

---

From the efficient trapping and  
accumulation of cooled antiprotons  
towards the formation of a pulsed beam  
of antihydrogen atoms

---

Dissertation

zur Erlangung des Doktorgrades

an der Fakultät für Mathematik, Informatik und  
Naturwissenschaften

Fachbereich Physik  
der Universität Hamburg

vorgelegt von

**Saiva Helga Rosalie Huck**

Hamburg

2025



Universität Hamburg

DER FORSCHUNG | DER LEHRE | DER BILDUNG



---

Betreuer/innen der Dissertation:	Prof. Dr. Erika Garutti Dr. Ruggero Caravita Dr. Michael Doser
Gutachter/innen der Dissertation:	Prof. Dr. Erika Garutti Dr. Ruggero Caravita
Zusammensetzung der Prüfungskommission:	Prof. Dr. Daniela Pfannkuche Prof. Dr. Erika Garutti Dr. Ruggero Caravita Prof. Dr. Gudrid Moortgat-Pick Prof. Dr. Christian Bressler
Vorsitzende der Prüfungskommission:	Prof. Dr. Daniela Pfannkuche
Datum der Disputation:	08.04.2025
Vorsitzender Fach-Promotionsausschuss PHYSIK:	Prof. Dr. Wolfgang J. Parak
Leiter des Fachbereichs PHYSIK:	Prof. Dr. Markus Drescher
Dekan der Fakultät MIN:	Prof. Dr.-Ing. Norbert Ritter



---

## Abstract

The work presented here comprises a number of significant contributions to the formation of a pulsed beam of antihydrogen atoms in the AEGIS experiment at CERN's Antiproton Decelerator facility. The main objective of AEGIS is a precise measurement of the gravitational acceleration of antimatter to probe the Weak Equivalence Principle and gain insight into possible explanations of our universe's matter-antimatter asymmetry. For this purpose, antihydrogen atoms are produced in a pulsed scheme through a charge exchange reaction of cold, trapped antiprotons and laser-excited positronium atoms. The anti-atoms are to be formed into a horizontal beam that passes through a deflectometer, enabling a determination of their vertical deflection due to the influence of gravity.

At the foundation of this work lies the development and implementation of a control system, CIRCUS, that has reliably run all experiments in AEGIS since 2021. CIRCUS has the capability of autonomous operation and comprises a core formed by the AERIALIST fast control unit that synchronizes experimental processes with nanosecond precision. It has been fundamental to any recent achievements of AEGIS, most prominently the first-ever successful laser cooling of positronium.

This work further includes the commissioning of the new electron system of AEGIS, which is a vital component of the experiment, both in the preparatory phase and for the implementation of the antiproton plasma routines.

A main result achieved in this work is the unprecedented accumulation of several hundred million cold antiprotons in an electromagnetic trap. It has become feasible thanks to the development of routines for the capture of a record 70% of the antiprotons available from CERN's ELENA decelerator as well as efficient electron cooling and plasma compression techniques and the stacking of multiple antiproton bunches inside the trap. Antiprotons being a core component of most bound antimatter systems, this achievement opens the door to a variety of antimatter research, including a strongly increased antihydrogen production in AEGIS and the formation and study of antiprotonic atoms.

The formation of the antihydrogen atoms into a horizontal beam, crucial to the gravity measurement, has been realized as part of this work through the forward-acceleration of the antiprotons towards the formed positronium via a parabolic potential, precisely synchronized to the positronium excitation. The functionality of the procedure has been successfully verified by an analysis of the observed signal on a scintillator/PMT detector, which has been implemented as part of this work as well. The same analysis also represents a first investigation of antihydrogen formation with AEGIS Phase II, following major upgrades to the apparatus, some of which are part of this work, and the transformation of the procedures to a collinear antihydrogen production. While evidence of an excess signal is observed in those runs expected to produce antihydrogen, compared to control runs without positronium excitation, in the relevant time window, the produced numbers are too low to be unambiguously significant. This impediment can be mainly attributed to the underperformance of the positronium line, which is currently being improved.

Thanks to the collectivity of the performed upgrades, many of which this work has contributed to, an antihydrogen formation efficiency boosted by several orders of magnitude is expected with respect to previous production runs. With the corresponding statistics and the beam formation procedure furthermore well in place, a determination of the gravitational acceleration of antimatter with a relative precision of 1% has become realistically achievable within a few months of antiproton beam time.

---

## Zusammenfassung

Die hier präsentierte Arbeit umfasst eine Anzahl maßgeblicher Beiträge zur Entwicklung eines gepulsten Strahls von Antiwasserstoff-Atomen im AEGIS Experiment in der Antiproton Decelerator-Anlage am CERN. Das Hauptziel von AEGIS ist eine präzise Messung der Gravitationsbeschleunigung von Antimaterie, um das Schwache Äquivalenzprinzip zu prüfen und Einblicke in mögliche Erklärungen für die Materie-Antimaterie-Asymmetrie unseres Universums zu erhalten. Zu diesem Zweck werden Antiwasserstoff-Atome in gepulster Form durch eine Ladungsaustausch-Reaktion zwischen kalten, gefangenen Antiprotonen und Laser-angeregten Positronium-Atomen hergestellt. Die Anti-Atome sollen zu einem horizontalen Strahl geformt werden, der ein Deflektometer durchquert, um eine Ermittlung ihrer vertikalen Ablenkung aufgrund der Gravitation zu ermöglichen.

Die Grundlage dieser Arbeit bildet die Ausarbeitung und Einführung eines Kontrollsystems, CIRCUS, das seit 2021 zuverlässig alle AEGIS Experimente steuert. CIRCUS ermöglicht einen autonomen Betrieb und beinhaltet als ein Kernstück die AERIALIST Schnell-Kontroll-Einheit, die experimentelle Prozesse mit Nanosekunden-Genauigkeit synchronisiert. Es ist fundamental für alle kürzlichen Errungenschaften von AEGIS, wobei herausragend die erste erfolgreiche Laser-Kühlung von Positronium zu nennen ist.

Diese Arbeit schließt außerdem die Inbetriebnahme des neuen Elektronen-Systems von AEGIS ein, das ein essenzieller Bestandteil des Experiments ist, sowohl in der Vorbereitungsphase als auch für die Implementierung der Antiprotonen-Plasma-Routinen.

Ein Hauptergebnis dieser Arbeit ist die beispiellose Akkumulation von mehreren hundert Millionen kalten Antiprotonen in einer elektromagnetischen Falle. Dies wird durch die Entwicklung von Routinen für das Einfangen von rekordbrechenden 70 % der vom ELENA Entschleuniger am CERN verfügbaren Antiprotonen sowie für effiziente Elektronenkühlungs- und Plasma-Kompressionstechniken und für das Ansammeln von mehreren Antiprotonen-Paketen in der Falle ermöglicht. Da Antiprotonen ein Kernbestandteil der meisten gebundenen Antimaterie-Systeme sind, öffnet dieser Erfolg die Tür zu einer Vielfalt an Antimaterie-Forschung, einschließlich einer stark erhöhten Antiwasserstoff-Produktion in AEGIS und der Herstellung und Untersuchung von antiprotonischen Atomen.

Die Formung der Antiwasserstoff-Atome zu einem horizontalen Strahl, entscheidend für die Gravitationsmessung, ist in dieser Arbeit durch die Beschleunigung der Antiprotonen in die Richtung des Positroniums mittels eines parabolischen Potentials umgesetzt, präzise synchronisiert mit der Positronium-Anregung. Die Funktionalität der Methode wurde erfolgreich durch eine ebenfalls im Zuge dieser Arbeit implementierte Analyse des Signals auf einem Szintillator/PMT Detektor verifiziert. Dieselbe Analyse stellt auch eine erste Untersuchung der Antiwasserstoff-Herstellung mit AEGIS Phase II dar, die auf die Transformation der Abläufe zu einer kollinearen Antiwasserstoff-Produktion und wesentliche Verbesserungen des Apparats folgt, von denen einige Teil dieser Arbeit sind. Während Evidenz eines Überschuss-Signals im relevanten Zeitfenster für solche Daten zu beobachten ist, in denen Antiwasserstoff-Herstellung erwartet wird, im Vergleich zu Kontroll-Daten ohne Positronium-Anregung, sind die erreichten Anzahlen zu gering um eindeutig signifikant zu sein. Diese Beeinträchtigung kann hauptsächlich auf die Leistungsschwäche der Positronium-Komponenten zurückgeführt werden, die aktuell verbessert werden.

Dank der Gesamtheit der durchgeführten Upgrades, zu vielen von denen diese Arbeit beigetragen hat, wird eine im Vergleich zu vorherigen Durchführungen um mehrere Größenordnungen gesteigerte Effizienz der Antiwasserstoff-Herstellung erwartet. Mit der resultierenden Statistik und der darüber hinausgehenden Prozedur für die Strahl-Formung wird eine Bestimmung der Fallbeschleunigung von Antimaterie mit einer relativen Genauigkeit von 1 % realistisch innerhalb einiger Monate der Antiprotonen-Strahlzeit erreichbar.

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Antimatter and the gravitational acceleration</b>	<b>3</b>
2.1	Fundamentals of antimatter physics . . . . .	3
2.1.1	The Standard Model of particle physics . . . . .	3
2.1.2	Antiproton properties . . . . .	7
2.1.3	CPT and the matter-antimatter asymmetry . . . . .	7
2.2	The Weak Equivalence Principle - Universality of Free Fall . . . . .	9
2.2.1	Testing the WEP . . . . .	10
2.2.2	WEP tests with antimatter . . . . .	10
2.3	Neutral antimatter systems . . . . .	11
2.3.1	Positronium . . . . .	11
2.3.2	Antihydrogen . . . . .	12
2.3.3	Protonium . . . . .	14
2.3.4	Antiprotonic helium . . . . .	14
<b>3</b>	<b>Introducing AEGIS</b>	<b>15</b>
3.1	The Antiproton Decelerator at CERN . . . . .	16
3.1.1	CERN's accelerator complex . . . . .	16
3.1.2	Antiproton production . . . . .	17
3.1.3	Antiproton deceleration . . . . .	18
3.1.4	ELENA . . . . .	19
3.2	Trapping antiprotons: Penning traps . . . . .	20
3.2.1	Thermal equilibrium in Penning-Malmberg traps . . . . .	23
3.2.2	The Rotating Wall technique . . . . .	23
3.3	AEGIS: A procedure for pulsed production of antihydrogen . . . . .	23
3.3.1	Antiproton manipulation . . . . .	24
3.3.2	Positronium formation and Rydberg excitation . . . . .	25
3.3.3	The number of produced antihydrogen atoms . . . . .	25
3.3.4	Collinear antihydrogen formation: AEGIS Phase II . . . . .	26
3.4	AEGIS setup . . . . .	27
3.4.1	Magnet setup and coordinate system . . . . .	28
3.4.2	Vacuum system . . . . .	29
3.4.3	Cryogenic system . . . . .	29
3.4.4	Positron system . . . . .	30
3.4.5	Positronium converter target . . . . .	30
3.4.6	Laser system . . . . .	31
3.4.7	Degraders . . . . .	33
3.4.8	Electromagnetic traps . . . . .	35
3.4.9	Detectors . . . . .	39
3.5	Antihydrogen gravity measurement . . . . .	44

3.5.1	The gravity module . . . . .	44
3.5.2	Sensitivity expectations . . . . .	47
3.5.3	Detector options . . . . .	49
<b>4</b>	<b>The CIRCUS control system</b>	<b>51</b>
4.1	Control system hardware . . . . .	52
4.1.1	Sinara hardware overview . . . . .	52
4.1.2	Using the external clock . . . . .	56
4.1.3	The AEGIS Pulsar . . . . .	58
4.1.4	Rotating Wall electronics . . . . .	58
4.2	Hardware communication . . . . .	59
4.2.1	ARTIQ . . . . .	59
4.2.2	Kasli Direct Memory Access . . . . .	60
4.2.3	Gateware . . . . .	61
4.2.4	The AEGIS library system . . . . .	61
4.2.5	Efficient programming . . . . .	67
4.3	Calibration of the amplifiers . . . . .	68
4.3.1	Calibration setup and procedure . . . . .	69
4.3.2	Calibration results . . . . .	71
4.4	The AERIALIST and TALOS: CIRCUS . . . . .	74
4.4.1	TALOS overview . . . . .	76
4.4.2	Integration of the AERIALIST and TALOS . . . . .	77
4.4.3	Error handling . . . . .	78
4.4.4	Multi-Kasli modes . . . . .	79
4.5	Timing requirements . . . . .	80
4.5.1	Master/satellite operation . . . . .	81
4.5.2	Kasli master synchronization . . . . .	82
4.6	Results and reliable operation . . . . .	83
4.7	Summary . . . . .	84
<b>5</b>	<b>Electrons in AEGIS</b>	<b>87</b>
5.1	The AEGIS electron gun . . . . .	87
5.1.1	Using the electron gun . . . . .	88
5.1.2	Re-conditioning of the electron gun . . . . .	91
5.2	Tests of the new antihydrogen production trap . . . . .	91
5.3	Experimental alignment procedures using electrons . . . . .	93
5.4	Detector calibration using electrons . . . . .	96
5.4.1	CMOS calibration . . . . .	96
5.4.2	MCP as an amplified Faraday Cup: an initial exploration . . . . .	96
5.5	Summary . . . . .	104
<b>6</b>	<b>Efficient capture and accumulation of cold antiprotons in AEGIS</b>	<b>105</b>
6.1	Efficient capture of antiprotons . . . . .	106
6.1.1	Tail-less electron cooling plasma preparation . . . . .	107
6.1.2	Beam steering . . . . .	110
6.1.3	Optimization of the trap closing time . . . . .	112
6.1.4	Degrader studies . . . . .	113
6.1.5	The antiproton capture efficiency . . . . .	117
6.2	Efficient antiproton cooling . . . . .	118
6.2.1	Cooling efficiency . . . . .	120

---

6.2.2	Rotating Wall compression . . . . .	124
6.3	Accumulation of cold antiprotons in the trap . . . . .	128
6.4	Axial energy components of antiprotons in the trap . . . . .	135
6.5	Applications for large numbers of cold antiprotons . . . . .	139
6.6	Summary . . . . .	142
<b>7</b>	<b>Antihydrogen developments in AEGIS</b>	<b>143</b>
7.1	Ballistic transfer of antiprotons . . . . .	145
7.1.1	Antiproton launch trap . . . . .	145
7.1.2	The parabolic potential . . . . .	147
7.1.3	Electron removal and recycling . . . . .	150
7.1.4	Antiproton launch and recycling . . . . .	151
7.2	Antihydrogen production . . . . .	154
7.2.1	Electric potentials at the antiproton launch . . . . .	155
7.2.2	Timing considerations and the target electrode potential . . . . .	156
7.3	ESDA analysis . . . . .	160
7.3.1	Peak tagging . . . . .	161
7.3.2	Event discrimination and selection . . . . .	163
7.3.3	Antiproton oscillations . . . . .	167
7.3.4	The antihydrogen signal . . . . .	171
7.4	Expected $\bar{\text{H}}$ production improvement . . . . .	173
7.4.1	The Ps Rydberg level reach . . . . .	173
7.4.2	Summary of the expected $\bar{\text{H}}$ number increase . . . . .	175
7.4.3	Considerations of the $\bar{\text{H}}$ axial velocity . . . . .	177
7.5	Laser cooling of positronium . . . . .	178
7.5.1	Setup for broadband laser cooling of positronium . . . . .	178
7.5.2	Experimental results of positronium laser cooling . . . . .	181
7.6	Summary . . . . .	184
<b>8</b>	<b>Discussion and outlook</b>	<b>187</b>
	<b>Appendix A List of acronyms</b>	<b>207</b>
	<b>Appendix B Examples of positron-producing processes</b>	<b>210</b>
	<b>Appendix C The minimum proton energy for proton pair production</b>	<b>211</b>
	<b>Appendix D The <math>\mathbf{E} \times \mathbf{B}</math> drift in Penning-Malmberg traps</b>	<b>213</b>
	<b>Appendix E List of the electrodes of the AEGIS trap system</b>	<b>214</b>
	<b>Appendix F Setup of the Sinara Kasli</b>	<b>216</b>
	<b>Appendix G Issues of the custom amplifier boards</b>	<b>218</b>
	<b>Appendix H Calibration results of the amplifier channels</b>	<b>221</b>
	<b>Appendix I CMOS images of electron and antiproton plasmas</b>	<b>223</b>
	<b>Appendix J Potential reshapings between two <math>\bar{\text{H}}</math> production cycles</b>	<b>226</b>
	<b>Appendix K Scan over the target electrode potential</b>	<b>228</b>



# Chapter 1

## Introduction

Modern physics, thanks to a combination of strong, ever-improving theoretical foundations and sophisticated experiments operating at the technological forefront, has gained a deep understanding of the origin, composition and evolution of our universe, much of which is expressed in the Standard Model of particle physics. However, some riddles remain. Prominently, the mystery of the impressive matter-antimatter asymmetry, which is consistently observed, is yet to be solved. Various techniques can be employed to investigate this discrepancy and search for differences in the behaviour of the two.

One approach is the direct production of antimatter systems in laboratory conditions and at low energies to investigate its properties, compare to regular matter and draw conclusions on the origin of the matter-dominated universe. This is what is uniquely done at CERN's Antiproton Decelerator facility, also referred to as Antimatter Factory.

The work presented here has been conducted as part of the AEGIS collaboration, which is based at the Antiproton Decelerator complex and whose main objective is the formation of a pulsed, horizontal beam of antihydrogen atoms. The influence of gravity on these anti-atoms is to be probed by letting them pass through a deflectometer to end up on a position-sensitive detector, determining their vertical deflection due to the gravitational acceleration. Insights into the gravitational interaction of antimatter could allow to draw conclusions on the validity of the Weak Equivalence Principle, a pillar of General Relativity whose violation could account for a difference in the present amounts of matter and antimatter in the universe.

Experiments mandated with the production of antimatter typically utilize procedures from a wide range of research areas and become very complex. Such experiments can only function efficiently if they are steered by a powerful control unit that can coordinate multiple independent subsystems and provide precise timing synchronization to the micro- or nanosecond, while ensuring a stable operation over several months. This is the **first main task** dealt with in this work: the development and implementation of a new, streamlined and reliable control system for the experiments performed in AEGIS (and, thanks to its versatility, in a variety of experimental setups). The outcome is CIRCUS, which has chaperoned all measurements of the collaboration over the past three years and has proven to fulfill all of the criteria outlined above, enabling particularly precise timing and proper control of the experimental parameters thanks to the AERIALIST framework of Sinara hardware and a software library based on ARTIQ software. The birth of CIRCUS forms part of a range of radical upgrades the AEGIS apparatus has undergone between 2020 and 2023 aimed at improving antihydrogen production and making the experiment more versatile, which have in part been topics of this work as well and are presented here. It has furthermore enabled the development of a procedure to perform

laser cooling of positronium atoms for the first time.

Antiprotons lie at the core of any experimental endeavours to produce bound antimatter systems, as they constitute a crucial building block. On the core line of AEGIS, they are brought together with laser-excited positronium atoms in an electromagnetic Penning-Malmberg trap to undergo a charge exchange reaction, combining with the positron to form antihydrogen in a pulsed manner. To perform a statistically significant, precise gravity measurement, the production of antihydrogen in sufficient amounts is essential, requiring a source of large numbers of readily available, cold antiprotons. The realization of such a source in AEGIS is the **second main task** reported on here. Relying on a procedure implemented in the AERIALIST and CIRCUS for the extremely efficient routine capture of 70% or more of the antiprotons from CERN's ELENA decelerator, sympathetic cooling on electrons from a well understood source, and the refined accumulation of the antiprotons between the electrodes of the trap system, several hundred million cold antiprotons can be confined in the tens of cubic centimeters small trap volume. These advances pave the way for a significantly improved antihydrogen production.

Finally, as the **third main task** of this work, a technique has been designed to form the produced antihydrogen atoms into a horizontal beam that travels towards the future gravity module. For this purpose, antiprotons are, prior to being combined with positronium, accelerated towards the antihydrogen production region on a trajectory following a parabolic potential, arriving at a time precisely synchronized to the formation and excitation of the positronium. In this way, a tuning of the resulting axial antihydrogen velocity is enabled. The implementation of these procedures in the AERIALIST/CIRCUS are reported on here, as is a newly implemented analysis of the first antihydrogen production data obtained with the new AEGIS setup in 2023.

Following the brief introduction in this Ch. 1, the subsequent Ch. 2 gives an overview of the relevant underlying physics and the motivation for performing antimatter experiments, with a particular focus on probing the influence of gravity on neutral antimatter systems. Ch. 3 introduces the AEGIS experiment, including its setup, procedures involved in the formation of antihydrogen, and the future gravity measurement. In Ch. 4, the newly developed CIRCUS control system is described in some detail, highlighting its hardware and "fast control" parts as well as the overall integration and importance for the techniques and developments reported on here. Ch. 5 focuses on the use of electrons in the experiment, particularly the incorporation of the upgraded electron gun in the control system and newly developed procedures for hardware tests and calibration techniques. Included in Ch. 6 are all developments involved in the efficient capture, cooling and accumulation of antiprotons in one of the electromagnetic traps of AEGIS, based on the capabilities of the control system. The steps taken for a record accumulation of cold antiprotons are outlined and applications for such a  $\bar{p}$  source are discussed. Ch. 7 describes all developments that have been achieved for the formation of a forward-boostered beam of antihydrogen, including in particular the antiproton launch along a parabolic potential and the integration of all procedures in the new apparatus and control system. The analysis of the  $\bar{H}$  signal of the 2023 run as well as a general discussion of the expected improved production rate is also presented. Furthermore, the successful implementation of positronium laser cooling is outlined. Ch. 4 through Ch. 7 all include individual brief chapter summaries. Finally, Ch. 8 offers a concluding discussion and an outlook to future developments and research.

# Chapter 2

## Antimatter and the gravitational acceleration

Investigations of the properties of antimatter can probe fundamental assumptions regarding the formation and composition of the universe, many of which are manifested in the Standard Model of particle physics. Any observations of basic symmetry violations, in particular CPT symmetry, or a difference in the influence of the gravitational acceleration on matter and antimatter can open a window to as of yet unknown physics.

Since antimatter does not naturally occur in significant quantities, it is created in laboratories through intricate experiments and then studied in detail.

This chapter gives an overview over the established Standard Model as well as some of its shortcomings, before describing the properties of known antiparticles, production mechanisms of neutral, bound antimatter systems, and their role in studies of fundamental symmetries.

### 2.1 Fundamentals of antimatter physics

The existence of antiparticles and antimatter is a well established concept that has been experimentally verified and is widely accepted. However, since antimatter is not readily available in nature and annihilates quickly once produced, many of its properties remain yet to be investigated in detail, and the results could yield insight into very foundational questions of the physical universe.

#### 2.1.1 The Standard Model of particle physics

The most successful relativistic quantum gauge theory to date describing the properties of the elementary particles forming all known matter (and antimatter) as well as three of the four established elementary forces acting between them is called the Standard Model (SM) of particle physics [1].

#### Particles and antiparticles

In the Standard Model, all particles are considered excited states of fundamental quantum fields and account for either the manifestation of the matter of the universe or for its interactions.

The elementary matter particles are fermions<sup>1</sup> with a spin of  $\pm 1/2$  and are divided into

---

<sup>1</sup>Fermions are particles with half-integral spin values.

quarks and leptons, each grouped into three generations with the mass increasing with the generation.

The quarks are further categorized as up-type quarks with an electromagnetic charge of  $+2/3$  or down-type quarks with an electromagnetic charge of  $-1/3$ . All quarks additionally carry a color charge, which enables them to interact via the strong force.

The leptons do not carry a color charge and consequently do not take part in strong force interactions. The group of the leptons consists of electrons, muons, and tau leptons, all with an electromagnetic charge of  $-1$ , as well as their corresponding neutrinos, which are electrically neutral.

All quarks and leptons in the Standard Model have a corresponding antimatter partner. These elementary antifermions have the same characteristics as their matter counterparts, except for the basic charges, which have an opposite sign. For example, the positron ( $e^+$ ), which was predicted by Paul Dirac in 1928 (see the paragraph below) and experimentally discovered from cosmic rays in a cloud chamber by Carl Anderson in 1933 [2] as the first antiparticle, is the antimatter partner of the electron ( $e^-$ ) and has an electromagnetic charge of  $+1$ . Positrons are produced naturally, for example in radioactive  $\beta^+$  decays or when a photon decays to an electron-positron pair (pair production). Example Feynman diagrams of these processes are shown in Appx. B. When a particle comes into contact with its antiparticle, the two annihilate and produce new particles. For example, the collision of an electron and a positron typically yields two photons.

Except for the electrically charged leptons and bosons, antimatter particles are denominated by a horizontal line above their identifying symbol.

In addition to the elementary fermions, there are six known elementary bosons<sup>2</sup>: the first five are vector bosons with a spin of  $+1$ , and the most recently discovered Higgs boson has a spin of  $0$ . These bosons are the mediators of the three fundamental forces between the fermions described in the Standard Model.

### The Dirac equation

To find a solution for the reconciliation of Quantum Mechanics and Special Relativity in the description of the properties of spin  $1/2$ -particles, while conserving Lorentz invariance and at the same time maintaining causality, in 1928, Paul Dirac proposed the relativistic wave equation given in Eq. 2.1 in its modern form [3]. Here,  $i$  is the imaginary unit,  $\hbar$  the reduced Planck constant,  $\gamma^\mu$  are the complex gamma matrices, and  $\partial_\mu$  the partial derivatives. The index  $\mu$ , over which a summation is implied, ranges from  $0$  to  $3$ , with  $0$  representing the time variable and  $1$  to  $3$  as space coordinates. The Dirac equation is compatible with Special Relativity in particular because the space and time derivatives enter at the same order.  $m$  is a parameter associated to the particles' mass and  $c$  the speed of light.  $\psi$  represents a four-component column vector, known as Dirac spinor, whose components are coupled by the gamma matrices. Eq. 2.1 is thus a construct of four coupled differential equations.

$$(i\hbar\gamma^\mu\partial_\mu - mc)\psi(x) = 0 \quad (2.1)$$

The solutions to this equation yield positive as well as negative energy eigenvalues, a feature that is incompatible with energy conservation and ultimately led to the postulation of *electrons with exactly opposite electric charge (but same spin and mass)*, i.e. positrons. The Dirac equation has become a fundamental equation of relativistic quantum mechanics

---

<sup>2</sup>Bosons are particles with integral spin values.

and quantum field theory, describing all spin-1/2 particles (and antiparticles) included in the Standard Model.

## Fundamental forces

The first force included in the Standard Model, the electromagnetic force, acts between all electrically charged particles and is mediated by the massless photon ( $\gamma$ ), which is itself electrically neutral. The interactions of the photon are theoretically explained by Quantum Electrodynamics (QED). The electromagnetic force is for instance responsible for the process of the electron-positron annihilation into photons.

The second fundamental force is the weak force. It has two different types of mediators: the electrically neutral Z boson with a mass of approximately  $91 \text{ GeV}/c^2$ , and two W bosons, which have an electromagnetic charge of  $\pm 1$  and a mass of about  $80 \text{ GeV}/c^2$ . All elementary SM fermions participate in interactions via the weak force, but only the electrically charged ones and neutrinos interact with the Z boson<sup>3</sup>. One example of an interaction via the weak force is the beta decay, transforming a proton into a neutron or vice versa.

The Standard Model assumes a unification of the electromagnetic and weak interactions, combining to the electroweak force at high enough energies.

Based on the theory of electroweak symmetry breaking, the generation of the masses of the Standard Model particles is explained by an interaction with the Higgs field via the Higgs boson, where the strength of the interaction is proportional to the masses of the particles [4, 5]. This process is called the Higgs mechanism. The Higgs boson was the last missing SM particle until the CMS and ATLAS experiments at the Large Hadron Collider at CERN discovered a Higgs candidate with a mass of approximately  $125 \text{ GeV}/c^2$  in 2012, whose properties are being investigated in great detail since [6, 7].

The strong force, as the third fundamental force included in the Standard Model, is mediated by eight different gluons, which carry combinations of three color and three anti-color charges, and are massless and electrically neutral. The gluon interactions are accounted for in the theory of Quantum Chromodynamics (QCD). The strong force acts only between these gluons and the color-charged quarks and is responsible for the confinement of the quarks into composite particles called hadrons, as detailed below.

The longest-known fundamental force is also the least understood one: gravitation, which acts on all particles via their mass-energy and, compared to the other three fundamental interactions, has a very small strength on a subatomic level. Gravitation is not included in the Standard Model, and, to date, there is no established theory explaining its mediation on a quantum level<sup>4</sup>.

## Hadrons and antihadrons

Hadrons are classified as baryons or mesons, depending on the number of valence quarks they are made of: baryons contain an odd number of primary quarks or antiquarks (at least three), while mesons are formed by an even number, usually a quark and its anti-quark.

In addition to the valence quarks, which contribute to the quantum numbers of a hadron, hadrons also contain so-called virtual sea quarks. The sea quarks are spontaneously created by the splitting of a gluon and typically annihilate again quickly inside the hadron

<sup>3</sup>Only left-handed particles and right-handed antiparticles interact with the W bosons; for interactions with the Z boson, the interaction strength depends on the particle's chirality.

<sup>4</sup>On a classical level, gravity is well described by General Relativity.

when encountering their antimatter counterpart, thus creating a constant flux of the total quark and gluon content.

The most common hadron is the proton (p, electromagnetic charge of +1), whose valence quark structure is formed by two up-quarks and one down-quark. The primary structure of the antimatter counterpart of the proton, the antiproton ( $\bar{p}$ , electromagnetic charge of -1), on the other hand, is made up of two anti-up-quarks and one anti-down-quark.

Just like electrons and positrons, hadrons typically annihilate when encountering their antimatter partner - or, in fact, any hadron containing at least one corresponding antiquark to their valence quark structure, as the respective quarks will be the ones undergoing the annihilation process<sup>5</sup>. Owing to the composite nature of the hadrons, their annihilation events typically do not only involve photons (as is the case for  $e^- - e^+$  annihilations), but a complex procedure of rearrangement due to strong interactions of the resulting gluon and the remaining constituents. This process is called hadronization and leads to the creation of new hadrons, mostly pions<sup>6</sup>, which then interact with the surrounding material, depositing energy and yielding further particles, or, ultimately, decaying.

### Shortcomings of the Standard Model

Although the SM describes – and has predicted – well the known elementary particles and most fundamental forces acting between them, it is incapable of providing a complete explanation of all characteristics of the universe and has several shortcomings, including the lack of explanations for the following processes and observations:

- Gravity as a fundamental force on a quantum level (as explained above)
- The matter-antimatter asymmetry of the universe (see Sect. 2.1.3)
- The existence of dark matter and dark energy, which, according to cosmological observations, account for over 90 % of the universe's mass-energy [8, 9]
- The so-called "hierarchy problem": A unification of all fundamental forces described above, in a theory that is universally (i.e. also up to the Planck scale at  $\mathcal{O}(10^{19}$  GeV), where quantum effects of gravity become relevant) valid, would require quantum loop corrections to the masses of particles such as the Higgs boson to be much higher than the electroweak scale, as found by experimental observations, or so precisely tuned that they are often considered "unnatural" [10, 11].

Several "beyond the Standard Model" (BSM) theories have been developed to extend the Standard Model and propose explanations for above concepts; for example, theories based on Supersymmetry or string theory are among the most well-known, suggesting the existence of additional, often heavy but weakly interacting particles that could provide dark matter or "quantum gravity particle" candidates as well as "natural" solutions to the hierarchy discrepancy of the fundamental forces [12–14]. None of these predictions have, as of yet, been experimentally verified but extensive research is done by a multitude of different experiments.

---

<sup>5</sup>Due to the instability of the sea quarks, annihilation processes between two hadrons involving sea quarks are very rare.

<sup>6</sup> $\pi^0$ :  $u\bar{u}$  or  $d\bar{d}$ ,  $\pi^+$ :  $u\bar{d}$ ,  $\pi^-$ :  $d\bar{u}$

### 2.1.2 Antiproton properties

The antiproton ( $\bar{u}\bar{u}\bar{d}$ ), as the most common anti-hadron, is generally present in the nucleus of any anti-atom<sup>7</sup> that is being produced and is therefore at the core of most antimatter research. It was first intentionally produced and observed in 1955 by Emilio Segrè, Owen Chamberlain et al. at the Bevatron accelerator, where antiprotons were created, similarly to the procedure now used at CERN's Antiproton Decelerator (AD), by directing a 6.2 GeV proton beam onto a copper target [15].

Since then, investments into the deceleration of the produced antiprotons at CERN's Low Energy Antiproton Ring (LEAR) and the AD complex have enabled precision studies of the properties of antiprotons. The most precise measurements of this kind to date have been performed by experiments at the AD:

In 2016, the ASACUSA collaboration has determined the antiproton-to-electron mass ratio as 1836.1526734(15), in agreement with the known value for protons within  $8 \times 10^{-10}$ , in a spectroscopic measurement of antiprotonic helium atoms, which are produced by mixing antiprotons with helium gas and are subsequently cooled to below 2 K [16].

The BASE collaboration has determined the ratio of the antiproton and proton's charge-to-mass ratios to be 1 within 16 parts per trillion through measurements of both particle's cyclotron frequencies inside a cryogenic (order of 1 K) Penning trap in 2022 [17].

BASE has also precisely measured, from the antiprotons' cyclotron and Larmor frequencies inside their trap, the antiproton magnetic moment in 2017 as 2.7928473441(42), to be compared to the most precise measurement of the proton magnetic moment by the same collaboration and compatible at a level of 0.8 parts per million, with the precision of the antiproton measurement being higher than that of the proton by an order of magnitude [18].

