and exclude the spurious tracks from the vertex determination. The following track exclusion algorithm, for an event with N tracks, attempts to minimize the mean distance of closest approach, d_{vertex} , while including as many tracks as possible:

1. Do not proceed if the event has $N \leq 2$ tracks, as at least two tracks are needed for form a vertex.
2. Reconstruct the vertex position using all the available tracks ($N > 2$) and calculate the mean distance of closest approach for this initial configuration, d_0 .
3. Reconstruct N new vertices, such that each new vertex excludes one of the original tracks.
4. Calculate the mean distance of closest approach for each new vertex configuration.
5. Determine which set of tracks has the smallest mean distance of closest approach, and call this value d_{min} .
6. Calculate $\delta = (d_0 - d_{\text{min}})/d_0$, which gives the fractional improvement in the mean distance of closest approach by excluding that specific track from the vertex.
7. If $\delta \leq 0.4$, exit the algorithm, keeping the configuration associated with d_0 as the final vertex determination.
8. If $\delta > 0.4$, the track configuration associated with d_{min} is promoted to the current accepted configuration, and relabeled as d_0 for the remainder of the algorithm.
9. If the track configuration now associated with d_0 has more than two tracks, return to Step 3 with this configuration.

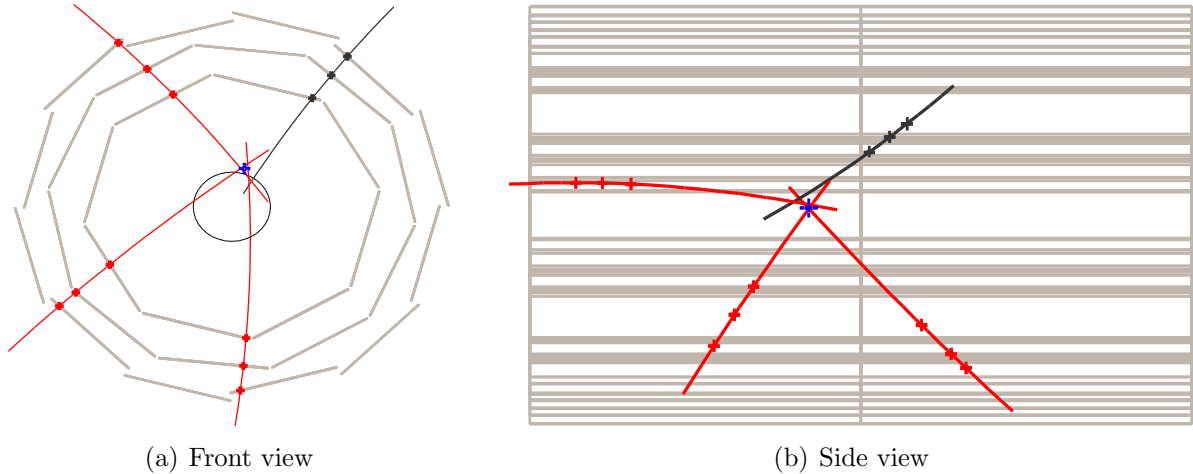


Figure 5.15: Example reconstruction of an annihilation event, where (a) shows the $x - y$ (front) view, and (b) shows the $x - z$ (side) view. Here, the grey track has been considered, but ultimately excluded according to the algorithm described in Sec. 5.6.4. The blue cross shows the reconstructed vertex position.

This algorithm will excluded tracks one at a time and look for an improvement in the mean distance of closest approach for each iteration. Figure 5.15 shows an example event containing four tracks, which combined return a mean distance of closest approach of $d_0 = 0.32$ cm. However, the configuration with the grey track excluded return $d_{\min} = 0.16$, with an associated $\delta = 0.5$. The algorithm then evaluates the two-track configurations, but finds the maximum improvement of $\delta = 0.32$, so the three-track configuration indicated by the red tracks in Figure 5.15 gives the final vertex location. Since this algorithm continues until there is no longer a large improvement through track exclusion, it can also handle the case where there are several spurious tracks.

5.7 Vertex distributions

5.7.1 Annihilation on background gas

There are several long pauses during the antihydrogen formation and trapping experiments (see Section 3.6). In particular, during the electron kick-out (Section 3.4.3), a

cloud of 3.8×10^4 antiprotons remain in a stationary potential well (the well does undergo fast potential manipulations, but these predominately affect the electron plasma, and the majority of antiprotons remain spatially confined) for close to 122 s. Combining a series of 335 experimental cycles gives a total dataset of 40780 s. During this time the antiprotons have the opportunity to interact with, and annihilate on, residual gas present in the trapping volume (vacuum conditions are estimated to be $\sim 10^{-14}$ mbar [157]). The annihilations between the antiprotons and the residual gas can then be reconstructed and the distribution of vertices is shown as Figure 5.16. In Figure 5.16a) the z -integrated $x - y$ cross-section shows how the antiproton cloud is located in the radial centre of the trapping volume. Similarly, Figure 5.16b) shows the axial distribution of antiproton annihilations, where the coordinate along z indicates the axial position relative to the center of the neutral trap. The unhatched region in 5.16b) represents the axial extent of the electrodes providing the confining potential. The axial position of the annihilation distribution corresponds to the location of the potential well during the e-kicking procedure.

5.7.2 Octupole-induced antiproton annihilation

The ALPHA apparatus includes an octupole magnet (Section 3.5) as part of the neutral atom trap. The radially-increasing magnetic field of the octupole magnet provides the confining potential for the neutral trap, in the radial direction. However, the addition of the octupole field breaks the azimuthal symmetry of the Penning trap field needed for the confinement of the charged particles. Charged antiparticles, outside a critical radius imposed by the octupole field strength [62, 158], will follow the octupole field lines into the electrode walls. In this case, the antiproton annihilations will form a distribution dictated by the octupole field line configuration. The annihilation distribution for antiprotons in ALPHA's octupolar field is dominated by eight spots where the field lines meet the

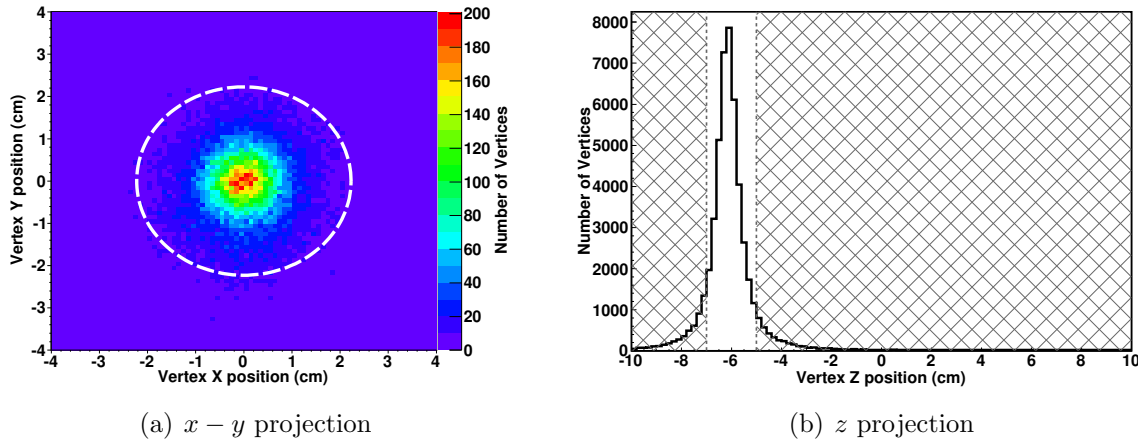


Figure 5.16: Vertex distribution for annihilations on background gas; (a) shows the $x - y$ projection of the vertex distribution, with the colour palette representing the number of vertices in each bin. The dashed white circle represents the position of the electrode walls. (b) shows the z distribution of vertices, where the unhatched region shows the extent of the electrodes providing the confining electric potential.

electrode wall (the number of spots corresponds to the number of turning points in the octupole winding). Four annihilations spots are located on both axial extents of the electrostatic confinement well, with the groups of four spots being shifted in azimuthal angle by 45 degrees from each other.

Figures 5.17 a) and b), along with Figure 5.18 show the antiproton annihilation distribution in the octupole magnetic field. These figures show the result of 6 cycles where 3.8×10^4 antiprotons (per cycle) were exposed to the octupole fields. The antiprotons were initially confined in a short electrostatic well (~ 4 cm in length). The well length is quickly (on the order of microseconds) expanded to about 10 cm, with a corresponding increase in the radius of the antiproton cloud. The octupole current is then ramped up, and the antiprotons beyond the critical loss radius [62] follow the octupole field lines into the electrode wall. Figure 5.17a) shows the z -integrated $x - y$ projection, with the eight octupole-induced annihilation points clearly visible. Likewise, Figure 5.17 shows the distribution of z vertex coordinates. The locations of the two peaks (at around ± 3




Figure 5.17: Vertex distribution for annihilations induced by the impinging octupole field lines on the confinement region; (a) shows the $x - y$ projection of the vertex distribution, with the colour palette representing the number of vertices in each bin. The dashed white circle represents the position of the electrode walls. (b) shows the z distribution of vertices, where the unhatched region shows the extent of the electrodes providing the confining electric potential.

cm) are set by the electrostatic confining potential, as it is at the axial extents of the antiproton orbits where the particles experience the greatest field line divergence.

Figure 5.18 shows a different vertex projection, with the z vertex coordinate along the horizontal axis, and the azimuthal angle of the vertex along the vertical axis. This projection is used to highlight the azimuthal features of the octupole-induced annihilation distribution. All eight annihilation points are present, and the 45 degree azimuthal phase separation between the the left and right bunches is clearly shown. This is a dramatic demonstration of the azimuthal symmetry breaking of the octupole field and a confirmation of the effect of the octupole field lines, predicted in [158]. Moreover, with the assumption that all of the antiprotons annihilate at the same point at each of the eight spots, experimental reconstruction resolutions of $\sigma_{z,\text{octupole}} \sim 0.7$ cm (which is in decent agreement with another estimate of the z vertex resolution in Sec. 5.7.4) and $\sigma_{\phi,\text{octupole}} \sim 12$ deg are estimated from the distribution. This example clearly shows that the reconstruction methods have the power to resolve interesting physical features in




Figure 5.18: $R - \phi$ projection of the octupole-induced vertex distribution (same distribution as shown in Fig. 5.17). The eight peaks correspond to the turning points in the octupole winding.

experimental annihilation distributions.

5.7.3 Antihydrogen formation in the neutral-atom trap field

The primary purpose of the ALPHA detector is the detection and study of antihydrogen. After formation, unconfined antihydrogen will quickly reach the electrode walls, where the bound antiproton produces the same annihilation signature as a bare antiproton. Although the same techniques can be used to reconstruct the events due to bare and bound antiprotons, the overall spatial vertex distributions formed are significantly different.

Figures 5.19 a) and b) show the z -integrated $x - y$ projection and z distribution, respectively, during the mixing of positrons and antiprotons in the magnetic field of the neutral atom trap. A total of 335 s of mixing is presented, resulting from 335 experimental cycles. Here, there are no electric or magnetic field manipulations as the particles interact, and the majority of vertices are the result of antihydrogen atoms




Figure 5.19: Vertex distribution during antihydrogen formation (with the neutral trap magnets engaged); (a) shows the $x - y$ projection of the vertex distribution, with the colour palette representing the number of vertices in each bin. The dashed white circle represents the position of the electrode walls. (b) shows the z distribution of vertices, where the unhatched region shows the extent of the electrodes providing the confining electric potential.

forming above the neutral trap depth, and as such being unconfined by either the electric or magnetic potentials. However, there is also a sizable component of vertices attributable to bare antiprotons following the magnetic field lines of the octupole. Specifically, the two peaks (at ± 2 cm) in Figure 5.19b) represent antiprotons that had formed weakly bound antihydrogen and travelled to large radii before being ionized by the trap fields (which are stronger at large radii) or the self-field of the positron plasma [50].

The spatial distribution of antihydrogen vertices is discussed at length in [157], which includes work that attempts to model the dynamics of produced antihydrogen in an effort to reproduce experimental distributions. Through careful analysis and modeling, distributions similar to Fig: 5.19b) are used to extract information about the initial antihydrogen conditions, specifically: positions, velocities, and binding energies.

5.7.4 Reconstructed vertex position resolution

The reconstructed vertex position resolution is an important quantity for evaluating the performance of the reconstruction routines. Due to finite hit resolution (Sec. 4.1.6.2), scattering in the apparatus material (Sec. 5.1), and energy loss of a charged particle as it passes through the apparatus material (Appx. A), the distribution of reconstructed vertices will be broadened. This spread reflects how well, statistically, the reconstructed vertex position is expected to determine the actual ('true') annihilation position.

One way to estimate the vertex position resolution is to generate a large number of annihilations using the Monte Carlo simulation. Since the generated annihilation positions are known (as they are given as input), the reconstructed vertex positions can be compared with the true positions. Similarly, an experimental estimate of the vertex position resolution can be made if the annihilations are known to all occur at a fixed location (and the spread in this location is much less than the spread in the reconstructed vertex positions). In this case, the distribution of reconstructed vertex positions directly reflects the reconstruction resolution, since broadening that is greater than the variance of the known annihilation position comes from reconstruction uncertainty.

For the ALPHA detector system, the distributions of reconstructed vertex positions can be fitted with a double-Gaussian resolution function (in this case, for the axial resolution),

$$f(z) = \frac{C_{z,1}}{\sqrt{2\pi\sigma_{z,1}^2}} \exp\left[-\frac{1}{2}\left(\frac{z - \mu_{z,1}}{\sigma_{z,1}}\right)^2\right] + \frac{C_{z,2}}{\sqrt{2\pi\sigma_{z,2}^2}} \exp\left[-\frac{1}{2}\left(\frac{z - \mu_{z,2}}{\sigma_{z,2}}\right)^2\right], \quad (5.14)$$

where $C_{z,1}$ and $C_{z,2}$ give the amplitudes, $\mu_{z,1}$ and $\mu_{z,2}$ the means, and $\sigma_{z,1}$ and $\sigma_{z,2}$ are the standard deviations (where the subscript 1 indicates the first term, and the subscript 2 indicates the second term). The first term of Eq. 5.14 represents the distribution of well-reconstructed vertices, which form a narrow Gaussian. Conversely, the second term represents poorly determined vertices (events with tracks that undergo large-angle

scattering, for example, will often return a poor determination of the vertex position) and contribute a broader Gaussian to the distribution.

The position resolutions are then taken as the weighted mean of the Gaussian widths. For example, the axial resolution is then:

$$\sigma_z = (C_{z,1}\sigma_{z,1} + C_{z,2}\sigma_{z,2}) / (C_{z,1} + C_{z,2}). \quad (5.15)$$

This averaging is done to account for the effect of the tails of the resolution function on the overall spread of the vertices. (The resolutions of the other coordinates are constructed analogously to Eqs. 5.14 and 5.15, with the coordinates interchanged.)


5.7.4.1 Monte Carlo resolution estimate

The Monte Carlo simulations result for the reconstructed vertex position resolution is estimated by simulating 10000 uniformly distributed annihilations on the electrode surface, and then comparing the reconstructed vertex positions with the known annihilation positions. Figure 5.20 shows the distributions of differences between the known annihilation location and reconstructed positions.


The z resolution (Fig. 5.20(a)) is found to be $\sigma_{z,MC} = (0.67 \pm 0.04)$ cm, while the resolution in the radial component (Fig. 5.20(b)) of the vertex position is found to be $\sigma_{R,MC} = (0.68 \pm 0.04)$ cm, and the azimuthal resolution (Fig. 5.20(c)) is $\sigma_{\phi,MC} = (21.0 \pm 0.9)$ deg (which, at the electrode radius of 2.2275 cm, corresponds to $\sigma_{\phi,MC} = (0.82 \pm 0.04)$ cm). The narrow and wide Gaussian widths are tabulated in Table 5.2.

5.7.4.2 Experimental resolution estimation


The background gas annihilation distribution for cold trapped antiprotons (Section 5.7.1) can be used to estimate the reconstruction resolution of the vertex position. This estimation assumes that the size of the antiproton cloud is much smaller than the reconstruction



(a) Simulated axial resolution



(b) Simulated radial resolution



(c) Simulated azimuthal resolution

Figure 5.20: Linear histograms illustrating the distributions of differences between simulated and reconstructed vertex distributions for the (a) axial, (b) radial, and (c) azimuthal coordinates. The dashed line shows a double-Gaussian fit to the distributions.


Type of estimate	Resolutions		
	Axial (cm)	Radial (cm)	Azimuthal (cm)
Monte Carlo simulation	0.67 ± 0.04	0.68 ± 0.04	0.82 ± 0.04
Monte Carlo simulation (narrow distribution)	0.27 ± 0.01	0.32 ± 0.02	0.38 ± 0.01
Monte Carlo simulation (wide distribution)	1.29 ± 0.05	0.99 ± 0.03	1.53 ± 0.06
Background gas estimate	0.75 ± 0.02	0.80 ± 0.02	-
Background gas estimate (narrow distribution)	0.427 ± 0.004	0.50 ± 0.05	-
Background gas estimate (wide distribution)	1.54 ± 0.02	1.38 ± 0.01	-
Octupole loss estimate	~ 0.7	-	~ 0.5

Table 5.2: Summary of the simulated and experimental reconstruction vertex position resolutions estimates.

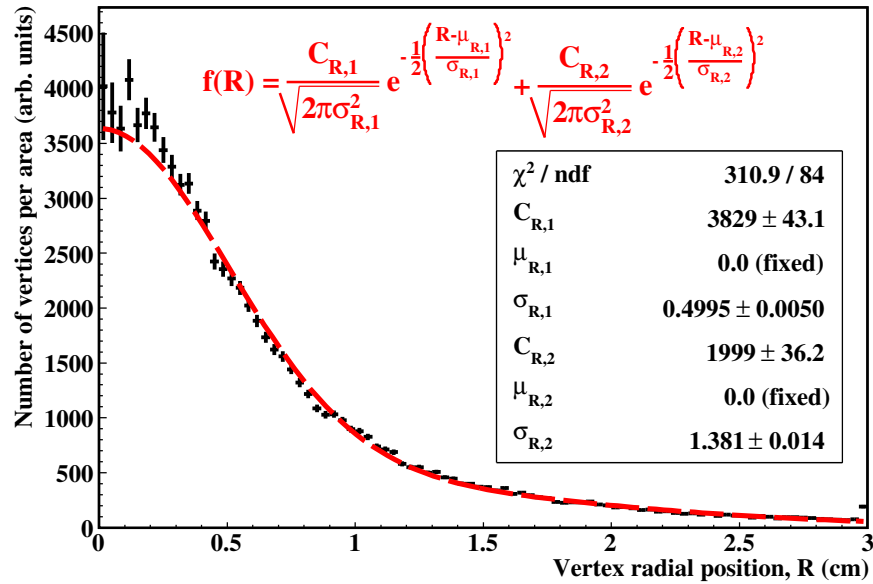
resolution, such that the antiproton cloud can be treated as a point source in the center of the confining potential. The extent of the annihilation distribution is then entirely due to the resolution of the event reconstruction. Figure 5.21(a)) shows the vertex z coordinates, which have been shifted such that the mean is centered at $z = 0$ cm. The radial reconstruction resolution can also be estimated by examining the antiproton annihilations on background gas. Here, because the radial distribution (Fig. 5.21(b)) is integrated over the azimuthal direction, each bin is divided by dR to properly scale for the increase in integration area. (Since the annihilations on background gas are symmetric about the origin in the azimuthal coordinate (Fig. 5.16(a), a distribution to apply the resolution function is unavailable.)

The experimental estimate of the axial resolution (Fig. 5.21(a)) is then $\sigma_{z,\text{Exp}} = (0.75 \pm 0.02)$ cm, and the radial resolution (Fig. 5.21(b)) estimate is $\sigma_{R,\text{Exp}} = (0.80 \pm 0.02)$ cm. The narrow and wide Gaussian widths are tabulated in Table 5.2. These estimates are consistent with the simulated resolutions to within better than 20%.

This determination of the reconstruction resolution is only valid if the originating



(a) Experimental axial resolution



(b) Experimental radial resolution

Figure 5.21: Plots of vertex frequency versus (a) vertex axial position and (b) vertex radial position. This data enables an experimental determination of the reconstruction vertex resolutions for the axial and radial coordinates. The distributions are taken from annihilations on background gas (Sec. 5.7.1), where the antiproton cloud is assumed to be well localized. The dashed red curves show the fitting of a double-Gaussian function to the distributions.

antiproton cloud is much smaller than the reconstruction resolution. The radius of the antiproton cloud is directly measured on the MCP (Section 3.2) to be $r = 0.08$ cm, which is much smaller than the 0.8 cm radial resolution found above. However, the axial dimension of the antiproton cloud is not directly measured. The background annihilations occur during the electron cooling and kick-out stage (Section 3.4.3), where the antiproton cloud cools, via collisions with the electrons, into the bottom of the potential well. It is difficult to estimate the axial length of the antiproton cloud, as the number of electrons varies over the course of the electron kick-out procedure and the kick-out manipulations will often heat the antiproton cloud. However, the simulated and measured distribution widths ($\sigma_{z,MC} = (0.67 \pm 0.04)$ cm vs. $\sigma_{z,Exp} = (0.75 \pm 0.02)$ cm) are roughly consistent, which suggests that using the distribution of annihilations on the background gas is perhaps an adequate proxy for the reconstruction resolution at the $\sim 10\%$ level.

Sec. 5.7.2 provides an additional estimation of the experimental z resolution from the octupole-induced annihilation distribution ($\sigma_{z,octupole} \sim 0.7$ cm), as well as a complementary estimate of the azimuthal resolution ($\sigma_{\phi,octupole} \sim 12$ deg). Like the other estimates in this section, the octupole measurement relies on the assumption that the antiproton annihilations all occur with a position spread much smaller than the reconstruction resolution. However, this assumption is more difficult to verify since the octupole field was changing over the course of the measurement (that is, the ramping of the octupole magnet to induce antiproton losses). A long antiproton cloud could account for poor goodness-of-fit seen in Fig. 5.21(a). At the very least, the estimates from the octupole loss experiment and the estimates based on the rest of the measurements and simulations are all roughly consistent with each other.

5.8 Summary

This chapter described the method used to reconstruct the vertex position for events in the ALPHA detector. First, a couple of challenges associated with the specifics of the ALPHA detector and apparatus were discussed. This was followed by a description of the algorithms for finding and reconstruction charged particle tracks, along with the criteria for rejecting poor track candidates. The optimization of these criteria using Monte Carlo simulation was also presented in detail. The vertex determination method, and the algorithm for excluding inappropriately chosen tracks, were then presented. Finally, a number of experimentally observed distributions were given as examples of the utility of the vertex reconstruction.

The vertex reconstruction routines presented above identify charged particle tracks through the silicon detector and interpolate their intersection as an estimation of the annihilation position. As shown, the distributions of reconstructed vertices can reveal the dynamics of the antiprotons and antihydrogen atoms under study. In the next chapter, the vertex distributions will be analyzed in order to reduce the number of background events accepted, while maximizing the antihydrogen annihilations events found.

Chapter 6

Rejection of background events¹

The ALPHA detector is sensitive to the passage of charged particles through its silicon volume. As demonstrated in Chapter 5, the annihilation position can be determined by reconstruction of the particle tracks through the three detector layers.

However, not all of the detector readout events are due to antihydrogen annihilation. That is, there are a number of background processes that result in undesired detector events that can obscure, or mimic, antihydrogen annihilation. This chapter will describe the various detector backgrounds, with a particular focus on background events due to cosmic ray particles. The method used to discriminate and suppress cosmic ray background events is described in detail. Moreover, the optimization of the background rejection criteria is also presented.

6.1 Backgrounds to the antihydrogen annihilation signal

Although this thesis is mainly concerned with the detection of trapped antihydrogen, it is important to consider alternate situations, other than antihydrogen annihilation, that result in the triggering and readout of the detector. From here on, ‘signal’ will refer exclusively to antihydrogen annihilation, and ‘background’ will refer to any other type of readout event. Several kinds of background events are described in this section: noise events, where no discernible particles are present; mirror-trapped antiproton events, where the adiabatic conservation of the magnetic moment of a gyrating antiproton can cause bare antiprotons to be confined in the non-homogeneous magnetic neutral trap

¹A section of this chapter (Sec. 6.3) has been submitted for publication: G. B. Andresen, et. al., *Antihydrogen Annihilation Reconstruction with the ALPHA Silicon Detector*, Submitted to Nucl. Instrum. Meth. A.

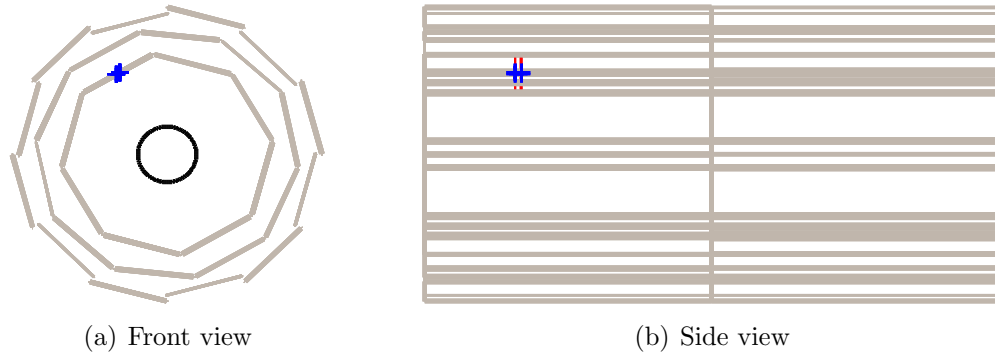


Figure 6.1: Detector read-out data representing a small amount of detector activity triggering the full readout. Recalling the description of the readout trigger in Sec. 4.1.5, two or more p-side TA signals are required to trigger the analog readout. The blue crosses indicate hit positions, while the red lines show the signal strips. The left view (a) shows the $x - y$ projection of the event, while the right view (b) shows the $z - y$ projection. The black ring shows the electrode radius.

field; and cosmic ray events, which are due to the constant flux of atmospheric muons.

6.1.1 Environment and hardware background

The detector readout trigger (described in Sec. 4.1.5) is intended as a very inclusive trigger. That is, since annihilation events during trapping experiments are rare, an overly selective trigger might fail to read out the detector for legitimate annihilation events. However, such an inclusive trigger will also accept unintended, non-annihilation events. For example, Fig. 6.1 shows an event which can trigger the detector readout, but does not contain much information. While this event contains the requisite two (or more) p-side TA signals, there is no indication of the presence of any charged particles. This event is likely the result of environmental or electronic hardware noise on the TA signal lines which induce a spurious detector readout. These events are trivial to recognize, given the small number of hits and absence of any hit patterns, and a simple solution is to reject any event with less than six hits (i. e. $N_{\text{hits}} < 6$, since two tracks are needed, with three hits each, to be able to form a vertex).

An example of another type of background is shown as Fig. 6.2. In this case, an




Figure 6.2: Detector read-out data representing a readout event with an overwhelming number of hits (possibly due to an electromagnetic particle shower within the apparatus or detector). The blue crosses indicate hit positions, while the red lines show the signal strips. The left view (a) shows the $x - y$ projection of the event, while the right view (b) shows the $z - y$ projection. Accurate reconstruction is unlikely, and events of this type are discarded. The black ring shows the electrode radius.

overwhelming number of hits makes proper event reconstruction difficult. Fig. 6.2 is likely the result of an electromagnetic particle shower, where a high-energy particle (usually a high-energy cosmic ray) produces high-energy photons via bremsstrahlung, which in-turn produce electron-positrons pairs. This cycle can continue until the particles no longer have enough energy to produce electron-positron pairs. These events are identifiable by a very large number of hits, and as such, events with greater than $N_{\text{hits}} > 100$ are discarded.

Additionally, there are several kinds of background events corresponding to operations in the AD. For example, to produce antiprotons, high-energy protons are collided with the AD target (see Sec. 3.3.1). While the produced antiprotons are magnetically selected and directed in the AD storage ring, a large number of other particles (mostly pions and muons) are also produced and can reach the detector. Similar to Fig. 6.2, these events are characterized by a large number of hits and their synchronicity with the AD production cycle. Similarly, a large particle flux occurs when the ALPHA experiment receives a pulse of antiprotons from the AD. As described in Sec. 3.4.2, most of the antiprotons are




Figure 6.3: Detector read-out data representing an example of a background event generated by electrical noise associated with the fast shut-down of the neutral atom trap magnets. The blue crosses indicate hit positions, while the red lines show the signal strips. The left view (a) shows the $x - y$ projection of the event, while the right view (b) shows the $z - y$ projection. The black ring shows the electrode radius.

not captured and a large number annihilate on the vacuum window between the ALPHA apparatus and the AD. Both of these types of events can be rejected both through the aforementioned $N_{\text{hits}} > 100$ filter, but also by discarding any events coincident with the AD production and extraction time-windows.

Figure 6.3 shows an example of background events related to the fast shutdown of the neutral-atom trap magnets (described in Sec. 3.5.3). During the fast shutdown the trap magnets are de-energized over tens of milliseconds. This operation creates a large change in the magnetic flux, which can interfere with nearby electronics, including the on-board readout circuits on the detector modules. Coincident within tens of microseconds with the fast magnet shutdown, both the digital and analog detector signal lines experience a large amount of induced noise, often enough to reach the trigger threshold and trigger the detector readout. These noise-induced events are also easily identifiable by the large number of hits. Also, as illustrated in Fig. 6.3, a distinct hit pattern can be seen where, for each ASIC, the first several strips all exceed the analog strip threshold (Sec. 4.1.6.1). This is caused by a large shift in the readout pedestals (the first several strips are readout immediately and experience the most analog noise), which the strip

thresholding algorithms are unable to correct. These events occur intermittently, but always within the first tens of microseconds after the magnet shutdown is initiated. As such, these can easily be rejected by discarding any event with a large number of hits within about 20 μs of the start of the magnet shutdown.

6.1.2 Mirror-trapped antiprotons

A very difficult background to separate from antihydrogen signal is due to bare antiprotons that have been magnetically confined in the inhomogeneous magnetic field of the neutral-atom trap. These are referred to as mirror-trapped or mirror-confined antiprotons. This background is difficult to address since the annihilation signature of bare antiprotons is identical to that of antihydrogen (since it is the antiproton annihilation that is reconstructed in either case, since the silicon detector is not capable of reconstructing positron annihilations).

Antiprotons (or other charged particles) can be ‘mirror confined’ as a consequence of the adiabatic conservation of their magnetic moment, μ [159]:

$$\mu = \frac{\frac{1}{2}mv_{\perp}^2}{B}, \quad (6.1)$$

where m is the mass of the particle, v_{\perp} is its speed in the plane perpendicular to the magnetic field (which has a magnitude B). An example of mirror confinement is an antiproton in the magnetic field configuration given in Sec. 3.5.1. Here, two ‘mirror coils’ provide regions of increasing magnetic field magnitude. For a particle to conserve its magnetic moment in the regions of increasing field, energy is transferred from the parallel degree of freedom to the perpendicular degree of freedom (and vice versa for particles traveling away from the region of increasing magnetic field magnitude). This conservation of energy can be expressed as the following equation for the parallel particle

speed, v_{\parallel} (in the absence of electric potentials),

$$v_{\parallel}^2 = v_0^2 \left(1 - \frac{v_{\perp,0}^2}{v_0^2} \frac{B}{B_0} \right), \quad (6.2)$$

where $v_0 = \sqrt{v_{\parallel,0}^2 + v_{\perp,0}^2}$, and $v_{\parallel,0}$ and $v_{\perp,0}$ are the parallel and perpendicular speeds of the particle at a point where the magnetic field magnitude is B_0 [53].

Turning points for the particle can be found from Eq. 6.2 when $v_{\parallel} = 0$ and the condition for charged particle confinement in the magnetic mirror field comes from $v_{\parallel} < 0$, or

$$\left(\frac{v_{\parallel}}{v_{\perp}} \right) < \sqrt{\frac{B_{\max}}{B_{\min}} - 1}, \quad (6.3)$$

where B_{\max} and B_{\min} are the maximum and minimum magnetic field magnitudes, respectively. This condition shows that antiprotons with large perpendicular energies can be confined in the magnetic field of the neutral-atom trap. This is an especially concerning background, as it would closely mimic the antihydrogen annihilation signal. Experimental identification and exclusion of mirror-trapped antiproton events will be discussed in detail in Chapter 7.

6.1.3 Cosmic Rays

Cosmic rays are high energy particles (predominately protons) of extraterrestrial origin, which, through collisions in the upper atmosphere, can produce secondary particles that may reach the Earth's surface. Specifically, the majority of cosmic ray particles that reach the surface are muons (μ^+/μ^-), with mean energy of about 4 GeV and mean rate of $\sim 1 \text{ cm}^{-2} \text{ min}^{-1}$ for horizontal detectors [9]. For the ALPHA detector, with a horizontal cross-section of $15 \text{ cm} \times 46 \text{ cm} = 690 \text{ cm}^2$, the expected rate of $Si > 1$ triggers (Sec. 4.1.5) due to cosmic muons is $\sim 11 \text{ s}^{-1}$. The background $Si > 1$ trigger rate is measured as $9.65 \pm 0.02 \text{ s}^{-1}$, which is consistent to about 12%.




Figure 6.4: Detector read-out data representing an example of a cosmic event, shown in the a) $x - y$ projection, and b) $y - z$ projection. Two tracks are found (which are segments of a single cosmic track), and shown as the blue curves. The black diamond gives the reconstructed vertex position, following the method described in Chapter 5. The blue crosses indicate hit positions, while the red lines show the signal strips. The black ring shows the electrode radius.

In the current experimental setup, there is no straightforward way of absolutely differentiating between annihilation and cosmic background events². As such, the passage of cosmic ray muons through the detector can result in a background event that, when the reconstruction algorithms are applied, erroneously returns a vertex position. Figure 6.4 illustrates such a background event, with the reconstructed ‘vertex’ position shown as the black cross. For this event, the contiguous cosmic trajectory resulted in two co-linear tracks on opposite sides of the detector. Since, to the reconstruction algorithms, these tracks are entirely valid, a vertex position can be calculated using the exact same method as with the annihilation events.