Antiprotons are so far assumed to be stable, in the same way as protons; however, they generally annihilate quickly after production when coming into contact with regular matter. The record time antiprotons have ever been kept alive, trapped in a vacuum Penning trap, is 405 days, yielding a lower limit on the  $\bar{p}$  lifetime of 10.2 years (68% confidence level) [19]. For certain decay modes, more stringent lifetime constraints of up to  $10^5$  years have been reported [20, 21].

### 2.1.3 CPT and the matter-antimatter asymmetry

According to the Charge, Parity, and Time reversal (CPT) theorem, which was proven to be a consequence of few, basic hypotheses (Lorentz invariance, locality, and unity) in the middle of the twentieth century by Gerhart Lüders and has not been contradicted by any observations since, CPT symmetry holds for any relativistic quantum field theory [22, 23]. This means that symmetry is conserved when at the same time applying charge conjugation, parity transformation, and time reversal to a system described by such a theory, while the individual symmetries are broken in the Standard Model.

As a result of the CPT theorem, the Standard Model assumes complete symmetry between matter and antimatter states. Therefore, bound systems of antiparticles are expected to possess the same characteristics as their matter counterparts, including having the same lifetime and energy levels. Furthermore, any known process that creates matter particles also creates antimatter counterparts.

---

<sup>7</sup>There are of course some exceptions of anti-atoms without antiprotons; for example, individual positronium systems (bound state of an electron and a positron) are also commonly referred to as "exotic" atoms. Therefore, from here on, the term "Ps atom" is also used to describe a single bound positronium system.

Accordingly, the primordial universe is predicted to have consisted of matter and antimatter in equal amounts and, more precisely, because matter particles annihilate upon meeting their antimatter counterparts, it should have quickly been devoid of any matter at all. However, present-day observations of the Cosmic Microwave Background (CMB) show an extensive excess of ordinary matter over antimatter, also referred to as the baryon asymmetry. In fact, the difference between the two is quantified by the cosmologically observed ratio of the baryon and photon number densities in the universe,  $n_{\text{baryon}}/n_{\gamma} = 6 \times 10^{-10}$  [24, 25].

Several hypotheses have been developed which are aimed at explaining the observed asymmetry. These include the theorized existence of an antimatter-universe (having existed before the Big Bang) to pair to the matter-dominated one, or that of antimatter-galaxies within distant areas of the known universe, separated by boundary regions from the matter one [26, 27]. Some theories also suggest "antigravity", i.e. gravitational repulsion between matter and antimatter, as a cause for a possible separation of the two [28–30]. In general, a difference in the effective gravitational interaction of matter and antimatter, which could extend the SM for example through additional gauge bosons, offers possible explanations for the current asymmetry. Such differences can be searched for in precision tests of the Weak Equivalence Principle (see Sect. 2.2) [31].

Most of these theories provide valid explanations for the baryon asymmetry as well as certain additional observed phenomena, even those not included in the Standard Model, but are in contrast to, or lack explanations for, others, and none of them have been experimentally confirmed [32, 33].

Already in 1967, Andrei Sakharov defined theoretical conditions to be met for a mechanism to yield a matter-antimatter asymmetry, i.e. to generate baryogenesis<sup>8</sup> [34]. These three so-called Sakharov conditions are the following:

- Baryon number violation
- Violation of C and CP symmetry
- The process is happening out of thermal equilibrium.

The third condition was formulated assuming exact CPT symmetry. If, on the other hand, one allows for CPT violation to occur, it is possible to generate a matter-antimatter asymmetry in thermal equilibrium, given also the first two conditions.

The Standard Model itself provides an explanation for a difference in behaviour of matter and antimatter particles, which is known as CP (Charge conjugation, Parity) violation. CP violation was first demonstrated in the decays of neutral K mesons by James Cronin and Val Fitch in 1964 and has been confirmed for the weak interaction in other systems since [35]. Symmetry is thus not conserved by a combined transformation of charge and parity (i.e. flip in the sign of one spatial coordinate) alone, and the additional reversal of time is needed to achieve symmetry conservation. The extent of the matter-antimatter asymmetry predicted by this phenomenon within the Standard Model and observed in experiments does, however, not suffice to explain the imbalance observed in our universe today [36].

It is possible that, if also CPT symmetry is observed to be violated in certain processes, the additional effect would be able to predict a sufficiently large imbalance. Such hypothetical violations of the CPT theorem can be directly probed by comparing the properties

---

<sup>8</sup>Since most of the existing matter in the universe consists of baryons (and the antibaryons are missing), the process for the generation of the imbalance of matter and antimatter is also referred to as baryogenesis.

of matter and antimatter systems. However, stringent tests of the CPT symmetry between matter and antimatter systems have been performed in recent years [37, 38], all of whose results are consistent with the CPT theorem [39, 40].

Some BSM theories, whose probing forms part of current high energy particle physics research, suggest the presence of extended or modified Higgs sectors as an explanation for the emergence of baryogenesis [41, 42]. Generally, a hypothetical creation of asymmetric baryon densities during the electroweak phase transition<sup>9</sup> (assumed to have been a first order process), so-called electroweak baryogenesis, could satisfy the Sakharov conditions. At the same time, electroweak baryogenesis models predict the existence of additional particles not included in the Standard Model and/or modified properties of the Higgs boson, which can be probed in collider experiments such as at CERN's LHC [42, 43].

Recently, extensive research has also been performed to probe the electric dipole moment (EDM), i.e. a difference in the center of charge with respect to the center of mass along the angular momentum axis, of fundamental leptons, in particular the electron and the muon. The magnitude of these is limited to very small values in the Standard Model (maximum  $1.0 \times 10^{-35}$  e·cm for the electron) because they are only non-zero due to rare quantum interactions with virtual quarks in the SM: They arise from CP violation processes involving only quarks [39, 44, 45]. The presence of larger EDM would violate both parity and time reversal symmetry, opening the door to a different behaviour of matter and antimatter particles [46]. However, all experiments performed so far yield results in agreement with the Standard Model predictions, placing limits on any theories beyond [39, 47], while experimental sensitivities still range a few orders of magnitude above SM predictions. The same is true for experiments on baryon EDMs, such as that of the neutron (expected to be of the order of  $1.0 \times 10^{-32}$  e·cm) [46, 48].

## 2.2 The Weak Equivalence Principle - Universality of Free Fall

In 1638, Galileo Galilei observed that the trajectory of a freely falling object is independent of its mass [49]. This principle, known as the Universality of Free Fall (UFF), requires, in terms of the physics of Isaac Newton [50], an exact equivalence of gravitational and inertial mass. Gravitational mass ( $m_g$ ) is defined as the charge to which gravitation couples (correspondent to the electromagnetic force coupling to electrically charged entities), while inertial mass ( $m_i$ ) is a measure of how slowly an object is accelerated by a given force. Accordingly, all objects dropped at the same location near the surface of the earth with the same velocity fall with the same free-fall acceleration of approximately  $9.8 \text{ m s}^{-2}$ , independently of their internal structure and mass.

Until today, there is no evidence to contradict the validity of the Universality of Free Fall. In the context of General Relativity, Albert Einstein postulates a more general objective, known as the Equivalence Principle (EP) [51]. It states that, locally, i.e. in a sufficiently small region of space-time, no experiment can distinguish between a homogeneous gravitational field and a uniformly accelerated frame of reference.

Universality of Free Fall is thus a direct consequence of General Relativity and manifests itself in one of the two forms of the Equivalence Principle that are conventionally distinguished: the Weak Equivalence Principle (WEP). Contrarily to the EP in its Strong form (SEP), the WEP neglects gravitational contributions to the binding energy of test masses, while the SEP requires the gravitational self-energy to obey the Equivalence Principle as

---

<sup>9</sup>The electroweak phase transition refers to the onset of the Higgs mechanism.

well (which can generally only be tested for self-gravitating bodies, i.e. those with a large enough mass of the order of astronomical objects).

### 2.2.1 Testing the WEP

Owing to the variety and diverse scaling of EP test setups, experimental results probing the validity of the Equivalence Principle are customarily compared by employing the dimensionless Eötvös parameter  $\eta$ , given in Eq. 2.2, which is defined as the differential acceleration of two test masses A and B in a given force field divided by their average acceleration in that field. Since the relation  $m_i \vec{a} = m_g \vec{g}$  (with  $\vec{a}$  as measured acceleration) should hold according to the WEP, the Eötvös parameter can be rewritten as shown, using the ratio  $\alpha_g = g_1/g_2$  to express differences in the gravitational interactions of the two test objects. The last conversion is valid for  $\alpha_g \ll 1$ .  $\eta$  quantifies the deviation from the UFF and, thus, any non-zero result can be interpreted as an indication of a violation of the Weak Equivalence Principle. The parameter is named after Loránd Eötvös, who performed the first precise comparison of inertial and gravitational mass of matter in a torsion balance experiment between 1885 and 1909 [52].

$$\eta = 2 \frac{|a_A - a_B|}{|a_A + a_B|} = 2 \frac{|(m_g/m_i)_A - (m_g/m_i)_B|}{|(m_g/m_i)_A + (m_g/m_i)_B|} = 2 \frac{|\alpha_g - 1|}{|1 + \alpha_g|} \approx 2|\alpha_g - 1| \quad (2.2)$$

The acceleration  $g$  of matter in the gravitational field of the earth is well established, and any WEP tests performed with matter are in agreement with above assumptions, the most stringent limit having been set by the MICROSCOPE collaboration in 2022 via a measurement of the acceleration of objects free-falling in a satellite in earth's orbit, with  $\eta$  of the order of  $1 \times 10^{-15}$  [53–55]. Measurements of the effect of  $g$  on antimatter, on the other hand, are scarce, mostly due to the shortage of available antimatter. Nonetheless, they can shed light on very fundamental questions such as the formation and composition of the universe.

### 2.2.2 WEP tests with antimatter

Certain measurements aimed at stringent CPT invariance investigations also yield results which can be interpreted as indirect WEP tests. Particularly interesting are observations of the cyclotron frequency of protons and antiprotons in Penning traps, such as those of the TRAP collaboration at LEAR and, more recently, the BASE collaboration at the AD [17, 56]. It is assumed that if gravitation coupled anomalously to antimatter, the oscillation frequencies of the two systems inside the trap would differ (WEP for matter/antimatter clocks) [57]. With this, BASE suggests a limit of the WEP-violating gravitational anomalies of  $|\alpha_g - 1| < 1.8 \times 10^{-7}$  [17], restricting the Eötvös parameter to  $\eta < 3.6 \times 10^{-7}$  according to Eq. 2.2. However, due to a strong dependency of this result on the exact local gravitational potential and difficulties in its determination, these absolute values are controversial [58]. For this reason, it is possible to state a differential result, obtained by repeating measurements at different times of the year and determining the differences in the oscillation frequencies according to the elliptical orbit of earth around the sun, which then depends instead on the gravitational potential of the sun. With this approach, BASE claims  $|\alpha_{g,\text{diff}} - 1| < 0.03$  [17] and the Eötvös parameter can be determined as  $\eta_{\text{diff}} < 0.06$ .

Already in 1968, F. Witteborn and W. Fairbank devised an experimental scheme to directly test the influence of gravity on a *charged* antiparticle: the positron. After attempting a free-fall experiment with both electrons and positrons, they concluded that

it was not possible to control the systematic effects of the experiment on the charges in a sufficient manner to obtain a meaningful result [59]. The influence of electromagnetic fields, which are generally present in an experimental environment, exceed that of gravity by several orders of magnitude for such small masses. Some efforts were made by the PS200 collaboration at LEAR to mitigate the obstacles of such field effects when working towards a gravity measurement with antiprotons [60]. However, LEAR was closed in favor of CERN's AD program in 1996 before the experiment commenced operation.

Because of difficulties such as these, the focus of direct gravity measurements has shifted more towards *neutral* antimatter systems, which are much less affected by the external fields. First direct observations of the effect of gravity on antimatter, obtained in a free-fall experiment with antihydrogen atoms by the ALPHA collaboration, are in agreement with the local gravitational acceleration of these anti-atoms being the same as for regular matter [61]. However, this measurement is mainly limited in its precision ( $\Delta g/g \approx 0.29$ ) by being conducted inside a magnetic field, whose modelling contributes essentially to the uncertainties. Still, this measurement constitutes the first direct gravitational measurement on antimatter and is widely accepted to have determined the sign of  $g$  for antimatter, confirming that its gravitational interaction points in the same direction as for regular matter. Most of the mass of bound systems is generated from their binding energy (almost 99% for the proton for example) and only a small portion from the constituent masses themselves. The indication, found from the WEP tests on matter outlined above, that gravity influences binding energy in general as it does regular matter, therefore limits the expected allowed deviation of  $g$  for antihydrogen atoms to around 1%, motivating the advancement of precision experiments.

## 2.3 Neutral antimatter systems

Currently, it is experimentally possible to produce four different neutral atomic systems containing antiparticles in significant amounts at low enough temperatures to be manipulated, i.e. 100 K or below:

- Positronium
- Antihydrogen
- Protonium
- Antiprotonic helium.

### 2.3.1 Positronium

Positronium (Ps) is an unstable exotic atom formed by an electron and a positron orbiting each other, which was first experimentally produced in 1951 [62]<sup>10</sup>. Being a purely leptonic system, Ps is very well suited for precision tests of Quantum Electrodynamics with antimatter, as no hadronic interactions are involved. On the other hand, thanks to the similarity of its composition to hydrogen atoms, the energy levels of positronium can also be approximated in the same way as for hydrogen, using Eq. 2.4 with a reduced mass  $\mu$  that differs from the one for hydrogen according to the mass difference of the

---

<sup>10</sup>While "positronium" itself can refer to an individual "atom", or bound system, it is also used to label the substance as such (as is for example done for (anti-)hydrogen). For this reason, the single instances of this substance are from here on also called "Positronium (or Ps) atoms".

proton and the positron: Since the electron and the positron have exactly the same mass ( $m_{e^+} = m_{e^-}$ ),  $\mu$  can be expressed as shown in Eq. 2.3, adding a factor of approximately 1/2 compared to the reduced mass of the electron and the heavier proton in hydrogen. Accordingly, the energy levels  $E_n$  of positronium are also about half those of hydrogen. The constants used in Eq. 2.4 are the electric charge  $q_e$  of the electron, the Planck constant  $h$  and the vacuum permittivity  $\epsilon_0$ .  $n$  denotes the principal quantum number of the positronium atom at the given energy level.

$$\mu = \frac{m_{e^+}m_{e^-}}{m_{e^+} + m_{e^-}} = \frac{m_{e^-}^2}{2m_{e^-}} = \frac{m_{e^-}}{2} \quad (2.3)$$

$$E_n = -\frac{\mu q_e^4}{8h^2\epsilon_0^2} \frac{1}{n^2} \approx -\frac{6.8 \text{ eV}}{n^2} \quad (2.4)$$

An efficient way to produce positronium atoms is the insertion of positrons into a medium, a so-called converter target, where, if energetically favourable, one type of interaction the positrons can undergo is combining with an electron to form Ps. Inside this medium, a large electron density to enable Ps formation is given and momentum conservation can be achieved more easily than when freely mixing electrons and positrons. The positronium production procedure in AEGIS is described in Sect. 3.3.2.

As the electron and the positron inside the Ps are each other's antiparticles, they annihilate after a short time, making positronium unstable. Ground state Ps can exist in two different configurations of the hyperfine structure, depending on the alignment of the spins of the electron and positron. The singlet state with antiparallel spins ( $S = 0$ ) is called para-positronium and has a mean lifetime of less than one nanosecond, primarily decaying into two photons. If the spins of both particles are parallel ( $S = 1$ ), i.e. Ps is formed in a triplet state, it is referred to as ortho-positronium, whose lifetime is of the order of 142 ns with a predominant decay to three photons [63].

It is possible to excite positronium to higher quantum levels (e.g. laser excitation to Rydberg levels as done in AEGIS), which is experimentally useful since it increases the Ps lifetime and additionally has a favourable impact on its reactivity. Both of these features are exploited in the AEGIS experiment, where a charge exchange reaction between positronium and antiprotons is used to produce antihydrogen atoms.

### 2.3.2 Antihydrogen

Being the only neutral bound system purely made of antimatter (antiproton and positron) that can currently be produced in reasonable amounts at low energies ( $\sim K$ ), and moreover being stable, antihydrogen ( $\bar{H}$ ) is uniquely qualified for use in precision experiments investigating the properties of antimatter.

#### Antihydrogen properties

According to the CPT theorem, antihydrogen is expected to share most of the well-known characteristics of regular hydrogen, for example the transition frequencies of the atoms' energy levels (and thus their emission spectra), the magnetic moment, its mass, and its behaviour in external fields (including the influence of gravity). With production of antihydrogen atoms now established (see below), precision tests of  $\bar{H}$  properties are a growing field of research.

For example, the ALPHA collaboration at CERN's AD performs spectroscopy on magnetically trapped  $\bar{H}$  atoms and has confirmed the agreement of the transition frequencies

for the 1S-2S, 1S-2P, and 2S-2P transitions and the fine structure Lamb shift for  $\bar{\text{H}}$  with the predicted and experimentally observed values for regular hydrogen [64–66]. The neutrality of the  $\bar{\text{H}}$  electric charge has also been confirmed to below 0.71 parts per billion from an investigation of the antihydrogen atom trajectory in electric fields [67]. Furthermore, several collaborations at the AD are working on measurements of the influence of gravity on antihydrogen atoms [61, 68, 69].

### Antihydrogen production mechanisms

The first antihydrogen atoms were produced and detected in the 1990s at CERN’s LEAR decelerator and at Fermilab by directing a beam of antiprotons onto a target, leading to a small fraction of them capturing a positron from the preceding pair production and forming  $\bar{\text{H}}$  atoms [70, 71]. Both experiments, while successful, were yielding only individual antihydrogen atoms at much too high energies (i.e. too hot) to allow for a precise investigation of their characteristics.

CERN subsequently constructed and commissioned the Antiproton Decelerator (AD) complex (see Sect. 3.1), paving the way for the production and analysis of low-energy antimatter systems in larger amounts. In 2002, both the ATHENA and ATRAP collaborations used a procedure of mixing antiprotons from the AD and positrons inside a Penning trap to produce thousands of colder ( $\sim$  a few 1000 K) antihydrogen atoms [72, 73]. The plasma mixing, whose principle is still used today by the ASACUSA and ALPHA collaborations [74, 75], produces antihydrogen via three-body (Eq. 2.5) and radiative (Eq. 2.6) recombination, with the former being predominant.



ATRAP has devised a second procedure for antihydrogen production, which, in their case, involves two charge exchange reactions [76]: First, cesium atoms from a thermal beam, previously laser-excited to Rydberg states, are collided with cold positrons to form positronium. The Ps atoms are then used for a second charge exchange with trapped antiprotons, yielding antihydrogen. The second reaction, Eq. 2.7, is still employed by the AEGIS and GBAR collaborations [68, 69].



One advantage of the charge exchange approach, while being orders of magnitude less efficient than the mixing, is the negligible plasma heating – and thus lower temperature of the produced  $\bar{\text{H}}$  – thanks to the static trap potential. Furthermore, and particularly important for the endeavours of AEGIS to measure the influence of gravity on the anti-atoms via their vertical deflection from the knowledge of their time of flight in a deflectometer, this procedure allows for a very precise determination of the instant of creation – and thus launch – of the  $\bar{\text{H}}$  atoms, without the need to trap them, thanks to a pulsed production mode. While GBAR produces non-excited Ps to form ground state antihydrogen for their  $\bar{\text{H}}$  ion formation, AEGIS directly excites the Ps to Rydberg states. Since the excited state of the product of a charge exchange reaction depends on that of the reactant, and the cross section for  $\bar{\text{H}}$  formation from Eq. 2.7 scales as  $a_0^2 n_{\text{Ps}}^4$  ( $a_0$  being the Bohr radius and  $n_{\text{Ps}}$  the principal quantum number of the used Ps), this direct Ps excitation can selectively target higher  $\bar{\text{H}}$  production rates than the ground state and two-stage charge exchange [77].

### 2.3.3 Protonium

Protonium (Pn) is the bound state of a proton and an antiproton and is, as such, also referred to as antiprotonic hydrogen. Pn, like Ps, consists of matter-antimatter partners (the (anti-) quarks inside the hadrons here being each other's antiparticles) and is, therefore, not stable, but can have lifetimes of up to a few microseconds [78].

Other than the relatively unregulated production in particle collisions, protonium can be produced by simply mixing antiprotons and protons inside the same particle trap, which was first demonstrated by the ATHENA collaboration in 2007 [79].

While some theoretical calculations and spectroscopy experiments have been performed with protonium, many of its properties have not yet been studied in detail and could contribute to the understanding of interesting processes, in particular those involved in QCD and internucleon forces.

### 2.3.4 Antiprotonic helium

Antiprotonic helium is the metastable (lifetime  $\sim \mu\text{s}$ ) Coulomb-bound state of a helium nucleus ( $\text{He}^{2+}$ ) with an electron and an antiproton orbiting around it. It was discovered in 1991 at KEK in Japan through the observation of unexpectedly long antiproton lifetimes inside helium gas and can be synthesized by mixing antiprotons with He gas, whereby a few percent of the introduced antiprotons will replace one of the electrons orbiting the nucleus [80, 81].

At CERN, the ASACUSA collaboration uses spectroscopy on cold antiprotonic helium atoms to precisely measure the antiproton mass [16].

# Chapter 3

## Introducing AEGIS

*This chapter presents the experimental setup of AEGIS as it is at the time of writing. Between 2019 and 2023, multiple core components of the AEGIS setup have been extensively upgraded to improve antihydrogen production and broaden the physics reach. My main contribution to this is the new control system, which includes almost all main experimental subsystems and is introduced in the next chapter. However, I also contributed to further hardware upgrades, such as the installation and commissioning of the new Penning-Malmberg trap, and the redesigned electron gun, the preparation and routing of all new cables for the trap electrodes, and the characterization of the new degrader setup. Details of the most significant upgrades to the system can be found in [82], which was written and submitted by me for the collaboration, and in the reports of AEGIS to the SPS Committee [83–85], which I contributed to together with other collaboration members. The new alexandrite laser system, which has been specifically designed and constructed for laser cooling of positronium and which utilizes the AERIALIST synchronization capabilities implemented by me, is reported on in [86], of which I am a contributing author. I have additionally contributed to the integration and data acquisition of the prototype detector planned to be used for the gravity measurement, which is reported on in [87] with me included in the co-authors list. I have also taken part in the acquisition and have performed the analysis of the data for the ESDA splitter calibration and SC1112-SC12 intercalibration. Furthermore, during the antiproton run campaigns in 2021, 2022 and 2023, I was part of the "cryo team", in charge of the continuous re-filling of liquid helium into the cryogenic apparatuses.*

AEGIS, along with five other experiments<sup>1</sup> based at the Antiproton Decelerator (AD) at CERN, aims at investigating the properties of antimatter, comparing them to regular matter properties, and drawing conclusions on fundamental concepts such as the CPT theorem. AEGIS, in particular, focuses on measuring, with high precision, earth's gravitational acceleration on antihydrogen, to probe the Weak Equivalence Principle. For this purpose, the AEGIS experiment produces bound antimatter systems – in particular neutral antihydrogen atoms.  $\bar{\text{H}}$  is formed by combining cold antiprotons, captured from the AD complex and confined in Penning-Malmberg traps, with positrons from laser-excited positronium through a charge exchange reaction. The atoms are to be produced in a pulsed, horizontal beam to infer the influence of gravity from their vertical deflection over a known travelling distance in a deflectometer [88].

This chapter introduces the AD complex and the AEGIS experiment itself, giving details

---

<sup>1</sup>The five other active AD experiments are: ALPHA, ASACUSA, BASE, GBAR, and PUMA.

on antiproton manipulation techniques as well as the setup of the experiment and the procedure to form antihydrogen in a pulsed way.

## 3.1 The Antiproton Decelerator at CERN

The Antiproton Decelerator complex is the only existing facility to provide large numbers of low-energy antiprotons, which enables a variety of unique research. It has been integrated in the main accelerator structure of CERN since its installation in the year 2000, with the major upgrade of the new ELENA decelerator in 2021 [89, 90], and the  $AE\bar{g}IS$  experiment commenced operation in 2012, having first demonstrated production of antihydrogen in a pulsed scheme in 2018 [68].

### 3.1.1 CERN's accelerator complex

The Large Hadron Collider (LHC), located at the European Laboratory for Particle Physics (CERN), is the world's most powerful particle collider experiment. In its main operation mode, protons are accelerated to energies of up to 7 TeV and then brought to collision with the purpose of investigating the resulting products of their interaction and drawing conclusions on the fundamental constituents of matter (and antimatter) and the forces acting between them.

In order to accelerate the protons (and, in other modes of operation, also heavy ions) to such high energies as to be instrumental to these investigations, a large complex of accelerator machines is employed into which the particles are injected. This complex is shown schematically in Fig. 3.1.

The acceleration of the particles is done by a combination of radiofrequency cavities and cooling mechanisms, which work in the same way as the deceleration procedures employed in the AD, simply with electric field oscillations used in the opposite way, and are thus explained in more detail in Sect. 3.1.3.

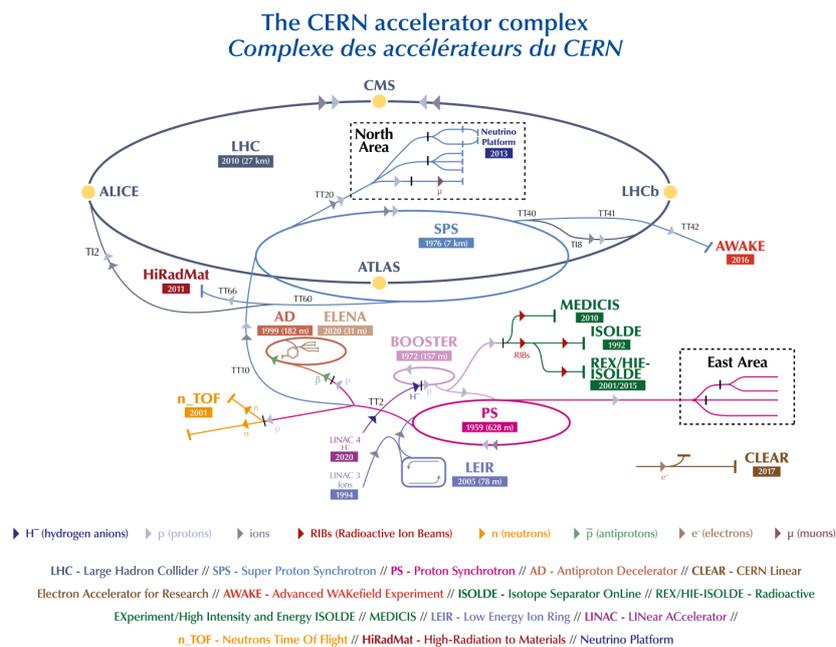


Figure 3.1: Schematic of the CERN accelerator (and decelerator) complex [91].

The origin of the protons is a source of negative hydrogen ions ( $\text{H}^-$ ) which are inserted into a linear accelerator (LINAC4), accelerating them to 160 MeV in several stages before transferring them to the Proton Synchrotron Booster (PSB) ring. During this transfer, the negative ions are stripped of their electrons through an electric field, leaving only protons. These are then accelerated to 1.4 GeV in the PSB before again being transferred, from the PSB to the Proton Synchrotron (PS), which increases the particle energy to 26 GeV and subsequently injects the protons into the larger Super Proton Synchrotron (SPS). The SPS marks the last stage of the pre-acceleration and delivers the protons to the LHC ring with an energy of 450 GeV. The LHC accelerates the particles to their final energy and brings them to collision inside four experimental locations along the tunnel.

### 3.1.2 Antiproton production

Slow enough antiprotons can be caught and trapped inside experimental setups, allowing to directly investigate their properties or to combine them with other particles, for example with positrons to form antihydrogen (see Sect. 3.3). As with all antiparticles, antiprotons annihilate when encountering their matter counterparts which, in the universe we know, happens basically immediately after their creation without the application of a high vacuum. Furthermore, there is no natural radioactive decay which produces antiprotons (as is the case for example for positrons).

At CERN, antiprotons are created by directing the 26 GeV proton beam from the Proton Synchrotron (PS, see Fig. 3.1) onto an iridium target. The impact of the protons onto the target initiates the creation of various secondary particles through the process of pair production by interaction with the target nucleons. A Feynman diagram of a simple pair production process, where an electron-positron pair is created from an incoming photon, is shown in Appx. B. In the same way, the creation of other particle-antiparticle pairs from a neutral boson is possible as well given the observance of energy and momentum conservation laws. Such processes are facilitated by the principle of mass-energy equivalence ( $E = mc^2$  in a system's rest frame, with  $E$  as energy,  $m$  as rest mass and  $c$  as the speed of light) whose discovery finds its origin with Albert Einstein in 1905 [92]. Therefore, in order for this conversion to take place, the incoming boson has to carry a total energy which is at least equivalent to the rest mass of the resulting particle pair ( $E_{\gamma, \min} = 2m_e c^2 = 1022 \text{ keV}$  for the production of an electron-positron pair), but in most cases significantly larger because of momentum conservation. If its energy is larger, the additional amount is transferred to the newly created particles as kinetic energy. In general, the cross section of pair production processes increases with the energy of the incoming particle.

For the momentum conservation to hold as well, a "collision partner" is required (for example a second target nucleon), which enables the interaction by experiencing a recoil, balancing the momentum of the system.

As described above, pair production can occur for leptons, but the creation of quark-antiquark pairs is just as well possible, enabling also the formation of hadron-antihadron pairs.

The minimum energy of the boson required for the creation of a proton-antiproton pair is significantly higher than for the above case due to their much larger mass ( $m_p c^2 = 938 \text{ MeV}$ ). When directing the highly energetic protons from the PS onto the dense

iridium target<sup>2</sup>, they are stopped very quickly via nuclear interactions with the target material, producing also photons with sufficient energy to create proton-antiproton pairs. It is possible to calculate the minimum required energy of an incoming proton hitting a fixed target for the creation of a proton-antiproton pair. Assuming the interaction takes place in the field of a proton or neutron of the target nucleus, one can employ four-momentum conservation laws (see Appx. C) to compute a minimum proton energy of 6.6 GeV ( $7m_p c^2$ ) or, in other terms, a threshold proton kinetic energy of approximately 5.6 GeV.

Thus, among the different types of resulting secondary particle pairs after the incidence of the 26 GeV proton on the iridium target, there are protons and antiprotons of various energies which move in different directions. According to their distinct charge-to-mass ratio, the antiprotons can be separated from the other particles by a magnetic horn [94] and are then focused until those with the selected energy of 3.5 GeV are inserted into the circular structure of the Antiproton Decelerator<sup>3</sup>.

### 3.1.3 Antiproton deceleration

The AD is a storage ring decelerator with a circumference of 188 m and is composed of various magnets used to focus the beam (quadrupole magnets) and store the antiprotons on a circular track inside the ring (dipole magnets). Inside the AD, strong electric fields from radiofrequency (RF) cavities are employed to slow down the antiprotons axially. This is done via oscillations (switches of direction) of the field inside each cavity according to the predicted current position of the antiproton bunches passing through them, so as to counteract their forward momentum. However, in accordance with phase space conservation laws (Liouville's theorem [95, 96]), this would lead to a radial divergence of the beam, increasing the momentum spread and eventually causing strong losses. To prevent this, a combination of stochastic cooling and electron cooling in several steps is used on the antiproton beam, alternating with the deceleration steps.

The first cooling steps after the injection of the beam are composed of stochastic cooling. This technique is based on the principle of a feedback loop: a detector on one side of the AD measures the deviation of a sample of particles from the center of gravity of the bunch and sends the corresponding signal diagonally across the ring, where the electric field of a kicker corrects for this deviation during the following passage of the bunch. After many cycles of these "random" corrections, the overall momentum spread and transverse emittance is significantly reduced.

For lower beam energies, electron cooling becomes more efficient. Electron cooling is a type of sympathetic cooling which makes use of the large mass difference between electrons and antiprotons, which have the same charge: a cold beam of electrons of the same velocity as the antiprotons is merged with the antiprotons, which then undergo Coulomb collisions with the electrons. During these collisions, the antiprotons transfer momentum to the much lighter electrons, thus effectively cooling. After many collisions, thermal equilibrium is reached when both particle types have the same momentum, and the electrons

---

<sup>2</sup>Iridium was chosen as a target material because of its high density: the target length should be kept as short as possible in order to maximize the phase space density and thus the collection efficiency of the antiprotons. The optimized length is finally a compromise of allowing the majority of protons to interact and preventing re-absorption of the antiprotons and is found to be of the order of one proton nuclear interaction length [93].

<sup>3</sup>The value of 3.5 GeV for the injection energy of the antiprotons into the AD was chosen for practical reasons including the existing infrastructure (e.g. the transfer region from PS to SPS) as well as the availability of space and power from the pre-cooling electrodes.

are steered away from the antiproton beam.

At the end of the deceleration and cooling steps, the antiprotons have an energy of 5.3 MeV.

### 3.1.4 ELENA

In order for the AD experiments to be able to effectively trap and manipulate the arriving antiprotons, they need to have an energy of no more than a few keV, as this is an energy whose voltage equivalent can be reached employing commercially available high-voltage electrodes. For this reason, the antiprotons arriving directly from the AD have to be further decelerated. Typically, the AD experiments have used relatively thick aluminium degrader foils (160  $\mu\text{m}$  in the case of AEGIS) to achieve this large difference in energy for a fraction of the antiprotons by their interactions with the foil material. The majority of the antiprotons passing through are, however, completely stopped in this thick degrader, and thus up to 99 % of them are lost [68,97]. Comparable trapping efficiencies of a maximum of a few percent are achieved with an alternative approach used by the ASACUSA collaboration, which employs a radiofrequency quadrupole decelerator to further reduce the  $\bar{p}$  energy [98].