The identification and rejection of cosmic background events is the topic of Sec. 6.3 and the subject of the majority of the remainder of this chapter.

²Two external cosmic identification systems are currently being considered: the first would rely on vertical Time-Of-Flight differentiation between annihilation and cosmic events, while the other system would use direction-sensitive Cherenkov detectors to identify cosmic events.

6.2 Importance of background rejection

Section 6.1 gave brief descriptions of a number of backgrounds to the antihydrogen annihilation signal. Currently, the expected antihydrogen trapping rate in ALPHA is low, especially when compared to the cosmic background (the experimental trapping rate will be discussed further in Chapters 7 and 8). As such, it is imperative to maximize the ratio of signal to background in order to have any chance of observing rare trapped antihydrogen events. The trapping experiments outlined in Sec. 3.6, for example, rely heavily on the clear identification of annihilation events. To this end, a large effort is made to identify and suppress signals from the background sources described in Sec. 6.1.

6.2.1 Blind analysis

As described in Sec. 3.5.3 and 3.6, the characteristic signature of trapped antihydrogen occurs immediately after the fast shutdown of the neutral-trap magnets is engaged. It is tempting to optimize the background rejection routines to accept as many events in this time window as possible. After all, since this is the region where the trapped antihydrogen signal is expected, the more signal events found would mean more trapped antihydrogen atoms. However, this kind of approach is prone to introducing experimenter bias, where the resulting signal is overstated, since the optimization is done on the signal itself [160].

Such unintentional bias can be minimized through the use of a blind analysis approach. In a blind analysis, the data of interest is set aside (placed in a ‘box’), and the decisive analysis is performed on a simulated or auxiliary (‘training’) dataset. Usually, a training dataset contains many more events than are found in the box. This is to ensure that the surrogate signal and background samples are well characterized and representative. Once the full analysis has been optimized and finalized on the training set, the same analysis is applied the box dataset and the final result revealed (‘opening of the box’). A blind analysis was performed for the ALPHA trapping experiments, where the

background rejection criteria were optimized on auxiliary annihilation and background datasets (described in Sec. 6.3.2). The final cuts were only applied to the trapping signal region when the background analysis was completed.

6.3 Cosmic background rejection

It is important to reduce the cosmic ray background in order to assist with the identification of antihydrogen annihilations, especially for the observation of magnetically trapped antihydrogen. Fortunately, the track geometries associated with the large majority of cosmic background events are identifiably different from those of annihilation events and can be rejected on an event-by-event basis.

6.3.1 Discriminating variables

Cosmic rays that graze the detector do not typically produce a vertex and are automatically rejected. A vertex is often reconstructed, however, when the cosmic ray particle passes through the center of the detector. Fortunately, the distinct topologies of annihilation and cosmic ray events can be used to classify events as signal (annihilation, Fig. 6.5(a)) or background (cosmics, Fig. 6.5(b)). There are several variables that can be used to quantify the different signal and background topology: the number of charged tracks, N_{tracks} ; the combined linear fit residual, δ ; and the vertex radial position, R .

6.3.1.1 Number of charged tracks, N_{tracks}

The majority of cosmic background events which successfully return a vertex position contain two charged tracks ($N_{\text{tracks}} = 2$), according to our definition that tracks must contain exactly three hits, with one hit in each detector layer. This follows from the fact that these events are generally produced by the passage of a single charged particle, such that the two tracks found in the event are just segments of a single charged track.




Figure 6.5: Example detector events for an (a) annihilation (accepted signal), and a (b) cosmic ray (rejected background). The blue crosses show the hit positions, the blue curves show the reconstructed tracks, and the red diamonds show the reconstructed vertex position.

Conversely, the average charged multiplicity from antiproton annihilation results in roughly three charged tracks per annihilation. This means that a large number of annihilation events will contain more than two charged tracks ($N_{\text{tracks}} > 2$). However, there is still substantial overlap between the signal and background N_{tracks} distributions, wherein many annihilation events contain only two charged tracks, while some cosmic background events are accompanied by particle showers or scattering resulting in larger numbers of tracks in those events.

6.3.1.2 Combined linear fit residual, δ

Cosmic ray particles passing through the ALPHA detector are expected to follow, to first order, straight-line trajectories. Thus, events consistent with a single, linear, particle track are likely to be the result of the passage of a cosmic particle. To test for this case, the hit positions in an event can be fitted to a line. The combined linear residual, δ , can be used to evaluate how closely an event conforms to a single straight line track. This

estimator is written as,

$$\delta = \min_{i \in \mathcal{F}} \{ \Delta_i \}, \quad (6.4)$$

where \mathcal{F} is the set of all combinations of two tracks in the event, and the term in the brackets is given by

$$\Delta_i = \sum_{j \in G} d_{\perp,j}^2 \quad (6.5)$$

where $d_{\perp,j}$ is the perpendicular distance, or residual, between the fitted line and the j^{th} hit in G , which represents all the hits in the two tracks under consideration. The minimization over all available combinations of tracks is done to ensure that, even if the event contains several tracks, the combination providing the best fit is chosen.

If the hits in an event fit a perfect line, δ will evaluate to zero. However, due to the magnetic field and multiple scattering, cosmic trajectories often deviate from the ideal, resulting in a broadening of the δ distribution. Annihilation events, on the other hand, are not expected to produce many co-linear tracks, and should return values of δ well removed from the cosmic distribution.

6.3.1.3 Vertex radius, R

Annihilations must originate from within the trapping region of the apparatus, either on the surface of the electrodes, or on background gas. This physical constraint restricts the possible locations of the reconstructed vertex. In particular, the radial coordinate of the annihilation vertex is expected to be, within the radial reconstruction resolution, at or inside the electrode radius.

A cosmic event with two co-linear tracks, however, will return a vertex that is unconstrained in the radial coordinate and often well outside the trapping volume (an example of such a vertex is Fig. 6.4). Thus, events where the vertex radius is much greater than the electrode radius are attributed to cosmic rays, and categorized as background.

6.3.2 ‘Training’ datasets

Representative sample sets are required for both the annihilation signal and cosmic background in order to place the cuts on (i. e., to ‘train’) the discriminating variables. In the case of the ALPHA experiment, the signal and background can be measured separately, and as such, dedicated data samples can be collected for each.

The annihilation signal sample was constructed from 335 cycles where positrons and antiprotons were mixed together to form antihydrogen in the magnetic field of the neutral atom trap [53]. Each mixing cycle lasted for 1 s, and a total of 165520 readout events, largely due to unconfined antihydrogen (discussed in Sec. 5.7.3), were collected. Over the combined 335 seconds of signal collection, about 3350 events are expected from cosmic background, which constitutes a contamination of $\sim 2\%$.

Conversely, the background sample set was collected by operating the detector with no antiparticles present within the apparatus. So as to best emulate the situation of interest, the neutral trap magnets were kept engaged throughout the background collection. Data were recorded over several periods, which totaled almost 3 hours, with 109824 readout events. With no antiparticles present, these events should all be background signals, from such sources as cosmics or electronic noise-induced detections.

6.3.3 Cut placement and optimization

The separation of signal from background is only effective with well-placed cuts on the discriminating variables. They are placed to maximize the number of cosmic events rejected, while retaining as many annihilation events as possible. This analysis is focused on optimizing the expected signal significance during attempts to trap antihydrogen.

An appropriate figure of merit for this optimization is the p-value for a Poisson distribution, α . The p-value represents the probability of observing n_0 events (or more), due to a Poisson background rate b . Here, the expected signal and expected background

can be parametrized as functions of cuts on the combined linear residual, δ_{cut} , and vertex radius, R_{cut} , as well as categorized according to the number of charged tracks, N_{tracks} , found in the event. The p-value figure of merit can then be written as

$$\alpha(n_0, b) = \sum_{n=n_0}^{\infty} \frac{b^n e^{-b}}{n!}, \quad (6.6)$$

where the number of observed events and background rate depend on the cuts used ($n_0 = n_0(R_{\text{cut}}, \delta_{\text{cut}}, N_{\text{tracks}})$ and $b = b(R_{\text{cut}}, \delta_{\text{cut}}, N_{\text{tracks}})$). The signal optimization then proceeds by finding the set of cuts which minimize Eq. 6.6.

In order to reduce the dimensionality of the cut space, events are separated into two categories: $N_{\text{tracks}} = 2$, and $N_{\text{tracks}} > 2$. To first order, as the vast majority of background events fall into the $N_{\text{track}} = 2$ category, and thus this division separates events into background-dominated and signal-dominated sets. Therefore, the $N_{\text{tracks}} = 2$ set contains primarily background events, with some signal events to be extracted, while the $N_{\text{tracks}} > 2$ category contains mostly signal events, with some background events to be rejected. Thus, cuts on the $N_{\text{tracks}} = 2$ events should be relatively stringent to reject as many background events as possible, while the $N_{\text{tracks}} > 2$ cuts should be set to much more inclusive levels.

Fig. 6.6 shows the parameter space of the figure of merit as a function of the discriminating variables. The p-value is expressed in terms of standard deviations for a one-sided normal distribution, such that maximizing the expected significance corresponds to minimizing the expected p-value. These distributions are constructed by considering an array of $(R_{\text{cut}}, \delta_{\text{cut}})$ cut value pairs. For each pair, the expected background rate, b , is extracted from the application of these cuts to the background sample. Similarly, the expected number of observed events, n_0 , is scaled according to the signal sample, with the overall normalization reflecting the number of observed events in previous measurements and expectations about the number of trapped atoms.




Figure 6.6: Contour plots for the expected signal significance as a function of the cuts on the vertex radius, R_{cut} , and combined linear residual, δ_{cut} , for a) events with $N_{tracks} = 2$ ($N_{tracks} > 2$ events held constant), and b) events with $N_{tracks} > 2$ ($N_{tracks} = 2$ events held constant). The final cut decisions are shown as the white crosses.

N_{tracks}	Vertex radius, R_{cut} (cm)	Combined linear residual, δ_{cut} (cm ²)
= 2	< 4	> 2
> 2	< 4	> 0.05

Table 6.1: Final parameter cut conditions. Events satisfying these conditions are classified as annihilation signal.

Each $(R_{\text{cut}}, \delta_{\text{cut}})$ pair in Fig. 6.6 is the result of 5000 pseudo-experiments. Each pseudo-experiment is performed as follows:

1. A Poisson distribution with mean λ , $\text{Pois}(\lambda)$, is sampled so as to obtain an pseudo-experimental number of observed events, $n_s = \text{Pois}(n_0)$.
2. The p-value for the pseudo-experiment is calculated $\alpha_i = \sum_{n=n_s}^{\infty} b^n \exp(-b)/n$.
3. The representative value for the ensemble of p-values is taken as the log-average $\alpha = \exp(1/n \sum_{i=1}^n \ln \alpha_i)$. This measure of central tendency takes into account the logarithmic nature of the distribution of pseudo-experimental p-values.

The set of cuts were then chosen so as to take into account the expected significance, as well as attempting to maximize the acceptance of signal events. The final choices of cuts are shown as the white crosses in Fig. 6.6, and enumerated in Table 6.1.

6.3.4 Results of the background rejection

The effectiveness of the cut placement optimization can be evaluated by applying these cuts to the datasets described in Sec. 6.3.2. Figure 6.7 shows the distributions of the discriminating variables for both signal and background datasets, along with the distributions after the application of the cuts. By applying the cuts to the background sample, $(99.54 \pm 0.02)\%$ of the events are rejected, corresponding to a background acceptance rate of $(47 \pm 2) \times 10^{-3}$ events/s. Likewise, $(64.4 \pm 0.1)\%$ of the events in the signal sample pass the cuts. Since the events included in these datasets were collected

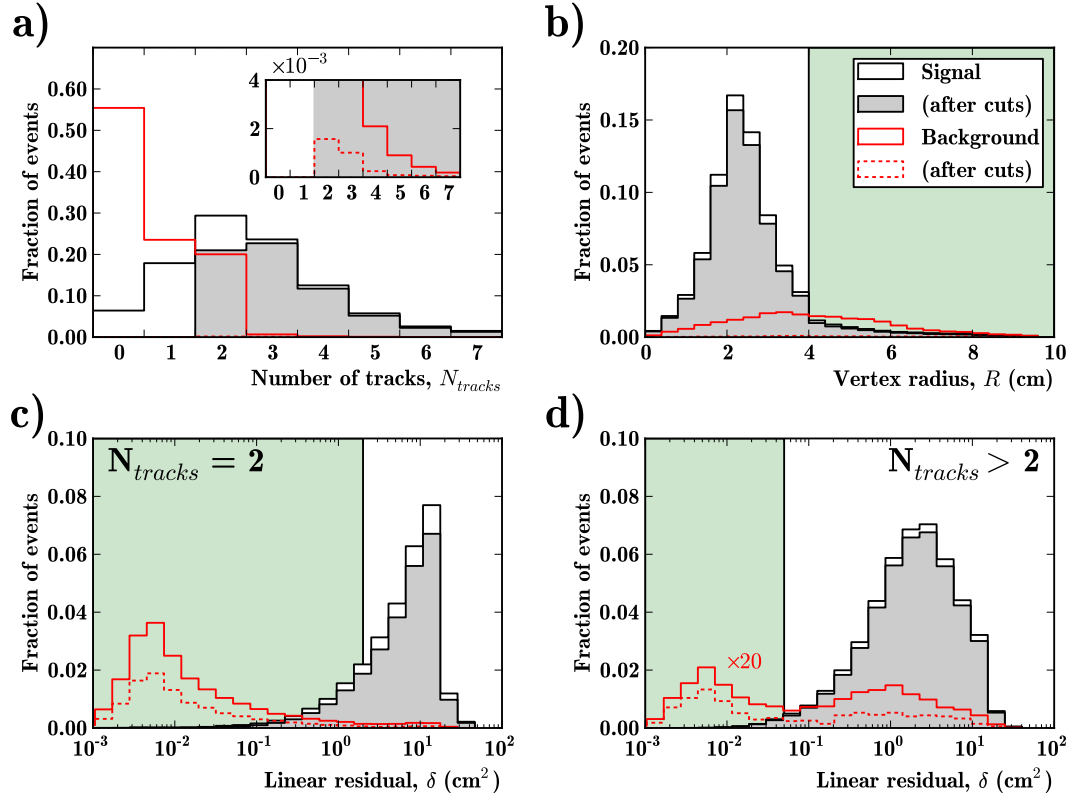


Figure 6.7: Linear histograms of the measured signal and background distributions for the discriminating variables. Shown are (a) the distribution of the number of charged tracks, N_{tracks} per event, (b) the radial component of the reconstructed vertex, R , and the combined linear residuals for the cases where (c) $N_{tracks} = 2$, and (d) $N_{tracks} > 2$. The annihilation signal sample is shown as the solid black trace, while the signal sample after the application of the cuts is shown as the solid grey filled trace. The cosmic background sample is shown as the solid red trace, and the red dashed trace shows the background sample after the application of the cuts. All distributions are normalized to the number of events in each sample set. The inset in (a) shows a magnification of the y-axis, highlighting the effect of the cuts on the background sample. Likewise, both background traces have been multiplied by a factor of 20 in (d) to appear for comparison. The green shadings represent the regions rejected by the cuts.

from *in situ* measurements (Sec. 6.3.2), the resulting background rejection and signal acceptance fractions correspond to total efficiencies, which include both the silicon and reconstruction efficiencies.

6.4 Summary

This chapter described the various detector backgrounds and methods for identifying and suppressing the majority of background events. First, the backgrounds affecting the silicon detector, especially those mimicking antihydrogen annihilation, were presented. The case was then made for the application of a blind analysis to minimize unintentional experimenter bias. The remainder of the chapter deals with the discrimination and rejection of cosmic ray background events. Discriminating variables related to the topology of the annihilation signal and cosmic background events were then introduced and the signal and background datasets described. Finally, the method for optimizing the background rejection criteria was presented and the resulting criteria evaluated.

The analysis presented in this chapter resulted in an efficient discrimination between annihilation and cosmic background events. By rejecting this class of background, the signal-to-noise is drastically increased for the identification of trapped antihydrogen annihilations. The next chapter, describing the experimental observation of trapped antihydrogen, will rely extensively on this analysis for the suppression of the cosmic ray background and determination of the annihilation vertex position.

Chapter 7

The trapping of antihydrogen

The overarching goal of the ALPHA collaboration is to magnetically confine antihydrogen atoms for spectroscopic studies. An outline of the trapping experiment procedure was already presented in Sec. 3.6. Moreover, Chapter 6 demonstrated that the cosmic background is under control, and the vast majority of cosmic events can be identified and rejected. However, cosmic rays are not the only background competing with the antihydrogen signal. As described in Sec. 6.1.2, mirror-trapped antiprotons can also be confined in the magnetic field of the neutral-atom trap. This chapter will start by describing the mirror-confined antiproton background, along with the methods used to attempt to remove the bare antiprotons and/or to discriminate out the signals associated with them. To assist the background identification and neutral-atom trapping effort, careful Monte Carlo simulations were performed, and will be described below. Several control and cross-check measurements, which complement the trapping experiments, are also detailed.

Finally, the results of two sets of trapping experiments are presented: first, the initial series of experiments performed in 2009, then the more recent experiments conducted in 2010. Backgrounds, such as cosmic rays and mirror-trapped antiprotons, are ruled out and it is concluded that the observed annihilation events are the result of the release of trapped antihydrogen. As will be seen, the observation of trapped antihydrogen relies critically on the silicon detector and the reconstruction routines presented earlier.

7.1 Mirror-trapped antiprotons

As discussed in Sec. 6.1, there are a number of background processes that can trigger the silicon detector readout. While Sec. 6.3 presented a method for efficiently identifying and suppressing cosmic ray background events through software analysis, discussion of the direct annihilation backgrounds (that is, annihilation signal that is not due to antihydrogen) was deferred. An important annihilation background comes from bare antiprotons (that is, antiprotons not bound with a positron) that are mirror-confined in the inhomogeneous magnetic field of the neutral-atom trap (Sec. 6.1.2).

Recalling Sec. 3.6, the antiproton cloud is already within the minimum-B region of the neutral-trap when the superconducting magnets are energized. This means that the antiprotons are inside the magnetic potential and, after antihydrogen formation is completed, it is necessary to remove any remaining antiprotons in order to eliminate competing annihilation events. Bare antiproton annihilations events are difficult (if not impossible) to separate from antihydrogen annihilation events using the software methods described in Chapter 6. This is because the detected antihydrogen annihilation signature comes from its antiproton, and is indistinguishable from the signature of a bare antiproton annihilation. As will be discussed in Sec. 7.3, the resolution of this dilemma is to modify the trapping experiment to spatially separate annihilations resulting from charged and neutral systems.

The vast majority of positrons and antiprotons are removed from the neutral-atom trap after mixing is complete through straight-forward potential manipulations, where the particles are dumped down the axis and their numbers recorded as they annihilate in distant regions of the apparatus. Additionally, four clearing pulses are applied across the trap in an attempt to remove any mirror-trapped particles. Each pulse consists of an axial electric field across the trap of about 2.5 V/cm for 10 ms (the on-axis electric


potential is shown in Fig. 7.1(a)). This field is intended to push any charged particles down the axis and away from the neutral-atom trap. The pulses alternate in polarity, with two pulses pushing antiprotons to the left (and positrons to the right) and two pulses pushing antiprotons to the right (and positrons to the left).

However, it is difficult to ensure that all of the antiprotons are cleared out of the trap, as particles with large transverse kinetic energies are harder to clear through the mirror-coils. This can be seen by examining the pseudopotential, U ,


$$U = E_{\perp,0} \left(\frac{B - B_0}{B_0} \right) + (-e)\Phi, \quad (7.1)$$

where $E_{\perp,0}$ and B_0 are the minimum transverse kinetic energy and minimum magnetic field magnitude, respectively, e is the fundamental charge, and Φ is the electric potential. Equation 7.1 combines the magnetic potential energy (given by $U_{\text{mag}} = \mu\Delta B$, where $\mu = E_{\perp}/B_0$) and electrostatic potential energy ($U_{\text{electric}} = (-e)\Phi$ for antiprotons) to give a single expression for the antiproton potential energy. Figure 7.1(b) shows the on-axis pseudopotential during the clearing pulses for several values of the transverse kinetic energy, $E_{\perp,0}$. As can be seen from the figure, a potential well forms for large transverse antiproton energies. This analysis leads to the conclusion that antiprotons need at least 20 eV of transverse kinetic energy to remain mirror trapped in these fields [53].

The pseudopotentials for off-axis antiprotons become complicated because of the octupolar magnetic field. Monte Carlo simulations (discussed further in Sec. 7.2) were performed to numerically examine off-axis trajectories and no simulated particles with less than 20 eV of transverse kinetic energy were found [53]. Another possibility is that a fraction of antiprotons is left with more than 20 eV of perpendicular energy after thermalizing with the positron plasma during mixing. However, the thermal spread of the antiprotons is limited to about 0.6 eV (the depth of the well that the antiprotons must overcome to enter the positron plasma). A 0.6 eV thermal distribution of antiprotons



(a) On-axis electric potential (solid) and magnetic field magnitude (dashed).



(b) Pseudopotential combining the electric and magnetic potential energies.

Figure 7.1: Electromagnetic potentials and fields versus position during the antiproton clearing pulses (reproduced with permission from [53]). Fig. (a) gives the on-axis electric potential (solid black line) and magnetic field magnitude (dashed red line) across the trap. The electric potential and magnetic field magnitude are combined via Eq. 7.1 to give a pseudopotential (b), for various transverse kinetic energies, E_{\perp} . Charged particles with large transverse kinetic energies can be confined by the magnetic field of the neutral trap.

will have a very small fraction ($\lesssim 10^{-10}$) of particles above 20 eV, so any mirror-confined antiprotons are likely the result of a non-thermal source. Some possible non-thermal mechanisms include: ionization of weakly bound antihydrogen which can result in the antiproton gaining parallel kinetic energy (which might then be transferred to the perpendicular degree of freedom through a hard collision with a residual gas atom or another antiproton), or a direct collision with a residual gas atom or another antiproton could impart a large amount of transverse kinetic energy. Rates for these processes are estimated to be very low [53, 157], meaning that it is unlikely that many mirror-trapped antiprotons survive the clearing pulses. However, the simulations and estimates alone are not enough to definitively rule out the possibility of mirror-confined antiprotons. Experimental verification is needed to validate the simulation assumptions and accuracy.

7.2 Simulation of particle trajectories during magnetic trap shutdown

The trapped antihydrogen signal window occurs during the fast shutdown of the neutral-atom trap magnets (Sec. 3.5.3). The changing magnetic fields during this time make it difficult to analyze the expected particle trajectories. To this end, careful Monte Carlo simulations of the bare antiproton and antihydrogen trajectories were performed, taking into account the effect of the clearing pulses (described in Sec. 7.1) and time-dependent magnetic fields [53, 54].

Particle trajectories are simulated immediately after most of the positrons and antiprotons have been removed from the neutral-trap region, and proceed throughout the clearing and magnet rampdown operations. Since the majority of charged particles have already been removed, a single-particle treatment of the dynamics is appropriate. As such, these simulations use classical force equations to propagate the particles through the apparatus. Two particle propagation methods were implemented: a full Lorentz-force

propagator and a guiding-center approximation. The two method approach allowed for cross-validation – where consistency between results helps gain confidence in the simulation.

The magnetic field was first analytically modeled using the combination of vector potentials to approximate the solenoid-octupole-mirror coils array. The model parameters were then fitted to a TOSCA/OPERA3D [161] calculation of the full magnetic field, which included effects of the magnet and apparatus structure and material. This approach resulted in an analytic vector potential model of the magnetic field, which avoids problems associated with using discrete magnetic field maps, and also agrees with the full magnetic field calculation to within 0.02 T.

Accurate modeling of the electric fields is very important when simulating the trajectories of bare antiprotons. A finite-difference method, along with a multi-grid relaxation technique, are used to calculate the discrete electric potentials at fixed radial and azimuthal steps. For positions between the grid spacing, a bicubic spline is used to interpolate the potential. The time-dependent response of the electrode amplifiers is also taken into consideration when calculating the electric potentials.

When simulating a distribution of antihydrogen atoms or antiprotons, the individual trajectories are calculated independently, with their initial positions and momenta taken from the underlying distribution. For example, a spatially flat distribution can be generated by assigning an equal probability to each point in the trap. Likewise, a thermal distribution of arbitrary temperature can be simulated by sampling the initial momenta from a Maxwell-Boltzmann distribution. Once initialized, the simulated particles are subjected to the clearing fields and the changing magnetic fields during magnet ramp-down. The simulation is concluded when the particle encounters the electrode surface, or a time limit of 50 ms after the beginning of the magnet shutoff is reached. These simulations are extremely useful in the modelling of the dynamics of mirror-confined an-

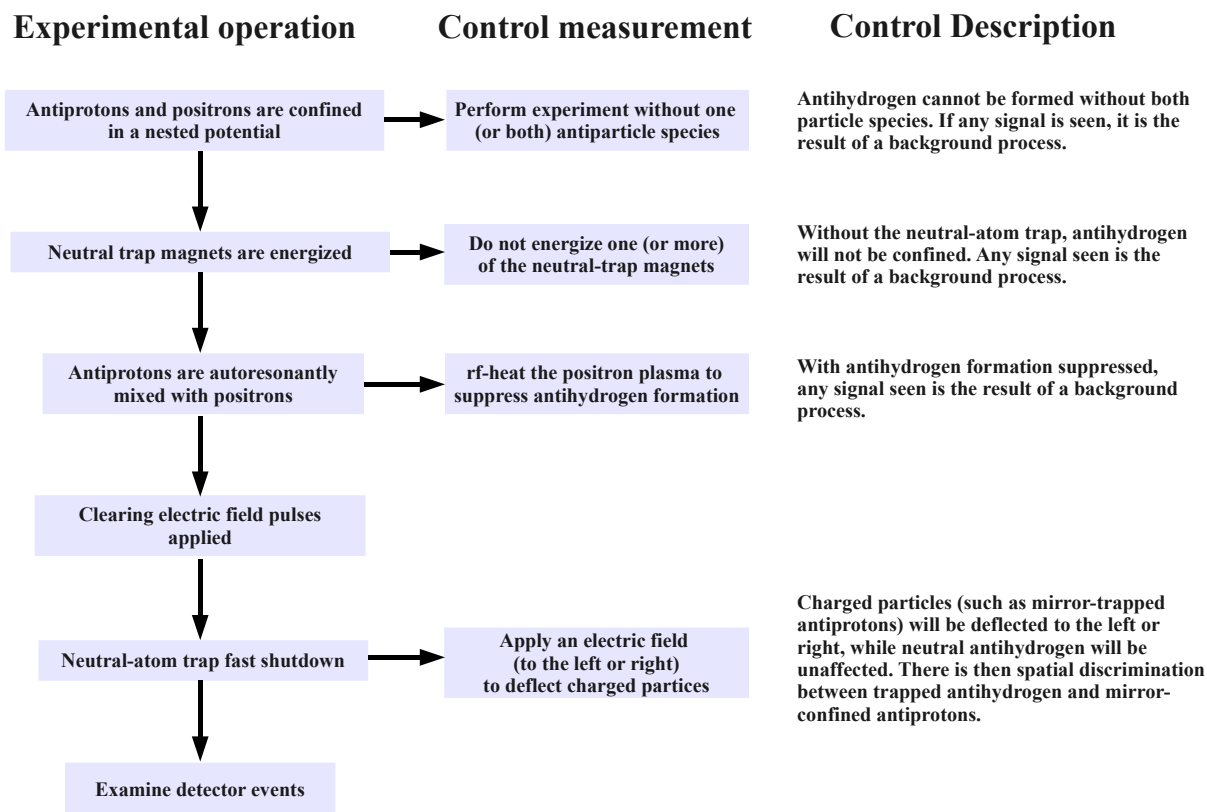


Figure 7.2: A flow chart outlining the various control and cross-check measurements and how they correspond to apparatus operations during the trapping experiments.

tiprotons (Sec. 7.1 and 7.3.3), and the release signature of trapped antihydrogen (Sec. 7.4 and Sec. 7.5).

7.3 Control measurements

With such a complicated experiment, it is necessary to implement a number of control measurements and cross-checks. These measurements involve modifying an operation in the trapping experiment in order to rule out backgrounds as the source of observed signal. For a cross-check measurement, one or more experimental components are disabled such that the magnetic trapping of antihydrogen is impossible. Therefore, any observed annihilations are due to procedural faults, which need to be resolved before an unambiguous trapping signal can be isolated.

Likewise, control measurements are needed to separate trapped antihydrogen from mirror-trapped antiprotons. A flow chart of the various control and cross-check measurements and how they relate to the trapping experiment operation is given as Fig. 7.2. Electric fields can be used to spatially separate the bare charged antiprotons from the neutral antihydrogen atoms. Both the control and cross-check measurements are described in the following sections.

7.3.1 Neutral trap magnets disengaged

For these cross-check measurements, either one of the mirror coils or the octupole magnet is not engaged during the trapping experiment. In this way, the neutral-atom trap does not cover the entire solid angle, and any trapped antihydrogen atoms will eventually find this hole and escape. Thus, any annihilation signal observed during these experiments indicates that antiprotons are being inadvertently confined and released. Moreover, observation of annihilation signal when only one of the mirror coils is energized is likely due to antiprotons being confined in an inappropriate electric potential, as only one mirror coil is inadequate to axially mirror-trap antiprotons.

7.3.2 Without antiprotons and/or positrons

Another cross-check measurement involves proceeding with the trapping experiment, but without either the antiprotons or the positrons (or proceeding with neither antiparticle species). Since both antiprotons and positrons are needed to produce antihydrogen, any annihilation signal during these experiments must be the result of some background process. For example, an annihilation signal in the detection window for the antiproton-only experiment is a good indication that there are mirror-confined antiprotons that are not being removed by the clearing pulses. Annihilation signals during positron-only experiments are likely the result of the cosmic ray background.

7.3.3 Electric potential biasing during fast magnetic shutdown

A powerful control, which can be used to discriminate between antihydrogen and bare antiproton annihilations, involves applying an axial electric field during the fast magnet shutdown. This bias field will deflect charged particles, while leaving neutral antihydrogen unaffected. In this way, mirror-confined antiprotons can be spatially separated from the neutral antihydrogen, and thus identified.

Figure 7.3 shows the result of modified trapping experiments designed to deliberately mirror-confine antiprotons. Here, the positron source is not engaged and only antiprotons are collected so that no antihydrogen atom are produced to interfere with our mirror-trapped antiproton signal. Instead, the antiprotons are pushed over a 40 eV potential drop and held while antiproton-antiproton collisions distribute the imparted energy between the longitudinal and transverse degrees of freedom. The redistribution of energy from the longitudinal to the transverse motion means that some of the antiprotons will be able to survive the pulsed clearing fields in mirror-trapped orbits. Additionally, fewer clearing pulses are applied, and as a result, several mirror-trapped antiprotons can be observed every experimental cycle. Moreover, the polarity of the applied bias can be inverted to push the bare antiprotons axially in either direction. The red triangles in Fig. 7.3 show the reconstructed z position of the annihilation vertex and event time after the start of the magnet shutdown (where every included point has satisfied the annihilation event selection criteria of Chapter 6). Similarly, the green triangles give the events recorded with the left bias field, and the blue events give the events during the right bias field. The coloured points in the background show the simulated antiproton distributions (where the point colours correspond to colours of the event markers). The simulated hit locations have been convolved with the detector resolution function (Eq. 5.14) to reflect the uncertainty in the z coordinate reconstruction.

It can be seen from Fig. 7.3 that the electric bias field effectively pushes the antipro-




Figure 7.3: Scatter plot of the reconstructed axial vertex coordinate, z , and time after the start of the fast magnet shutdown, t , for intentionally mirror-confined antiprotons. Simulated antiproton distributions are shown as the small dots (the simulated distributions have been convolved with the reconstruction resolution function to reflect the uncertainty in z). The measured no-bias events are shown as the upside-down triangles, the left-bias events as the green rightside-up triangles, and the right-bias events as the blue circles. The simulated no-bias events are shown as the red dots, the simulated left-bias events as the green dots, and the simulated right-bias events as the blue dots.

tons either axially to the left or to the right. Also to note is the good correspondence between the simulated and measured antiproton distributions, both in position and in time¹. This provides a validation of the simulation of Sec. 7.2, which are shown to adequately predict antiproton trajectories in the time-dependent electric and magnetic fields.

7.3.4 Heated positrons

Another powerful control measurement comes from the heating of the positron plasma during antihydrogen formation. Here, the positron plasma is heated to ~ 1100 K through the application of an rf-drive to the dipole plasma mode. This effectively suppresses antihydrogen formation, since the kinetic energy of the positrons becomes much greater than the antihydrogen binding energy. The antiproton temperature will come to equilibrium with the heated positrons (which act like a heat reservoir). Any antihydrogen that does form, when the relative velocity of the positrons and antiprotons match, will have a large amount of kinetic energy and will likely be unconfined by the neutral trap.

The main advantage of this control measurement is that it does not affect (to first order) the electrostatic or magnetic environment. While the other measurement in Sec. 7.3 can change the particle dynamics within the neutral trap, this measurement only affects the temperature parameter. As such, it can be considered a null measurement and counterpart to the standard trapping experiment. Moreover, since an annihilation signal during the positron-heated experiments is unlikely to be due to trapped antihydrogen (and the situation is identical to the standard experiment aside from particle temperature), any observed annihilation events can be considered direct background to

¹Note that the electrodes at the ends of the trap decrease in radius. This radial step occurs at ± 13 cm, and is coincident with the four red triangles in the upper right-hand section of Fig. 7.3. The antiproton simulations take into account this feature, and a small number of simulated antiproton annihilations can be seen at ± 13 cm. It is also of note that all four late-time annihilations occur to the right, which may be the result of small field errors not included in the simulations.

the trapping signal. The positron heating technique can also be combined with the bias during magnet shutdown (Sec. 7.3.3) to measure their backgrounds directly.