A new decelerator ring, the Extra Low ENergy Antiproton ring (ELENA), has been commissioned inside the AD hall in 2021 [90]. A schematic of the setup inside the AD hall is shown in Fig. 3.2, including the AD and ELENA "rings" and the locations of the different experiments. The slowed-down antiprotons are transferred from the AD ring to the smaller ELENA ring with a circumference of 30 m, where their energy is further reduced, again employing both radiofrequency deceleration cavities and electron cooling analogously to the AD procedure<sup>4</sup>, to finally reach 100 keV. The momentum spread stated by the ELENA collaboration is of the order of 0.05 %. From ELENA, the slow antiprotons are conveyed to the different AD experiments via magnetic transfer lines. This upgrade enables an improvement of the possible antiproton capture efficiencies of the different AD experiments by up to two orders of magnitude through the use of much thinner, optimized degrader materials.

---

<sup>4</sup>Stochastic cooling is not efficient for antiprotons with energies as low as those of the ones in the ELENA ring, making electron cooling the only used cooling method here.

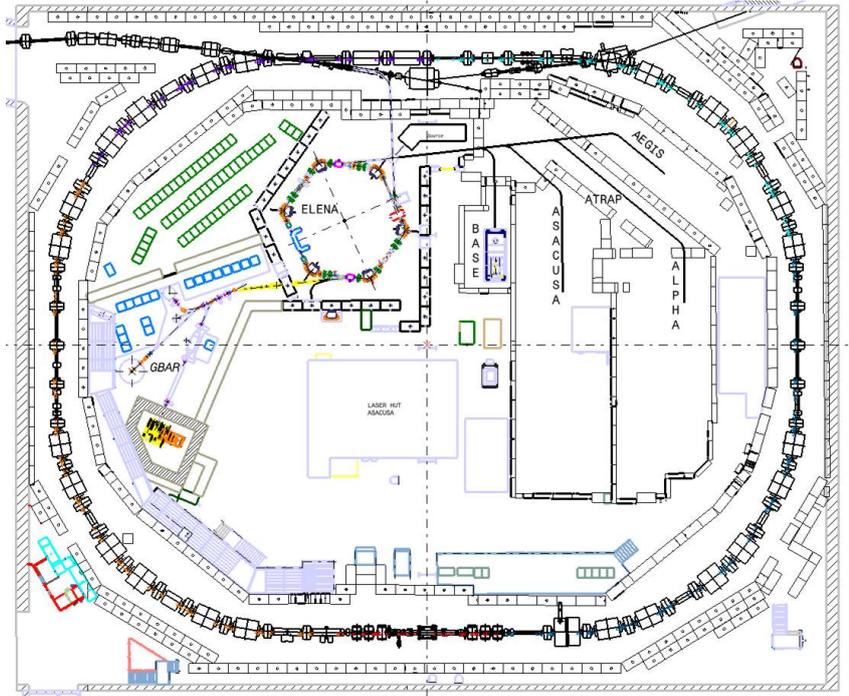


Figure 3.2: Schematic of the AD hall at CERN, showing the large ring of the AD around the inside border of the hall together with the smaller ELENA decelerator (hexagonal structure) within, as well as the transfer lines to the different AD experiments [89]. The operation of the ATRAP experiment ended in 2018 and the available space is taken over by the BASE-STEP collaboration, whose apparatus is currently being installed, as is the PUMA experiment in a new area close to GBAR.

In addition to allowing experiments to much more efficiently capture the supplied antiprotons, another advantage of the deployment of ELENA is the accompanying change of the antiproton provision scheme: Previously, experiments were supplied with antiprotons in rotation for eight-hour slots, i.e. depending on how many experiments were requesting the antiproton beam, each one had to wait its turn for up to 40 hours without beam. With ELENA, antiprotons are available 24 hours per day every day because four experiments can be supplied simultaneously. If more than four experiments request beam at the same time, each experiment is left out for one shot every so often (again depending on the number of involved experiments), in rotation. This further increases the number of antiprotons the experiments can use and thus greatly benefits the statistical development. On the other hand, this non-stop antiproton provision also poses a challenge for the experiments to be operated in a way not to miss beam while it is available for several months at a time. This functionality is one of the core pillars of the new experimental control system developed in AEGIS since 2021, which is introduced in Ch. 4.

ELENA provides antiprotons to the connected experiments in bunches of up to  $7 \times 10^6 \bar{p}$  (bunch lengths of the order of 125 ns) approximately every 110 seconds, mainly defined by the cycle length of the AD; the cycle length of ELENA itself is below 20 seconds [90].

### 3.2 Trapping antiprotons: Penning traps

The devices typically employed by the AD experiments to capture and confine large numbers of charged (anti-) particles in a small volume are called Penning-Malmberg traps.

Generally, these traps have high-voltage (up to few tens of kV) endcap electrodes to stop and capture the antiprotons from ELENA and a bulk of lower-voltage electrodes for efficient manipulation of the confined charges. If the density of the particles in these traps is high enough and they are kept at low temperatures of a few K, the particles form plasmas.

According to Earnshaw's theorem, which follows directly from the Poisson equation, the motion of charges can only be restricted in two spatial directions at the same time by electrostatic fields [99]. It is therefore not possible to construct stable particle traps using only such static electric fields.

Penning traps rely upon the combination of a strong homogeneous axial magnetic field and an inhomogeneous electric quadrupole field to enable trapping in all three spatial dimensions [100, 101]. A schematic of the basic operation principle of a Penning trap is shown in Fig. 3.3.

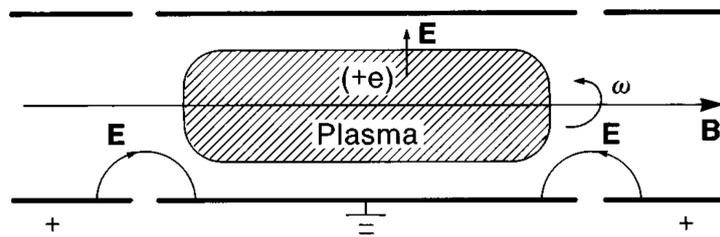


Figure 3.3: Drawing (by F. Penning, 1936) of the operation principle of a Penning trap to confine a plasma, from [102]. The trap consists of three cylindrical electrodes (two endcaps biased oppositely to the plasma charge and one central electrode kept at ground potential) for axial confinement by the electric field  $E$ . The plasma shown here has a positive charge; reversal of signs yields validity of the principle for negative charges. A uniform magnetic field  $B$  along the trap axis provides radial confinement. Due to the interactions of the electric and magnetic fields, the plasma forms a spinning (frequency  $\omega$ ) rod in the central part of the trap.

The magnetic field confines the charged particles radially: due to the Lorentz force, they oscillate on a circular path around the magnetic field lines with the cyclotron frequency  $\omega_c$  that depends only on the ratio of charge  $q$  and mass  $m$  of the particles and the magnetic field  $B$ , a relation that can be derived from an equalization of Lorentz force and centripetal force for the circular motion of a charged particle in a magnetic field:

$$\omega_c = \frac{q}{m}B. \quad (3.1)$$

In the axial direction, the particles' motion is limited by the electrostatic repulsion from the quadrupole field: In order to confine negatively charged particles axially, a setup of two negatively charged end-caps and one (relative to the end-caps) positive central (ring) electrode is in principle sufficient, creating a potential saddle-point in the center, where the particles are trapped.

In comparison to the ideal Penning design, realistic Penning-Malmberg traps possess several coaxial cylindrical electrodes instead of one ring and two end-caps to generate the (not accurately quadratic) electric field, optimized for the large numbers of confined particles as well as the typically elongated shape of the plasmas and yielding more freedom for the manipulation of the particles.

The addition of the electric quadrupole field modifies the motion of the charged particles

inside the trap from the simple cyclotron frequency [103]. To be precise,  $\omega_c$  is slightly reduced to the so-called modified cyclotron frequency  $\omega'_c$ , while a second type of motion in the radial plane, a slower drift around the trap center with the magnetron frequency  $\omega_m$ , is introduced:

$$\omega'_c = \frac{\omega_c}{2} + \sqrt{\frac{\omega_c^2}{4} - \frac{\omega_z^2}{2}} \quad (3.2)$$

$$\omega_m = \frac{\omega_c}{2} - \sqrt{\frac{\omega_c^2}{4} - \frac{\omega_z^2}{2}}. \quad (3.3)$$

$\omega_z$  in above equations denotes the frequency of the trapped particles' motion in the axial direction, i.e. along the axis between the end-cap electrodes. This frequency also depends again on mass  $m$  and charge  $q$  of the particles as well as on the potential difference  $V_0$  between end-caps and center and a geometric parameter  $d$  of the trap<sup>5</sup>, and can be expressed as

$$\omega_z = \sqrt{\frac{qV_0}{md^2}}. \quad (3.5)$$

While it follows from Eq. 3.2 and Eq. 3.3 that  $\omega_c = \omega'_c + \omega_m$ , this relation is, in fact, only true for ideal Penning traps. On the other hand, the invariance theorem given in Eq. 3.6 can be shown to remain valid also for real traps which possess misalignments and are subject to imperfections of the applied potential [103].

$$\omega_c^2 = \omega'^2_c + \omega_m^2 + \omega_z^2 \quad (3.6)$$

In order to obtain a stable, confining trap, the experimental parameters have to be chosen in such a way to meet the trapping condition, Eq. 3.7. Otherwise, the outwardly pushing, radial electric field starts to overcome the confining magnetic field and the particles drift away from the trap center.

$$\omega_c^2 > 2\omega_z^2 \quad (3.7)$$

Additionally utilizing Eq. 3.1 and Eq. 3.5, the trapping condition can be expressed in terms of the parameters of an experimental setup as follows:

$$\frac{qB^2}{m} > \frac{2V_0}{d^2}. \quad (3.8)$$

Another kind of trap for charged particles is a radio frequency (RF) trap, also called Paul trap [104]. Paul traps do not make use of magnetic fields but employ dynamic electric quadrupole fields, whose potential is made to change quickly enough (RF rates) for the contained particles to oscillate between the electrodes without actually reaching them.

One advantage of Penning traps compared to Paul traps is the exclusive use of static fields, avoiding the micro-motion and resulting heating that would be caused by dynamic fields. For this reason, most AD experiments use variants of Penning-Malmberg traps, as it is essential to keep the trapped antiparticles at the lowest possible temperatures in order to be able to investigate their properties.

---

<sup>5</sup>For a trap with hyperbolic electrodes as described above,  $d$  can be determined from the distance  $z_0$  between end-caps and trap center and the trap radius  $r_0$  according to Eq. 3.4.

$$d^2 = \frac{z_0^2}{2} + \frac{r_0^2}{4} \quad (3.4)$$

### 3.2.1 Thermal equilibrium in Penning-Malmberg traps

Non-neutral plasmas confined in Penning-Malmberg traps generally arrange themselves in a cylindrical rod, producing a radial electric field  $E_{\text{plasma}}$  which pushes the particles outward. As a result of the interplay with the confining magnetic field, the plasma rod experiences a Lorentz force  $E_{\text{plasma}} \times B$  drift and comes to thermal equilibrium, spinning around the azimuthal axis as a rigid body with a frequency  $f_{E_{\text{plasma}} \times B}$  that is linearly proportional to the plasma density  $\rho$ , as given in Eq. 3.9 [102]. Here,  $\epsilon_0$  is the vacuum permittivity constant.

$$f_{E_{\text{plasma}} \times B} = \frac{q\rho}{4\pi\epsilon_0 B} \quad (3.9)$$

Eq. 3.9 is derived in Appx. D.

The linear dependence of the plasma rotation frequency on  $\rho$  allows to influence the density of the plasma by controlling its rotation via external forces, one example of which is explained below.

### 3.2.2 The Rotating Wall technique

One possibility to manipulate charged particles and plasmas stored in electromagnetic traps, in particular to influence their compression and density, is the use of the Rotating Wall (RW) technique [105–107].

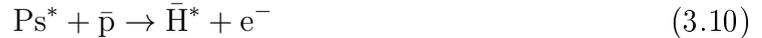
RW relies on the application of a rotating electric field around the confined plasma, perpendicular to its axis of symmetry. This field couples to the plasma and induces an electric dipole which creates a torque. This feature can be exploited to increase or decrease the rotation speed of the plasma (rotating the external electric field faster or more slowly than the natural plasma rotation respectively). Consequently, the plasma compression, and its density, can be modified according to the Lorentz force (see Eq. 3.9).

Typically, such a rotating electric field is achieved by azimuthally segmented ring electrodes installed as part of the confining trap. In the case of AEgIS, such four-fold sectorised electrodes form part of both of the main plasma traps [108]. Phased, sinusoidal voltages at a given frequency are applied to the segments, creating the spinning electric field and influencing the plasma properties as described above. See Sect. 4.1.4 for details on the AEgIS setup.

## 3.3 AEgIS: A procedure for pulsed production of anti-hydrogen

The AEgIS (Antimatter Experiment: Gravity, Interferometry, Spectroscopy) collaboration, based at CERN’s AD/ELENA complex, aims at a precise, direct measurement of the gravitational acceleration of antimatter by means of the position measurement of a horizontal beam of antihydrogen atoms guided through a Moiré deflectometer [88, 109]. The production of antihydrogen has been successfully demonstrated using several different production mechanisms (see Sect. 2.3.2), all starting from antiprotons and combining them with positrons in different ways. For measurements relying on the time of flight of the formed atoms over a known distance, like the AEgIS gravity measurement, precise knowledge of the moment of launch is essential. A possible way to tackle this is to directly produce the  $\bar{\text{H}}$  atoms in a pulsed procedure. The only known such procedure is the one developed for the first time in AEgIS [68]: Formation of antihydrogen atoms in Rydberg states ( $\bar{\text{H}}^*$ ) is facilitated by the charge exchange reaction Eq. 3.10 between positronium

in excited states ( $\text{Ps}^*$ ) and cold antiprotons ( $\bar{p}$ ), leaving as a side product the electrons ( $e^-$ ) released from the positronium atoms. This procedure allows for a precise (within a few hundred nanoseconds) determination of the excitation instant of the Ps and, thus, the time of formation of  $\bar{\text{H}}$ . If the formation happens at a well-defined location in-flight, this time also defines the "launch".



For this procedure to work, both the antiprotons and the positronium atoms need to undergo a series of preparatory steps before finally being combined inside the antihydrogen production trap of the experiment. Details on the exact implementation of these steps in the upgraded version of AEGIS are given in Ch. 7.

### 3.3.1 Antiproton manipulation

For an efficient production of antihydrogen atoms travelling at velocities that facilitate a gravity measurement (of the order of  $1000 \text{ ms}^{-1}$  or below), the antiprotons themselves also need to be cold enough and controlled for sufficient amounts of time. The steps performed to achieve this are the following:

- Preparation of an electron plasma in a low-voltage ( $\pm 200 \text{ V}$ ) potential well between high-voltage electrodes inside a Penning-Malmberg trap within a 5 T magnetic field, and application of Rotating Wall frequencies to the included segmented electrode(s)
- Further reduction of the energy of the antiprotons arriving from ELENA (from 100 keV to 14 keV or below) by letting them pass through thin foils of material called "degraders"
- In-flight capture of the slowed-down antiprotons between high-voltage electrodes reaching up to  $-14 \text{ kV}$  with the use of a nanosecond-switching mechanism
- Cooling of the caught antiprotons by Coulomb collisions with the electron plasma and accumulation in the low-voltage trapping region
- Dumping of the fraction of antiprotons that has not cooled from the trap by ramping down the high voltage ("hot dump")
- Further cooling and Rotating Wall compression of the mixed electron-antiproton plasma
- Careful modification of the trap potentials to transfer antiprotons from the capture trap (in the 5 T part of the experiment) to the antihydrogen production region (where a 1 T field is present).

The last step can be done in different ways, either optimized for re-capture of the antiprotons in the downstream trap to essentially consider them at rest when the  $\bar{\text{H}}$  production occurs or for a ballistic  $\bar{\text{H}}$  production with the antiprotons accelerated in the desired direction when meeting the Ps atoms. Details on the implementation of the second option are given in Ch. 7; the first option is the one used for the previous  $\bar{\text{H}}$  production [68]. It is furthermore possible to perform the capture and hot dump procedures several times in a row in order to stack several ELENA shots in the trap and accumulate a larger number of cold antiprotons. Further details on this are given in Ch. 6.

### 3.3.2 Positronium formation and Rydberg excitation

The positronium atoms used for the charge exchange reaction, Eq. 3.10, have to first be produced from a source of positrons. As the charge exchange antihydrogen production rate increases proportionally to the fourth power of the principal quantum number  $n_{\text{Ps}}$  of the positronium, as given by the cross section  $\sigma_{\text{Ps}^*+\bar{\text{p}}\rightarrow\bar{\text{H}}^*+\text{e}^-}$  in Eq. 3.11 (with  $a_0$  as the Bohr radius) [77,110], the Ps is subsequently excited to Rydberg levels. This manipulation has the added benefit of extending the lifetime of the produced Ps by several orders of magnitude, depending on the excitation level [110].

$$\sigma_{\text{Ps}^*+\bar{\text{p}}\rightarrow\bar{\text{H}}^*+\text{e}^-} \propto n_{\text{Ps}}^4 \pi a_0^2 \quad (3.11)$$

The necessary steps for the preparation of the positronium, as implemented in AEGIS, are briefly summarized here, with further details given in Sect. 3.4:

- Accumulation and cooling of positrons from a  $^{22}\text{Na}$  source
- Magnetic transfer of the positrons into the main apparatus and onto the positronium converter target
- Positronium production by implantation of the positrons in a nanochannel Si target and recombination with the electrons of the target material, with subsequent re-emission of the Ps
- Laser excitation of the Ps atoms in two steps [111]:
  - Excitation to  $n = 3$  by a broadband UV laser
  - Excitation to Rydberg levels by a broadband IR laser.

### 3.3.3 The number of produced antihydrogen atoms

The number  $N_{\bar{\text{H}}}$  of antihydrogen atoms produced through the charge exchange process, Eq. 3.10, can be estimated, as derived in [112], on the basis of the definition of the cross section, Eq. 3.12, where  $\sigma$  is  $\sigma_{\text{Ps}^*+\bar{\text{p}}\rightarrow\bar{\text{H}}^*+\text{e}^-}$ , the charge exchange cross section introduced in Eq. 3.11, which is obtained from Classical Trajectory Monte-Carlo calculations and depends heavily on the ratio of the Ps center of mass velocity to the velocity of the positron in its classical orbit and on the magnetic field [77], as well as on the Ps excitation level to the fourth power. Typical Ps velocities in AEGIS are of the order of  $10^5 \text{ m s}^{-1}$  [113]. Generally, lower Ps temperatures have a favourable effect on the cross section.  $\rho_{\bar{\text{p}}}$  denotes the density of the antiproton plasma, which can be approximated from their number  $N_{\bar{\text{p}}}$  by assuming a cylindrical plasma equilibrium distribution in the trap, as given in Eq. 3.13 for cylinder radius  $r_{\bar{\text{p}}}$  and length  $L_{\bar{\text{p}}}$ .  $N_{\text{Ps}^*}$  corresponds to the number of positronium atoms that is Rydberg-excited and reaches the antiproton cloud.  $v_{rel}$  is the relative velocity of positronium and antiprotons, which can be well approximated by the average center of mass velocity of the Ps atoms when assuming the antiprotons to be at rest in comparison. This can be assumed to be the case in the current conditions with typical antiproton velocities up to  $10^3 \text{ m s}^{-1}$ .  $v_{rel}$  is the only quantity in Eq. 3.12 that depends on the interaction time  $dt$ , i.e. the average time the Ps atoms spend within the  $\bar{\text{p}}$  cloud.

$$\frac{dN_{\bar{\text{H}}}}{dt} = \sigma \rho_{\bar{\text{p}}} N_{\text{Ps}^*} v_{rel} \quad (3.12)$$

$$\rho_{\bar{p}} \approx \frac{N_{\bar{p}}}{\pi r_{\bar{p}}^2 L_{\bar{p}}} \quad (3.13)$$

To obtain the number of produced  $\bar{H}$  atoms, Eq. 3.12 is integrated over the interaction time. The result is well approximated by the interaction length of a linear jet and a planar, elliptical cloud and can therefore be expressed as  $\pi \rho_{\bar{p}} / 2 \sin(\theta_{Ps})$ , when allowing for a non-horizontal entering of the Ps atoms into the  $\bar{p}$  cloud.  $\theta_{Ps}$  is the angle between the Ps trajectory and the vertical direction. Using this result and plugging in Eq. 3.13, the number of produced antihydrogen atoms is finally given by Eq. 3.14.

$$N_{\bar{H}} = \sigma \frac{N_{\bar{p}} N_{Ps^*}}{2 \sin(\theta_{Ps}) r_{\bar{p}} L_{\bar{p}}} \quad (3.14)$$

### 3.3.4 Collinear antihydrogen formation: AEGIS Phase II

Fig. 3.4 shows schematics of the antihydrogen production via the charge exchange reaction, Eq. 3.10, in the AEGIS apparatus according to the described procedures. In AEGIS Phase I (left hand side of Fig. 3.4), the setup used for the first  $\bar{H}$  production in 2018 [68], positronium was produced off-axis and entered the experiment from the top. This meant that the Ps atoms had to travel orthogonally through the magnetic field of the superconducting magnets to reach the antihydrogen production trap, which had a small, central opening in the top of one of its electrodes for this purpose. This setup caused two main issues:

- Firstly, the hole in the electrode represented an azimuthal asymmetry which caused expansion and heating of the confined plasmas, i.e. relevant particle losses.
- Secondly, the highly excited Ps, travelling perpendicularly to the magnetic field lines, experienced dynamical field ionization, also called Stark ionization [110, 111, 114]. These losses furthermore increase with the principle quantum number of the Ps atoms, thus limiting the usable excitation level  $n_{Ps} < 20$  to maintain an efficient compromise between the losses and the possible antihydrogen production according to its dependence on  $n_{Ps}$ .

A more detailed discussion of the second point can be found in Sect. 7.4.

For these reasons, the AEGIS collaboration has, as part of its Phase II upgrade between 2020 and 2023, redesigned the entire setup for antihydrogen production from an orthogonal to a collinear scheme, with the positronium conversion target installed on the same axis as the particle traps, parallel to the magnetic fields lines [82]. A corresponding schematic is shown on the right hand side of Fig. 3.4. The incoming positrons now pass through the trap holding the antiprotons towards the Ps target and are emitted back on-axis from the target. This furthermore allows to tune the distance between the target and the antiprotons to reduce solid angle losses of the Ps atoms.

Additionally, in the original proposal for the AEGIS experiment, the formation of the horizontal  $\bar{H}$  beam was envisioned via Stark acceleration, by applying an accelerating electric field on the electrodes behind the production region and utilizing the sensitivity of the highly excited atoms, which are essentially produced at rest, on the field gradients [109]. This idea has been abandoned in favor of a new approach based on the ballistic acceleration of the antiprotons prior to  $\bar{H}$  formation. The development of this procedure is reported on in Ch. 7 and is only realistically implementable thanks to the on-axis  $\bar{H}$  production.

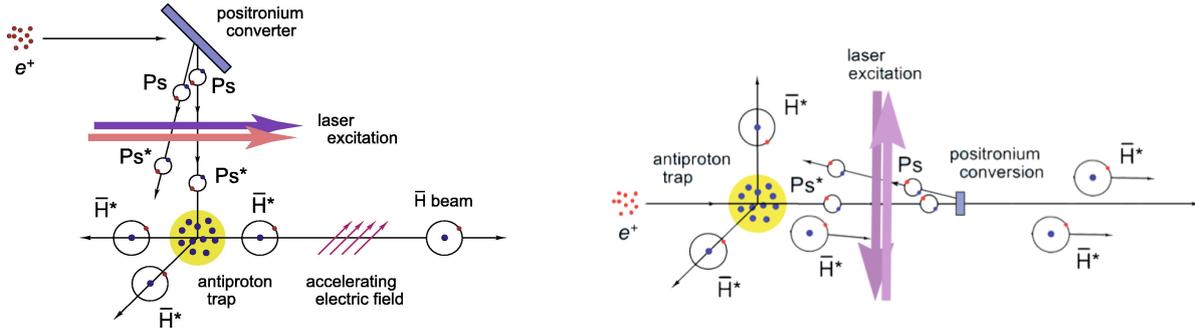


Figure 3.4: Schematics of antihydrogen production with the AEGIS experiment. Positrons enter the conversion target and the produced Ps is emitted to be two-step excited (pink/purple arrows). They encounter the antiprotons, which have previously been captured in an electromagnetic trap (yellow circle), and form excited  $\bar{H}$  via charge exchange, which can be formed into a horizontal beam. Left: AEGIS Phase I with Ps entering the  $\bar{H}$  production trap orthogonally [115]. Right: AEGIS Phase II with collinear  $\bar{H}$  production [84].

### 3.4 AEGIS setup

In 2018, AEGIS has already successfully produced individual antihydrogen atoms with the procedure outlined in Sect. 3.3, with a precisely known formation time of the order of 250 ns [68]. However, this represents only a feasibility study (AEGIS Phase I) of this pulsed production mechanism: the production efficiency was extremely low (of the order of 0.05 to 0.1 atoms per formation cycle) and thus infeasible for any kind of statistically relevant measurement of the  $\bar{H}$  properties (see Sect. 3.5.1). Therefore, AEGIS has undergone several significant upgrades since, aimed at improving the antihydrogen production by a few orders of magnitude and entering the experiment into its Phase II towards a gravity measurement [82].

A schematic of the setup of the AEGIS experiment is shown in Fig. 3.5 in simplified blocks, not to scale. The different components are briefly explained below. Generally, at the core of AEGIS lies a structure of two main Penning trap setups (small black dashes in Fig. 3.5), kept at vacuum and located inside 5 and 1 T magnetic fields, which are supplied by superconducting magnets (yellow blocks), whose cryogenic apparatus (light and dark blue areas) surrounds the entire trap region. For antihydrogen production, antiprotons ( $\bar{p}$ ) from ELENA enter the experiment and, after being energy-degraded (by the "degrader" structures), captured and cooled in the traps, meet positronium atoms that emerge from the conversion target (red spot) after implantation of previously prepared positrons and are subsequently laser-excited. A vacuum test chamber dedicated to positronium experiments also forms part of the setup and can be supplied with positrons as an alternative to their injection into the  $\bar{H}$  beam line. This is the location where successful laser-cooling of Ps atoms has been demonstrated for the first time in 2023 [116]. An MCP detector is located at the downstream end of the trap axis, and the cryostat is surrounded by a detector formed by scintillator slabs (orange blocks). The behaviour of the different components is governed and monitored by a novel experimental control system, which has been developed in AEGIS and is introduced in Ch. 4.

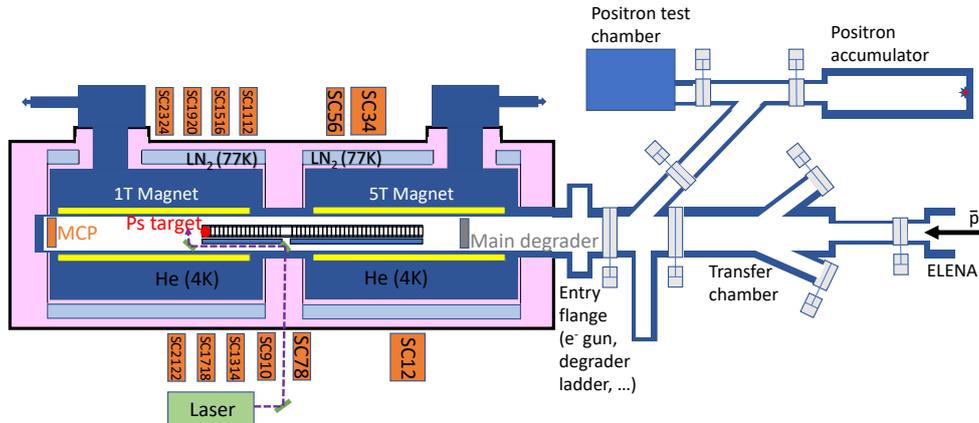


Figure 3.5: Simplified schematic setup of the AEGIS experiment. Not to scale. The connection to ELENA for the antiproton supply is on the right hand side (upstream end) of the experiment. The different components are described in the main text.

### 3.4.1 Magnet setup and coordinate system

The setup of the AEGIS magnet system is the result of a compromise between two contradictory requirements of the different subsystems of the experiment. On the one hand, a magnetic field  $B$  as high as possible is needed in the capture trap region for two main purposes: to maximize the radial confinement of the delivered antiproton bunches, and to facilitate a strong cyclotron radiation ( $\propto B^2$ ), i.e. loss of energy, of the electrons used to cool the antiprotons by Coulomb collision. On the other hand, in the antihydrogen production region, the magnetic field needs to be kept small enough to limit the mixing of atomic levels of the Rydberg atoms and (antihydrogen and positronium).

For this reason, the antihydrogen production region has been spatially separated from the capture region of the AEGIS trap system, and the magnet setup has been designed to supply each with a homogeneous magnetic field optimized to the respective requirements. A strong magnetic field, which can reach up to 5 T, is applied in the capture trap area, while the antihydrogen production trap is positioned inside a 1 T magnetic field.

For the generation of each magnetic field, a superconducting solenoid of a Niobium Titanium alloy is employed, which is cooled down to 4.2 K using liquid Nitrogen and Helium. The experimental axis and the magnets are aligned to one another such that the magnetic field lines run along the axis of the experiment, i.e. along the same axis as the centres of the Penning-Malmberg traps (see Sect. 3.4.8). This axis is referred to as  $z$ -axis, with the upstream end starting at the entrance region of the experiment, where the connection to ELENA is installed, and the downstream end located behind the positronium target and MCP detector. The center position ( $z = 0$ ) is defined to be between the two magnetic field regions. The direction from here towards the entrance region of the experiment is referred to as "upstream", while the direction towards the MCP detector is called "downstream". The simulated magnetic flux density distribution in the experiment is shown in Fig. 3.6, and further details can be found in [117]. The magnetic fields in the plateau regions shown

there are homogeneous to 0.2 mT (in the 5 T trap region) and 0.05 mT (in the 1 T trap region).

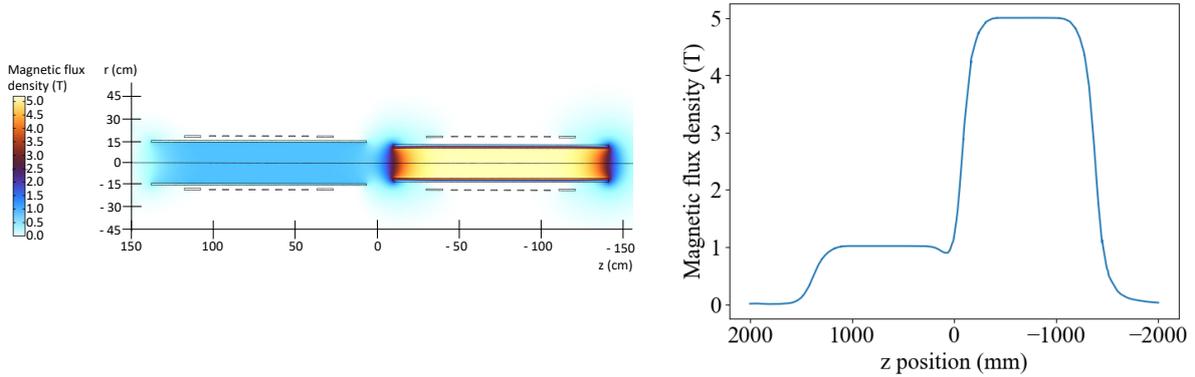


Figure 3.6: Magnetic flux density in the AEGIS experiment. Left: Magnetic field map around the electromagnetic trapping regions in axial and radial direction. Right: Plot of the total axial magnetic field over the horizontal length of the experiment.

### 3.4.2 Vacuum system

The limiting factor for the lifetime of the antiprotons is ultimately their annihilation probability, which is directly dependent upon the amount of residual gas (mostly hydrogen gas) present in the trapping region of the experiment. In order to maximize the  $\bar{p}$  lifetime to macroscopic values, an ultra-high vacuum (UHV) of up to  $10^{-13}$  mbar is therefore applied around the Penning-Malmberg traps confining the (anti-) particles.

Even in the case of a leak or when opening the flange keeping the innermost vacuum chamber closed, there is no direct transition between UHV and atmospheric pressure. Instead, an Outer Vacuum Chamber (OVC) is installed around the central magnets, which has a pressure of the order of  $10^{-7}$  mbar.

While the cryo-pumping described below removes much of the residual gas in the core apparatus, hydrogen and helium are generally not affected by this and remain inside because their condensation temperatures are below the reachable levels. A combination of a set of getter pumps (for hydrogen) and an ion pump (for helium) are used to minimize the presence of these gases [118]. Furthermore, the cryogenic apparatuses are heated before running experiments to release residue left in their surfaces that would otherwise slowly outgas and contaminate the vacuum, a process that is referred to as baking.

### 3.4.3 Cryogenic system

The core elements of the experiment, i.e. the superconducting magnets and the electromagnetic traps, are installed inside a cryogenic vessel filled with liquid helium to keep it at a stable temperature of 4 K. This temperature is needed to maintain the superconductivity of the magnets and has the added benefits of avoiding the heating of confined particles and plasmas and provoking a cryo-pumping of the surfaces, i.e. adsorption of molecules on the apparatus walls, in the trap regions, thus improving the vacuum [119]. A second, outer layer cryogenic container is placed around this central one, inside the OVC. This one is filled with liquid nitrogen, providing a temperature of 77 K.

While the nitrogen only needs to be refilled every few months during experimental cam-

paigns, the helium reservoirs require filling every couple of days to maintain a level that prevents magnet quenching and keeps the superconductivity alive.

### 3.4.4 Positron system

The positrons used to produce positronium in AEGIS originate from the  $\beta^+$  decay of a  $^{22}\text{Na}$  source (200 MBq at the time of the most recent  $\bar{\text{H}}$  run of 2023, being replaced in 2024 by a 1850 MBq source), releasing them with a continuous energy distribution of up to 546 keV and a mean energy of the order of 216 keV [120, 121]. They pass through a solid neon moderator, which is grown on the sodium. A low-voltage potential, applied to the entire source region, slows down a fraction of the positrons to an energy of the order of eV, allowing them to enter a buffer gas trap; the too hot ones (e.g. the ones emitted in the opposite direction and experiencing the opposite effect from the electric potential) are instead filtered out<sup>6</sup>. The trap, like the main Penning-Malmberg traps of AEGIS, confines the charged particles through a combination of a (in this case mT) magnetic field and electric potentials from cylindrical electrodes. For efficient confinement, a fast cooling mechanism is needed for the light positrons, which is realized by the buffer gas at high pressure (of the order of  $10^{-4}$  mbar). Once sufficiently cooled, the positrons are transferred from this first trap to a second one, whose buffer gas pressure is significantly lower (of the order of  $10^{-8}$  mbar) to allow for  $e^+$  lifetimes of several minutes (while providing less efficient cooling than in the first trap). This second trap accumulates the cold positrons and can release pulses of the order of  $10^6$   $e^+$  upon the arrival of a trigger signal. The released positrons are then accelerated to some keV (to prevent an axial broadening) and magnetically directed towards the positronium formation target [122], either inside the main apparatus for  $\bar{\text{H}}$  production or in the additional test chamber for dedicated Ps studies. In 2022, a transport of the positrons using only the electrostatic field has been achieved for the first time as well, opening the door for Ps measurements in a magnetic field-free environment.

Further details on the positron system as well as on dedicated positronium experiments performed with the AEGIS apparatus can be found in [121].

### 3.4.5 Positronium converter target

A piece of material inside which positronium can be formed by insertion of positrons and recombination with the material electrons is called a positronium converter (or conversion) target. The Ps targets used in AEGIS are made of nanoporous silicon dioxide ( $\text{SiO}_2$ ): Nanochannels (i.e. pores with a diameter of several nanometres, here up to 20 nm) are electrochemically etched into silicon wafers and subsequently heated in air, provoking the formation of a thin ( $\sim 10$  nm) layer of  $\text{SiO}_2$  on the surface. Through multiple etching and heating cycles, the final depth of the nanochannels is of the order of 1 to 2  $\mu\text{m}$  [123, 124]. These parameters have been optimized for the energies of the incoming positrons (a few keV) and for a high ortho-Ps yield in cryogenic temperatures, as the Ps target used in AEGIS forms part of the Penning-Malmberg trap installed inside the He-cooled cryostat. Inside the target material, the arriving positrons lose energy through different processes, including ionization and scattering at the relevant energies, as well as phonon excitation, and can thermalize enough to annihilate, either directly or subsequently to Ps formation,

---

<sup>6</sup>Only the cold positrons are able to adiabatically follow curved magnetic field lines provided by a saddle coil around a tungsten target, while the hot ones, i.e. the ones with a too high momentum, end up on the obstacle on a straighter path.

or form positronium by charge exchange with surrounding atoms or molecules or undergoing a bound with a quasi-free electron [125,126]. The chosen  $\text{SiO}_2$  setup can have Ps conversion efficiencies of up to over 80 %, and the positronium is produced at energies of several eV [127,128]. Routinely, efficiencies of 5 to 10 % are observed in  $\text{AE}\bar{\text{g}}\text{IS}$ . Then, the Ps atoms undergo inelastic collisions with the walls of the nanochannels, essentially cooling, before being emitted into the vacuum. Therefore, the cooling efficiency is defined by the geometry of these channels and, given sufficient optimization, the lower bound of the Ps temperature is set by the surrounding environment. In the cryogenic setup of  $\text{AE}\bar{\text{g}}\text{IS}$ , temperatures of a few tens of meV have been achieved in this way [123,129].

The geometry used in  $\text{AE}\bar{\text{g}}\text{IS}$  results in a re-emission of the produced Ps in the direction from which the positrons enter, which is referred to as reflection production. A production in transmission is also generally possible, has been investigated in  $\text{AE}\bar{\text{g}}\text{IS}$ , and would have advantages such as additional flexibility for experimental setups, however, the production of such converters is quite challenging and the targets themselves are more fragile. Further details on such discussions can be found in [121].

As explained below, the  $\text{AE}\bar{\text{g}}\text{IS}$  Ps converter is now mounted on the axis of the trap, incorporated on the last ring electrode of the antihydrogen production trap, which is installed on a movable actuator. To reduce solid angle losses of the formed antihydrogen atoms having to travel past the target, the dimensions of the new target are reduced to  $(4.6 \times 4.6) \text{ mm}^2$ , with a thickness of 1 mm, and the entire setup, including a heater and thermal sensor, is held by a thin, 3D-printed titanium support structure. The target and the surrounding structure are shown in Fig. 3.7. In 2023, the Ps conversion efficiency achieved by the target has been significantly lower than previously, reaching less than 3 % and thus strongly limiting the  $\bar{\text{H}}$  yield. Possible reasons are discussed in Ch. 7. Upgrades in the planned production technique are expected to improve the efficiency to approximately 28 % and reduce the most probable Ps temperature from 1.5 to below  $1 \times 10^5 \text{ m s}^{-1}$  through deeper positron implantation [129].

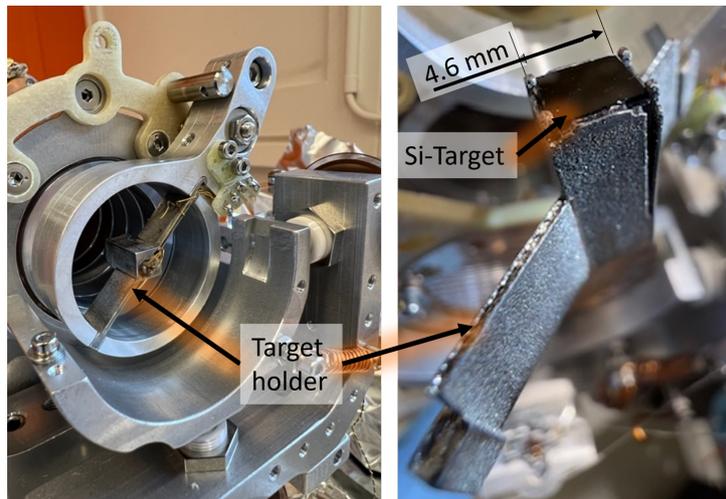


Figure 3.7: Photographs of the Ps converter target and its holder, mounted in the center of the last trap electrode.

### 3.4.6 Laser system

The laser system of  $\text{AE}\bar{\text{g}}\text{IS}$  consists of two individual setups, one (the "EKSPLA") for the purpose of exciting positronium and photo-ionization of Rydberg atoms, and another

one (the "Alex") for spectroscopy and laser cooling experiments. Both systems are briefly introduced below; detailed information can be found in [86, 112, 113, 130]. The two lasers are operated individually and have to be extremely well synchronized to the nanosecond despite their difference in frequency when used together (e.g for measurements of positronium laser cooling). The beam parameters of the used laser systems are summarized in Table 3.1.

### The "EKSPLA"

The origin of the excitation laser is an EKSPLA model NL303 Q-switched Nd:YAG laser pumping module that produces a beam at 1064 nm and also includes crystals for Second Harmonic Generation (SHG) and Fourth Harmonic Generation (FHG) stages for frequency doubling, providing laser light at 532 nm and 266 nm as well. Through a combination of Optical Parametric Generation (OPG), Optical Parametric Amplification (OPA), and Sum Frequency Generation (SFG), the latter two are used to produce a broadband ultra-violet (UV) (tunable between 204.9 nm and 205.2 nm) laser beam. It is used to excite positronium to the  $n = 3$  level, under standard conditions with an efficiency of the order of 15 % [113, 131]. The 1064 nm beam can be directly passed into the Ps test setup (for photo-ionization) or, again through OPG and OPA, be converted into the broadband infrared (IR) (tunable between 1670 nm and 1720 nm) beam used for subsequent Rydberg excitation of the Ps [130].

Using the  $n = 3$  as intermediate level, the IR laser wavelength range corresponds to the positronium excitation transitions to Rydberg levels between  $n = 14$  and  $n = 22$ . This can be determined using the equation for the positronium energy levels, Eq. 2.4, and the photon energy-wavelength relation, Eq. 3.15.

$$E = \frac{hc}{\lambda} \quad (3.15)$$

In 2022, the crystal responsible for the IR wavelength generation has been replaced to allow for higher Ps Rydberg levels (which are usable thanks to the new collinear  $\bar{H}$  production, see section Sect. 3.3.4). However, due to a manufacturing error, the lowest possible wavelength is of the order of 1665 nm, corresponding to  $n_{Ps} = 24$ , instead of the planned  $n_{Ps} = 40$ . With  $n_{Ps} = 24$ , the increase in the  $\bar{H}$  production cross section according to Eq. 3.11 is a factor 3.9 with respect to the orthogonal production limited to  $n_{Ps} = 17$ . This can be compared to an increase by a factor 12.6 for  $n_{Ps} = 32$  and a factor 30.7 for  $n_{Ps} = 40$ . An upgrade is ongoing.

It is possible to use the entire EKSPLA laser in the default way with the two components for Ps excitation but inhibit the IR pulse arrival in the experiment through the use of an IR filter, which is installed on a movable flipper that is controlled by a dedicated part of the control system ( $\mu$ Service, see Ch. 4). This feature is regularly exploited to prevent Rydberg excitation of the produced positronium and, thus, antihydrogen formation, when performing background runs, alternating with  $\bar{H}$  production, without modifying the overall procedure.

### The "Alex"

A Q-switched alexandrite-based laser system (729 nm), which has been modified to match the requirements of AEGIS, forms the basis of the second laser setup [86]. Through non-linear frequency conversion (SHG, dual wave plate for polarization rotation, Third Harmonic Generation (THG, mixing of the fundamental and SHG beams)), light at

243.024 nm is produced, corresponding to the resonance wavelength of the Ps  $1^3S-2^3P$  transition. This is the transition which is saturated for the laser-cooling of positronium in AEGIS (see Ch. 7). Especially relevant to the Ps laser cooling are (see Table 3.1) its broadband nature (to cover the broad Ps velocity distribution before cooling), its high intensity and its long pulse duration (to allow for a long interaction time with Ps and multiple cooling cycles). The wavelength can be selected to 10  $\mu\text{m}$  precision, while ensuring the broad spectral bandwidth, thanks to a Volume Bragg Grating (VBG). Furthermore, a connected high-voltage switch can open or close the cavity to allow for the laser pulse generation or to inhibit it with nanosecond precision through a sharp falling edge, enabling a precise definition of the interaction end time to ensure a relaxation of the Ps to the ground state prior to performing the velocity probing.

With the CIRCUS control system (see Ch. 4), the laser timing can be synchronized with nanosecond precision. Furthermore, the settings of the crystals involved in the optical chain can be set without the need for manual re-calibration, and it is possible to exactly calibrate the firing time for each shot, adapted to temperature and humidity variations that would otherwise cause slight wavelength modifications to the sensitive laser, with an active feedback loop included in the system [86, 132].

	EKSPLA			Alex
	Nd:YAG	IR	UV	Ps $1^3S-2^3P$
Wavelength	1064 nm	1665 - 1720 nm	204.9 - 205.2 nm	243 nm
Energy	700 $\mu\text{J}$	1.1 mJ	75 $\mu\text{J}$	2.3 mJ
Pulse length	6 ns	3 ns	1.5 ns	215 ns
Bandwidth	30 GHz	150 GHz	179 GHz	101 GHz
Repetition rate	10 Hz	10 Hz	10 Hz	4 Hz

Table 3.1: Summary of the parameters of the laser systems currently used in AEGIS. The exact energies and bandwidths that are reachable depend on the current state of the beam steering, and the most recent values from the 2023 runs are given here. The values stated for the pulse lengths correspond to the FWHM. Uncertainties of the parameters are not stated here and can be found in [130], [113] and [86]. Where relevant for the work presented here, they are also stated in the corresponding sections of the main text.

The Ps excitation efficiency using the current EKSPLA laser system is mainly limited by its spectral bandwidth and pulse length, i.e. suboptimal overlap of the laser pulse with the positronium atoms. The fundamental pulse of the "Alex" laser (tunable between 728 and 742 nm) can also be used for two-photon resonant process Ps Rydberg excitation (and Ps photoionization for  $n_{\text{Ps}} = 2$ ), covering a larger fraction of the emitted positronium and thereby increasing the excitation efficiency, expectedly from 15 to 30% [83]. Its properties are a bandwidth of 130 GHz, a pulse length of up to 215 ns FWHM and an energy up to 40 mJ. This laser system is installed as a separate beam line of the "Alex", whose commissioning is currently being finalized.

### 3.4.7 Degradors

The 100 keV antiprotons arriving from the AD/ELENA are too highly energetic to be manipulated and trapped by conventional high-voltage electrodes, which can typically reach a maximum of few tens of kilovolts. In order to further reduce their energy before entering the trapping region, they are passed through degraders, i.e. thin foils of material

optimized for  $\bar{p}$  energy loss through collisions with the material's electrons, while minimizing multiple scattering.

To cope with the requirements of the lower-energy antiprotons now provided by ELENA, a dedicated degrader configuration is needed that is carefully optimized with regard to the used materials and thicknesses. In AEGIS, a set of two degrader units has been developed, both of which are shown in Fig. 3.8: The first one is the so-called "main degrader" (photograph on the left), which is located just in front of the antiproton capture trap, inside the 5 T magnet at cryogenic temperatures. This degrader consists of two Mylar<sup>7</sup> foils (1300 and 100 nm) that are (partially) coated with aluminium, which makes it furthermore usable as a beam position monitor, as explained in Sect. 3.4.9. Prior to entering the experiment, the antiprotons already pass through the second degrader unit: A so-called "degrader ladder" (photograph on the right hand side of Fig. 3.8) has been placed just outside the entrance region of the experiment ( $\sim 40$  cm away from the main degrader) on a movable actuator, in a specifically designed vacuum chamber that can be opened (e.g. to replace individual degrader foils) without having to warm up the entire experiment. The ladder accommodates five ultra-thin foils made of Parylene-N with a diameter of 15 mm and varying thicknesses (in steps of 100 nm between 100 and 500 nm), allowing to fine-tune the overall degrader thickness for an optimized antiproton capture efficiency. The foils starting from a thickness of 300 nm are formed by a stack of individual, thinner foils, as detailed in Table 3.2. The thicknesses given here are the ones originally foreseen; however, due to damage from mechanical stress, internal re-organization and unclear statements from the manufacturing company, the actual thicknesses at a given time cannot be determined with certainty. This is briefly discussed in Sect. 6.1.4 and has contributed to the decision taken in 2023 to discontinue the use of the thin foils and rely on the implementation of stronger antiproton capture voltages in order to be able to efficiently use the thick degrader alone.

Degrader thickness	Used foil thicknesses
100 nm	100 nm
200 nm	200 nm
300 nm	100 nm + 200 nm
400 nm	100 nm + 100 nm + 200 nm
500 nm	100 nm + 200 nm + 200 nm

Table 3.2: Composition of the individual foil stacks of the degrader ladder shown on the right hand side of Fig. 3.8.

---

<sup>7</sup>Tests with Parylene-N, which has a similar density and atomic number as Mylar, as the main degrader material have also been performed in AEGIS, but Mylar is chosen as the final candidate thanks to its superior mechanical properties at cryogenic temperatures [83, 84].



Figure 3.8: Photographs of the two degrader units installed in the AEGIS apparatus. Left: Main degrader of Mylar foil and segmented Al layer coating, with amplifier PCB mounted directly in front. Right: Degrader ladder with variable Parylene-N foil thickness.

### 3.4.8 Electromagnetic traps

The Penning-Malmberg traps of the AEGIS apparatus are formed by stacks of in total 66 coaxially arranged cylindrical electrodes with varying lengths between 7.5 and 62 mm and uniform diameters of 30 mm. Being cylindrical, the electrodes are expected to produce potentials that are radially symmetric around the central axis of the experiment. They are separated into two individually constructed units: one in the 5 T region of the experiment, which is used to capture the antiprotons from ELENA and sympathetically cool them, and one in the 1 T part, where the antihydrogen production takes place. A schematic of the combined trap region is shown in Fig. 3.9, with the depicted dimensions to scale and the electrode names and distances labelled.

The different trapping regions are formed by the involved electrodes as listed in Table 3.3 and referred to as "C Trap", "P Trap", and "A Trap" accordingly. The "inlet" electrode is the one located on that entrance side of the trap which is closer to the origin of the arriving particles, while the "outlet" is on the further away side. Both inlet and outlet are also referred to as "endcap" electrodes and can be opened to release the trapped particles. In a typical Penning-Malmberg configuration, the endcaps are at an absolute higher repulsive potential (negative/positive for negatively/positively charged particles) and those electrodes in between, i.e. the "floor", are biased to be more attractive. The C and P Traps are part of the Penning-Malmberg trap inside the 5 T region of the experiment, closer to the entrance region, while the A Trap is located further downstream in the 1 T part, where antihydrogen production takes place. According to specific requirements (e.g. larger or smaller traps, optimized stability or symmetry), other trapping regions can be formed as subsets or combinations of the listed default traps. The T and B electrodes also shown in the schematic serve as transfer electrodes, to move the contained plasmas between the two trap regions. The passage area between them (i.e. between electrodes T6 and B0) is also the location of the transition of the magnetic fields and the defined axial center of the experiment ( $z = 0$ ).

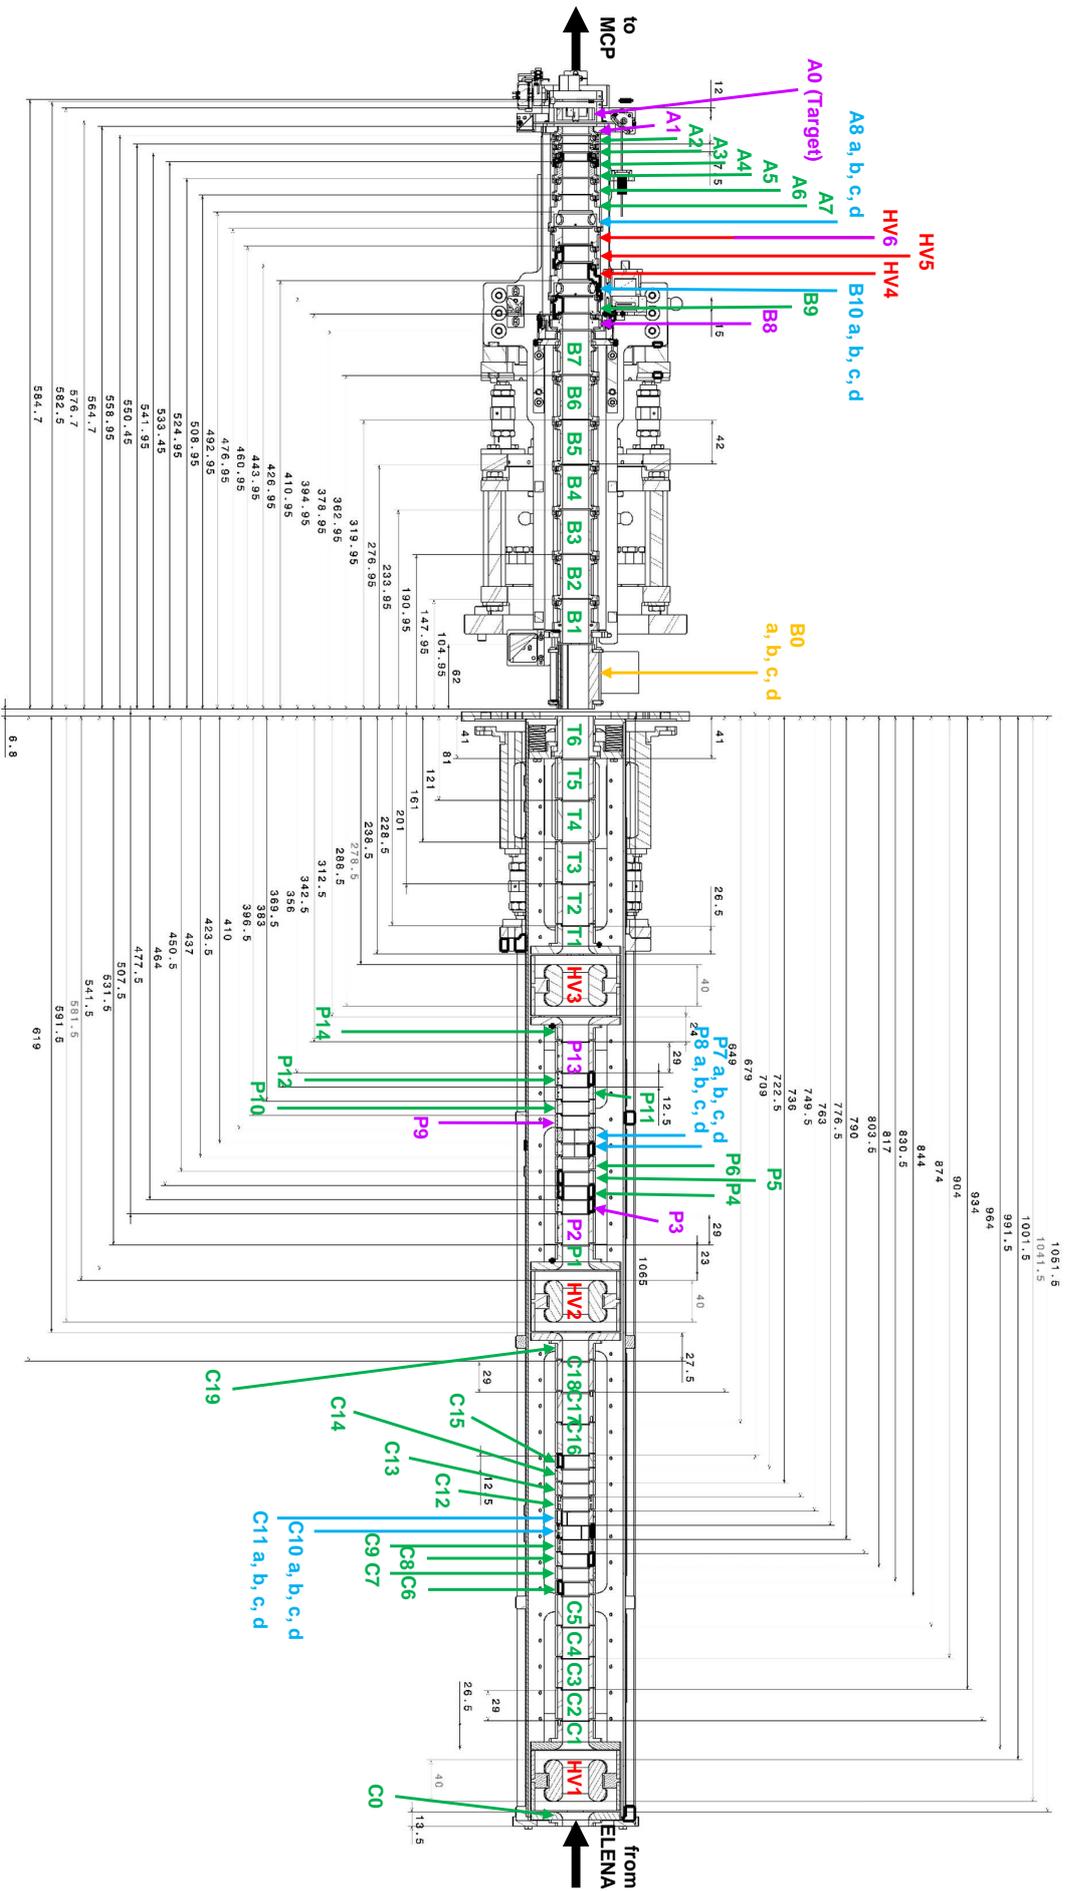


Figure 3.9: Schematic of the Penning-Malmberg trap regions in the AEgIS apparatus. The electrodes are labelled and their lengths and distances (in mm) are stated. The antiprotons enter from the right from the AD/ELENA. The gap between the T and B electrodes marks the transition from the 5 T (towards the ELENA side) to the 1 T (towards the MCP side) magnetic field regions and is considered the center of the experiment. The colors of the electrode labels indicate special functionalities: standard electrodes (green), high-voltage electrodes (red), electrodes connected to a Pulser channel (purple), Rotating Wall electrodes (blue), steering electrode (yellow).

Trap name	Inlet electrode	Floor electrodes	Outlet electrode(s)	Magnetic region
C Trap	C5	C6 - C15	C16	5 T
Long P Trap	P2	P3 - P12	P13	5 T
Short P Trap	P3	P4 - P8	P9 - P11 (*)	5 T
A Trap	A7	A3 - A6	A1, A2	1 T

Table 3.3: List of the default trapping regions of the AEgIS electromagnetic trap system, formed by different electrodes. The inlet, floor and outlet electrode(s) used for each trap are given. Additionally, the magnetic field region the trap is located in is stated. (\*): For logistics reasons, only P9 alone is sometimes used as outlet of the Short P Trap, if it is less important to be sure to prevent spillovers from the too low potential of a single electrode, for example when the other electrodes are used for a different purpose.

The electrodes denoted with the letters a, b, c, and d in Fig. 3.9 (blue labels, except for B0) are four-fold sectorized around their center, i.e. around the central axis of the trap. This enables the application of sinusoidal phase-shifted radio frequency signals on the sectors, facilitating the use of the Rotating Wall (RW) technique for radial plasma manipulation (see Sect. 3.2.2). The base voltage for all sectors is given by the same supply and the RW frequency is added to this (refer to Ch. 4 for details). One exception to this is the B0 electrode (yellow label), whose four sectors have individual voltage supplies, allowing it to be used to steer the particles travelling through the transfer region between the magnets via a bipolar gradient.

Electrodes HV1 to HV6 (red labels) are high-voltage electrodes, which can receive the standard  $\pm 200$  V from the low-voltage control system part but have an alternative connection to a ns high-voltage switch via special cryogenic cables optimized for these high voltages, allowing them to provide potentials of up to  $-15$  kV. These are the electrodes usually used to capture the antiprotons arriving from the decelerator complex.

The A0 electrode incorporates, in its center, the Ps converter target with its holding structure and is installed on a movable actuator that can be moved in and out of the beam line within a few seconds.

Some electrodes (purple labels) are connected to one of the channels of the Pulser (see Sect. 4.1.3), enabling the fast application of short voltage pulses to quickly modify their potentials for specific operations. HV6 is a high-voltage electrode that is pulsable as well (in the low-voltage range).

All electrodes are connected inside the apparatus by cryogenic coaxial cables (limited to 1 m length for the 1 T electrodes to reduce the heat load) and, via air flanges, outside by new 10 m long double-shielded coaxial BNC cables, carefully arranged for easy connection with the apparatus both open and closed and for space efficiency, while complying with the cables' minimum bending radius.

A list of all electrodes included in the AEgIS trap system can be found in Appx. E. Also stated there are their axial locations and lengths. The length of an electrode has a large influence on the potential that it actually produces, given an applied voltage. As a reference, the potential produced by the application of a voltage of  $-150$  V on a single electrode of the P Trap is plotted in Fig. 3.10. The potential is obtained from a finite element model determining the potential along the z-axis with 1 mm binning for a voltage of 1 V applied on a specified electrode, the results of which are saved in a matrix that is convoluted with that of the applied voltages on all electrodes to determine the actual potential. Shown in Fig. 3.10 is a scan over all electrodes from P1 to P14, assuming

for each that all other electrodes are kept at 0 V. The horizontal axis represents the location along the  $z$  axis in the experiment. The dashed black line indicates the  $-150$  V that are applied. It is clear that the maximum amplitude, found at the center of a given electrode, deviates drastically from the applied voltage for short electrodes, e.g. only reaching around  $-80$  V  $\approx 55\%$  for the 12.5 mm electrodes, while a negligible difference is observed for example for the 42 mm B electrodes. The longer an electrode is, the closer the amplitude of the produced potential approaches the set voltage. The achieved voltage fraction is consistent for different absolute values. It is possible to increase the amplitude at a given point by additionally applying a voltage on the neighbouring electrodes. These observations should be kept in mind, since all potential configurations plotted here generally include the influence of the electrode lengths and neighbours, unless otherwise stated. As a side note, Fig. 3.10 also serves as an eye guide for the center locations of the individual P electrodes in the experiment, as they are involved in most of the operations reported on here but not pointed out explicitly any more from here on.

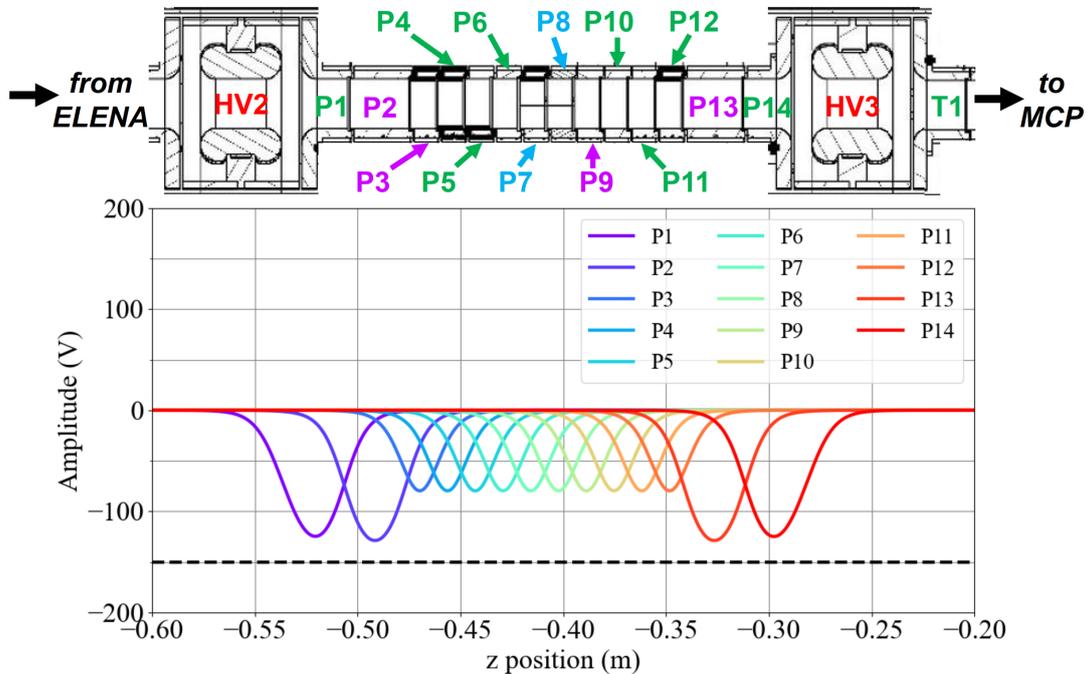


Figure 3.10: Potentials produced by the application of a voltage of  $-150$  V (dashed black line) on a single electrode along the horizontal axis of the experiment. All other electrodes are kept at 0 V. The different colors represent a scan over all electrodes in the P Trap. Above the plot, a schematic of the involved range of trap electrodes is included, scaled to match the dimension of the x-axis, with the colors of the electrode labels having the same meaning as in Fig. 3.9.

In the course of 2022, the trap system of the antihydrogen production region of the experiment has been fully replaced by a newly constructed setup optimized for on-axis  $\bar{\text{H}}$  production, which has been commissioned and fundamentally tested with electrons, as described in Ch. 5. Fig. 3.11 shows a schematic of the design as well as a photograph of the final trap. This trap incorporates, in addition to the cylindrical electrodes, cryogenic actuators to allow for an in-situ alignment, grids for the ionization of excited atoms, the Ps converter target, and optical elements for the positioning and diagnostics of the entering

laser light (a 1300 Thorlabs FG200AEA multi-mode silica fiber bundle, which is led out of the cryostat, where its light is separated into the different laser components by dichroic mirrors and then imaged by cameras). For electrical insulation and to reduce thermal noise, the entire trap is surrounded by a small Faraday cage, and the last eight electrodes, closest to the Ps target, have cryogenic low-pass filter boards installed.