7.4 Initial search for trapped antihydrogen (2009)

Between late October and early November 2009, ALPHA performed its first systemic effort to magnetically confine antihydrogen [53]. Although very similar to the trapping experiment described in Sec. 3.6, these initial experiments did not include: evaporative cooling of the positron plasma, experiments where the electric potential bias was applied during magnet shutdown control, or experiments where the positrons were rf-heated. Also, due to a signal routing error, the detector readout trigger was not as described in Sec. 4.1.5 (and shown as Fig. 7.4). Rather, the modules included in the readout trigger are shown as Fig. 7.5. This erroneous mapping drastically reduced the solid angle coverage of triggerable modules. However, the reduction in solid angle was mitigated by the track multiplicity for annihilation events. That is, since each annihilation will, on average, produce about three charged pions, there are several tracks that can satisfy the the readout trigger. Overall, this mismapping resulted in a decrease in the trigger efficiency by about 10%, from the intended $\sim 98\%$ to $\sim 86\%$. This mapping also decreased the unsuppressed cosmic background rate from the nominal $(9.65 \pm 0.02) \text{ s}^{-1}$ (Sec. 6.1.3) to $(4.35 \pm 0.01) \text{ s}^{-1}$. After the application of the background rejection cuts (which were optimized for this dataset, and determined to be the same as given in Sec. 6.3), the expected cosmic background rate for the initial trapping experiments is $(2.2 \pm 0.1) \times 10^{-2} \text{ s}^{-1}$.

During this trapping series, 212 ‘standard’ experiments were conducted. In each cycle, around 4.5×10^4 antiprotons (with mean temperature of $358 \pm 55 \text{ K}$) were autoresonantly mixed with about 2.2×10^6 positrons (with mean temperature of $71 \pm 10 \text{ K}$) for 1 s. In

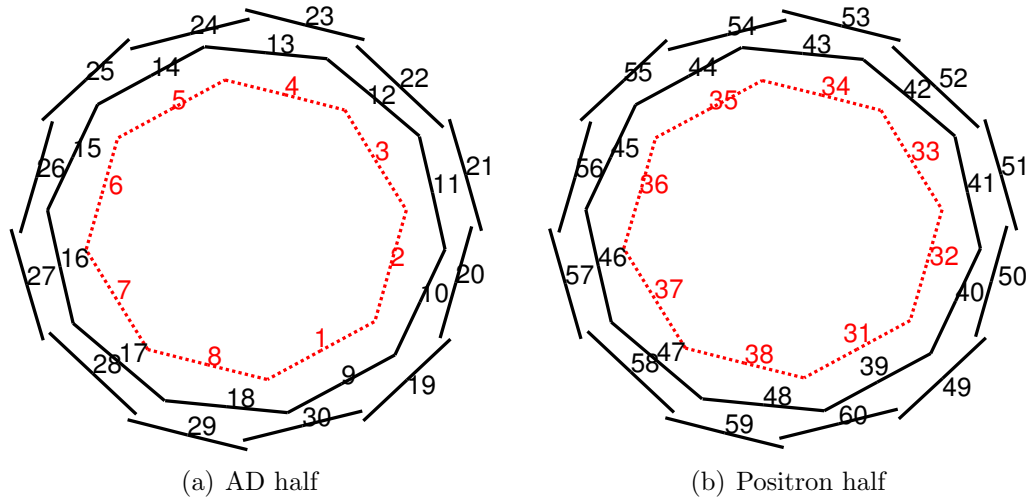


Figure 7.4: Diagram illustrating the intended detector $Si > 1$ trigger module radial and azimuthal locations for use during the 2009 experimental run. The red dashed modules were intended to be included in the primary detector trigger algorithm.




Figure 7.5: Diagram illustrating the actual detector $Si > 1$ trigger module radial and azimuthal locations for use during the 2009 experimental run. The red dashed modules were the actual modules included in the primary detector trigger algorithm.

total, about 10^7 antiprotons were mixed with 1.3×10^9 positrons. After applying the vertex reconstruction and background rejection routines, 7 events² satisfied all antihydrogen annihilation selection criteria. Figure 7.6 shows the $z - t$ distribution of observed events, including a comparison to the simulated antihydrogen distribution (Fig. 7.6a) and simulated mirror-confined antiproton distribution (Fig. 7.6b). As shown, the identified annihilation events are entirely consistent with the expected antihydrogen distribution, but fall outside of the expected distribution for mirror-trapped antiprotons. The number of observed events is also inconsistent with the measured rate of cosmic ray events. Given a total observation time of 6.36 s for the entire 212 experiments and the known (suppressed) cosmic rate, 0.14 ± 0.01 cosmic events are expected. Thus, the probability that all 7 observed events were due solely to statistical fluctuations of the cosmic background (the Poisson p-value, Eq. 6.6) is 1.9×10^{-10} , corresponding to a significance of 6.3 standard deviations.

The trapping experiments and cross-check measurements are summarized in Table 7.1. The number of events in the 30 ms detection window are tabulated (along with tallies for the extended period of 50 ms after magnet shutdown, which demonstrates the robustness of the time-cut for the events satisfying all the cut). An important feature to note is the large reduction between the total number of detector readouts and the number of events returning a vertex. This reduction is due to the presence of noise-induced background events resulting from the fast shutdown of the superconducting neutral trap magnets (see Sec. 6.1.1, and Fig. 6.3). Also to note is that none of the cross-check measurements (described in Sec. 7.3) resulted in any annihilation events. This is informative, as any observed annihilation events would indicate a major flaw in the trapping sequence. However, none of the cross-check measurements serves as a good null experiment (as they

²The initial analysis gave 6 events that satisfied all selection criteria [53]. The new number is based on the reconstruction routines described in Chapter 5, which are updated and improved from the routines used in the initial analysis. It should be noted that the author performed both the original and improved analyses, as well as implemented the improvements to the reconstruction routines.

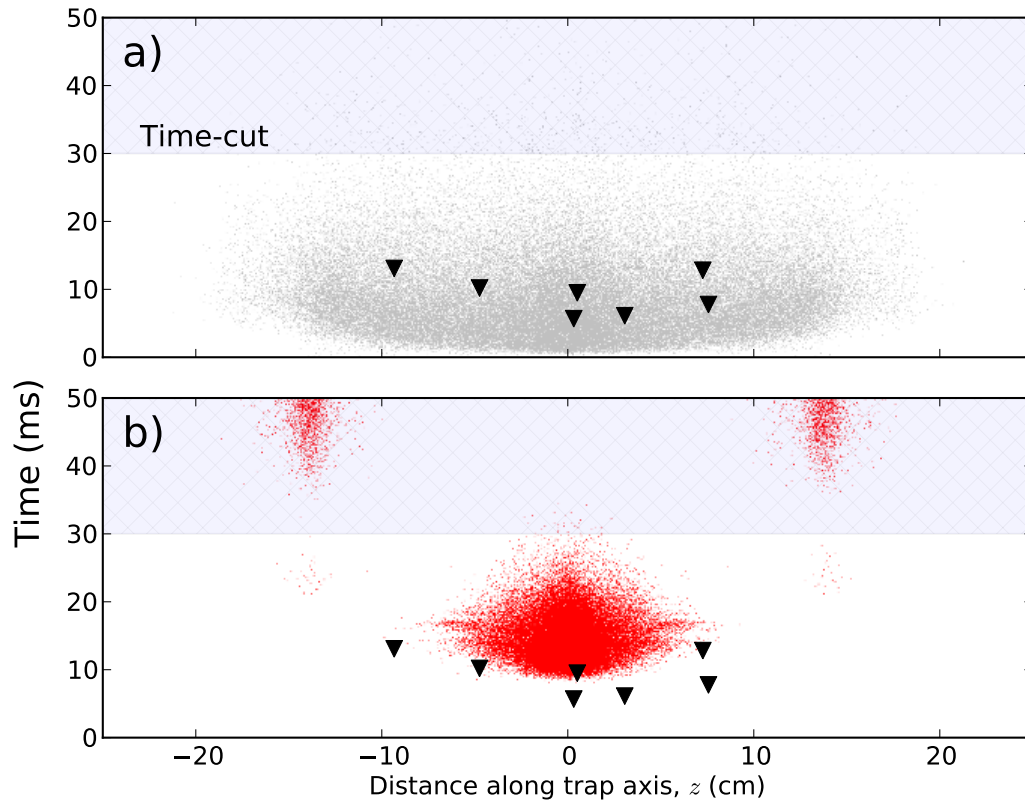


Figure 7.6: The $z - t$ distributions for a) simulated antihydrogen atoms (grey dots) and b) simulated mirror trapped antiprotons (red dots). The simulated distributions have been convoluted with the reconstruction resolution function to reflect the uncertainty in z . The observed events (satisfying all cuts) are shown as the black triangles. The region enclosed by the grey rectangle is excluded by the time-cut.

all affect the dynamics of the system), so they do not rule out the possibility that the observed events are due to background processes.

Without adequate null measurements or mirror-trapped antiproton discrimination, the 2009 trapping series was very encouraging, but not quite conclusive. Since the Monte Carlo simulations lacked experimental validation³, the distributions shown in Fig. 7.6 were not enough to claim that the observed events were definitively trapped antihydrogen.

³The validation of the Monte Carlo simulations shown in Sec. 7.2 for mirror-confined antiprotons was completed in the Summer of 2010.

Type of experiment	Number of experiments	Total number of readouts [†]	Number of events returning vertices [†]	Number of events satisfying all cuts (annihilation events) [†]
'Standard' trapping experiment	212	116 (138)	11 (15)	7 (7)
No octupole magnet	55	33 (39)	3 (4)	0 (0)
No left mirror coil	8	2 (2)	1 (1)	0 (0)
No antiprotons	39	7 (9)	0 (2)	0 (0)
No particles	121	41 (53)	0 (1)	0 (0)
No particles and no octupole magnet	23	3 (4)	0 (0)	0 (0)
No particles and no left mirror coil	24	2 (2)	2 (2)	0 (0)

[†]The first value quoted is for the 30 ms detection window. The value in parenthesis gives the total number of events for 50 ms after the magnet shutdown.

Table 7.1: Summary table of the relevant experimental parameters for the 2009 trapping experiment events.

Type of experiment	Number of experiments	Total number of readouts [†]	Number of events returning vertices [†]	Number of events satisfying all cuts (annihilation events) [†]
No bias ('standard experiment')	137	142 (179)	30 (36)	20 (20)
Left bias	101	69 (92)	20 (21)	14 (14)
Right bias	97	96 (120)	19 (24)	14 (14)
No bias, heated positrons	132	128 (157)	9 (11)	1 (1)
Left bias, heated positrons	60	18 (28)	1 (5)	0 (1)
Right bias, heated positrons	54	16 (25)	1 (2)	0 (0)

[†]The first value quoted is for the 30 ms detection window. The value in parenthesis gives the total number of events for 50 ms after the magnet shutdown.

Table 7.2: Summary table of the relevant experimental parameters for the 2010 trapping experiment events.

7.5 Observation of trapped antihydrogen (2010)

Almost a year after the initial experiments, the next set of systematic trapping attempts were completed. A total of 335 trapping experiments were completed between mid-August and mid-September 2010, combining 10^7 antiprotons with 7×10^8 positrons [54]. Unlike the initial attempts (Sec. 7.4), these experiments included evaporative cooling the positron plasma to about 40 K (the average antiproton temperature is also improved to about 200 K). Also, of the 335 experiments, 101 included a ‘left bias’ electric field during the fast magnet shutdown and 97 included a ‘right bias’ electric field. These modified experiments allow for definitive discrimination of mirror-trapped antiprotons. The correct detector trigger (Fig. 7.4) was also used for these measurements.

Table 7.2 summarizes the results of these trapping experiments. In total, 48 events⁴ survived all of the annihilation selection cuts (including the 30 ms time-cut). In contrast, 246 control measurements were performed, where the positron plasma was heated to 1100 K (Sec. 7.3.4), and only 1 annihilation event was found. This event (or 1.4 ± 1.4 , when scaled to match the 335 experiments conducted with cold positrons) constitutes the direct background for these trapping measurements. This should be compared to the measured $(4.6 \times 10^{-2}) \text{ s}^{-1}$ cosmic background rate (after suppression), which is expected to contribute 0.46 ± 0.01 counts resulting from misidentified cosmic rays. This background rate, either observed or inferred, is clearly insufficient to account for the observed signal.

The $z - t$ distributions of these events are shown as Fig. 7.7, and compared with the simulated antihydrogen (Fig. 7.7a) and mirror-trapped antiproton (Fig. Fig. 7.7b) distributions. Additionally, Fig. 7.7 is separated into several categories, depending on the variation of experiment performed (no bias, left bias, right bias). Recalling Fig. 7.3, with

⁴A total of 38 events were reported in [54]. The results presented here follow the reconstruction routines described in Chapter 5, which are updated and improved from the routines used in the original publication. It should be noted that the author performed both the original and improved analyses, as well as implemented the improvements to the reconstruction routines.

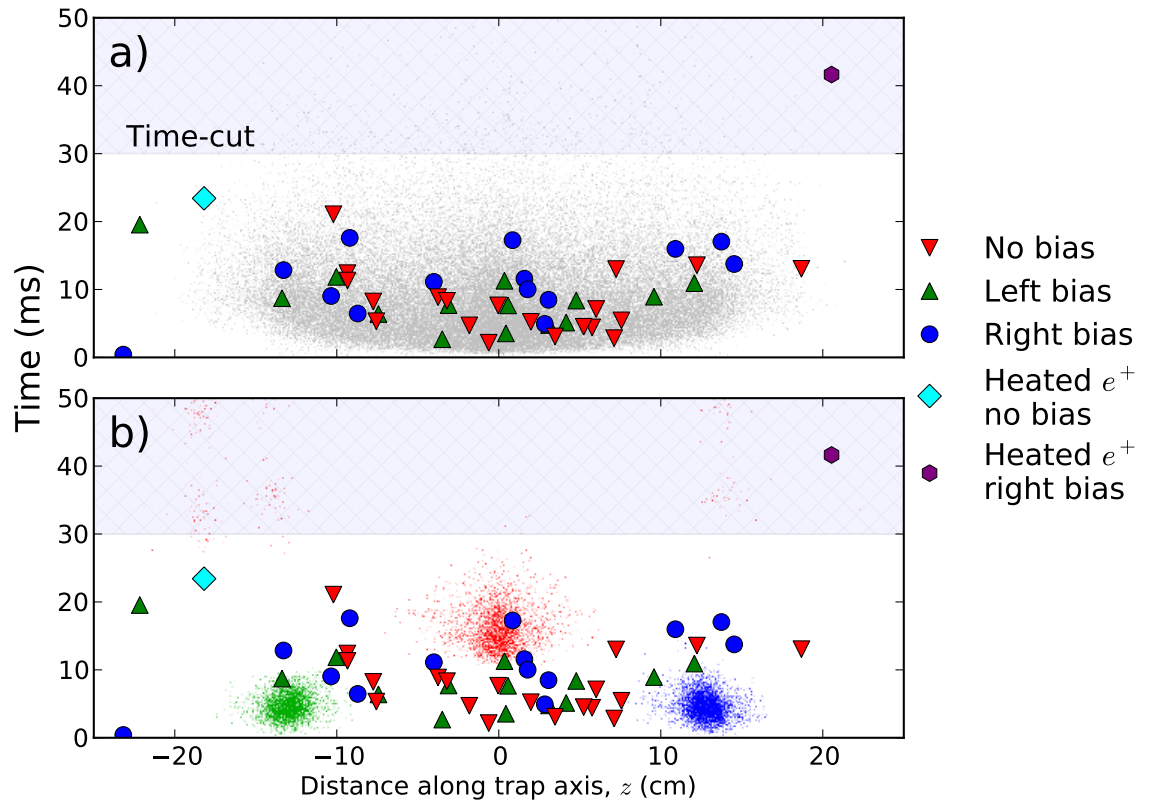


Figure 7.7: The $z-t$ distributions for a) simulated antihydrogen atoms (grey dots), and b) simulated mirror trapped antiprotons (coloured dots). The simulated distributions have been convoluted with the reconstruction resolution function to reflect the uncertainty in z . The measured no-bias events are shown as the upside-down triangles, the left-bias events as the green right-side-up triangle, and the right-bias events as the blue circles. Events during heated-positron control experiments are shown as cyan(purple) for the no-bias(right-bias) cases. The simulated no-bias events are shown as the red dots, the simulated left-bias events as the green dots, and the simulated right-bias events as the blue dots. The region enclosed by the grey rectangle is excluded by the time-cut.

no bias field, mirror-trapped antiprotons escape from the center of the trap. Conversely, when the left or right bias fields are applied, the antiprotons escape from either extent of the trap. This behaviour is not seen for any of the trapping series event distributions, for which the events are spread out over the entire volume of the trap. This indicates that these annihilation events are insensitive to electric fields, characteristic of antihydrogen.

The combination of the heated-positron null measurement and the insensitivity to the bias potentials rules out background processes being responsible for the events observed in the trapping measurements. This leads to the conclusion that the annihilation events in Table 7.2 and Fig. 7.7 are due to the release of trapped antihydrogen [54]. The time between the end of mixing and the fast shutdown of the neutral trap magnets is 172 ms, which sets the lower limit for the confinement time for this set of experiments (that is, if an antihydrogen atom remain trapped from the beginning of mixing onward, it would be confined for just over a second, see Sec. 3.6).