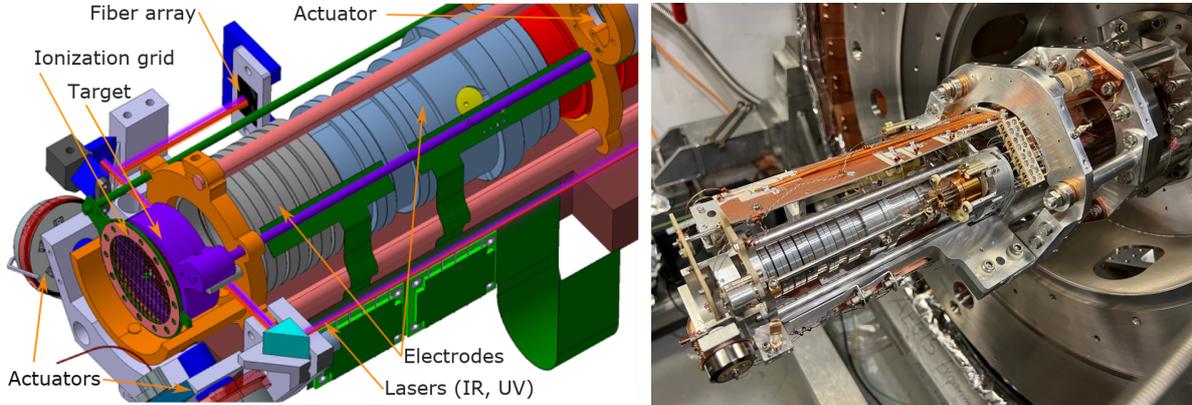


Figure 3.11: Images of the Penning-Malmberg trap installed in the 1 T apparatus of the AEGIS experiment. Left: Schematic showing the different components. Right: Photograph of the trap during installation.

### 3.4.9 Detectors

A set of different detector setups is employed in the AEGIS apparatus, which are each optimized to specific requirements. The different detectors of the main apparatus are briefly described here. Their locations are indicated in Fig. 3.5. Some additional detection devices, which are mainly used for the positron system, are detailed in [121].

#### Micro-channel plate detector

A Micro-channel plate (MCP) stack in chevron configuration (Hamamatsu model F2225-21PGF<sup>8</sup>) is installed at the downstream end of the experiment, i.e. approximately 30 cm behind the electrode holding the positronium conversion target, inside the cryogenic apparatus.

On the output side, the MCP is coupled to a phosphor screen, whose luminescence signal (peak emission wavelength: 545 nm) can be recorded by a CMOS camera (ORCA-Flash4.0 V2, installed in combination with a telescopic magnifying system) located outside the apparatus. In this way, the distribution of incoming charged particles on the MCP can be imaged: Higher recorded pixel intensities correspond to more (same-charge) particles arriving in a given location. This feature is used, for example, to investigate the shapes of different plasmas trapped by the electrodes. Additionally, the phosphor screen signal can also be used as the current source of an RC circuit by adding a resistor and a capacitor (in case of the current AEGIS setup a high-speed NI 5152 oscilloscope with an 8-bit ADC) to filter out the direct current components and provide a high time resolution for the

<sup>8</sup>According to the corresponding technical sheet, the diameter of the effective area of the Hamamatsu F2225-21PGF MCP is 42 mm and the channel diameter is 12  $\mu\text{m}$ . The gain is stated as  $10^6$  and the dark count rate as  $3 \text{ s}^{-1} \text{ cm}^{-2}$  at 1 kV operating voltage. The standard operation voltage is 2 kV.

measurement of passing fast pulses.

The input side of the MCP has a metallized front face and is also used as a separate detector; in fact, it is connected to a relay that allows to either polarize it (currently negatively) by a CAEN high-voltage power supply to guide arriving particles or digitize and record the arriving charge (in the case of AEGIS using a NI 6133 data acquisition device, featuring a 14-bit ADC, used with a 10 k $\Omega$  resistor), essentially using the MCP front face as a Faraday Cup (FC) charge collector [133]. This second option is therefore also referred to as "MCP-FC" detector. It allows to read the DC component of the arriving signal but has no gain element in the readout chain and therefore relatively low sensitivity and is used for the determination of the arriving amount of charge. It is possible to acquire this data simultaneously with the phosphor screen signal described above.

## ESDA

A set of twelve 1 cm thick and  $\sim 1$  m long curved plastic scintillator (EJ200) slabs covers a 120 $^\circ$  arc around the magnets and cryostats, i.e. at a distance of approximately 70 cm to the trap axis, and is referred to as External Scintillating Detector Array (ESDA). Each scintillator has a width of 10 or 20 cm (see Table 3.4) and is coupled to a pair of independent, magnetically shielded Photomultiplier Tubes (PMT, Philips-Photonis XP2020 phototubes and Thorn-EMI 9954B), one at each end. The locations of the scintillator slabs are given in Table 3.4; as indicated in Fig. 3.5, some slabs are installed above the magnets and some below. The scintillators are named according to which PMTs read them out (e.g. "SC12" is read out by PMTs 1 and 2).

Name	$z$ Position	Width	Magnetic region	Vertical location
SC12	-96 cm	20 cm	5 T	Below
SC34	-76 cm	20 cm	5 T	Above
SC56	-45 cm	10 cm	5 T	Above
SC78	-35 cm	10 cm	5 T	Below
SC910	+39 cm	10 cm	1 T	Below
SC1112	+46 cm	10 cm	1 T	Above
SC1314	+59 cm	10 cm	1 T	Below
SC1516	+64 cm	10 cm	1 T	Above
SC1718	+75 cm	10 cm	1 T	Below
SC1920	+79.5 cm	10 cm	1 T	Above
SC2122	+89 cm	10 cm	1 T	Below
SC2324	+96 cm	10 cm	1 T	Above

Table 3.4: List of the axial locations (distances from  $z = 0$ ) and widths of the ESDA scintillator slabs used in AEGIS. The magnetic field region the slabs are located in are also stated as well as information on whether they are installed above or below the magnets.

The ESDA is mainly used to detect pions produced by the annihilation of antiprotons and antihydrogen inside the apparatus, and its signals are recorded in two different ways: A splitter on each PMT splits the signal in half, with one half directly used as discriminated signal to determine the number of counts and monitor continuously the antiproton annihilations [134], and the second half getting ten-fold attenuated<sup>9</sup> and then digitized

<sup>9</sup>The attenuation, using Phillips Scientific model 804, is needed to prevent a saturation of the dynamic range of the digitizer.

(using the 12-bit CAEN V1720 ADC at 250 MHz, 2 V input range, one channel per PMT) for fast detection of the antihydrogen signal.

A calibration of the ESDA system has been performed in the past using cosmic ray muons in a dedicated QDC measurement for an equalization of the PMT gains and a dedicated Monte Carlo simulation to determine the conversion of number of ESDA unit counts to number of annihilating antiprotons, whose results are used here [134, 135].

The acquisition chain of the discriminated signal involves the following elements: CAEN N413, LeCroy 623A, and LeCroy 622 for discrimination, with a 50 mV threshold, LeCroy 622C for coincidence determination, and SIS3820 for counting. The coincidence requirement is needed to avoid the counting of PMT dark counts, whose rate is above 100 Hz [134]. To obtain a clean signal in the discriminated ESDA acquisition chain, it is necessary to subtract the expected background caused by cosmic muons arriving on the scintillators and surrounding radioactive sources with measured rates of 40 to 60 Hz. This procedure is discussed in more detail in Ch. 6.

An individual tuning of the threshold cuts on the PMT coincidence spectra to maximize agreement in the number of count ratios between data and the simulation for the different ESDA units to SC12 has yielded a conversion factor for the number of counts on SC12 to number of antiprotons of  $f_{\text{SC12} \rightarrow \bar{p}} = 16.0 \pm 0.2$ . SC12 is taken as a reference thanks to its favourable solid angle, as its location is closest to the interaction point of the antiprotons when letting them annihilate on the degrader structure, which was done in the simulation and which is also the standard procedure for measuring antiproton numbers in the experiment. This value had been determined prior to the installation of above-mentioned splitters on the PMTs. Since the thresholds have been left as they were even once the splitters were installed, the resulting reduced amplitudes of all hits necessitate an adaptation of this conversion factor, as more hits are filtered out than during the calibration. For this purpose, a dedicated measurement has been performed, leaving the splitter installed on one of the other ESDA units (SC56 shown here) and taking alternating measurements with and without the splitter installed on SC12. The measurement consists of trapping antiprotons in the P Trap for half a minute and then releasing them (hot dump (HD), compare Sect. 3.4.8 and Ch. 6), observing the number of counts within a time window defined by the length of the dumping procedure on the ESDA units. Fig. 3.12 shows the number of observed counts on SC12 for the two sets of measurements plotted against the number of counts on SC56. A clear difference can be observed in the slopes of the data points. The average ratio of counts between the two units is then calculated for each set and the ratio of the ratios is finally determined to yield the calibration factor needed to take into account the reduction of the number of counts on SC12 caused by the splitter installation. This technique allows the determination of the correction factor removed of systematic effects such as the ELENA intensity. The value found in this way is  $f_{\text{splitter}} = 1.321 \pm 0.017$ . A systematic uncertainty of  $\pm 0.016$  is obtained from comparisons of factors obtained by using different reference scintillators (both in the 5 T and the 1 T regions) for the determination of the calibration factor; the statistical error on the mean corresponds to 0.001. When converting the number of counts on SC12 to the number of annihilating antiprotons, the conversion factor previously used without the splitters is multiplied by this correction factor.

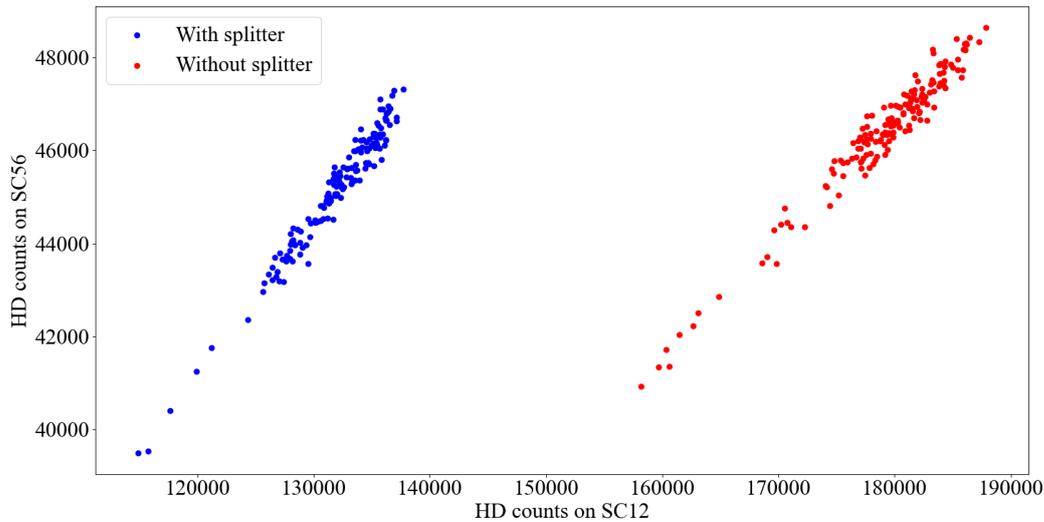


Figure 3.12: Plot of the number of counts observed on SC12 with (blue markers) and without (red markers) a splitter installed on the corresponding PMTs, for a set of measurements releasing captured antiprotons from the P Trap, plotted against the number of counts on SC56. The PMTs of SC56 have the splitters installed in both cases.

Normally, all ESDA units acquire data simultaneously during an experiment and their results can be compared. However, when accumulating large numbers of particles in the traps and dumping them within a relatively short time, the data generated by the annihilations in the ESDA exceed the limitations of the data acquisition system and cannot be reliably stored. For this reason, it is important to intercalibrate the different units so as to be able to use only those ones that are best suited for a given measurement. For example, since SC12 is used as a reference for the conversion to antiproton numbers but it is very close to the degrader structure, on which the antiprotons typically annihilate, and SC1112 closer to the center of the experiment is less prone to saturation effects, the number of counts on the two are intercalibrated so that it is then possible to perform particularly long antiproton accumulation runs with only the use of SC1112 and still convert to the number of antiprotons having been present in the trap. A set of 700 intercalibration runs yields a conversion factor of  $f_{\text{SC1112} \rightarrow \text{SC12}} = 30.04 \pm 0.19$  (mean value and standard error). This is the factor that the number of counts on SC1112 is multiplied by when converting to the number of counts on SC12, for example to determine the absolute number of annihilating antiprotons in Ch. 6. The same set of runs gives an intercalibration factor of  $f_{\text{SC56} \rightarrow \text{SC12}} = 3.22 \pm 0.02$  for SC56, which is briefly used in Sect. 6.1.5 as well.

For the digitized ESDA acquisition chain, a hit is recorded if the two PMTs of one scintillator produce signals in coincidence (typically within 50 ns) with an amplitude above a given threshold, and the average of their signal amplitudes is used. This procedure is needed to avoid the background from dark counts of the PMTs. In practice, to discriminate  $\bar{p}$  and  $\bar{H}$  events (mainly charged pions) from the background induced by the positrons inserted for positronium formation (photons releasing electrons through Compton scattering or the photoelectric effect, i.e. particles with significantly lower energies), a software amplitude cut of the order of 250 to 350 mV is introduced, assuming, below saturation, a direct proportionality between the energy released in the scintillators, their

light yield and the PMT signal [68, 135]. This value is motivated by a Monte Carlo simulation of antiprotons annihilating on the AEGIS trap walls and a comparison to data acquired with and without antiprotons in the trap. A plot of the corresponding results is shown in Fig. 3.13. The main plot compares the distribution of signal amplitudes of the trapped antiprotons (red line) to that originating from positrons, shortly after their implantation in the Ps target (blue line). The inset plot shows the simulated data from  $\bar{p}$  annihilations. Further details on the procedures involved in the analysis of the digitized ESDA signals are given in Ch. 7.

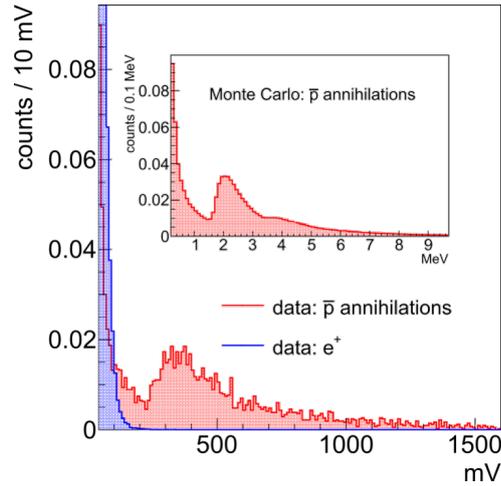


Figure 3.13: Experimental and simulated ESDA amplitude spectra used for the determination of useful threshold cuts. In the main plot, spectra from trapped antiproton annihilations (red) and  $e^+$  signals following their implantation in the Ps target (blue) are compared. The inset plot displays the results from a Monte Carlo simulation of the  $\bar{p}$  signal [68].

### Beam position monitor/Beam counter

The NI 6133 device used to digitize the charge arriving on the MCP front face is also employed to digitize the signals obtained from the main degrader beam position monitor alluded to previously: The 100 keV antiprotons delivered to experiments from ELENA cannot be monitored by silicon pixel detectors with thicknesses of tens of micrometers (contrarily to the higher-energy direct AD antiprotons). The AEGIS main degrader therefore has a dedicated design that still allows to record the  $\bar{p}$  arrival position and intensity: It is separated into two layers of foils, the first (1500 nm) one of which is coated with a 10 nm aluminium layer that is segmented into four pads with a non-coated 1 mm cross in the center, while the second (100 nm) foil is fully Al-coated. The segmented coating can be used as beam position monitor: Any deviation of the incoming beam from the center causes a detectable signal in the corresponding pad. The full homogeneous aluminium layer, on the other hand, can act as a beam counter to record the intensity of the incoming antiprotons. At the same time, this layer can also be used as a Faraday Cup charge detector for other particles passing through or being trapped in the experiment. The signals from these foils are ten-fold amplified by cryogenic amplifiers on the PCB directly behind (see Fig. 3.8), and both the amplified and the un-amplified signals are passed outside the experiment and read out.

### 3.5 Antihydrogen gravity measurement

The gravity measurement envisaged by the AEGIS collaboration is based on the observation of the vertical deflection of antihydrogen atoms travelling horizontally over a known distance in a precisely determined time. For this purpose, the  $\bar{\text{H}}$  atoms are passed through a set of gratings (matter or optical), defining possible parabolic trajectories, and end up on a detector with a high spacial resolution [109].

The gravity module, including the gratings and the detector, is currently being developed, to be finalized in 2024/2025.

#### 3.5.1 The gravity module

Fig. 3.14 shows the original design idea for such a module, a device based on a Talbot-Lau interferometer in a classical regime without interference: the Moiré deflectometer [136–139]. A Moiré deflectometer consists of two parallel grating structures and a position-sensitive detector installed at a known axial distance  $L$  from each other. The gratings filter out a narrow range of trajectories from an originally divergent atomic beam and the detector records the atomic density modulation, i.e. the fringe pattern, to determine the vertical displacement  $\Delta y$  due to the influence of gravity. From  $\Delta y$ , the gravitational acceleration  $a_g$  can be deduced according to Eq. 3.16, with  $\Delta t$  as the atoms' time of flight between the two gratings (and, equivalently, between the second grating and the detector plane) over the distance  $L$ . In the absence of a detector with a high enough resolution, a third grating, which is vertically movable, can be used as an alternative in combination with a simple counting detector, albeit necessitating much longer data acquisition periods (at different vertical settings) to yield the same sensitivity.

$$\Delta y = a_g \Delta t^2 \quad (3.16)$$

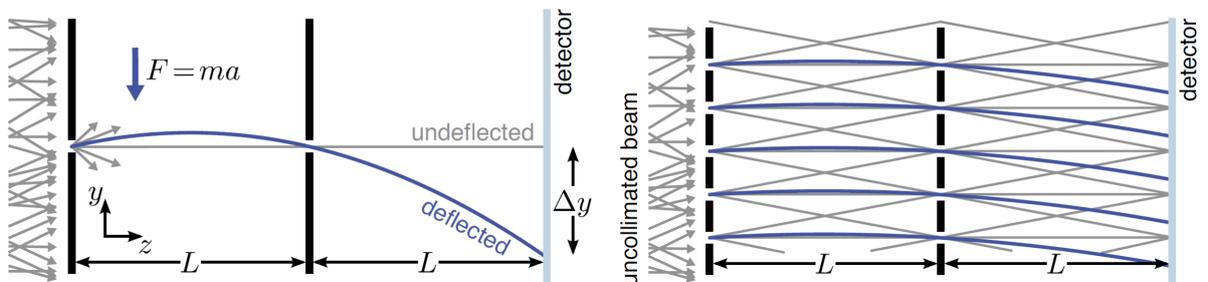


Figure 3.14: Schematic functionality of a Moiré deflectometer, as designed for the gravity measurement of AEGIS [138]. Not to scale. Two gratings and one position-sensitive detector are placed at an axial distance  $L$ . A set of atom trajectories (dark blue parabolic lines) from the initially divergent antihydrogen beam is selected by the two gratings and ends up on the detector plane, creating a fringe pattern to determine their vertical displacement  $\Delta y$  due to the influence of gravity. The grey lines represent trajectories not affected by gravity for comparison. Left: Outline of the technique, showing a single grating opening and only one trajectory each. Right: Multiple openings in the grating lead to the creation of a fringe pattern on the detector from the trajectories.

A reference measurement for a determination of the particle arrival positions without a gravitational deflection is planned using a photon beam in the same device, i.e. essentially

employing the deflectometer as a Talbot-Lau interferometer. A discussion of Talbot-Lau interferometry and first corresponding measurements in AEGIS, both with light and with antiprotons in a Moiré deflectometer prototype, can be found in [138]. Fig. 3.15 shows the resulting signal obtained in the prototype, using silicon gratings with a  $40\ \mu\text{m}$  periodicity and a  $12\ \mu\text{m}$  opening together with an emulsion detector [138, 140]. The measurements were performed in the 1 T magnetic field region of AEGIS. On the left hand side, the patterns produced by a light source (red/white fringe pattern) and annihilations from antiprotons being led through the deflectometer (blue dots) are shown in a restricted range. The intensity distributions obtained from these patterns are plotted on the right hand side, together with curves representing the distributions expected from a simulation. The observed shift of  $9.8\ \mu\text{m}$  between the two distributions corresponds to a force acting on the antiprotons. Both the uncertainties of the result itself ( $6.4\ \mu\text{m}$ , mainly due to the required software alignment) and the origin of the force (due to the surrounding magnetic field and in general stray electromagnetic fields) remain very large: This campaign constitutes the first ever interferometry experiment performed with antimatter but does not yet represent a gravity measurement.

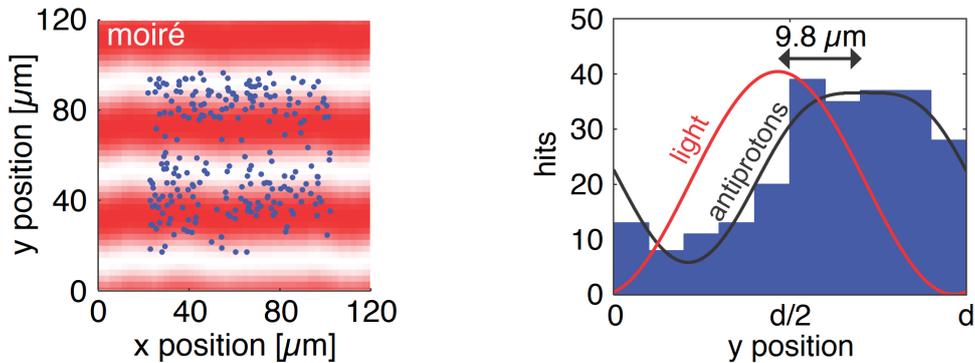


Figure 3.15: Results of light and antiproton measurements in a Moiré deflectometer prototype [138]. Left: The patterns produced by both particle types (red/white fringe pattern from photons, blue dots from antiproton annihilations). Right: The resulting intensity distributions for both measurements, with a shift due to a force acting on the antiprotons but not on the light.

While progress is being made towards a de-excitation of Rydberg antihydrogen [141], the  $\bar{\text{H}}$  atoms currently available for a gravity measurement are not in the ground state. Therefore, such studies have to be performed in very homogeneous magnetic field regions to limit the influence of the magnetic dipole force, which can mimic the deflective influence of gravity already at gradients of a few  $\mu\text{T}/\text{cm}$ , depending on the excitation level of the atoms [121]. To guarantee this condition and for practical reasons related to the AEGIS setup, a horizontal distance of approximately 30 to 35 cm is available for the gravity module, i.e. axial grating distances  $L$  of the order of 15 cm are possible.

Reasonable values for achievable axial velocities of the  $\bar{\text{H}}$  atoms in the current experiment, allowing for a sufficient axial acceleration with achievable  $\bar{\text{H}}$  temperatures, are of the order of several hundred to a few thousand  $\text{m s}^{-2}$  (see Sect. 7.4.3), leading to a time of flight in the deflectometer of a few 100  $\mu\text{s}$ , given above considerations concerning the available space. According to Eq. 3.16, for the expected gravitational acceleration  $g = 9.8\ \text{m s}^{-2}$ , this would produce a vertical displacement of below  $1\ \mu\text{m}$ , beyond the resolution of most

realistic detectors.

Currently, developments are ongoing to perform the gravity measurement outside the magnetic field regions of AEGIS, i.e. in an intrinsically homogeneous field, through the extension of the vacuum apparatus behind the downstream MCP setup by an additional tube that holds the gravity module (see Fig. 3.16, right hand side, where an extension to the left of the apparatus is envisioned). Axial grating distances  $L$  of the order of 50 cm or more will thus be possible (and a longer distance from the  $\bar{H}$  source to the first grating required), allowing the atoms to travel long enough for an expected deflection of several micrometers. This development requires the redesign of the detectors in that area and the installation of an additional vacuum flange on the main apparatus as well as a focusing apparatus for those  $\bar{H}$  atoms that are affected by magnetic lensing when leaving the magnetic field region.

Studies regarding a setup involving, instead of matter grids, laser light gratings in an optical time-domain ionizing matter wave (OTIMA) interferometer have also been performed [142,143]. Optical gratings allow for a smaller periodicity (down to below 100 nm), limiting the phase shift and thus giving better resolution for the same grating distance, and enable high-precision ( $\mu\text{m}$ ) placement of the grating with ns-precise laser pulses. However, due to the needed timing overlap of the laser pulse and the  $\bar{H}$  atoms, the needed flux of cold enough atoms for a sensitivity to the influence of gravity of 10% or below (e.g. 1100 atoms at 1 K in every production cycle for a 6-month measurement campaign) is not currently feasible.

The AEGIS collaboration is thus currently focusing on the construction of a Moiré deflectometer, which has an adjustable grating distance and is rotatable around the axis of the experiment by  $90^\circ$  to allow for systematic control studies also in the horizontal plane. Fig. 3.16 shows a preliminary design of the gravity module and its planned installation location in the apparatus. The illustration shown on the left hand side incorporates three gratings with a dimension of  $40 \times 40 \text{ mm}^2$  and a thickness of  $150 \mu\text{m}$ . The grating periodicity is  $100 \mu\text{m}$  and the opening size is  $40 \mu\text{m}$ . The gratings are mounted on rails to facilitate a variation of  $L$  and the entire structure is in turn installed on two rotating frames. As shown in the schematic on the right hand side, the deflectometer is planned to be attached to the downstream end of the experimental apparatus, with the MCP detector removed and a new flange added, extending the vacuum tube. An iron coil could act as a shield of the magnetic field gradients. This location outside of the magnetic field region implies a distance between the  $\bar{H}$  source and the gravity module of the order of 1 m. Three scintillation detectors are envisaged on the outside of the module to monitor  $\bar{H}$  annihilations on the gratings and the detector. Instead of a third grating, this schematic includes the currently favoured design with the time- and position-sensitive detector at the end of the deflectometer.

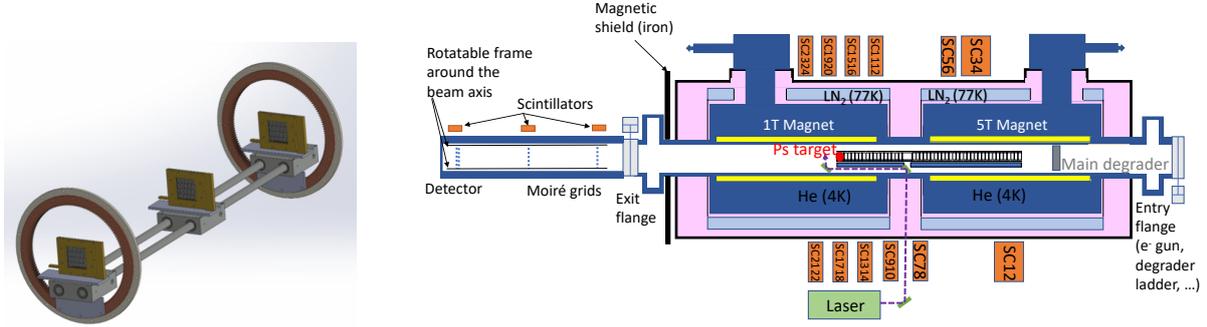


Figure 3.16: Schematic illustration of the preliminary design of the AEGIS gravity module. Left: Technical sketch of the Moiré deflectometer/Talbot-Lau interferometer. Right: Schematic of the main cryostat parts of the AEGIS apparatus, incorporating the gravity module on the downstream end. Not to scale.

### 3.5.2 Sensitivity expectations

The minimal detectable acceleration  $a_{g,\min}$  depends on a variety of experimental parameters, as given in Eq. 3.17 [136, 138, 144]. Here,  $d$  denotes the grating periodicity.  $\nu$  is the visibility of the periodic fringes and is defined as the ratio of the difference between maximum and minimum of the obtained intensity modulation to their sum. It depends on the geometrical features of the module, including the open fraction  $\eta$  of the gratings,  $d$ , and the coherence in space of the stream of incoming particles. Example values of  $\nu$  obtained from simulations are given in [138] and [112]. A good compromise for an incoherent beam ranges around values of 0.3 for  $\eta$  and 0.8 for  $\nu$ , reaching minimal values of  $a_{g,\min}$ .  $N_{\text{det}}$  represents the number of detected atoms and  $N_{\text{det}}$  can be estimated from the number of originally produced atoms  $N_{\text{prod}}$  as given in Eq. 3.18, taking into account the relevant contributions: From considerations of the  $\bar{H}$  beam divergence and the geometry of the deflectometer, the grating radius and, equivalently, the detector radius  $r$  together with the distance between the antihydrogen source and the first grating  $L_{1\text{st}}$ , the grating distance  $L$  (as introduced above) and  $\eta$  determine the geometrical acceptance. The term  $\beta = v_{\parallel}/v_{\perp}$ , where  $v_{\parallel}$  and  $v_{\perp}$  denote the mean longitudinal and transversal velocities, takes into account the influence of the beam boosted in the direction of the gravity module (with respect to an isotropic production).  $\beta$  is referred to as the boost factor. Plugging Eq. 3.18 into Eq. 3.17, the final expression for  $a_{g,\min}$  can be obtained [138]. Since the atoms are only axially accelerated,  $v_{\perp}$  is assumed to be dominated by thermal velocity with a Maxwell-Boltzmann distribution, i.e.  $v_{\perp} = \sqrt{k_{\text{B}}T/m}$  (see Eq. 7.18 and Sect. 7.4.3).  $k_{\text{B}}$  is the Boltzmann constant,  $T$  the temperature of the  $\bar{H}$  atoms and  $m$  the mass of one antiproton. The flight time  $\Delta t$  can also be expressed in terms of  $L$  and  $v_{\parallel}$ ,  $\Delta t = L/v_{\parallel}$ . The factors in Eq. 3.17 are grouped according to the experimental aspect they are associated with.

$$a_{g,\min} = \frac{d}{2\pi\nu\Delta t^2\sqrt{N_{\text{det}}}} = \underbrace{\frac{d}{2\pi\nu\eta r}}_{\text{gratings}} \cdot \underbrace{\frac{L_{1\text{st}} + 2L}{L^2}}_{\text{geometry}} \cdot \underbrace{\beta}_{\text{boost}} \cdot \underbrace{\frac{k_{\text{B}}T}{m\sqrt{N_{\text{prod}}}}}_{\bar{H} \text{ source}} \quad (3.17)$$

$$N_{\text{det}} = \frac{r^2}{(L_{1st} + 2L)^2} \frac{v_{\parallel}^2}{v_{\perp}^2} \eta^2 N_{\text{prod}} \quad (3.18)$$

The comparison of Eq. 3.17 and Eq. 3.18 visualizes the need for a compromise:  $\beta$  contributes significantly (quadratically) to the number of detected  $\bar{\text{H}}$  atoms, which makes sense intuitively: The more atoms are boosted toward the gravity module, the more of them are detected. However, at the same time, larger  $\beta$  values also increase the minimal detectable acceleration as a result of the reduced vertical displacement of the atoms according to their higher velocity. Depending on the desired precision of the gravity measurement,  $\beta$  should therefore be chosen as low as possible to guarantee the required statistics. Similar considerations also apply for the grid distance  $L$ , which limits the number of atoms reaching the detector (due to solid angle losses) for larger values but at the same time enables bigger  $\Delta y$ .  $L$  should therefore be chosen as large as possible, while maintaining a sufficient  $\bar{\text{H}}$  flux.

Of course, for some parameters the optimization capabilities are apparent: For example, the grating and detector radius  $r$  only contributes to the number of detected atoms and should be maximized (which is mainly limited by cost and available space), while the grating periodicity  $d$  enters only the reachable sensitivity and should be minimized (constrained by manufacturing options). Importantly, the  $\bar{\text{H}}$  temperature should be kept as low as possible, as it enables equal  $\beta$  (and thus  $N_{\text{det}}$ ) values, while reducing at the same time  $a_{\text{g,min}}$ . Trivially, if the flux of produced  $\bar{\text{H}}$  atoms is high enough to yield sufficient  $N_{\text{det}}$  anyway, an increased  $\beta$  and reduced  $L$  can be omitted in favour of better sensitivity. Obviously, to achieve a 1% precision on the gravity measurement,  $a_{\text{g,min}}$  should be below  $9.8 \times 10^{-2} \text{ m s}^{-2}$ . Table 3.5 summarizes the realistic experimental parameters and available gratings outlined above for the gravity module under construction. Using a boost factor  $\beta \approx 3$  and assuming  $\bar{\text{H}}$  temperatures of the order of 50 K, under these circumstances, even a minimal sensitivity of the order of  $g$  itself,  $a_{\text{g,min}}/g \approx 1$ , would require the detection of approximately 1000  $\bar{\text{H}}$  atoms. In reality, this number and the resulting statistics are planned to be accumulated over a few months of measurement time during ELENA  $\bar{\text{p}}$  provision. Obviously, with the number of  $\bar{\text{H}}$  atoms produced in AEgIS Phase I, such undertakings are impossible [68] and it is vital to increase the flux of produced antihydrogen by several orders of magnitude to 10 or more atoms per production cycle, keeping their temperature as low as possible. Given the upgrades performed on AEgIS for Phase II, such numbers are becoming realistic. A detailed discussion on the expected improved number of  $\bar{\text{H}}$  atoms can be found in Sect. 7.4.

If it is true that the  $N_{\text{prod}}$  flux of antihydrogen can be improved to 50-100 atoms per production cycle (see Sect. 7.4.2), this number would still necessitate a continuous operation and data acquisition for three to six months straight, around the clock every day of the week. In theory, this is certainly possible within the AD/ELENA campaign times. However, since many of the assumed parameters (notably  $\bar{\text{H}}$  numbers, temperature and acceleration) are still being studied in detail, systematics (e.g. detection efficiency, support structure effects) are not taken into account in these estimations and experimental procedures hardly ever work flawlessly, especially on the first attempt, this may appear quite ambitious. Limiting the forward boost to  $\beta \approx 2$  has a slightly beneficial influence on  $a_{\text{g,min}}$  (only a few  $\text{m s}^{-2}$ ) but drastically reduces  $N_{\text{det}}$  for these parameters of the  $\bar{\text{H}}$  source, approaching the limit for which above estimations are valid (a few hundred detected atoms [138]).