From Sec. 6.3.4, the measured acceptance of the reconstruction algorithms is $(64.4 \pm 0.1)\%$, and when combined with the estimated trigger efficiency of $(90 \pm 10)\%$ gives an overall detection efficiency of $(58 \pm 7)\%$. This detection efficiency can then be used to estimate the number of antihydrogen atoms trapped, N_{trapped} , from

$$N_{\text{detected}} = N_{\text{trapped}} \times f_{\text{detection}}, \quad (7.2)$$

where N_{detected} is the number of observed antihydrogen annihilation events, and $f_{\text{detection}}$ is the detection efficiency. Using Eq. 7.2, the number of atoms trapped is $N_{\text{trapped}} = 83 \pm 19$, or 0.25 ± 0.06 trapped atoms per attempt.

The rate of trapped atoms per attempt is an important measure when designing or attempting spectroscopic measurements. Obviously, a trapping rate of zero is incompatible with attempting spectroscopic measurements, since trapped atoms are required to observe any result. Conversely, a large trapping rate helps facilitate spectroscopic mea-

surements, as weaker incident radiation can then be used to measure the atomic spectra. For this reason, maximizing the trapping rate is an important near-term experimental goal for the ALPHA collaboration.

7.6 Summary

This chapter described the experimental search and eventual observation of magnetically confined antihydrogen. First, the mirror-trapped antiproton background was examined along with the employment of pulsed electric fields to clear the neutral trap of charged particles. Monte Carlo simulations of bare antiprotons and antihydrogen atoms during the fast magnet shutdown were then described. Before presenting the results of the trapping experiments, several control and cross-check measurements were outlined. Finally, the initial trapping series conducted in 2009 was described along with the experiments performed in 2010. The cross-checks and control measurements are also presented, which ultimately rule out the background possibilities and conclude that the observed events are, in fact, trapped antihydrogen.

This achievement marks the first time any group has successfully confined antihydrogen. Moreover, this success crosses a major milestone on the road to high-precision studies of antihydrogen and tests of fundamental physical symmetries.

Chapter 8

Future prospects and conclusion


The magnetic confinement of antihydrogen makes possible a number of interesting experiments. Generally, these experiments require the antihydrogen atoms to remain in the neutral-atom trap for long enough to complete the given measurement. Indeed, the study of antihydrogen has been, so far, mainly concerned with the production of antihydrogen, as the atoms vanished almost immediately after formation.

This chapter will first describe another recent ALPHA result: the confinement of antihydrogen for as long as 1000 s. This result then motivates discussion of possible mid- and long-term measurements with trapped antihydrogen, specifically, microwave and laser spectroscopy of antihydrogen. Finally, some remarks to conclude this dissertation will be given.


8.1 Recent measurement: long-time confinement of antihydrogen

Immediately after the ‘proof of concept’ trapping experiments, where antihydrogen was confined for more than 172 ms (but not much more than a second), ALPHA began work on extending this confinement time. When designing the initial trapping experiments, it was not known how long the antihydrogen atoms would remain in the neutral-atoms trap. For this reason, the neutral trap magnets were set to shut off as soon as operationally possible. As a follow-up, longer confinement times were then investigated by simply adding more wait-time, with the atom trapping fields on, before the fast magnet shutoff.

Figure 8.1 shows the results of several sets of trapping experiments with various waiting-times [55]. Both the trapping rate (defined as the number of atoms trapped per



(a) Trapping rate



(b) Signal significance

Figure 8.1: The long-time confinement of antihydrogen (reproduced with permission from [55]). Figure (a) shows the trapping rate as a function of confinement time. The errors bar give the counting uncertainty only. Figure (b) shows the statistical significance of the observed events with respect to the cosmic ray background. The point for 0.4 s is off the scale ($> 20\sigma$), and not shown.




Figure 8.2: The $z - t$ scatter-plot, showing the results of trapping experiments with various confinement times, as indicated in the upper right scale in the figure (reproduced with permission from [55]). The simulated antihydrogen distribution is shown as the grey dots.

attempt, Fig. 8.1a) and the signal significance over the cosmic background (Fig. 8.1b) are given. Although the trapping rate appears to decrease for longer confinement times, the number of observed events is statistically significant for experiments up to 1000 s (an event was observed for the 2000 s experiments, but had only a 2.6σ signal significance).

In total, there were 309 events for the entire 2010 AD run (all shown in Fig. 8.2). With these additional statistics (and feedback from the simulations), the kinetic-energy distribution of the trapped antihydrogen atoms can also be inferred [55]. Just as importantly, the calculations indicate that most antihydrogen atoms appear to have de-excited to the ground state within ~ 1 second. Thus, for extended confinement times ($\gg 1$ s), the released antihydrogen atoms will likely be in the ground state, which is preferable for precision spectroscopic measurements.

8.2 Upcoming measurement: microwave spectroscopy

Certain ground state antihydrogen hyperfine transitions can be measured by inducing transitions that force the atom from a trapped to an untrappable state. Figure 8.3 shows




Figure 8.3: Breit-Rabi diagram for ground state hyperfine levels for (anti)hydrogen. The states shown in red correspond to the low-field seeking spin states, which are anti-parallel to the magnetic field, while the blue states correspond to the high-field seeking spin states, which are parallel to the magnetic field.

how the energy states in a ground state (anti)hydrogen atoms vary with magnetic field. The two high-field seeking (untrappable) states are labeled as $|a\rangle$ and $|b\rangle$, while the low-field seeking (trappable) states are labeled as $|c\rangle$ and $|d\rangle$. An example measurement would then proceed as follows: a population of trapped antihydrogen atoms in state $|c\rangle$ ($|d\rangle$) are subjected to microwave-range photons resonant with the $|c\rangle \rightarrow |b\rangle$ ($|d\rangle \rightarrow |a\rangle$) transition. The atom undergoes a positron spin-flip, leaving it in the untrapped state, and it quickly leaves the trap and annihilates on the electrode surface. The measurements of these transitions can then be combined to determine:

$$f_{|d\rangle \rightarrow |a\rangle} - f_{|c\rangle \rightarrow |b\rangle} = a/h, \quad (8.1)$$

where $f_{|d\rangle \rightarrow |a\rangle}$ and $f_{|c\rangle \rightarrow |b\rangle}$ are the two transitions frequencies, h is Planck's constant, and a is the zero-field hyperfine splitting constant (in hydrogen, $a/h = 1420$ MHz is the famous '21 cm' transition and is determined to parts in 10^{12} [20]).

An advantage of such a microwave measurement is that the current ALPHA apparatus would not require major modifications. With only minor alterations to aid the injection of the microwaves into the trap region, it should be possible to measure a/h in antihydrogen to better than parts in 10^5 [162]. The limiting factor with the current apparatus is the magnetic field inhomogeneity, which strongly affects the transition intervals.

8.3 Future measurement: new apparatus and laser spectroscopy

In hydrogen, the ‘gold-standard’ of high precision metrology is the $1S - 2S$ two-photon transition. This transition is particularly attractive, as it has a natural linewidth of 1.3 Hz. Moreover, the first order Doppler shifts can be canceled using two counter-propagating lasers. This transition has been measured, using a phase-coherent comparison with an atomic cesium fountain clock to parts in 10^{14} [11]. The $1S - 2S$ line provides an excellent yardstick to compare hydrogen and antihydrogen, which, according to CPT symmetry, should have identical transition frequencies.

However, the current ALPHA apparatus has no means for laser access, so such a measurement must wait for a major upgrade to the experimental setup (the ALPHA collaboration is currently pursuing a major apparatus redesign with the goal of accommodating laser access). Even with such an upgrade, the spectroscopic measurement might rely heavily on the silicon detector and event reconstruction routines for sensitive antihydrogen detection. A possible scheme might proceed as follows:

1. The ground state antihydrogen is first excited to the $2S$ level via two 242 nm counter-propagating pump lasers.
2. The excited atom is then removed from the magnetic trap, either through photoionization or by Stark-coupling the $2S$ and $2P$ states (thus, when the atom spontaneously decays from the $2P$ to the $1S$ state, it undergoes a $\Delta L = 1$ transition

and can end up in an unconfined angular momentum state).

3. The disassociated antiproton (or unconfined atom) annihilates on the electrode surface and is registered by the silicon detector (a similarity shared by the microwave measurement outlined in Sec. 8.2).

This scheme has the advantage of efficient detection of the unconfined antihydrogen atoms (or disassociated antiprotons). Moreover, with the long confinement times (Sec. 8.1) and good background rejection, it may be possible to spectroscopic measurements on very few trapped atoms. Some systematics involved with this measurement are described in [163].

8.4 Conclusion

This dissertation has described the first-ever trapping of antihydrogen, with an emphasis on its detection using the ALPHA silicon detector (the analysis for which the author was responsible for). ALPHA's demonstration of magnetically trapped antihydrogen marks the first time that atomic antimatter has been captured and confined. As outlined earlier in this chapter, this result paves the way for precision studies of the atomic spectra of antihydrogen (both microwave and laser), possible gravitational tests, and ultimately, the comparison with its hydrogen counterpart.

The silicon vertex detector played a crucial role in this result, along with the software routines used to reconstruct the antihydrogen annihilation position. Both the detector and the reconstruction routines have been described in detail, along with the background rejection methods using the silicon detector. This work has enabled the unambiguous detection of trapped antihydrogen, and will continue to be an important part of future spectroscopic studies in ALPHA.

Bibliography

- [1] G. Blanford *et al.*, Phys. Rev. D **57**, 6649 (1998).
- [2] Y. Enomoto *et al.*, Phys. Rev. Lett. **105**, 243401 (2010).
- [3] M. S. Fee *et al.*, Phys. Rev. A **48**, 192 (1993).
- [4] R. J. Hughes and B. I. Deutch, Phys. Rev. Lett. **69**, 578 (1992).
- [5] R. S. Van Dyck, P. B. Schwinberg, and H. G. Dehmelt, Phys. Rev. Lett. **59**, 26 (1987).
- [6] A. Angelopoulos *et al.*, Phys. Lett. B **471**, 332 (1999).
- [7] A. Kreissl *et al.*, Z. Phys. C **37**, 557 (1988).
- [8] M. Hori *et al.*, Nature **475**, 484 (2011).
- [9] K. Nakamura *et al.*, (Particle Data Group), J. Phys. G **37**, 075021 (2010).
- [10] G. Lüders, Ann. Phys. **2**, 1 (1957).
- [11] M. Niering *et al.*, Phys. Rev. Lett. **84**, 5496 (2000).
- [12] T. Pohl, H. R. Sadeghpour, Y. Nagata, and Y. Yamazaki, Phys. Rev. Lett. **97**, 213001 (2006).
- [13] C. L. Cesar, F. Robicheaux, and N. Zagury, Phys. Rev. A **80**, 041404 (2009).
- [14] S. Wu, R. C. Brown, W. D. Phillips, and J. V. Porto, Phys. Rev. Lett. **106**, 213001 (2011).
- [15] R. Bluhm, V. A. Kostelecký, and N. Russell, Phys. Rev. Lett. **82**, 2254 (1999).

- [16] B. Juhász and E. Widmann, *Hyperfine Interactions* **193**, 305 (2009).
- [17] M. Fujiwara *et al.*, *Hyperfine Interactions* **172**, 81 (2006).
- [18] M. D. Ashkezari *et al.*, (ALPHA Collaboration), Progress toward microwave spectroscopy of trapped antihydrogen, to be published in the proceedings of the 2011 LEAP conference.
- [19] C. Schwob *et al.*, *Phys. Rev. Lett.* **82**, 4960 (1999).
- [20] L. Essen, R. W. Donaldson, E. G. Hope, and M. J. Bangham, *Metrologia* **9**, 128 (1973).
- [21] E. G. Adelberger, B. R. Heckel, C. W. Stubbs, and Y. Su, *Phys. Rev. Lett.* **66**, 850 (1991).
- [22] E. G. Adelberger and B. R. Heckel, *Phys. Rev. Lett.* **67**, 1049 (1991).
- [23] T. Goldman *et al.*, *Phys. Rev. Lett.* **67**, 1048 (1991).
- [24] M. Fischler, J. Lykken, and T. Roberts, arXiv:0808.3929 [hep-th] (2008).
- [25] T. Goldman and M. M. Nieto, *Physics Letters B* **112**, 437 (1982).
- [26] T. W. Darling, F. Rossi, G. I. Opat, and G. F. Moorhead, *Rev. Mod. Phys.* **64**, 237 (1992).
- [27] R. Hughes, *Nucl. Phys. A* **558**, 605 (1993).
- [28] P. A. M. Dirac, *Proc. R. Soc. Lond. A* **133**, 60 (1931).
- [29] C. D. Anderson, *Phys. Rev.* **43**, 491 (1933).
- [30] O. Chamberlain, E. Segrè, C. Wiegand, and T. Ypsilantis, *Phys. Rev.* **100**, 947 (1955).

- [31] G. Baur *et al.*, Phys. Lett. B **368**, 251 (1996).
- [32] G. Blanford *et al.*, Phys. Rev. Lett. **80**, 3037 (1998).
- [33] G. Gabrielse, S. L. Rolston, L. Haarsma, and W. Kells, Phys. Lett. A **129**, 38 (1988).
- [34] G. Gabrielse *et al.*, Phys. Rev. Lett. **57**, 2504 (1986).
- [35] G. Gabrielse *et al.*, Phys. Rev. Lett. **63**, 1360 (1989).
- [36] M. Amoretti *et al.*, (ATHENA Collaboration), Nature **419**, 456 (2002).
- [37] G. Gabrielse *et al.*, (ATRAP Collaboration), Phys. Rev. Lett. **89**, 213401 (2002).
- [38] M. Amoretti *et al.*, (ATHENA Collaboration), Phys. Lett. B **578**, 23 (2004).
- [39] G. Gabrielse *et al.*, (ATRAP Collaboration), Phys. Rev. Lett. **89**, 233401 (2002).
- [40] M. Amoretti *et al.*, (ATHENA Collaboration), Phys. Lett. B **583**, 59 (2004).
- [41] G. Gabrielse *et al.*, (ATRAP Collaboration), Phys. Rev. Lett. **93**, 073401 (2004).
- [42] M. Amoretti *et al.*, (ATHENA Collaboration), Phys. Lett. B **590**, 133 (2004).
- [43] M. C. Fujiwara *et al.*, (ATHENA Collaboration), Phys. Rev. Lett. **101**, 053401 (2008).
- [44] C. H. Storry *et al.*, (ATRAP Collaboration), Phys. Rev. Lett. **93**, 263401 (2004).
- [45] M. Amoretti *et al.*, (ATHENA Collaboration), Phys. Rev. Lett. **97**, 213401 (2006).
- [46] M. C. Fujiwara *et al.*, (ATHENA Collaboration), Phys. Rev. Lett. **92**, 065005 (2004).
- [47] N. Madsen *et al.*, (ATHENA Collaboration), Phys. Rev. Lett. **94**, 033403 (2005).

- [48] G. Andresen *et al.*, (ALPHA Collaboration), Phys. Rev. Lett. **98**, 023402 (2007).
- [49] G. B. Andresen *et al.*, (ALPHA Collaboration), J. Phys. B: At. Mol. Opt. Phys. **41**, 011001 (2008).
- [50] G. Andresen *et al.*, (ALPHA Collaboration), Phys. Lett. B **685**, 141 (2010).
- [51] G. Gabrielse *et al.*, (ATRAP Collaboration), Phys. Rev. Lett. **98**, 113002 (2007).
- [52] G. Gabrielse *et al.*, (ATRAP Collaboration), Phys. Rev. Lett. **100**, 113001 (2008).
- [53] G. B. Andresen *et al.*, (ALPHA Collaboration), Phys. Lett. B **695**, 95 (2011).
- [54] G. B. Andresen *et al.*, (ALPHA Collaboration), Nature **468**, 673 (2010).
- [55] G. B. Andresen *et al.*, (ALPHA Collaboration), Nat. Phys. **7**, 558 (2011).
- [56] M. Amoretti *et al.*, (ATHENA Collaboration), Phys. Rev. Lett. **91**, 055001 (2003).
- [57] L. V. Jørgensen *et al.*, (ATHENA Collaboration), Phys. Rev. Lett. **95**, 025002 (2005).
- [58] R. Funakoshi *et al.*, (ATHENA Collaboration), Phys. Rev. A **76**, 012713 (2007).
- [59] G. Gabrielse *et al.*, (ATRAP Collaboration), Phys. Lett. B **548**, 140 (2002).
- [60] G. Gabrielse *et al.*, (ATRAP Collaboration), Phys. Rev. Lett. **106**, 073002 (2011).
- [61] A. Speck *et al.*, (ATRAP Collaboration), Phys. Lett. B **650**, 119 (2007).
- [62] G. B. Andresen *et al.*, (ALPHA Collaboration), Phys. Plasmas **15**, 032107 (2008).
- [63] G. B. Andresen *et al.*, (ALPHA Collaboration), Phys. Plas. **16**, 100702 (2009).
- [64] G. B. Andresen *et al.*, (ALPHA Collaboration), Phys. Rev. Lett. **105**, 013003 (2010).