A future upgrade of the module is planned to use gratings with 20  $\mu\text{m}$  periodicity and 50 mm radius, and it is expected that a reduction of the  $\bar{\text{H}}$  temperature to below 10 or even 1 K will be feasible. Assuming the same  $\bar{\text{H}}$  flux and measurement time,  $a_{\text{g,min}} \approx 0.098 \text{ m s}^{-2}$ ,

i.e.  $a_{g,\min}/g \approx 0.01$ , is finally achievable already at temperatures of 10 K with  $\beta \approx 2$ . For lower temperatures, the same boost factor can produce similar results already in less time with much fewer  $\bar{\text{H}}$  atoms (e.g. 1.5 months at 5 K) or yield even better resolution.

Overall, it is clear that investing in a reduction of the  $\bar{\text{H}}$  source temperature seems lucrative, especially if the production flux cannot be further increased. This should be feasible, as antiproton temperatures of a few Kelvin (basically liquid helium temperature) or even below are routinely achieved in the cryogenic environments of antimatter and plasma physics experiments through sympathetic cooling [145–147]. Using electron cooling in combination with resistive or laser cooling can allow to reduce their temperature down to 100 mK [68, 103, 109, 148]. AE $\bar{\text{g}}$ IS is also investigating the possibility of sympathetic cooling of antiprotons with Doppler-laser-cooled ions, with  $\text{C}_2^-$  as a promising candidate, which also opens the door to sub-Kelvin temperatures [85, 149]. With a temperature of 1 K and  $\beta \approx 5$  (i.e.  $v_{\parallel} \approx 450 \text{ m s}^{-1}$ ), considering the flux of 50  $\bar{\text{H}}$  atoms produced per cycle, almost 1400 atoms could be detected within two weeks (four weeks assuming a detection efficiency of only 50%), yielding the desired 1% resolution on the  $g$  measurement.

Denotation	Parameter	Current value	Possible upgrade
$d$	Grating periodicity	100 $\mu\text{m}$	20 $\mu\text{m}$
$\nu$	Visibility	0.8	–
$\eta$	Grating open fraction	0.4	–
$r$	Grating and detector radius	20 mm	50 mm
$L_{1st}$	Distance from $\bar{\text{H}}$ source to first grating	1 m	–
$L$	Grating distance	0.5 m	–

Table 3.5: Summary of the parameters of the Moiré deflectometer currently under construction in AE $\bar{\text{g}}$ IS. Possible future upgrades are included in the last column.

### 3.5.3 Detector options

Thanks to the pulsed production modality of AE $\bar{\text{g}}$ IS, the instance of formation of the  $\bar{\text{H}}$  atoms is well defined (within a few hundred ns). Therefore, knowledge of their arrival time on the detector plane also yields their time of flight  $\Delta t$ . This arrival time can be precisely determined in a straightforward way through their annihilation products, for example using the Fast Annihilation Cryogenic Tracking (FACT) detector of AE $\bar{\text{g}}$ IS [150], which is currently being refurbished. FACT is formed by scintillating fibers that are arranged in two concentric double-layer cylinders around the  $\bar{\text{H}}$  production trap and are read out by arrays of silicon photomultipliers.

The second component for determining the influence of gravity on the antihydrogen atoms according to Eq. 3.16 is their vertical deflection in space  $\Delta y$ .

Recently, a high-resolution (1-1.2  $\mu\text{m}$  expected spacial resolution) detector has been developed for the purpose of imaging low-energy antiparticles, which is based on an array of CMOS camera sensors (Sony IMX219, 8 megapixels,  $(1.12 \times 1.12) \mu\text{m}^2$  pixel size,  $(3.67 \times 2.76) \text{mm}^2$  sensitive area) used in commercial smartphones. A prototype of a detector using one such sensor has been successfully tested with a positron beam in 2023 to achieve a resolution of 0.97  $\mu\text{m}$  [151]. It is sensitive both to charged particles and to photons.

The same prototype has also been installed in the AE $\bar{\text{g}}$ IS setup during the antiproton campaign of 2023 [85, 87]. A routine developed using the AERIALIST control system

on a dedicated branch (see Ch. 4) together with a newly installed upstream extraction beam line [85] has enabled the ejection of captured and confined antiprotons "backwards" from the 5 T trap towards the entry region, in front of which the detector has recorded the annihilation events of the arriving  $\bar{p}$ . These development runs were performed in alternation with the antihydrogen production during the last days of available antiprotons from ELENA and were therefore strongly limited in statistics; however, over 2500 useful events were recorded, yielding an overall spacial resolution of 1-2  $\mu\text{m}$  through the reconstruction of the annihilation vertices [87]. A schematic of the detector assembly as well as a selection of antiproton annihilation events imaged on it are shown in Fig. 3.17. The next step is the assembly of the array structure (around 50 sensors, arranged in a compact rectangular tessellation, yielding a sensitive area of the order of  $5.8 \times 5.7 \text{ cm}^2$  and an expected detection efficiency of approximately 56%) and readout chain to cover the entire plane required for the detection of  $\bar{\text{H}}$  annihilations during the gravity measurement.

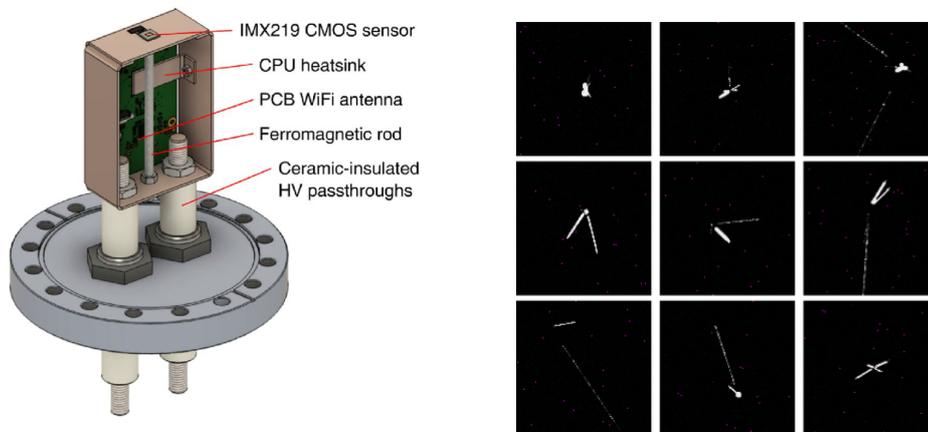


Figure 3.17: A prototype high-resolution detector for reconstruction of antiproton annihilations. Left: The flange and assembly holding the CMOS sensor, which has been attached to the upstream end of AEgIS to detect antiproton annihilations. Right: Selection of antiproton annihilation events imaged by the sensor.

An alternative detector design based on nuclear emulsions has also achieved a resolution of the order of 1  $\mu\text{m}$  [140] and would be applicable in the AEgIS setup. However, contrary to the camera sensor device, it is unable to provide real-time feedback on experimental results and is therefore not currently further pursued.

One option for a hybrid imaging/timing detector has been developed within the AEgIS collaboration recently as well, based on the combination of an MCP with a TimePix3 pixel detector [152]. This detector, while achieving a timing resolution of the order of few tens of nanoseconds, has reached a maximum spatial resolution of 12  $\mu\text{m}$  in test runs with positron beams. Some possible upgrades to make such a detector utilizable in antihydrogen (or positronium) gravity measurements with the needed precision are also being evaluated.

# Chapter 4

## The CIRCUS control system

*This chapter introduces the new experimental control system, CIRCUS, that has been developed to manage the AEGIS experiment with all its subsystems. This system is reported on in [153], of which I am the corresponding author. As stated in the authors' contributions, M. Volponi and I are the main contributors to the building up, commissioning, and maintenance of the control system: My focus includes the acquisition, assembly, installation, calibration, and operation of the hardware and electronics as well as the direct programmatic control of these, i.e. the software library infrastructure written in ARTIQ, responsible for the ns-precise synchronization and fast control elements of the experiment. Furthermore, I have developed the procedure for and conducted the calibration of the amplifier units. While I have also developed, and taken part in the development of, some of the components of the TALOS LabVIEW™ architecture (notably the Electron Gun  $\mu$ Service), the main parts of this work are done by M. Volponi and J. Zieliński. I am a contributing author of the corresponding article, [132].*

In order to be able to operate an experiment as complex as AEGIS, which requires micro- or even nanosecond synchronicity for particular timing aspects and involves various intricate subsystems, a powerful control system is required. This system should be flexible enough to allow for an integration of a diverse set of subsystems as well as easy modifications and extensions, and able to reliably and reproducibly control the different aspects of the experiment, for extended periods of time but at the same time with a very precise and stable timing resolution, limiting the amount of needed human intervention.

The previous control system used by the AEGIS collaboration was in large part homemade, and, due to the fluctuations of the responsible members and the unavailability of replacements for specific components common in research collaborations, had become rather difficult to maintain. It has therefore been replaced by a more compact, open-source system, whose architecture is formed by Sinara hardware using the ARTIQ (Advanced Real-Time Infrastructure for Quantum Physics) control software for communication [154, 155], embedded in the TALOS (Total Automation of LabVIEW™ Operations for Science) software framework, which has also been developed in AEGIS. This new control system, CIRCUS (Computer Interface for Reliably Controlling, in an Unsupervised manner, Scientific experiments), has a strong focus on autonomy, giving the possibility to run in a completely unsupervised manner for long periods [153]. In addition to consisting of reliable, modular electronics optimized for limited available spaces, Sinara is able to provide precise control of time-critical, intricate experimental processes and allows for an easy integration into the new overall software control system.

The objective of this part of the work presented here has been the development of the

fast-control unit, involving software and hardware, with precise synchronization capabilities to steer the time-critical aspects of an experiment, as well as the interconnection with the slow-control architecture to form a broad, powerful control system. This chapter gives an overview over CIRCUS, its main functionalities and its performance. The hardware of the AE $\bar{g}$ IS control system and the software to directly program it form the AERIALIST (Antimatter Experiment (or AE $\bar{g}$ IS) Realtime Integration of Artiq Libraries and Sinara Technology) part of CIRCUS and are described in detail below. The AERIALIST has also been linked as a central component with TALOS, which coordinates all integral parts of the experiment on a broader, less time-critical level and is briefly introduced here as well. Details on the LabVIEW<sup>TM</sup> architecture can be found in [153], [132] and [156].

## 4.1 Control system hardware

At the core of the control system electronics is hardware from the Sinara portfolio, which is used for the control of the trap electrodes as well as most relevant triggers. Instances of the Sinara core controller, Kasli, are thus responsible for the entire fast synchronization of all subsystems of AE $\bar{g}$ IS. This system has been built up between 2021 and 2023 and is now running reliably. In addition, some components of the previous trap control system have also been modified and integrated. Most of this hardware is now cleanly organized in an electronics rack in the experimental area.

The power cables of many of the hardware components of the control system are connected to a power switch, which facilitates remote rebooting without access to the experimental area<sup>1</sup>.

### 4.1.1 Sinara hardware overview

The Sinara hardware of the new AE $\bar{g}$ IS control system is assembled in compact  $50 \times 20 \times 35 \text{ cm}^3$  rack-standard Eurocard 84 HP electronics crates, an example of which is shown in Fig. 4.1. Each crate contains one main carrier, nicknamed Kasli, which comprises an Artix-7 FPGA (Field Programmable Gate Array) and can control up to twelve of the extension modules of the Sinara portfolio, each designed for a specific purpose. Since a precise synchronization of the experimental subsystems is vital to an efficient trapping and manipulation of the particle plasmas inside AE $\bar{g}$ IS, a high-speed communication between the devices is indispensable. This is facilitated by the use of the DRTIO (Distributed Real Time Input/Output) protocol, enabling Gbit/s information transfer.

Kasli can be used either as a stand-alone core device or as a satellite or repeater of the DRTIO communication when combined with additional carrier devices. In the second case, both the control sequences and the synchronization clock signal (which can be generated internally or adopted from an external source) are quickly transferred via optical fibres from the master to one or more satellite devices. This option also offers the possibility for straightforward extensions and adaptations of the experimental setup.

In addition to the power supply connector, Kasli's front face also comprises a JTAG microUSB connector for flashing of the FPGA gateway and for reading the log output as well as an SMA connector used to provide an external clock signal if needed, and four SFP

---

<sup>1</sup>The possibility for remote rebooting of key devices is both extremely useful during ongoing measurement campaigns when access to the experiment is restricted for radiation safety reasons and to achieve a higher degree of automation for the control system in general.

ports to receive the programming via an Ethernet connection<sup>2</sup> (port 0) and to transmit the clock signal and communication via optical transceivers to satellite Kaslis (ports 1-3). Internally, the connection to the extension modules is realized via 30-pin IDC ribbon cables. A schematic of the Kasli connectors and the typical setup used in AEGIS can be found in Appx. F.

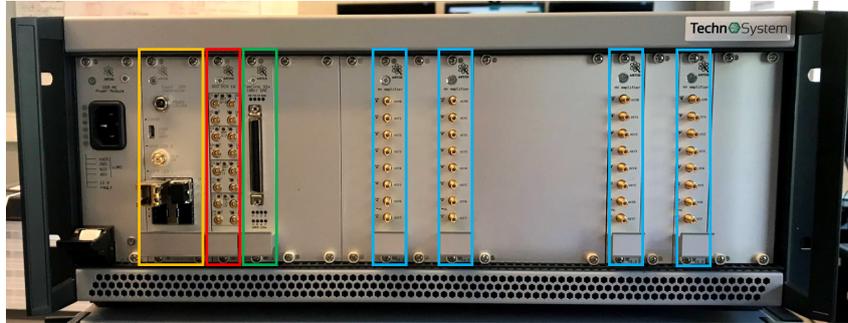


Figure 4.1: Fully equipped electronics crate of the AEGIS experiment composed of Sinara hardware. From left to right: power module, Kasli carrier (yellow), DIO module (red), Fastino DAC (green), HV amplifier boards (blue).

Fig. 4.1 shows the Kasli controller (bordered in yellow) as well as the typical Sinara extension modules used for AEGIS.

The module marked in red is the Digital Input/Output (DIO) unit, which features 16 MCX channels able to send and receive TTL signals with a precision below 1 ns. It is possible to individually configure these channels in batches of four to act as "input" (receiving incoming TTL signals) or "output" (sending out TTL signals). In AEGIS, these units are used for the trigger signals synchronizing the experimental subsystems (laser system, positron system, electron gun, detectors, etc.) among each other as well as with external procedures such as the antiprotons being ejected from ELENA towards the experiment and the data acquisition system. Essentially, they are thus responsible for the timing of the entire experiment. Over fifty such trigger connections have been realized in AEGIS.

The green border marks the Digital to Analog Converter (DAC) unit, referred to as Fastino. It provides output voltages on 32 channels individually in the range between  $-10\text{ V}$  and  $+10\text{ V}$  with a 16-bit precision, i.e. one step between two adjacent programmable voltage settings (so-called machine unit) corresponds to a step of approximately  $0.3\text{ mV}$ , with the exact value depending on the individual configuration of the channels. The voltages are provided simultaneously to the 68-pin SCSI front face output and, in batches of eight channels, to four IDC connections on the board. In AEGIS, Fastinos are mainly employed to provide the potentials for the over sixty electrodes of the Penning-Malmberg traps.

Since voltages of the order of up to  $\pm 190\text{ V}$  are used in AEGIS for an efficient trapping and manipulation of the confined particles and plasmas, the Fastino channels are used in combination with custom-designed amplifier boards (marked in blue in Fig. 4.1) with

<sup>2</sup>The communication to Kasli via Ethernet needs a Gigabit connection, for example via a network switch. In AEGIS, the use of such a switch, which is connected to an external power switch, has the added benefit of allowing for a remote reboot of the connection when needed, e.g. after a clock failure due to a power cut.

a 1 MHz bandwidth, a circuit schematic of which is included in Appx. G. One of these boards comprises eight amplifier channels, each benefiting from an individual OptoMOS isolation to prevent noise propagation and yielding a voltage amplification by a factor of 20, thus reaching a range of  $\pm 200$  V for the final output when connected to a Fastino channel, as done in AEGIS. Each electrode used in the AEGIS traps is directly, or through the Pulser or Rotating Wall electronics described below, connected to one such amplifier channel.

The voltage on the electrodes can thereby be ramped up (or down) simultaneously within microseconds, with the ramping time increasing with the exact voltage difference to be achieved. The plot on the left hand side of Fig. 4.2 gives an indication of these times. Plotted is the rise time within which the voltage settles to the desired value as a function of the voltage difference  $\Delta V$  to be covered. The shown data points are the mean values and corresponding standard errors on the mean of two reference measurements on one Fastino/amplifier channel, ramping to a given voltage from 0 V, both in positive and negative direction, as measured by an oscilloscope. This measurement is meant only to give a broad indication of the ramping speed and does not represent a systematic, precise statement. Due to a characteristic steep rise and slow flattening of the voltage ramp, an offset rise time of several microseconds is already observed for small voltage differences. From around  $\Delta V = 10$  V, the increase in the rise time per volt to be covered is approximately linear,  $\sim 0.35$   $\mu\text{s}/\text{V}$ . These findings are in good agreement with the rise times observed on any of the channels that have been randomly selected for verification tests. Example oscilloscope images of such voltage ramps can for example be found in Sect. 4.5 and in Appx. G.

Four amplifier boards are integrated in each of the AEGIS Sinara crates, thus providing amplification for all 32 channels of the one Fastino installed per crate. The noise of the channels is of the order of 5 mV, corresponding to an equivalent noise temperature of around 50 K. However, the individual OptoMOS relays of each channel facilitate an output disconnect within microseconds, giving the possibility to reduce the equivalent noise temperature to a few Kelvin. For this to be useful in practice, it is important to note that the applied voltage remains for an extended amount of time after disconnecting, with an RC time constant  $\tau$  of the order of several minutes. As a verification measurement, a voltage  $V_0$  has been applied on one of the channels before disconnecting the output and waiting a certain amount of time  $\Delta t$ . Then, using a probe and an oscilloscope, the amplitude  $V$  of the voltage left is measured. The result is plotted on the right hand side of Fig. 4.2 for  $\Delta t$  between 0 and 300 s. Each data point is the mean value obtained from three independent measurements, with the standard error on the mean. Of course, this manual readout method is subject to relatively large systematic uncertainties and this measurement is only used as a verification, not as a precision study. The lowest  $\tau$  that can be determined from this measurement, using Eq. 4.1, is of the order of 5 minutes, for  $\Delta t = 10$  s. Although this value is easily enough for any of the complex operations involved in antihydrogen production, the amplitudes still observed for longer  $\Delta t$  suggest a significantly larger  $\tau$  of up to 20 minutes.

Since the amplifier boards have been custom-designed for AEGIS, a few issues have been observed during the first operation tests, which have since been solved and are reported on in Appx. G as well.

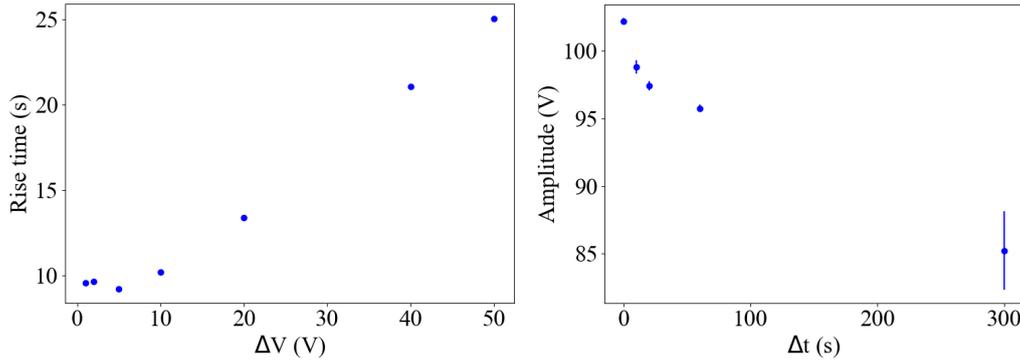


Figure 4.2: Rough verification measurements of amplifier characteristics. Left: Scan over the rise time as function of the voltage difference to be covered. Right: Scan over the voltage left on an amplifier channel with the output disconnected as a function of the time passed before the evaluation.

$$\tau = \frac{-\Delta t}{\ln\left(\frac{V}{V_0}\right)}. \quad (4.1)$$

The Sinara hardware responsible for the core control of the AEGIS experiment is formed by three of the crates depicted above: one ("1TC1") used in stand-alone mode and two ("5TC1", "5TC2") set up in a master/satellite configuration.

While the Fastino/amplifier channels of the stand-alone crate (32 channels) provide the voltages for the electrodes of the antihydrogen production trap (1 T region of the experiment), those of the two 5 T crates (in total 68 channels) are responsible for the potentials on the electrodes in the antiproton capture trap. Each set includes at least ten spare channels which can be readily used by the swap of one cable and the modification of one line of code in case of failure of individual channels. The electrodes in the transfer region between the two are split between both for logistics reasons. One channel of both the 1 T and the 5 T crates can also provide the voltage of the electron gun extractor introduced further below, and the taking charge of one or the other is managed by a relay controlled by one of the DIO output channels of the 1 T unit (see Sect. 5.1). Additionally, one such amplified DAC channel can also be used to provide a low bias voltage to the front face of the MCP installed at the downstream end of the experiment without employing the high-voltage system.

As for the DIO boards, there are in total two units, i.e. 32 TTL channels, per cluster of crates (two in the 1 T stand-alone crate and one each in the 5 T master and satellite). Some of these channels have the same purpose on both crates, e.g. the instruction for the control software to start a new run can be given from either one, and most detectors as well as the Pulser can be triggered from both. Contrarily, the external triggers from the AD/ELENA relevant for adapting the timing of the experiment accordingly are handled only by the DIO channels of the 5 T crate because its Fastino/amplifier channels are also responsible for capturing and trapping the arriving antiprotons. On the other hand, the 1 T DIO channels give the triggers for the laser and positron system, for example for the combination of the antiprotons with Ps in the 1 T trapping region. Both units also include channels to send back and forth TTL signals between the 5 T and the 1 T side to coordinate and synchronize the experiment for those routines involving both, such as antihydrogen production.

The Sinara portfolio also offers an option for the DDS-based frequency synthetization

needed for the use of the Rotating Wall technique, which has been successfully tested in AEGIS and is described in section Sect. 4.1.4.

### 4.1.2 Using the external clock

All time-critical operations in the AEGIS experiment are synchronized to an external clock signal with a frequency of 10 MHz, which is provided by a National Instruments™ PXI-6682 Timing Module. The output clock signal of this device is a square wave that has a peak-to-peak amplitude  $V_{pp}$  of approximately 3.0 to 3.3 V. However, while the Sinara Kasli allows for the use of an external clock with the given frequency and a sine or square wave signal and adapts its internal synchronization accordingly, the input connection for this external clock is optimized for use with a  $V_{pp}$  signal between 500 and 600 mV. On the one hand, continuous operation with a too high input voltage can eventually cause damage to the electronics, while on the other hand, a too low input signal can destabilize the system and lead to an unreliable timing in the experimental procedures. For this reason, it is important to attenuate the voltage of the clock provision circuit appropriately.

#### Home-made voltage dividers

Fig. 4.3 shows a simple circuit diagram of a resistive voltage divider. The output voltage is reduced according to Eq. 4.2 by a factor depending on the value of the resistors  $R_1$  and  $R_2$ . For designing such a component in reality, one has to take into account both the output impedance of the clock provision (here:  $50\ \Omega$ ), which forms part of  $R_1$ , and the input impedance of the external clock connection of the receiving device (here:  $50\ \Omega$ ), which serves as  $R_2$ . In order to reduce the provided clock amplitude to the ideal  $V_{pp}$ , a factor of  $1/6$  is needed, which, with  $R_2$  fixed to  $50\ \Omega$ , translates into a required  $R_1$  value of  $250\ \Omega$ . Since the output impedance of the clock module already amounts to  $50\ \Omega$  as well, additional resistors with a value of  $200\ \Omega$  soldered into BNC boxes are installed in series in the clock provision lines of all used Sinara crates.

$$V_{out} = \frac{R_2}{R_1 + R_2} \cdot V_{in} \quad (4.2)$$

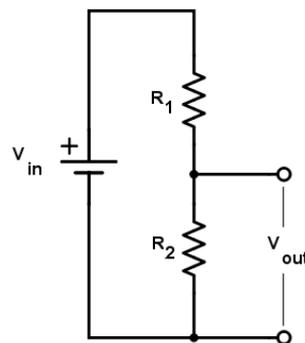


Figure 4.3: Simple electronics circuit of a resistive voltage divider.

In Fig. 4.4, oscilloscope images are shown of the clock signals going to the Kaslis with and without the addition of the  $200\ \Omega$  resistor. Both square functions have a frequency of 10 MHz but the additional resistor reduces the amplitude by  $1/3$ . As the output impedance of the clock ( $50\ \Omega$ ) and the input impedance of the receiving device (here:

the oscilloscope, with  $50\ \Omega$  termination) already form a voltage divider themselves, where  $R_1 = R_2$ , the visible amplitude in the image on the left is already approximately halved according to Eq. 4.2 compared to the provided value.

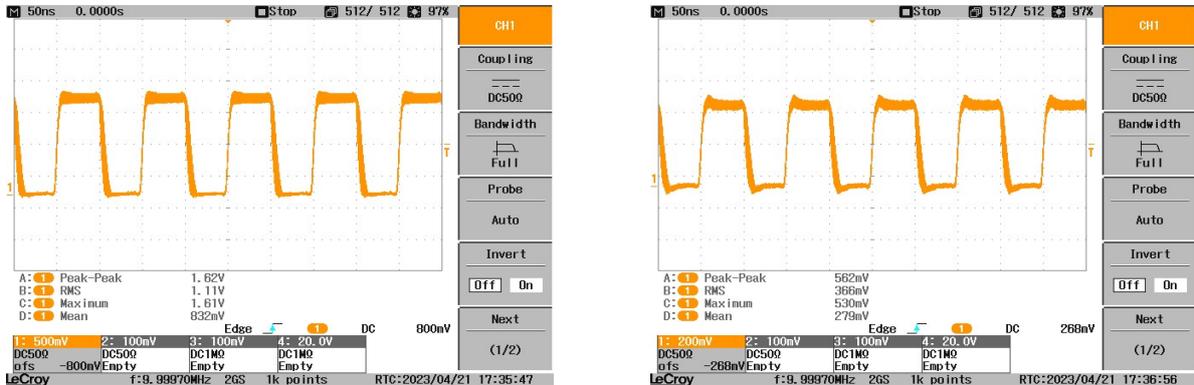


Figure 4.4: Oscilloscope views of the external clock input signals to the Sinara Kaslis without and with installed voltage dividers. Left: Without additional external resistors. Right: With the addition of a  $200\ \Omega$  resistor in series.

The installation of such simple voltage dividers can solve the acute issue of a too high input clock voltage and, for the most part, yields a reasonably stable operation. However, as can be seen in Fig. 4.4, the additional resistor not only decreases the amplitude of the signal but furthermore slightly modifies its shape, which also varies somewhat in every cycle. It can be observed that, when running experiments that need to be synchronized to the nanosecond over several minutes, hours or days (no systematic problematic amount of time has been obvious), individual clock signals sometimes do not seem to be registered by the Kasli because TTL triggers are skipped and the system eventually loses synchronization fully. Rebooting the system generally solves this issue so that the run can be restarted and goes on without problems (for a somewhat random amount of time) until the next de-synchronization event. Since one of the main features of the control system is the possibility to leave it to run the experiment in an unsupervised way, especially during nights and weekends, this issue becomes quite important. In addition, some measurement campaigns in AEgIS and, certainly, other experiments, require up to several days of continuous data acquisition without interruption.

## Standard electronics

Having identified the issue described above, a more sophisticated solution has eventually been implemented: the clock signals for the different Kaslis are now not passed through home-made voltage dividers but through an industrially produced  $50\ \Omega$  attenuator module (NIM Model 804 from Phillips Scientific), whose attenuation level can be adapted to match the needs of the Kasli connection. As depicted in the oscilloscope image in Fig. 4.5, the orange line, which shows the signal after attenuation, closely resembles the pink line (no attenuation) in shape and offers a well reduced amplitude of approximately 600 mV, as needed.

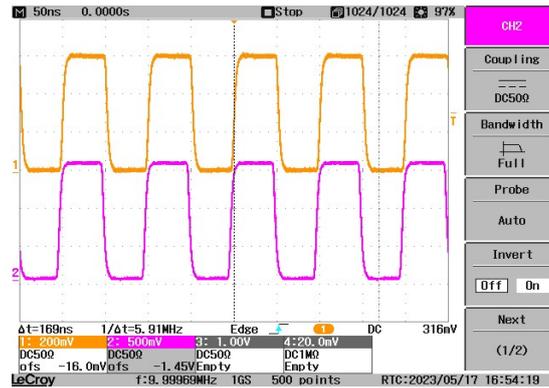


Figure 4.5: Oscilloscope image of the external clock input signals to the Sinara Kaslis without (pink line) and with (orange line) attenuators included in the circuit.

Having implemented this solution for the most recent antiproton campaign of AEGIS, the Sinara system has not once lost synchronization in this way over the entire time of six months of being run. Some reboots have been necessary for other reasons (e.g. power cuts, amplifier failures and replacements) but the synchronization has been maintained over the entire active time.

### 4.1.3 The AEGIS Pulser

A few of the electronics of the original AEGIS trap control system continue to be used, as they are serving relevant purposes while a future-proof replacement is being devised and are ideally suited for the use in the current experiments. These components have been modified slightly to be compatible with and have then been integrated into the new electronics setup of the control system.

A major such component is the AEGIS Pulser. The Pulser can provide fast voltage pulses in the amplitude range between  $-200$  and  $200$  V to up to nine trap electrodes, connected through filter boxes, via internal DAC units. These pulses can be synchronized on the nanosecond level and have a variable length between 10 and 30 000 ns. The pulse delay as well as the number of given pulses of each individual channel can also be varied as needed. The rise time of the pulses is of the order of a few to a few tens of nanoseconds, depending on the amplitude.

This feature has proved useful for various purposes in the experimental routines but is most prominently used to quickly move and transfer confined particles and plasmas from one of the employed Penning-Malmberg traps to another or simply remove a subset of them from a trap. Such procedures are vital for the antihydrogen production in AEGIS. The trigger signal for the Pulser to generate the pre-defined pulses on all channels is now provided by dedicated channels of the Sinara digital I/O units: both the 5 T and the 1 T Sinara units have full control over the Pulser triggers. This is particularly important when performing experiments involving both sides and furthermore yields additional flexibility for the utilization of the Pulser. The Pulser also has a dedicated TALOS  $\mu$ Service (see Sect. 4.4) and the settings of its different channels are steered by the AERIALIST code.

### 4.1.4 Rotating Wall electronics

As described in Sect. 3.2.2, it is possible to influence the radial extension of plasmas in Penning traps with the Rotating Wall (RW) technique, i.e. the application of rotating

electric fields on the confining electrodes. The electronics used in AEGIS to implement this technique are also part of the control system and are described here briefly.

### The custom RW crate

The second main component remaining from the original control system is a custom-made waveform synthesizer capable of adding phase-shifted sinusoidal signals of up to 5 V in a frequency range of 0 to 30 MHz on up to eight so-called sectorised electrodes. These electrodes themselves, of which six are integrated in the AEGIS traps, are separated into four sectors around their centre, i.e. around the central axis of the trap. The base voltage to be produced on them is, within the new system, supplied to the RW crate by Sinara Fastino/amplifier channels. The RW electronics, when instructed by the new control system, add the sinusoidal phase shift on each sector as needed, creating a rotating electric field perpendicular to the axis of symmetry of the contained plasma and thus producing a torque on the particles to manipulate the radial dimension.

### Sinara Urukuls

As an alternative to the custom RW crate and to future-proof the use of the Rotating Wall technique in AEGIS, a Sinara crate including DDS-based frequency synthesizers from the Sinara portfolio, the so-called Urukuls, has been commissioned and tested to provide the phased sinusoidal voltages as required.

The output frequency of the Urukuls can be tuned between 1 and 4000 MHz with a precision of 0.25 Hz, corresponding to a 32-bit resolution.

One Urukul board comprises four radio frequency output channels (SMA connectors), all of which can be controlled individually with respect to the frequency, amplitude, and phase provided.

## 4.2 Hardware communication

The Sinara hardware is optimized for use with the ARTIQ (Advanced Real-Time Infrastructure for Quantum physics) control software framework [155]. A library system built up on the basis of ARTIQ is used, as part of the AERIALIST, to directly control the used Kaslis and thus the connected core electronics of the control system.

### 4.2.1 ARTIQ

ARTIQ comprises a high-level programming language, which is based on Python and features a library of pre-generated, specialised functions for communicating control elements to the hardware. A device database file is needed for working with ARTIQ, which is essentially a list (Python dictionary) of all available devices as well as their corresponding drivers, controllers and connectors, and is stored in the memory of the used core device. With ARTIQ, procedures can be executed on the Kasli FPGAs and passed to the connected hardware with nanosecond precision on the timing and latency of the order of a few microseconds. These procedures, the so-called kernel, constitute the time-critical parts of the experimental routine codes and can be interfaced with more general Python code run on the host computer using Remote Procedure Calls (RPC).

## RTIO concepts

The kernel functions work based on Real Time Input or Output (RTIO) events that are to be executed on the hardware. These events make up a timeline which is used to guarantee the synchronization of all procedures involved in an experiment: An array of First In-First Out (FIFO) buffers (eight parallel lanes per default) holds the timestamps and data of when output events (e.g. a sent trigger pulse or the provision of voltages on Fastino channels) are scheduled and when input events (e.g. an incoming trigger pulse) happen according to the cycles of a provided (internal or external) clock. Essentially, a cursor is moved along the timeline according to the duration of scheduled events. When events occur with timestamps that are in the past with regard to the current clock value, i.e. timestamps that do not leave enough time for all other events required to have finished beforehand, or when the submitted cursor times cannot be strictly increasing, errors are thrown and the corresponding code is not executed<sup>3</sup>. When multiple kernel functions are invoked consecutively, the timeline cursor is maintained, allowing to cluster functionalities in useful blocks. Pre-defined functions also exist to force the kernel to wait for all previous events to finish executing and to clear the FIFO buffers.

### 4.2.2 Kasli Direct Memory Access

In some cases, it is necessary to perform a sequence of operations on the hardware in a very fast and synchronized way that would not be feasible in the normal operation mode. One example is the ramping of a voltage on multiple electrodes with nanosecond synchronicity between them and less than a microsecond between the ramping steps. This kind of procedure poses two problems because the setting of a voltage on a Fastino channel is an operation that takes time, i.e. advances the event cursor position on the timeline: Firstly, this means that before being able to set a voltage on another channel, a certain time delay is needed before the hardware can perform the task. This causes asynchronicities between the channels of the order of microseconds. Secondly, even on a single channel, a certain delay has to be introduced between the steps (around 4  $\mu$ s for one channel but increasing to hundreds of microseconds for multiple channels) when going from one voltage to another in multiple steps, as is needed in some cases as a hardware precaution or to minimize the disturbance of the particles or plasmas contained in the traps. This limits the flexibility and leads to unwanted delays, which can cause losses or prevent an efficient operation of the experiment. The first issue, when working with only one ramping step on multiple channels, can be overcome by a built-in functionality of ARTIQ: It is possible to put the entire Fastino output "on hold", then program all needed voltages to its respective channels without actually supplying them, and finally issue an update request to the Fastino for all involved channels. In this way, the voltages can be supplied on all channels with nanosecond synchronicity. The second issue remains,

---

<sup>3</sup>When an event is placed on the timeline where the current time cursor has not arrived after the execution of the existing events, i.e. the cursor is in the past with regard to the desired event timestamp, an "RTIO Underflow Exception" is generated.

Instead, if a FIFO holding input events becomes over-filled because the CPU has not read them out (with corresponding user functions) quickly enough, an "RTIO Overflow Exception" happens.

Events submitted with a decreasing or equal timestamp to the existing ones are first moved to the next of the available parallel FIFO lanes. Once such issues arise for all buffers, a "Sequence Error" is thrown. If more than one event is submitted on the same hardware channel with the same timestamp, a "Collision Error" is given.

"Busy Errors" correspond to the inability of a given channel to execute an event because it is already busy with a previous execution.

as several microseconds are still needed before the Fastino can again change its output voltages. The solution is a feature called Direct Memory Access (DMA). DMA allows to store event sequences directly in the SDRAM of the used Kasli. These pre-defined sequences can then be played back by the FPGA at much higher speeds than is possible for the CPU (e.g. nanoseconds between changes of electrode potentials instead of up to tens of microseconds without DMA). This feature is heavily exploited in AEgIS for the generation of the low-potential reshaping operations detailed in chapters Ch. 5 through Ch. 7.

### 4.2.3 Gateware

The gateware of an FPGA refers to the bitstream carrying the information on the configuration of its logic gates. In the presented setup, using ARTIQ and the Sinara hardware, it can be generated and compiled according to the specific needs of a setup in the form of binary image files and flashed into the FPGA board using the JTAG connector on Kasli's front face. A restart of the device is subsequently required. In this way, the setup can be modified according to changing needs of the experiment.

Generally, a JSON file is created which contains the information on the core device version used, whether it acts as master or satellite device, a variety of parameters including for example whether an external clock is used and which frequency is needed there, and the use of the peripheral connections, i.e. which extension modules are connected to which connectors of the main board. A typical example JSON file used for the gateware of a master Kasli used in AEgIS is shown in Appx. F. From this file, the gateware image is compiled, which can then be flashed into the Kasli.

### 4.2.4 The AEgIS library system

Every sequence of procedures to be performed by the hardware is written in the form of an ARTIQ/Python script that consists of a class with, at minimum, a `build` and a `run` function definition. Both of these functions are executed when the command `artiq_run <SCRIPT.py>` is issued within an ARTIQ environment, for example in a computer shell with ARTIQ activated. The `build` function defines all hardware components to be used in the given sequence of procedures (also referred to as "experiment" from here on) and allows the master core device to instantiate the corresponding drivers by interfacing with the device database. It is also possible to define additional class attributes and perform further preparations, such as the loading of configuration files, in `build`. The `run` function, then, defines the sequence of operations that are needed to perform the given task. This should, in most cases, start with a reset of the core device (clearing also the RTIO buffers) and an initialization and configuration of the hardware as needed (initialization of the DAC, setting the input/output direction of trigger channels, putting all channels in the expected state, etc.). Subsequently follows a list of function calls that are to be executed either on the core device or on the host, or both.

In principle, it would be possible, for every single experimental script (i.e. code block defining a continuous routine in the experiment) that is created, to define the `build` and `run` functions from scratch, using only the basic available methods. However, for experiments that are more complicated than for example the simple switching on and off of an LED, such scripts would become extremely unhandy, lengthy and more and more unreadable. Additionally, for most experiments, certain procedures, such as triggering and reading out a detector, are needed in various contexts and it would be quite unprofitable

to re-define these sequences every time.

It is no problem, when working in the ARTIQ environment, to define additional functions, which can, essentially, become infinitely complicated and long, and to determine whether these functions are to be executed on the kernel (by using the `@kernel` decorator to translate into machine code) or on the host (default). These functions can be defined inside the given experiment class or imported/inherited. For code executed on the host, the entire range of Python, including extensions, is of course available. On the other hand, the kernel code is somewhat limited in order to maintain timing efficiency: for example, the use of dictionaries and empty lists is not possible in kernel functions.

As a general rule, it is possible to call kernel functions from those run on the host and vice versa (via RPCs). However, when interlacing the two too much, an error is thrown because the sequential calls `host -> kernel -> host -> kernel` are not supported. For this reason, the `run` functions of all experiment scripts in AEGIS are defined on the host, with the used functions that need to be run on the kernel individually declared to do so.

### Class libraries

For above reasons and thanks to the possibilities of custom function definitions and multiple inheritance, a library system has been created to organize and standardize the used code and to avoid unnecessary duplication. The entirety of this library system is maintained as a git repository (called `kasli-code`), which enables version control and allows for an easy parallel development and usage of functionalities for different experimental programs, having proved extremely useful for example during the highly charged ions and precision detector campaigns in 2023 (see Sect. 6.5 and Sect. 3.5). A schematic of its structure is shown in Fig. 4.6. Each library defines a Python class with dedicated function definitions governing specific parts of the experimental procedures. All of these classes are used as parent classes of one main experiment class (here named `AEGIS Class`), which collectively imports from and inherits all of them. Each individual experiment script is then created as a child of this main class and consequentially has all functionality of all library classes available for use in its `build` functions and `run` routines. Some of the classes also import others, as indicated in the diagram, because their own function definitions need functions already defined higher up to perform their tasks. The majority of the base class (`TCP Library`, `Build & Init Library`, `Error Library`) functionality and many of the `Utility Library` and `Analysis Library` functions are experiment-agnostic and could be used in the same way - with different configuration files and some adaptations - for other experiments (and, more generally, any kind of task), while the `Trap Library` and the `Physics Libraries` contain functionality specific to the AEGIS apparatus and program:

- **TCP Library:** This class contains the functionality for interfacing with the TALOS LabVIEW™ framework via TCP tunnels and communication with the various VIs ("`µServices`", see Sect. 4.4). Details are given in section Sect. 4.4.1. As explained there, many of the functionalities of the other libraries that involve non-Sinara hardware rely on the `TCP Library` to communicate the needed action to the corresponding LabVIEW™ code which actually transfers it to the device. This possibility gives great flexibility to the overall system and is at the heart of connecting TALOS and the AERIALIST to form the core of the CIRCUS control system.
- **Build & Init Library:** The main purpose of this library is the provision of standardized `build` functions for the different units of Sinara electronics used (1 T crate,

full 5T system, 5TC1 crate in stand-alone mode), including all relevant hardware as well as the integration of configuration information, channel mappings and global attributes with default values (most of which are generated by dedicated functions also included in this library class). This is also where the relevant calibration data of the amplifier channels is loaded and lists of the used channels (DAC/amplifier and DIO) are created to be accessible in the following (often time-critical) experiments without delay. Instead of listing all these functionalities individually in the `build` function of an experiment script, these individual `build` functions thus contain only one line<sup>4</sup> corresponding to a call to one of these standardized functions. Furthermore, this library contains standardized initialization functions for each of the systems, the corresponding one of which can generally be included as the first line in the `run` function<sup>5</sup> instead of a list of individual steps, as well as general reset functions, for example to bring all DAC channels to 0V or to enable or disable all amplifiers.

- **Error Library:** This library takes care of the error handling, both for kernel and for host functions, and thereby contributes extensively to the stability of the control system: Unsupervised operation is only feasible if the system is resistant to occurring errors. When using the functionality of this class, a criticality code is assigned to each error, which is propagated to TALOS to evaluate how to proceed. For details on this, see Sect. 4.4. The functions of this class can be included at any point in the other class functions or directly in scripts where error handling is needed, which is why the other library classes all import the `Error Library` if needed.
- **Utility Library:** This library contains a large range of general functionality that is needed for many of the experimental routines. Aside from the configuration and triggering of detector electronics, the propagation of actuator movements to the corresponding `μServices`, and the preparation of relays as well as various subsystem components, among others, this also includes the main functions for starting and stopping experimental runs (called at the beginning and end of each experiment script, respectively), intricately involving the data acquisition (DAQ) system and TALOS. Furthermore, functions handling the timing, such as delays and the synchronization between multiple core devices are also defined here.
- **Analysis Library:** This library has been created for the purpose of defining functions needed for a quick analysis of incoming data, which can be used to return feedback to the overall system that is then applied to modify parameters in subsequent runs for optimized conditions. Such functionality gives the capability of self-optimization to the control system and renders it more autonomous. Recently, the data analysis done in AEGIS has been centralized in the broader ALPACA (All Python Analysis Code of AEGIS) framework, which is also introduced in [153] and [132].
- **Trap Library:** This class combines all the functionality to control the over 60 electrodes of the AEGIS apparatus providing the trapping potentials for the different involved particles and plasmas. For this reason, functions handling the determination of channels mappings and the application of calibration constants, using the

<sup>4</sup>If specific configuration steps are exceptionally needed in addition, corresponding functions can of course simply be added to the individual `build` function.

<sup>5</sup>Due to an issue concerning the initialization of the Fastino, this part is handled slightly differently. See below for details.

loaded data from the configuration files, are also included here. Furthermore, templates for setting specific voltages on specific electrodes are provided, as well as functionality to set up an entire potential trap shaped as chosen by any number of electrodes and quickly (or slowly, if needed) reshape the trap from one potential configuration to another. To guarantee synchronicity between the electrodes, the DMA functionality described above is used here extensively. Additionally, the control elements of the Pulser are also implemented in this library.

- **Physics Libraries:** These are a set of four class libraries dedicated to handling specialized parts of the overall experiment which are sometimes finally combined to perform more complex experiments, such as antihydrogen production, but are also frequently used individually, for example for positronium physics studies. These functionalities have been separated to keep the library files to a manageable length and maintain a good overview, and because they differ drastically in their nature such that it would not be beneficial to keep them together.

The **Plasma Library** takes care of all procedures involved in the provision and trapping of electrons from the electron gun. These can be individual measurements involving only electrons or their preparation for antiproton cooling. Furthermore, specific instances of the trap potential reshaping and electrode pulsing functions provided in the **Trap Library** are defined here and used in different combinations, together with detector and hardware functions from the **Utility Library** to manipulate plasmas of different particles as required and collect corresponding data. Over fifty reshaping operations between trap configurations are already defined here. Reshapings can be done fast (i.e. immediately according to the limits of the electronics) to be time-efficient or slowly (using a ramping procedure that goes from one voltage to the next on all electrodes in a defined time and number of steps) for operations that require care in not disturbing the trapped plasmas to prevent losses. As in the **Pbar Library** and **Ps Library**, defined functions are often blocks of more basic functions whose actions need to happen in the exact same order with the exact same timing for different measurements.

The **Pbar Library** itself handles all functionality to bring antiprotons from the AD/ELENA into AEGIS. This includes the control of the high-voltage capture electrodes, communication of beam steering parameters to the corresponding  $\mu$ Services, control of all precise timing triggers involved in the  $\bar{p}$  capture, and antiproton arrival detection, among the most prominent functionalities.

The **Ps Library**, on the other hand, hosts the functions governing the control of the positron and laser setups of the apparatus, mainly used to accumulate and bring positrons into the apparatus to form positronium on the conversion target, to excite the formed atoms, and to control the relevant detectors. This library is also the basis for the experiment runs for laser cooling of positronium [116]. An upgrade of the Sinara system is planned to include an additional crate dedicated to the positron system, which is currently still mostly controlled by custom electronics and integrated into the overall system via Sinara TTL lines.

The **Ion Library** has been created specifically for an independent branch of research in AEGIS: the formation of highly charged ions from the interaction of incoming antiprotons with the gas present inside the apparatus (and, in the future, antiprotonic atom formation via co-trapping). Details on these measurements can be found in Sect. 6.5. These measurements require special trap reshaping operations that are implemented here, combined with the needed detector functionality into function

blocks that differ from those used in the rest of the experiment and are therefore developed independently.

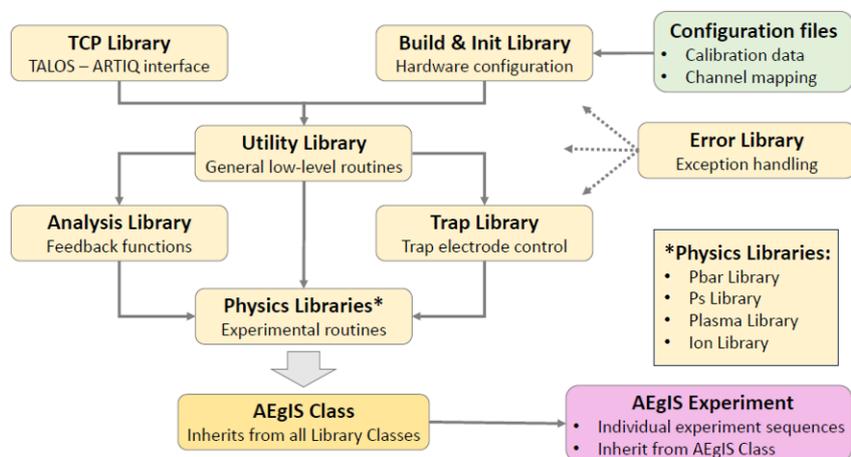


Figure 4.6: Schematic of the ARTIQ/Python library structure developed in AEGIS.

## The AEGIS Class

The `AEgIS Class`, which is the parent of each experimental script and thus forms the basis of all routines, defines a `run` method (see Fig. 4.7) that is executed every time a script is called. Included in this function are the `Simplified_Init` function (introduced below), the `StartRun` command (initialization of the TCP and DAQ functionality, start of a new run), the `experiment` function (see below), relevant error handling (see Sect. 4.4.3), the `closure` function (see below), and the `StopRun` command (conclusion of the TCP and DAQ functionality, transmission of the return code to the LabVIEW™ system (see Sect. 4.4)). A `try...except...finally` structure is chosen to ensure a safe termination of the script in case of any failures in the experiment.

Additionally, the `AEgIS Class` also provides empty default `build`, `experiment`, and `closure` methods. These are to be replaced in the individual scripts by the desired procedures: the `build` method by one of the pre-defined hardware setups described above (or individual configurations if needed), the `experiment` method by the actual sequence of operations to run the experiment, and the `closure` method by those processes that are required to happen at the end of the experiment, specific to the performed hardware operations, independent of its outcome.

```

def run(self):
    try:
        # Start procedure
        self.Simplified_Init()
        self.StartRun()
        # Experiment routine
        self.experiment()
    # Exception handling
    except Banana as error:
        self.ExceptBanana(error)
    except RTIOOverflow as error:
        self.ExceptKernelBanana(error)
    except Exception as error:
        self.ExceptException(error)
    finally:
        # Post-measurement cleanup
        print("Starting closure procedure...")
        self.closure()
        # Closing procedure
        self.StopRun(self.RetCode)
    return

```

Figure 4.7: `run` method of the AEGIS Class. The functions included here are executed every time an experimental script is launched.

## Configuration files

A set of configuration files (JSON files which can be loaded into Python dictionaries) has been created, which are read in according to the used `build` function from the `Build & Init Library` and decoded/used by the library classes (or, if needed, experiment scripts) where appropriate. Each of these files has a specific purpose and contains the corresponding information:

- `electrode_info.json`: information on the position (horizontal position in the apparatus relative to the front of the C0 electrode) and length of each electrode used in AEGIS (compare Appx. E);
- `HV_amp_calibration_data.json`: calibration data (slope and offset) of each Fastino/amplifier channel, as extracted from the procedure described in Sect. 4.3 and used to convert a desired output voltage to the closest suitable machine unit of the Fastino;
- `sinara_channels_HVamps.json`: mapping of each electrode name to an amplifier channel, according to the hardware connection;
- `sinara_channels_ttl.json`: mapping of each Sinara DIO channel to a hardware trigger purpose it serves, as well as information on its expected setting (input/output) and whether it is located on a master (or stand-alone) device or on a satellite;
- `pulser.json`: for each pulsable electrode (and the electron gun), information on which channel on which device is used for it;
- `trap_setup.json`: configurations of different trap setups (i.e. potential shapings using the trap electrodes), including information on the different parts of each trap (e.g. "inlet", "floor", "outlet") and a list of the involved electrodes.

Some configuration files contain different data for the 5 T and 1 T systems (e.g. different used trigger channels for the same purpose, different triggers for the pulser), and the `build` function of the `Build & Init Library` ensures that only the corresponding parts are loaded.

If one of the connected hardware lines changes (e.g. because of a failure replacement or for logistics reasons) or a connection is added anywhere, it suffices to change the one corresponding line in the dedicated configuration file for the entire system to be ready to run with the new setup. This renders adaptations to the system extremely easy, thus making it very flexible, and minimizes the time lost for such modifications.

### Special scripts and files

One additional python file, `AEgIS_imports.py`, is imported into every single experiment script. It has been created to automatically include all relevant external functionality: Special Python extensions (such as `numpy` and `json`) as well as the required ARTIQ packages are available to all scripts via one `import` line.

Furthermore, a special initialization script is created for every used control system unit (e.g. `Init_Script_5T.py`, `Init_Script_1T.py`), whose `run` function only contains the corresponding initialization function from the `Build & Init Library`. This script is run once after every reboot of the system so as to fully prepare the electronics. In principle, this initialization could happen at the beginning of every `run` function individually, however, a design choice in the setup of the Fastino DAC causes all of its channels to default to the zero setting in machine units upon initialization. Since the range of Fastino voltages goes from -10 to 10 V and the starting point of this setting in machine units is at the lower limit, all Fastino channels provide a constant voltage of  $-10$  V (which is then amplified to  $-200$  V) when the device is initialized, causing large disturbances in the potentials inside the apparatus and posing a risk to surrounding electronics in the case of too frequent occurrences. In addition, the `Init` script procedure also avoids unnecessary, frequent resets of the entire hardware, possibly extending its overall lifetime. With this approach, the initialization part of each defined `run` procedure is then reduced to a general `Simplified_Init` function included in the `Utility Library`, which only resets the RTIO buffers and prepares the voltages on the electrodes in the expected way.

### 4.2.5 Efficient programming

The library-based programming approach significantly improves the code efficiency, while at the same time rendering it more readable and approachable. Fig. 4.8 compares an experimental script to produce a well potential between three electrodes (two endcaps, one floor electrode) subsequently to an arriving external trigger without (left) and with (right) the use of a small fraction of the library structure.

Since such operations commonly form part of actual experimental routines in AEGIS, a standard function (built up again by sub-functions) has been defined, only requiring as parameters the denotation of the desired input trigger and its gate time as well as the set of electrodes (pre-defined as a trap configuration) and the voltages to be applied on them. Furthermore, all procedures involved in the default setup and initialization of the hardware are already included in the generalized `build` function or in the `AEgIS Class` itself. Thanks to these standardizations, the code of such simple functionalities can be reduced to a few lines of code and becomes quite intuitively usable.

For readability reasons, the routine on the left hand side, which uses the ARTIQ infrastructure directly and disregards the library system, does not include the application of the

amplifier calibration constants (see Sect. 4.3) and therefore applies uncalibrated voltages. To include the calibration, it is necessary to apply the voltage in "machine units" of the Fastino setting, which requires an additional conversion. This conversion is included in the configuration files of the library structure, allowing a straightforward integration in the routine on the right hand side.

```

from artiq.experiment import *
from artiq.coredevice.kasli_i2c import port_mapping

class HVAamp_Trigger(EnvExperiment):

    def build(self):
        self.setattr_device("core")
        self.setattr_device("fastino0")
        self.setattr_device("ttl0")
        self.setattr_device("ttl_hvamp0_sw1")
        self.setattr_device("ttl_hvamp0_sw2")
        self.setattr_device("ttl_hvamp0_sw3")
        self.setattr_device("dio_mcx_dir_switch")
        self.setattr_device("i2c_switch0")
        self.dio_mmcx_i2c_port = port_mapping["EEM0"]

    @kernel
    def set_dio_outputs(self):
        self.i2c_switch0.set(self.dio_mmcx_i2c_port)
        self.dio_mcx_dir_switch.set(0b00000001)
        self.core.break_realtime()
        self.ttl0.input()
        self.core.break_realtime()

    @kernel
    def SignalAtTrigger(self):
        t_gate = self.ttl0.gate_rising(120*s)
        t_trig = self.ttl0.timestamp_mu(t_gate)
        at_mu(t_trig)
        delay(10*us)
        self.fastino0.update(1<<3|1<<2|1<<1)

    @kernel
    def SetVoltages(self):
        self.fastino0.set_dac(1, 30.0/20.0)
        self.core.break_realtime()
        self.fastino0.set_dac(2, 5.0/20.0)
        self.core.break_realtime()
        self.fastino0.set_dac(3, 30.0/20.0)
        self.core.break_realtime()
        self.SignalAtTrigger()

    @kernel
    def run(self):
        self.core.reset()
        self.fastino0.init()
        self.core.break_realtime()
        self.fastino0.set_hold(1<<3|1<<2|1<<1)
        self.SetVoltages()

```

```

import sys
sys.path.insert(1, 'C:\kasli-code\Libraries')
from AEgIS_imports import *
from AegisExperiment import _AegisExpOfficial

class HVAamp_Trigger(_AegisExpOfficial):

    def build(self):
        self.Build5T()

    def experiment(self):
        self.FastReshapeAtTrigger("Trigger",
        120*s, "TestTrap", 30.0, 5.0, 30.0)

```

Figure 4.8: Experimental scripts in the ARTIQ environment defining the routine to produce a well potential trap formed by three electrodes upon the arrival of an external trigger signal. Left: Without the use of the AEgIS library structure and without the amplifier channel calibration. Right: Implementation of library-based programming, which also incorporates the amplifier calibration constants.

### 4.3 Calibration of the amplifiers

To efficiently manipulate plasmas confined in the electromagnetic traps of AEgIS, careful operation of the electrode potentials is essential, so as to not disturb the particles and

cause losses. For this reason, it is important to be able to control the applied voltages with an accuracy as high as possible. The voltages on the Fastino can be controlled in steps of 0.3 mV per machine unit, yielding approximately 6 mV steps on the amplifier outputs. Ideally, the accuracy of the produced voltage for each channel is thus also of the order of a few millivolts to efficiently use this property to tune the potentials.

When using the Fastino/amplifier channels directly without a calibration, the achieved accuracy varies strongly between different channels, according to their individual internal electronics setup, as well as over the voltage range for an individual channel and, for some, is as poor as 100 mV (see Sect. 4.3.2 for details). To improve this situation, a dedicated calibration procedure has been developed for the amplifier channels, which is described below and aims at determining, for every voltage in the producible range, the Fastino machine unit setting that is closest to the desired output. Since the produced voltage increases by a certain voltage with every one-step increase in the machine unit setting (approximately 6.1 mV on the amplifier output for the 16-bit Fastino DAC, but a characteristic individual property of each channel), a linear relationship between the two is expected, which can be used to extract the needed calibration constants.

### 4.3.1 Calibration setup and procedure

For the calibration of the Fastino/amplifier channels controlling the AEGIS trap electrode potentials, two programs have been developed, which act on the involved electronics. They are employed together and are easily usable in case the calibration of further channels becomes necessary. The calibration is performed channel-wise because, at the time of the calibration, no electronic component is available to connect the multimeter used for the data acquisition to multiple channels at the same time. However, once such an element is introduced, one additional `for` loop in each program would allow for the automated subsequent calibration of several channels as well.

The first program is an ARTIQ/Python script which instructs the used Kasli to scan through the range of the available machine units of the corresponding Fastino channel and, at every step, send a TTL pulse on one of the DIO channels, which provides the trigger signal for a multimeter connected to the amplifier channel that reads the produced voltage. The step size in machine units is a fixed number over the entire range and is determined by a definable number of measurement points. For the calibration performed here, 202 points are taken per scan, corresponding to a step size of approximately 100 mV for the Fastino setting. The number of times this scan is performed for every channel can also be determined by the user and is set to five here. Both the number of steps and the number of runs are a compromise between the statistical precision of the calibration procedure and time efficiency.

The second program is a LabVIEW™ VI that is responsible for the configuration of and data acquisition from a multimeter. The device used for the current calibration is a Keithley 2100 6 1/2 digit multimeter. The multimeter is connected to the amplifier channel that is being calibrated and reads the voltage produced on it. It allows to program the corresponding number of expected data points and the parameters identifying the channel. At every voltage step, the multimeter is triggered as described above, and the VI reads the obtained voltage and stores it in a JSON formatted text file together with the information on the channel and measurement number.

A third program has been developed, in regular Python code, to use the data obtained from the calibration measurements and extract from it the needed calibration constants: For every channel, the mean and standard deviation of the five measurements (or, more

generally, a definable number of measurements) are determined for every voltage step. This data, versus the programmed Fastino machine unit at every step, is then fitted with a linear function to extract the slope (i.e. the output voltage increase per machine unit step) and offset (i.e. the output voltage at the lowest machine unit setting) together with their corresponding standard errors. The calibration data of all channels of one example amplifier board used in AEGIS (board number 1 installed in the 5TC1 Sinara crate) is plotted on the left side of Fig. 4.9 together with the used linear fits. The inset plot shows a zoom to a smaller range of machine units for visibility as an example. The linear function fits well to the average data. This fact is also supported by the observed difference between the data points and the fit function, which is plotted on the right side of Fig. 4.9 for all channels at every used Fastino machine unit step and amounts to a maximum of around 6 mV, comparable to the Fastino precision.

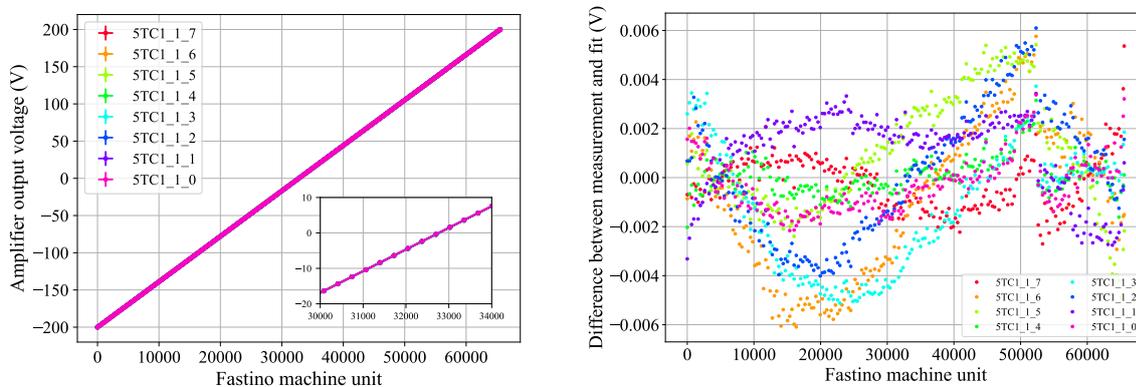


Figure 4.9: Calibration fit for the trap electrode amplifier channels. Left: Amplifier output voltage plotted against the used Fastino machine unit, together with a linear fit of the data, for all channels of one example amplifier board installed in AEGIS. Right: Difference between the calibration data points and the fit for each channel plotted against the used Fastino machine unit.

From the described fit, the calibration data for each channel is then saved in a JSON file, which is included in the ARTIQ/Python library of the CIRCUS control system of AEGIS, such that the corresponding calibration constants are automatically applied to find the most suitable Fastino setting when controlling the voltage of a trap electrode: The machine unit  $u_m$  to be used is determined in the dedicated Python function according to Eq. 4.3 from the desired voltage  $V_{\text{des}}$ , the offset  $V_{\text{offset}}$  and the slope  $V_{\text{slope}}$ .

$$u_m = \text{int}((V_{\text{des}} - V_{\text{offset}})/V_{\text{slope}} + 0.5) \quad (4.3)$$

During the channel characterization, it has become evident that the calibration does not only depend on the used amplifier or the used Fastino channel but on their combination and varies significantly for different combinations. This means that when one of the two needs to be replaced, a re-calibration of the channel is required. This is also the reason why the channels are calibrated as a whole and not for example only the Fastino channel individually.

The calibration of one channel in this way currently takes of the order of 30 minutes, mainly determined by the timeout given to the multimeter and LabVIEW™ program for the data acquisition and storage (a couple of seconds for good stability); the voltages of the Fastino/amplifier channels can be changed more quickly (see Sect. 4.5).

### 4.3.2 Calibration results

All 88 Fastino/amplifier channels used in AEgIS (including spares) have been calibrated with the described procedure. Fig. 4.10 shows the distributions of the extracted slope and offset values of all channels, grouped according to the Sinara electronics crate they belong to. No obvious dependence on the crate is observed, and the parameters are narrowly distributed around the expected values: All slopes are found to be between 6.098 and 6.109 mV, compatible with the 6.1 mV machine unit steps of the 16-bit DAC; all offset values are close to the quoted  $-200$  V minimum achievable voltage, between  $-200.14$  and  $-199.79$  V.

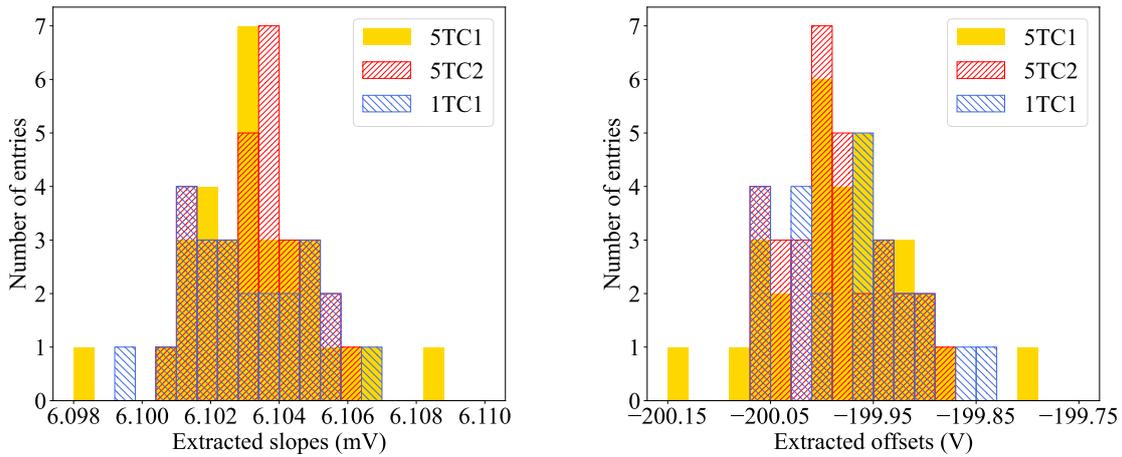


Figure 4.10: Distribution of extracted calibration constants for the amplifier channels used for the AEgIS trap electrode voltages. Left: Slope values, corresponding to the increase in output voltage per Fastino channel machine unit. Right: Offset values, corresponding to the lowest achievable voltage at the minimum machine unit setting.

A set of dedicated verification measurements for this calibration procedure have been performed, whose results are plotted in Fig. 4.11. Here, the extracted calibration constants have been applied as intended to convert the desired output voltages to the closest Fastino machine unit, which is then set on the given channel. To ensure stability of the calibration procedure over time and for varying environmental conditions, the shown measurements have not been taken directly after the calibration measurements but at random times of the day between a few days and one month later and have been verified since by individual checks up to two years after the initial implementation.

Fig. 4.11 shows the comparison of the produced voltage accuracy before and after the application of the calibration procedure for all channels of the example board whose calibration fit is shown above. Plotted is the difference between the obtained output voltage and the desired voltage programmed to the respective Fastino/amplifier channel for a number of voltages between the minimum ( $-200$  V) and maximum ( $200$  V) of the achievable range of each channel. Each data point is measured three times. The plotted error bars correspond to the statistical errors from these three measurements.