- [65] G. B. Andresen *et al.*, (ALPHA Collaboration), Phys. Rev. Lett. **106**, 025002 (2011).
- [66] G. B. Andresen *et al.*, (ALPHA Collaboration), Phys. Rev. Lett. **100**, 203401 (2008).
- [67] N. Kuroda *et al.*, Phys. Rev. Lett. **100**, 203402 (2008).
- [68] G. Gabrielse *et al.*, (ATRAP Collaboration), Phys. Rev. Lett. **105**, 213002 (2010).
- [69] G. B. Andresen *et al.*, (ALPHA Collaboration), Phys. Rev. Lett. **106**, 145001 (2011).
- [70] M. Doser, J. Phys.: Conf. Ser. **199**, 012009 (2010).
- [71] R. Hydomako, Modelling of antihydrogen formation and the commissioning of the alpha antihydrogen apparatus, Master's thesis, University of Calgary, 2007.
- [72] M. H. Holzscheiter, M. Charlton, and M. M. Nieto, Phys. Rep. **402**, 1 (2004).
- [73] F. Robicheaux, J. Phys. B **41**, 192001 (2008).
- [74] A. Müller and A. Wolf, Hyperfine Interactions **109**, 233 (1997).
- [75] F. B. Yousif *et al.*, Phys. Rev. Lett. **67**, 26 (1991).
- [76] M. L. Rogelstad, F. B. Yousif, T. J. Morgan, and J. B. A. Mitchell, J. Phys. B: At. Mole. Opti. Phys. **30**, 3913 (1997).
- [77] M. L. Wall, C. S. Norton, and F. Robicheaux, Phys. Rev. A **72**, 052702 (2005).
- [78] E. A. Hessels, D. M. Homan, and M. J. Cavagnero, Phys. Rev. A **57**, 1668 (1998).
- [79] M. E. Glinsky and T. M. O'Neil, Phys. Fluids B **3**, 1279 (1991).

- [80] F. Robicheaux and J. D. Hanson, Phys. Rev. A **69**, 010701 (2004).
- [81] F. Robicheaux, Phys. Rev. A **73**, 033401 (2006).
- [82] F. Robicheaux, Phys. Rev. A **70**, 022510 (2004).
- [83] W. Paul, Rev. Mod. Phys. **62**, 531 (1990).
- [84] E. L. Raab, M. Prentiss, A. Cable, S. Chu, and D. E. Pritchard, Phys. Rev. Lett. **59**, 2631 (1987).
- [85] M. A. Levine, R. E. Marrs, J. R. Henderson, D. A. Knapp, and M. B. Schneider, Physica Scripta **1988**, 157 (1988).
- [86] L. S. Brown and G. Gabrielse, Rev. Mod. Phys. **58**, 233 (1986).
- [87] D. Wineland, P. Ekstrom, and H. Dehmelt, Phys. Rev. Lett. **31**, 1279 (1973).
- [88] R. S. Van Dyck and P. B. Schwinberg, Phys. Rev. Lett. **47**, 395 (1981).
- [89] R. S. Van Dyck, S. L. Zafonte, S. Van Liew, D. B. Pinegar, and P. B. Schwinberg, Phys. Rev. Lett. **92**, 220802 (2004).
- [90] J. H. Malmberg and C. F. Driscoll, Phys. Rev. Lett. **44**, 654 (1980).
- [91] M. Kretschmar, Eur. J. Phys. **12**, 240 (1991).
- [92] R. Davidson, *Physics of Nonneutral Plasmas* (Allied, 2009).
- [93] T. M. O'Neil, Phys. Fluids **23**, 2216 (1980).
- [94] D. H. E. Dubin and T. M. O'Neil, Rev. Mod. Phys. **71**, 87 (1999).
- [95] J. Notte and J. Fajans, Phys. Plasmas **1**, 1123 (1994).
- [96] T. M. O'Neil, Phys. Fluids **26**, 2128 (1983).

- [97] T. M. O'Neil and P. G. Hjorth, *Phys. Fluids* **28**, 3241 (1985).
- [98] W. H. Wing, *Prog. Quant. Elect.* **8**, 181 (1984).
- [99] W. Bertsche *et al.*, (ALPHA Collaboration), *Nucl. Instr. Meth. Phys. Res. A* **566**, 746 (2006).
- [100] M. Charlton and J. Humberston, *Positron physics* (Cambridge University Press, 2001).
- [101] L. Michel, *Il Nuovo Cimento* **10**, 319 (1953).
- [102] C. Regenfus *et al.*, *Nucl. Instr. Meth. Phys. Res. A* **504**, 343 (2003).
- [103] M. Amoretti *et al.*, *Nucl. Instr. Meth. Phys. Res. A* **518**, 679 (2004).
- [104] E. Klempt, C. Batty, and J.-M. Richard, *Phys. Rep.* **413**, 197 (2005).
- [105] G. Bendiscioli and D. Kharzeev, *La Rivista del Nuovo Cimento* **17**, 1 (1994).
- [106] C. Ghesquière, An inclusive view on $\bar{p}p \rightarrow \pi$ at rest, in *Symposium on Antinucleon-Nucleon Interactions*, CERN, 1974.
- [107] D. Polster *et al.*, *Phys. Rev. C* **51**, 1167 (1995).
- [108] P. Lubiński *et al.*, *Phys. Rev. C* **66**, 044616 (2002).
- [109] S. J. Orfanidis and V. Rittenberg, *Nucl. Phys. B* **59**, 570 (1973).
- [110] G. Lütz, *Semiconductor Radiation Detectors: Device Physics* (Springer, 2007).
- [111] P. Burger, *Nucl. Instr. Meth. Phys. Res. A* **226**, 112 (1984).
- [112] L. Rossi, *Pixel detectors: from fundamentals to applications* (Springer, 2006).
- [113] G. Batignani *et al.*, *Nucl. Instr. Meth. Phys. Res. A* **277**, 147 (1989).

- [114] *Product Guide*, Micron Semiconductor, UK.
- [115] T. J. Murphy and C. M. Surko, *Phys. Rev. A* **46**, 5696 (1992).
- [116] R. G. Greaves, M. D. Tinkle, and C. M. Surko, *Phys. Plasmas* **1**, 1439 (1994).
- [117] G. B. Andresen *et al.*, (ALPHA Collaboration), *Rev. Sci. Instrum.* **80**, 123701 (2009).
- [118] D. L. Eggleston, C. F. Driscoll, B. R. Beck, A. W. Hyatt, and J. H. Malmberg, *Phys. Fluids B: Plasma Phys.* **4**, 3432 (1992).
- [119] G. B. Andresen, *Evaporative cooling of antiprotons and efforts to trap antihydrogen*, PhD thesis, Aarhus University, 2010.
- [120] D. Möhl, *Hyperfine Interactions* **109**, 33 (1997).
- [121] S. Maury, *Hyperfine Interactions* **109**, 43 (1997).
- [122] D. Möhl and A. M. Sessler, *Nucl. Instr. Meth. Phys. Res. A* **532**, 1 (2004).
- [123] S. van der Meer, CERN Report No. ISR-PO-72-31, 1972 (unpublished).
- [124] D. Möhl, G. Petrucci, L. Thorndahl, and S. van der Meer, *Phys. Rep.* **58**, 73 (1980).
- [125] G. I. Budker, in *Proc. Int. Symp. on Electron and Positron Storage Rings*, 1966.
- [126] H. Poth, *Phys. Rep.* **196**, 135 (1990).
- [127] M. Charlton *et al.*, *J. Phys.: Conf. Ser.* **262**, 012001 (2011).
- [128] J. D. Jackson, *Classical Electrodynamics* (Wiley, 1999).
- [129] S. Rolston and G. Gabrielse, *Hyperfine Interactions* **44**, 233 (1989).

- [130] L. Spitzer, *Physics of Fully Ionized Gases* (Wiley, 1962).
- [131] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell, *Science* **269**, 198 (1995).
- [132] W. Ketterle and N. V. Druten, *Adv. At. Mol. Opt. Phys.* **37**, 181 (1996).
- [133] G. Gabrielse *et al.*, (ATRAP Collaboration), *Physics Letters B* **507**, 1 (2001).
- [134] J. Fajans and L. Frièdland, *American Journal of Physics* **69**, 1096 (2001).
- [135] M. C. Fujiwara *et al.*, *AIP Conf. Proc.* **793**, 111 (2005).
- [136] M. C. Fujiwara *et al.*, *AIP Conf. Proc.* **1037**, 208 (2008).
- [137] *VAITA Manual*, Ideas ASA, Hovik, Norway.
- [138] S. Korpar, P. Križan, and S. Fratina, *Nucl. Instr. Meth. Phys. Res. A* **511**, 195 (2003).
- [139] J.-P. Martin and P.-A. Amaudruz, *IEEE-NPSS Real Time Conference* **0**, 85 (2005).
- [140] M. Goossens, *GEANT: Detector Description and Simulation Tool, Long Writeup W5013*, CERN Program Library (CERN, Geneva, 1993).
- [141] R. Brun and F. Rademakers, *Nucl. Instr. Meth. Phys. Res. A* **389**, 81 (1997).
- [142] I. Hřivnáčová *et al.*, The virtual monte carlo, in *Conference for Computing in High-Energy and Nuclear Physics*, 2003.
- [143] M. R. Clover, R. M. DeVries, N. J. DiGiacomo, and Y. Yariv, *Phys. Rev. C* **26**, 2138 (1982).
- [144] F. E. James, *Monte Carlo phase space* (CERN, Geneva, 1968).

- [145] H. Spieler, *Semiconductor detector systems* (Oxford Univ. Press, Oxford, 2005).
- [146] A. Strandlie and R. Frühwirth, *Rev. Mod. Phys.* **82**, 1419 (2010).
- [147] W. Kucewicz *et al.*, *Nucl. Instr. Meth. Phys. Res. A* **518**, 411 (2004).
- [148] W. J. Burger, *Nucl. Instr. Meth. Phys. Res. A* **435**, 202 (1999).
- [149] S. Straulino *et al.*, *Nucl. Instr. Meth. Phys. Res. A* **530**, 168 (2004).
- [150] M. C. Fujiwara *et al.*, *Phys. Rev. Lett.* **92**, 65005 (2004).
- [151] M. C. Fujiwara, *AIP Conf. Proc.* **793**, 111 (2005).
- [152] K. Kleinknecht, *Detectors for particle radiation* (Cambridge University Press, 1998).
- [153] R. Frühwirth and M. Regler, *Data analysis techniques for high-energy physics* (Cambridge University Press, 2000).
- [154] A. Strandlie, *Nucl. Instr. Meth. Phys. Res. A* **535**, 57 (2004).
- [155] P. Avery, *Applied Fitting Theory IV: Formulas for Track Fitting* (CBX 92-45, 1992), <http://www.phys.ufl.edu/~avery/fitting/fitting4.pdf>.
- [156] F. James, *MINUIT: Function Minimization and Error Analysis, Long Writeup D506*, CERN Program Library (CERN, Geneva, 1994).
- [157] E. Butler, *Antihydrogen formation, dynamics and trapping*, PhD thesis, Swansea University, 2011.
- [158] J. Fajans, N. Madsen, and F. Robicheaux, *Phys. Plasmas* **15**, 032108 (2008).
- [159] F. Chen, *Introduction to Plasma Physics and Controlled Fusion: Plasma physics* (Plenum Press, 1984).

- [160] P. F. Harrison, Blind analysis, in *Proceedings of the Conference on Advanced Statistical Techniques in Particle Physics*, 2002.
- [161] *OPERA-3D Reference Manual*, Cobham, Vector Fields Ltd, UK.
- [162] M. Ashkezari, Progress toward microwave spectroscopy of trapped antihydrogen, in *Proceedings of the LEAP 2011 Conference*, 2011, In preparation.
- [163] C. L. Cesar *et al.*, (ALPHA Collaboration), *Can. J. Phys.* **87**, 791 (2009).
- [164] R. M. Sternheimer and R. F. Peierls, *Phys. Rev. B* **3**, 3681 (1971).
- [165] R. Turton, *The physics of solids* (Oxford University Press, 2000).

Appendix A

Charged particles passing through matter

The average energy loss per unit distance $\langle dE/dx \rangle$ (also referred to as the stopping power) for moderately relativistic ($0.1 \lesssim \beta\gamma \lesssim 1000$) charged particles passing through the bulk of a material is given by the Bethe formula [9],

$$-\left\langle \frac{dE}{dx} \right\rangle = \left(\frac{e^2}{4\pi\epsilon_0} \right)^2 \frac{4\pi z_c^2}{m_e c^2 \beta^2} \frac{N_{\text{Av}} Z}{A} \left[\ln \left(\frac{2m_e c^2 \beta^2 \gamma^2 T_{\text{max}}}{I_{\text{avg}}^2} \right) - \beta^2 - \frac{\delta(\beta\gamma)}{2} \right], \quad (\text{A.1})$$

where

E	particle energy	z_c	charge number of the particle
x	the distance travelled	e	elementary charge
β	particle velocity ($= v/c$)	m_e	mass of the electron
γ	$1/\sqrt{1-\beta^2}$	N_{Av}	Avogadro's number
c	speed of light	Z	Atomic number
ϵ_0	vacuum permittivity	A	atomic weight
I_{avg}	mean excitation energy	T_{max}	maximum kinetic energy that can be transferred to an electron in a single collision.
$\delta(\beta\gamma)$	density effect correction		

For a particle of mass M , T_{max} can be written as [9],

$$T_{\text{max}} = \frac{2m_e c^2 \beta^2 \gamma^2}{1 + 2\gamma m_e/M + (m_e/M)^2}, \quad (\text{A.2})$$

and $\delta(\beta\gamma)$ represents a correction to the distant-collision contribution due to the polarization of the medium for large particle momenta [164].

A feature of Eq. A.1 is that the stopping power function has a global minimum (typically for particles in the several GeV momentum range). Particles with momenta near the stopping power minimum are referred to as Minimum Ionizing Particles (MIPs).

Charged particle passing through matter will also undergo many small-angle deflections, which is usually referred to as ‘multiple scattering’. The mean multiple scattering

angle, θ_0 , is given by the generally applied Molière formula [9]:

$$\theta_0 = \frac{13.6 \text{ MeV}}{\beta c P} z_c \sqrt{x/X_0} [1 + 0.038 \log(x/X_0)], \quad (\text{A.3})$$

where βc , P , and z_c are the particle velocity, momentum, and charge number, and x/X_0 is the thickness of the scattering material in units of radiation length. The radiation length, X_0 , is the mean distance over which, on average, a high-energy electron loses, due to bremsstrahlung, all but $1/e$ of its energy. The small-angle scattering is roughly Gaussian, but for at larger angles (more than a few θ_0) the distribution follows closer to Rutherford scattering, which has much larger tails than the Gaussian distribution [9].

Appendix B

Particle detection using silicon p-n junctions

Silicon, germanium, and diamond can all be used as particle detectors [145]. However, since silicon retains its semiconducting properties at room temperature and is abundant, it is a popular choice for particle detection devices. The general principle of silicon particle detectors is as follows:

1. An electric field depletes a region of the silicon bulk of charge carriers.
2. Energetic charged particles passing through this region liberate electron-hole pairs.
3. The electric field causes the electrons and holes to drift towards collection electrodes, where the liberated charge is collected and read-out.

A sample of intrinsic (pure) silicon has a full valence band and an empty conduction band. That is, the silicon atoms are arranged in a diamond lattice with four valence electron sites that are all occupied. Because of the empty conduction band, electrons that are promoted across the band gap acts as mobile charge – and likewise, the unoccupied valence state (or ‘hole’) can also act as a mobile charge carrier. Electron-hole pairs can be generated by passing charged particles imparting energy according to Eqn. A.1 (the energy needed to generate an electron-hole pair in silicon is about 3.6 eV at room temperature, which is about three times larger than the band gap). The generation of electron-hole pairs forms the basis of the signal generating mechanism for silicon semiconductor detectors.

The electrons in silicon obey Fermi-Dirac statistics. As such, $f(E)$ gives the probability for an electron to occupy an energy state E ,

$$f(E) = \frac{1}{e^{(E-E_f)/k_B T} + 1}, \quad (\text{B.1})$$


where E_f is the Fermi level. However, a non-negligible number of electrons will occupy the conduction band due to thermal broadening. This can be shown by assuming that the Fermi level is located in the middle of the band gap ($E_f = E_g/2$, where $E_g = 1.1$ eV is the silicon band gap), and with $e^{(E-E_f)/k_B T} \gg 1$, the occupancy function can be approximated as $f(E) = e^{-(E-E_g/2)/k_B T}$. Likewise, the density of states is given as $g(E)dE = \frac{1}{2\pi^2} \left(\frac{2m_e}{\hbar^2}\right)^{3/2} \sqrt{E - E_g} dE$, so the number of electrons in the conduction band is then given as [165],

$$N_e = \int_{E_g}^{\infty} f(E)g(E)dE = \int_{E_g}^{\infty} \frac{1}{2\pi^2} \left(\frac{2m_e}{\hbar}\right)^{3/2} \sqrt{E - E_g} e^{-(E-E_g/2)/k_B T} dE, \quad (\text{B.2})$$


$$= 2 \left(\frac{m_e k_B T}{2\pi\hbar^2}\right)^{3/2} e^{-E_g/2k_B T}, \quad (\text{B.3})$$

and with $E_g = 1.1$ eV and $T = 300$ K, the number of electrons in the conduction band is $N \sim 10^{10} \text{ cm}^{-3}$, which is non-negligible.

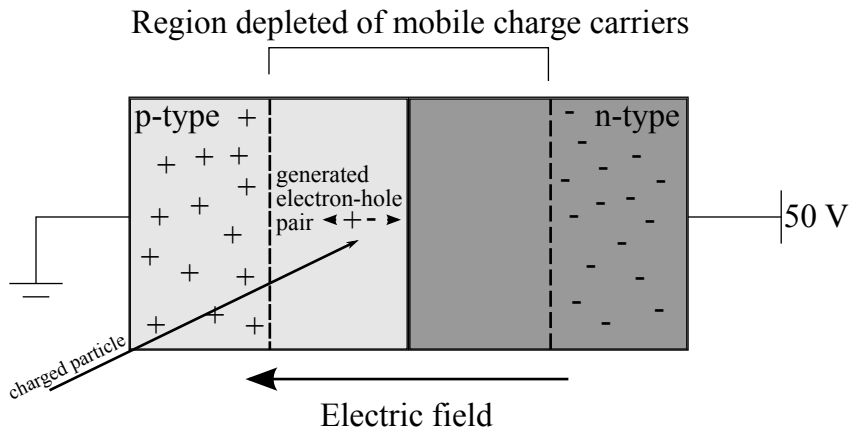
The semiconductor properties can be modified by doping the silicon with donor atoms (such as phosphorus, arsenic, or antimony) which have five weakly bound valence electrons, or acceptor atoms (such as boron, aluminum, gallium, or indium) have three valence electrons. Both types of dopants modify the band gap structure. For example, the introduction of donor atoms (known as n-type doping) to the silicon lattice adds a weakly bound valence electron (the Coulomb force between the electron and donor atom is reduced by the dielectric constant of the medium [145]) in an energy state just below the conduction band. Similarly, including acceptor atoms (known as p-type doping) leaves a silicon bond unpaired. The acceptor atom can then accept a neighboring electron to complete the bond, introducing a hole state just above the valence band.



(a) Unbiased p-n junction before thermal diffusion of mobile charge carriers.



(b) Unbiased p-n junction after thermal diffusion of mobile charged carriers.



(c) Reverse-biased p-n junction. A charged particle is shown passing through the depletion region and generating an electron-hole pair.

Figure B.1: Diagram of a p-n junction: (a) before thermal diffusion of mobile charge carriers (unbiased); (b) after thermal diffusion, with a net positive charge on the n-side and a net negative charge on the p-side, and a depletion region in between; (c) in the case of a reverse-biased junction with an enhanced depletion region and a passing charged particle generating an electron-hole pair.

By adjoining p-type and n-type doped semiconductors, a ‘p-n junction’ is created. In this configuration, a region with an abundance of valence electrons is adjacent to a region with an abundance of hole states (Fig. B.1(a)). Thermal diffusion will drive the electrons across the junction to the acceptor states, leaving a net positive charge in the n-type region; similarly, holes will diffuse into the n-type region, leaving a net negative charge in the p-type region. This separation of charge results in an electric field across the p-n junction, leaving the region free of charge carriers (the ‘depletion’ region), and resisting further diffusion of electrons and holes (Fig. B.1(b)). The p-n junction then functions as a diode: when a positive bias is applied to the p-type side and a negative bias to the n-side (‘forward bias’), the potential barrier across the junction is reduced and electrons and holes can flow freely; conversely, if a negative bias is applied to the p-side and positive bias to the n-side (‘reverse bias’), the potential barrier is increased and the depletion region widened (Fig. B.1(c)). The total depletion width, w , is given as (assuming that potential due to charge diffusion is small) [110],

$$w = \sqrt{\frac{2\epsilon V_b}{e} \frac{N_a + N_d}{N_a N_d}}, \quad (\text{B.4})$$

where V_b is the reverse bias applied, ϵ is the dielectric constant of the medium (11.9 for silicon), and N_a and N_d are the acceptor and donor concentrations, respectively. For example, with doping concentrations of $N_a = N_d = 10^{13} \text{ cm}^{-3}$ and an applied reverse bias of $V_b = 65 \text{ V}$, the depletion width is $w \simeq 300 \text{ }\mu\text{m}$.

Silicon particle detectors are commonly constructed with many reverse-biased p-n junctions. By fully depleting the silicon wafer through reverse-biasing of the junction, the bulk of the silicon will be empty of charge carriers (apart from a leakage current from a small number of electrons tunneling through the potential barrier). In this situation, energetic charged particles passing through the silicon bulk will generate electron-hole pairs (about 80 electron-hole pairs are generated per μm during the passage of a minimum

ionizing particle). The electrons and holes drift in opposite directions in the bias field, with drift velocities set by their respective mobilities ($1350 \text{ cm}^2/(\text{V}\cdot\text{s})$ for electrons and $450 \text{ cm}^2/(\text{V}\cdot\text{s})$ for holes in silicon). Then for an example silicon wafer of $300 \mu\text{m}$ thickness with a 65 V applied bias, it takes electrons $\sim 10 \text{ ns}$ and holes $\sim 30 \text{ ns}$ to fully cross the wafer. The induced charge can then be collected and converted into an electric signal for readout.

Appendix C

Particle propagation in GEANT3

The GEANT3 routines handle the propagation of the particles through the apparatus and detector geometry. Every particle is followed until a) it decays into daughter particles, b) it falls below the energy cut-off threshold, c) it leaves the simulation volume, or d) a maximum number of steps is exceeded. The propagation of each particle is summarized as follows [140]:

1. The particle to be simulated is taken from the ‘particle stack’, which is a Last In, First Out (LIFO) buffer, such that the last particle added is the next particle to be simulated. To set the distance to be propagated before undergoing some process, the number of interaction lengths (mean free paths), N_λ , for each particle process (photoelectric effect, bremsstrahlung, decay in flight, etc.) is sampled from,

$$N_\lambda = -\log(\eta), \quad (\text{C.1})$$

where η is an uniformly distributed random number between zero and one.

2. The distance to the interaction point is calculated for each process. This involves the evaluation of the mean free path, λ_{process} for the various particle processes

$$\lambda_{\text{process}} = \frac{A}{N_{\text{Av}}\rho\sigma(E, Z, A)_{\text{process}}} \quad (\text{C.2})$$

where A is the atomic weight, N_{Av} is Avogadro’s number, ρ is the material density, and σ_{process} is the total cross-section of the process, dependent on the particle energy, E , atomic number, Z , and atomic weight. The distance to the interaction point is then given as $s = N_\lambda\lambda_{\text{process}}$.

The step size is chosen to be the minimum between:

- (a) the distances to the various interaction points,
 - (b) the distance to the next geometric volume (a particle will never cross into a new material over the course of a step),
 - (c) a user-define maximum step limit,
 - (d) the distances that define the maximum fraction of energy loss, maximum angular deviation, and maximum and minimum step sizes for the present volume material (specified by the tracking material, see Sec. 4.2.2),
 - (e) and the predefined energy and time cuts (the step size should not propagate the particle past these limits).
3. The particle is transported according to the chosen step size (in a helical path if in a magnetic field, in a straight line, otherwise).
 4. The particle's energy is updated according to energy loss tables for the particle in the specific material. (Note that GEANT3 does not calculate the electron-hole pair generated in the reverse-biased silicon modules (Sec. B), rather, the signal digitization (Sec. 4.2.4) takes the total energy deposited in the geometry volume as input.)
 5. If the step involves undergoing a process, the final state of that process is sampled (adding new particles to the stack, if necessary). If the particle survives the process, resample the number of interaction lengths for that same process.
 6. Update the number of interaction lengths remaining for all processes according to the formula:

$$N'_\lambda = N_\lambda - \frac{\Delta x}{\lambda}, \quad (\text{C.3})$$

where Δx is the step size.

7. Return to 2), unless the particle falls below the energy threshold, progresses beyond the time cut, leaves the simulation volume, or disappears because of an interaction.

This prescription is followed for all particles added to the stack.

Appendix D

Simulated particle momenta through the silicon detector

Several heuristics used in the event reconstruction algorithms (Sec. 5.5) rely on the assumption that the transverse momenta of charged pions as they pass through the silicon detector are greater than that of other charged particles that may create tracks (most importantly, pair-produced positrons and electrons). This assumption is used to justify the decision to accept ‘straighter’ (large transverse momentum) tracks over tracks with smaller radii of curvature.

To investigate this assumption, 10 000 annihilation events were simulated. The annihilation locations were generated with a Gaussian distribution along the axial coordinate (with the distribution mean in the axial center of the detector, and width of 2.5 cm), uniformly in the azimuthal coordinate at the electrode radius ($R = 2.2275$ cm). The GEANT3 particle propagator was used to track and record the momentum of each particle. Only particles that passed through all three detector layers were considered, as this is also a condition for track acceptance used by the reconstruction routines.

Figure D.1 gives the generating processes for all of the particle types. It is interesting to note that the vast majority of charged pions come from ‘Primary particle emission’, which simply means that they are the initially generated particles (Sec. 4.2.3). Similarly, the majority of electrons and positrons come from pair-production, presumably through the $\pi^0 \rightarrow \gamma \rightarrow e^-e^+$ channel.

Table D compiles the number of number of particles of each type that passes through all three detector layers, along with their most probable momenta (taken from the momenta distributions, Figs. D.2-D.8). As can be seen, the vast majority of tracks come from π^\pm and e^\pm particles, and the charged pions clearly have larger transverse momenta




Figure D.1: Generation proces vs. particle type illustration for the simulated processes that resulted in tracks through all three layers of the detector for 10 000 simulated annihilations.

Particle type	Number of tracks	Most probable momentum	
		Transverse (MeV/c)	Axial (MeV/c)
π^-	6399	200	100
π^+	6486	200	100
e^-	4228	50	< 20
e^+	4167	50	< 20
μ^-	44	~ 150	~ 70
μ^+	42	~ 150	~ 70
p	163	400	150

Table D.1: The number of tracks through the silicon detector, probable momenta (from Figs. D.2-D.8) for the various particle types.




Figure D.2: The (a) transverse momentum, and (b) absolute axial momentum for simulated π^- particles that passed through all three detector layers.

(although protons generated through hadronic processes appear to generally have higher transverse momenta, they contribute very few tracks, and can be ignored). This simulation validates the heuristic decision that, when given the choice between two tracks, the tracks with higher transverse momentum should be accepted. Large radius of curvature tracks are more likely to be the result of the passage of a charged pion, which, the majority of the time, were produced by the antiproton annihilation (and not by a secondary process).




Figure D.3: The (a) transverse momentum, and (b) absolute axial momentum for simulated π^+ particles that passed through all three detector layers.

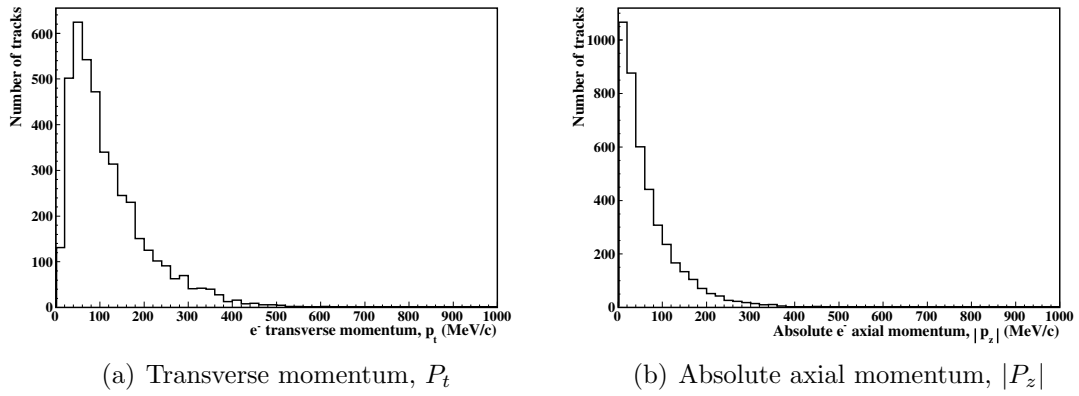


Figure D.4: The (a) transverse momentum, and (b) absolute axial momentum for simulated e^- particles that passed through all three detector layers.

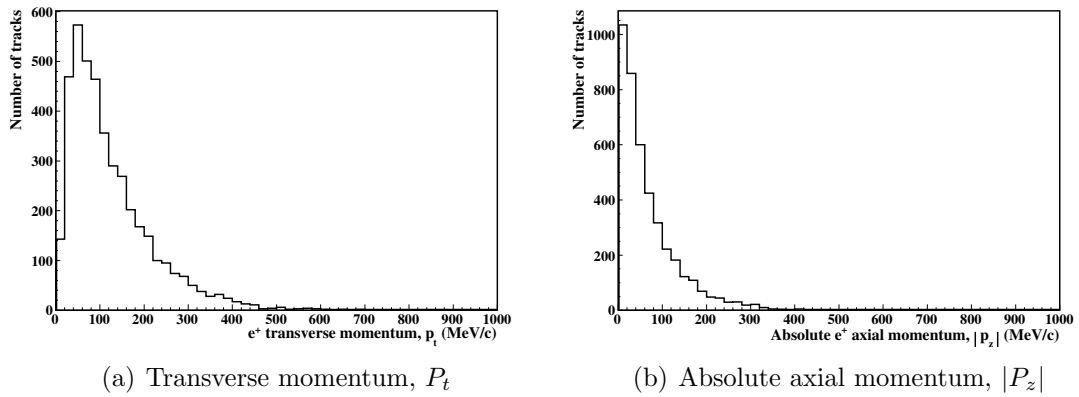


Figure D.5: The (a) transverse momentum, and (b) absolute axial momentum for simulated e^+ particles that passed through all three detector layers.




Figure D.6: The (a) transverse momentum, and (b) absolute axial momentum for simulated μ^- particles that passed through all three detector layers.

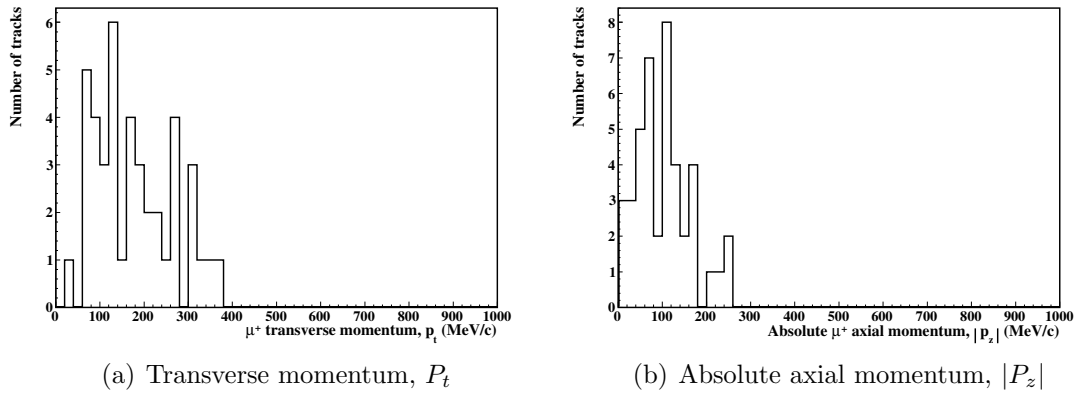


Figure D.7: The (a) transverse momentum, and (b) absolute axial momentum for simulated μ^+ particles that passed through all three detector layers.

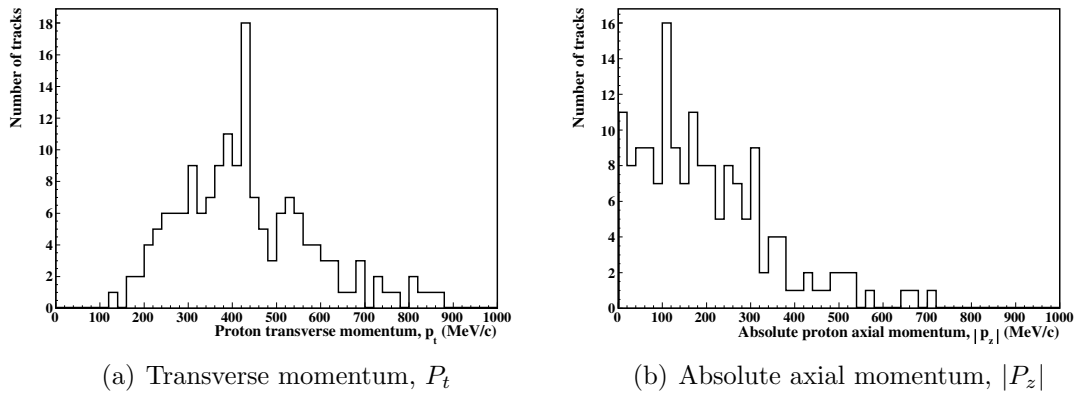


Figure D.8: The (a) transverse momentum, and (b) absolute axial momentum for simulated protons that passed through all three detector layers.