For an observation of the relative accuracy over the voltage range, and because for some applications, the accuracy at very low absolute (i.e. close to zero) voltages is particularly important, the number of data points taken is defined per order of magnitude step in this range. This means that the same number of data points is obtained between each of

the defined absolute intervals ( $(200, 20]$ ,  $(20, 2]$ ,  $(2, 0.2]$ ,  $(0.2, 0.02]$ ,) and the step size is adapted accordingly, resulting in the clustering of data points at low absolute voltages in Fig. 4.11.

It can be observed that the relatively large deviations from the desired voltage obtained without calibration (up to 100 mV) as well as the extremely varying voltage dependence of the accuracy per channel are removed by the calibration, yielding an accuracy of the order of a few millivolts for each channel over the entire voltage range. This value corresponds well to the 6 mV precision reachable by the electronics of the channel.

As described above and shown in Fig. 4.10, like the voltage change corresponding to one machine step, also the exact reachable minimum and maximum voltage is a characteristic of each individual channel and depends on the internal configuration of the Fastino. This causes more significant deviations from the desired voltages at the very boundaries of the range. However, the obtainable value is in no case further away than 0.2 V from the expected  $\pm 200$  V. For the use in AEGIS, this voltage reach suffices easily, as voltages beyond  $\pm 190$  V are never required for the application of the trap potentials.

The step structure of the measurement points in Fig. 4.11 is a consequence of the 16-bit precision of the voltage settings, rendering some of the tested voltages closer to the upper limit of the machine unit step and some closer to the lower one.

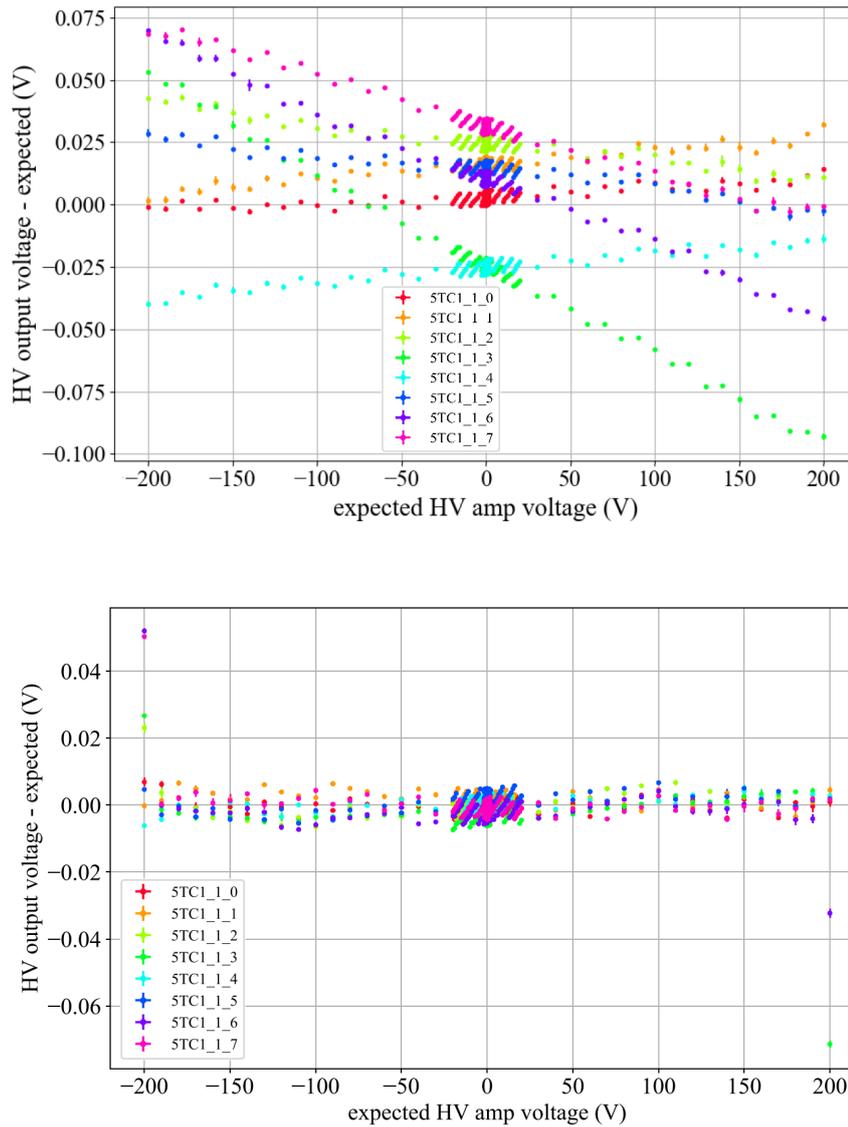


Figure 4.11: Difference between the desired voltage on the amplifier channels and the measured output voltage versus the expected voltage before (top) and after (bottom) the amplifier calibration for all eight amplifier channels of one example board. The legend identifies the channel numbers of the given board.

In between each change of the order of voltage magnitude of the verification measurements, one data point is also taken at 0 V to ensure the stability of the produced voltage independent of the previous setting. The result of this test is shown in Fig. 4.12 (left side) for two of the 5TC1 crate boards as an example: For all channels, the variation of the voltage read on the given channel with respect to its previous setting is far below the millivolt level and therefore insignificant, if any systematic effect can be observed at all. This means that no relevant timeout is needed between changes of potential in the trap to avoid instabilities.

A second systematic check is also shown on the right side of Fig. 4.12: To investigate the inter-channel influence, a scan through the Fastino machine unit settings is performed on one channel, while the output voltage is read on a neighbouring channel. A slight increase of the output voltage of the tested channel can be observed for higher machine units of its neighbour, however, this effect is limited to below 100  $\mu$ V and thus also insignificant for

the given accuracy. Fig. 4.12 (right side) shows the result for channels 5TC1\_1\_5 (tested channel) and 5TC1\_1\_4 (programmed channel) and the data points are the averages of four measurement runs; the error bars are given by the corresponding standard deviations. This test has not been performed for all possible neighbouring channels due to timing constraints but for a set of random combinations taken as examples, which all yield similar results.

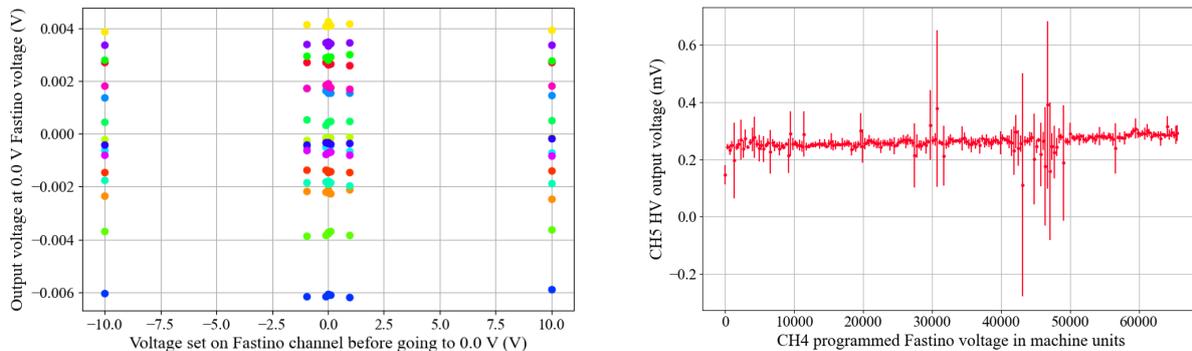


Figure 4.12: Systematic checks of the output stability of the calibrated amplifier channels. Left: Produced output voltages at 0 V programmed Fastino channel voltage plotted versus the previously (a couple of seconds before) applied voltage. A legend is not shown in this plot to maintain visibility; the different colors of the data points represent different channels of the amplifier boards with numbers 0 and 1 of the 5TC1 Sinara crate. Right: Plot of the voltage observed on one exemplary amplifier channel for a variation of the potential applied on a neighboring channel.

Unless otherwise stated, the example measurements shown here represent the results obtained with all Fastino/amplifier channels involved in the AEGIS system. Calibration results of all channels are shown in Appx. H.

## 4.4 The AERIALIST and TALOS: CIRCUS

The AERIALIST forms one of the two core pillars of the new control system developed in AEGIS, which is called CIRCUS [153]. While the AERIALIST takes care of those operations that need to be performed on the electronics with high synchronicity and timing precision down to the nanosecond (e.g. trap electrode control, trigger signals) using real-time code, it would be extremely inefficient to interconnect and control all experimental subsystems in this way. Furthermore, CIRCUS has been designed as a very general control system with a strong focus on autonomy and convenient user interfacing, which would not be easily achievable by the AERIALIST alone. In addition, the multitude of instruments involved in experiments and their different interfaces require a more diverse approach.

For these reasons, the second core component of the novel control system is a slow-control LabVIEW™ framework, named TALOS (Total Automation of LabVIEW™ Operations for Science), that unifies all involved computers in a single, distributed entity and thus allows for high degrees of automation. An overview over this framework is given below.

The strength of CIRCUS, thus, lies in the interplay between the AERIALIST and TALOS, which swap control back and forth according to changing requirements and, in this way, reliably steer the experiment over the diverse involved time scales (nanoseconds

for time-critical operations, minutes to several days for entire measurement sequences). This relationship is visualized in the schematic shown in Fig. 4.13. TALOS and its LabVIEW™  $\mu$ Services ( $\mu$ S) do the high-level slow control and coordination of the different hardware (HW) subsystems, scheduling experiment runs, informing the AERIALIST on which scripts to run when, and determining the next steps once completed, while the AERIALIST is in charge of the creation of the actual run routines and the running of fast procedures (ARTIQ/Python code) on the connected Sinara components, receiving and providing external triggers to additional hardware and requesting higher level (hardware or computer) operations from TALOS as needed.

The DAQ system and the ALPACA analysis framework are also incorporated in CIRCUS through dedicated  $\mu$ Services. With ALPACA, it is possible to perform automatized parameter optimization and quality assessments, a further step towards making CIRCUS autonomous: ALPACA features a Bayesian optimizer that can determine, based on the analysis of previous runs, the next useful point in a parameter space that is being explored, even in multiple dimensions. The AERIALIST allows for the external provision of parameters when running an experimental script. Therefore, when a completed run is analyzed by ALPACA and the next step is suggested to the control system, TALOS can provide this information to the next script in the overall program, essentially modifying the schedule while it is ongoing. This procedure renders parameter optimizations more efficient (and, thus, much faster) than manual scans and removes the need to assume the orthogonality of parameters for multi-dimensional tasks, allowing for improved settings. In AEGIS, this functionality is for example exploited for the laser calibration and for antiproton beam steering [132]. Similarly, ALPACA can also determine whether a completed run is "good", i.e. satisfies given quality standards, or needs to be retaken, which can then be used by TALOS to adapt the schedule accordingly.

The AEGIS DAQ system used to save and store the acquired data is briefly introduced in [153] as well. Information is stored in bundles containing a characteristic name, a timestamp and the data itself in JSON-formatted files, locally on a dedicated computer as well as on the the long-term tape/disk storage of CERN, EOS [157]. The incorporated DAQ system is optimized for the use in AEGIS, but CIRCUS can be operated with any data acquisition system supporting `Start`, `Stop` and `Send` commands by a simple adaptation of the corresponding  $\mu$ Service (see below).

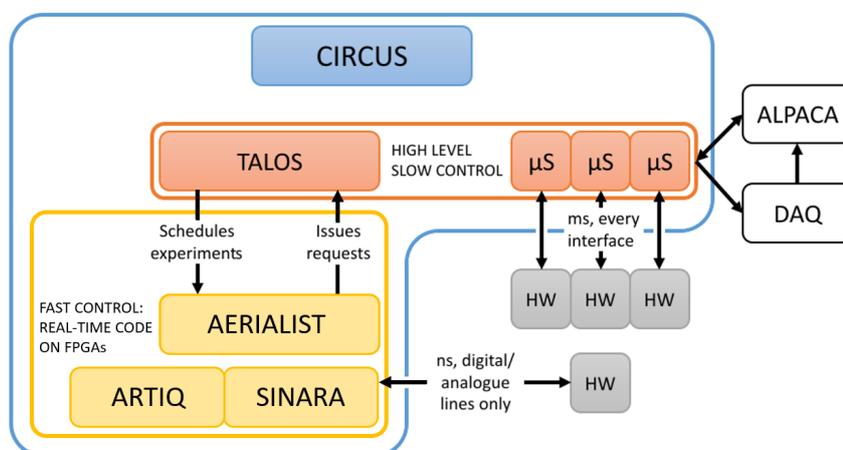


Figure 4.13: Visualization of the constituents of the CIRCUS control system (AERIALIST and TALOS) as well as its relationship to experimental subsystems.

### 4.4.1 TALOS overview

The TALOS framework is based on the NI<sup>6</sup> LabVIEW™ Actor Framework and is founded on two main pillars: the "everything is a  $\mu$ Service" approach and a distributed system architecture.

Just like the electronics, also the software setup of a powerful control system is required to be modular so as to allow for an independent development of the subsystems and, thus, render the system easily extendible and adaptable. This is achieved with the first concept: the LabVIEW™ code is divided into independent parts, called  $\mu$ Services, that are each in charge of a specific task within the experiment (e.g. control of different detectors, scheduling of run sequences, provision of instructions to the AERIALIST) and run asynchronously on one of the computers selected for this function. The base of all  $\mu$ Services is a template class (Father Of All  $\mu$ Services (FOAM)) that they all inherit from, ensuring a uniformity in the code (e.g. also making it mandatory to define sequences of actions for certain situations such as a restart of the system or entering the safe mode) and an option to modify common functionalities without having to adapt each component individually.

The second pillar is the key for making the system autonomous: copies of an identical process, called *Guardian*, run on every computer involved in the system and monitor at the same time the status of the local  $\mu$ Services (which are also launched by the *Guardian* according to a configuration file) and that of the other *Guardians* present in the network. A built-in TCP messaging system lets the  $\mu$ Services interact with each other in case of needed collaboration (e.g. to exchange information on the state of a given hardware component). In this way, the multitude of computers are unified as a single, distributed, entity, enhancing the reliability and stability of the system and allowing for global decisions to be taken in case of errors.

A schematic of the TALOS architecture is shown in Fig. 4.14. Further details are given in [132].

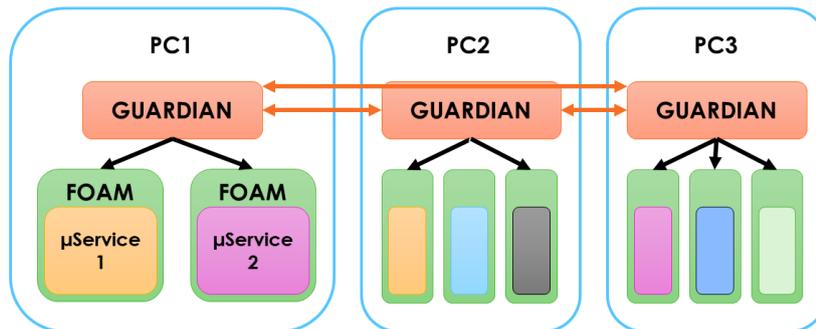


Figure 4.14: Schematic of the TALOS structure. Each computer (PC) in the network runs a *Guardian* process that launches the defined  $\mu$ Services and monitors them as well as the *Guardians* of the other computers.

<sup>6</sup>National Instruments Corp.

#### 4.4.2 Integration of the AERIALIST and TALOS

Generally, when controlling hardware that is not part of the core Sinara electronics from ARTIQ/Python scripts (i.e. using the corresponding library functions), the direct communication is actually done from a LabVIEW™  $\mu$ Service with the corresponding driver. In order to propagate the needed action from the experimental AERIALIST script to the  $\mu$ Service and be able to receive input from TALOS, double, asynchronous TCP tunnels are used, allowing for a communication without loss of the nanosecond-precise synchronization once the tunnel connection has been opened. In the scripts/libraries themselves, then, it is enough to call the dedicated TCP Library function (`TCP_Send`) with the name of the corresponding  $\mu$ Service and parameters defining the desired action to be taken by it at the correct moment of the timing sequence to initiate an operation. Another function (`TCP_Read`) is used to receive external input.

Naturally, ARTIQ/Python scripts are run on Kaslis from command line terminals. Therefore, to facilitate an integration on the TALOS side, a dedicated  $\mu$ Service, called *Kasli Wrapper*, has been developed that is launched when an experiment script for a given Kasli arrives in the overall schedule and encapsulates the ARTIQ shell, allowing to run experiments from the global system and at the same time log the standard output of the FPGA, thus intercepting low-level exceptions raised by the FPGA itself. These points are vital for a full, safe overall error management. Another  $\mu$ Service, the *Monkey*, provides the previously scheduled experiment scripts to the Kaslis in the given order<sup>7</sup>.

A screenshot of the running Electron Gun  $\mu$ Service (see Ch. 5) is shown as an example in Fig. 4.15. The base components of the Graphical User Interface (GUI) are the same for all  $\mu$ Services: The panels at the top left show the states of all computers involved in the network, i.e. their *Guardians*, as well as those of the  $\mu$ Services running on the viewed computer. In the top left, any relevant errors are displayed, with details of the most recent error in the box below. The run number and status as well as the selected  $\mu$ Service are displayed next to a secured "STOP" button, which can terminate the instance of TALOS on the given computer. On the right hand side, the "Kasli Log" of all `print` statements in the AERIALIST code together with relevant Kasli and TALOS messages is also shown. The central GUI components can be designed in standard LabVIEW™ code, following a set of templates, and form the actual  $\mu$ Service itself: the control and visualization of the power supply settings and the display of the emitted electron current together with additional safety and information features in the case of the Electron Gun  $\mu$ Service.

---

<sup>7</sup>The order of the experiment scripts in the *Scheduler*  $\mu$ Service can be modified while running, and this happens automatically in certain cases, for example when a run needs to be repeated.

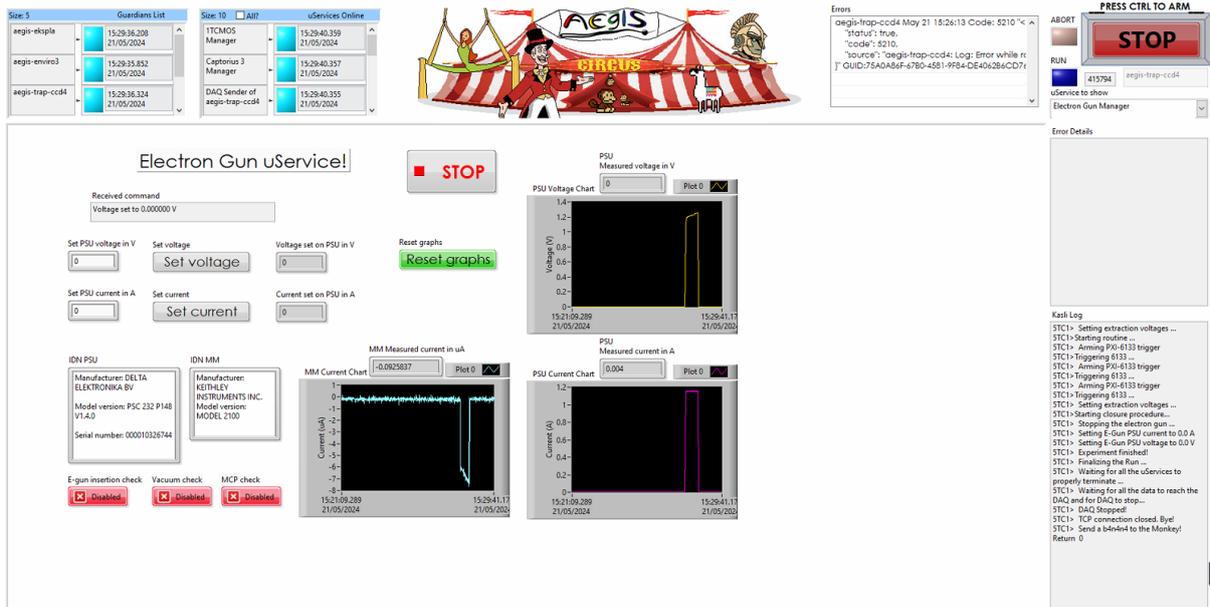


Figure 4.15: Screenshot of the Electron Gun pService running in the CIRCUS control system. Shown is the corresponding TALOS GUI on one of the computers in the network, providing an overview over the status of the entire system as well as the controls and displays of the given task.

### 4.4.3 Error handling

For an autonomous operation of the system, it is essential that occurring errors are handled without user intervention in a reliable way. This is facilitated by return codes that are generated by the AERIALIST according to obtained exceptions during the execution and suggested from ALPACA. These return codes are passed through the *Kasli Wrapper* and provoke specific actions to be initiated by the *Monkey*, determining how to proceed in the schedule of experiment scripts. The possible return codes and corresponding subsequent actions are, in increasing order of criticality<sup>8</sup>:

- **OK**: No errors have been encountered during the execution; the next script in the schedule will be launched.
- **RETRY**: A minor issue (e.g. the DAQ not having saved the complete dataset) has been encountered which is assumed to be unrelated to the specific script; the same script will be re-run.
- **SKIP**: A problem has prevented the completion of the run which is likely caused by the script itself (e.g. invalid parameters provided); the current script will be skipped and the next one in the schedule will be launched.
- **STOP**: A more global issue has been encountered during the run which is not critical (e.g. some hardware having become unresponsive in a safe state); the entire remaining schedule will be skipped and the system will enter the idle mode, i.e. ready to start the next schedule in the current state.

<sup>8</sup>If no or a wrong return code is given, *SKIP* is assumed.

- **ABORT**: A critical error has occurred (e.g. failure of the high-voltage system which could possibly cause damage to detection devices through over-voltages); the entire remaining schedule will be skipped and the system will go into safe mode, i.e. powering down any critical electronics (executed by the responsible  $\mu$ Services upon instruction) and potentially alerting experts<sup>9</sup>.

Any scripts that are being skipped are saved in a separate list to be edited and/or re-run at a later time.

The return codes are provided to the *Monkey* upon occurrence and it decides on a subsequent action, taking into account also a second set of error codes not generated by the AERIALIST but by TALOS itself via a *Error Manager*  $\mu$ Service, for example in case of one of the *Guardians* having become unresponsive. The final decision is based on the most critical return value obtained during the run. The user can observe the occurring errors in the *Error Manager* and obtain further information in the *Kasli Wrapper*, depending on their origin.

As introduced in Sect. 4.2, on the AERIALIST side, the **Error Library** contains the functionality needed to enable the global error handling. Depending on whether an error occurs in a kernel function or on the host, different exceptions are defined which are raised by the methods importing them and include criticality values corresponding to the return codes described previously. Specific exceptions that are needed as part of CIRCUS but are not generally included in ARTIQ or Python are defined here, such as the *Banana* needed to propagate the outcome of an experiment to TALOS.

#### 4.4.4 Multi-Kasli modes

To enable a simultaneous operation of multiple Kasli master devices, a  $\mu$ Service called *Tamer* has been developed. It instructs the creation of one instance of the *Monkey* (except for the sequential mode, see below) and of the *Kasli Wrapper* each for every Kasli master used in the system, based on the schedule defined by the user. The *Tamer* is then responsible for the distribution of the scripts in the schedule to the *Monkey* corresponding to the correct device as well as for the reception and distribution of the return codes from the *Kasli Wrappers* to take global decisions.

The user can decide to run the schedule of scripts for multiple Kaslis in one of three modes:

- *sequential mode*: One combined *Monkey* sends the scheduled scripts to the Kaslis one by one, waiting for the previous one to finish before starting the subsequent one, independently of whether they are run on the same device or on different ones.
- *asynchronous parallel mode*: Multiple *Monkeys* receive individual schedules for their corresponding Kaslis, which are run in parallel, going through the scripts until finished on the given Kasli, independently of what happens with the others.
- *synchronous parallel mode*: The *Tamer* acts as a coordinator between the multiple *Monkeys* sending scripts to their corresponding Kaslis: The execution of the next scripts for all Kaslis is initiated at the same time, and no new script is sent until the previous ones of all Kaslis have finished.

For simplicity, the *Tamer* is also included in the operation flow when using only one Kasli.

---

<sup>9</sup>Every  $\mu$ Service responsible for hardware in the system is required to have a defined routine of actions to be taken when ordered to enter the safe mode.

## 4.5 Timing requirements

To be able to efficiently coordinate the entire experiment and guarantee quick enough reactions to incoming triggers, latencies of no more than  $100\ \mu\text{s}$  are needed, allowing to correct for them within time scales that are compatible (i.e. lower by an order of magnitude) with experimental routines. Depending on the specific operations, the ramping of the trap electrode voltages to reshape the potentials confining the plasmas needs to happen within timescales of  $100\ \mu\text{s}$  to seconds so as to not cause unnecessary losses, and in a synchronous way on multiple channels so as to not disturb the confined plasmas excessively. Fig. 4.16 shows the voltage ramp from 0 to 20 V on three of the Fastino/amplifier channels of the AERIALIST, subsequently to ( $10\ \mu\text{s}$  after) the arrival time of an external trigger pulse on one of the digital I/O channels, as recorded by an oscilloscope. As indicated there, the potentials on the different channels can be produced very synchronously within no more than a few tens of microseconds, with the exact ramping time depending on the voltage difference to be covered (for details see Sect. 4.1.1), satisfying above requirements. Furthermore, reactions to given triggers are, as needed, easily possible within a few microseconds. Such fast operations are possible thanks to the use of Kasli's DMA (see Sect. 4.2.2).

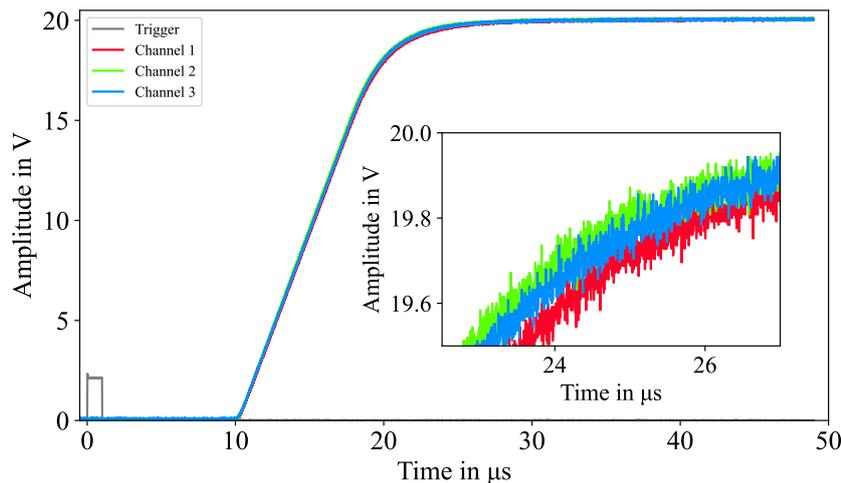


Figure 4.16: Synchronous voltage ramp from 0 to 20 V on three Sinara Fastino/amplifier channels subsequent to a common trigger signal (arriving at  $0\ \mu\text{s}$  in the figure). A zoomed plot of the rising shoulder region is shown in the inset for a better visualization of the synchronous ramp.

More stringently, the timing resolution for precise synchronization of certain routines is required to be of the order of nanoseconds. These include for example the synchronization of the different involved lasers with the moment of implantation of positrons on the Ps target for antihydrogen production and that of the trigger signals of the antiproton bunch from ELENA and certain detector acquisitions so as not to miss the arrival of the particles. The implementation of these "fast synchronization" aspects in the AERIALIST and CIRCUS are briefly discussed below.

Finally, while many procedural preparations of the "slow control" are limited to the speed of typical computer operations, the communication between involved subsystems should happen within 100 ms to ensure a smooth flow of the overall experimental routines. Direct

interfacing between computers and the real-time FPGAs is, within CIRCUS, possible on the order of a few milliseconds through the use of TCP tunnels. Since such communication generally happens over a network, it is inherently limited to the ms scale.

In the context of AEGIS, the broadest frame for the execution of experimental routines is given by three components: Firstly, the antiproton cycles of the AD/ELENA complex facilitate a  $\bar{p}$  provision approximately every 110 s. Therefore, in order to efficiently use the provided particles, all operations of one run should complete within this time. Because calls from the host to the kernel entail significant delays, which can amount to several seconds when going back and forth thousands of times, it starts to become relevant for complex procedures involving many potential reshaping (e.g. antihydrogen production) to clearly separate those actions to be taken on the host and those on the kernel. For example, the functions controlling the reshaping are separated into a first part (host) that decodes the sequence of voltages in machine units and the used electrodes into lists<sup>10</sup>, which are then recorded by the use of Kasli's DMA (host) in one run as a second part instead of alternating multiple times, even though that makes the procedure look less computationally efficient on the first glance.

Secondly, the safe ramping of the high-voltage electrodes for the antiproton capture requires of the order of 15 s, a time that needs to be included in all experiments involving  $\bar{p}$  capture. Thirdly, the data acquisition system needs a certain amount of time to process the acquired data, which depends on the intensity of the signal obtained on the detectors (up to 100 s for some of the accumulation experiments presented here but usually limited to a few seconds) and needs to be included as a timeout in the closure procedure to ensure that the DAQ is ready to acquire new data when a new run starts. Some of the used detectors also require themselves timeouts of a few seconds between subsequent armings and acquisitions.

#### 4.5.1 Master/satellite operation

As alluded to previously, while Kasli controllers can be used individually as stand-alone devices, it is also possible to connect several of them to each other, using one as the master, which is directly connected to the host computer and external clock and receives the communication, and the others as satellite devices, with the RTIO and time information forwarded via optical transceivers. This can be done both in a star topology (default for minimized latency; all satellites connected directly to a master channel) and as a daisy chain (devices connected in series). The configuration is done with a routing table binary file that is then stored in the flash storage of the core device. In AEGIS, the star topology is used wherever the master/satellite setup is needed, as currently no more than three satellites are involved and each Kasli can host three direct links, but daisy chains have also been successfully tested in the laser setup<sup>11</sup>, allowing for straightforward extensions of the system in the future.

Synchronization of events is maintained across the satellites thanks to the integration of the Gbit/s DRTIO system in ARTIQ. This is depicted on the left hand side of Fig. 4.17, which shows an oscilloscope acquisition of two TTL pulses<sup>12</sup>, one given by a master device and the other one originating from a satellite setup. The ramping of the voltages to produce these trigger pulses happens very synchronously, to below 1 ns, on both channels

<sup>10</sup>The kernel part is unable to handle more complex structures than lists, such as Python dictionaries.

<sup>11</sup>Parts of the AEGIS laser system are located at a further away location inside a new laser hut on a different level of the AD hall, making daisy chaining more convenient with fewer long cables.

<sup>12</sup>Only two channels are shown here but the synchronicity does not change with the number of channels.

despite their different origins. In the software, it is necessary to introduce a delay for the master channels according to the length of the optical transceiver when requiring synchronicity with satellite outputs. In the used setup, this delay is of the order of 280 ns.

### 4.5.2 Kasli master synchronization

As explained above, the trap electronics in the 5 T and 1 T regions of the experiment are controlled independently from different Sinara crates and ARTIQ/Python folders because they are in charge of independent procedures (e.g. antiproton preparation in the 5 T trap, positronium preparation in the 1 T trap). They are also separated from the point of view of TALOS. However, for certain procedures, such as antihydrogen production, it is necessary to operate the two systems together, in a synchronized way. It would, of course, be possible to turn one of the two masters into the other one's satellite for such occasions, which would simply require to change the gateway of both, undoing the change again afterwards. Because both modes of operation are frequently used, sometimes even alternating every few minutes, this would be a very inefficient technique, especially since it necessitates a restart of the electronics every time. For this reason, a different approach has been developed, leaving both systems as they are set up and synchronizing procedures via the software. Two techniques have been introduced to achieve this goal.

The first one relies on the use of a barrier function that is available to both systems and two additional hardware TTL connections between the two, one with input setting and one in output mode on each, to facilitate the exchange of trigger signals. For combined procedures, the barrier function is included in the experiment scripts of both systems ("Kasli A" and "Kasli B" from here on for readability) wherever a synchronization is needed. Once the code execution of Kasli A reaches this function, a trigger is sent from it to Kasli B. Kasli A then opens a gate time (with the timeout as a parameter), during which it waits for an incoming return trigger from Kasli B. If the trigger is not received within the gate time, a timeout exception is raised. If the trigger does arrive, a second gate is opened by Kasli A to wait for a second trigger signal. If this one is not received, Kasli A itself sends out two triggers, with a delay matching the gate timeout in between, and the function returns immediately after, leaving the code execution to continue. If the second trigger does arrive, the function returns after a delay that is necessary to ensure synchronization (of the order of 620 ns in the current setup to allow for the production, transfer and reception of the trigger signal through the digital line). In this way, the double trigger acts as the signal to both systems that they have both arrived at the synchronization point and are ready to continue the procedure together. The image on the right side of Fig. 4.17 shows the result of this technique: Synchronous (better than nanosecond level) pulses are produced on two digital I/O channels, one each chosen randomly on the 5 T setup and on the 1 T side. This procedure works independently of what happens on each setup previously, as long as appropriate timeouts are given to the barrier function.

The second part of the approach is the transformation of the 5 T setup into a *Particle Server*: The main coordination of the experimental routines is done from the 1 T side and the 5 T one, after performing all the needed preparatory tasks, remains in an infinite loop of waiting for instructions from it. These instructions essentially correspond to specific functions defined for the 5 T side, including, among others, the procedure for antiproton capture from ELENA, the loading of electrons into the trap, several reshaping of the trap potentials and Pulser operations to prepare and move the antiprotons, and resetting the system. One instruction is also the application of the barrier function described

above to synchronize the two systems. Essentially, at the corresponding point in the code procedure, a TCP tunnel message with a code word for the required action is sent from the 1T side to a *Kasli Server* `µService`, which transfers it to the 5T system, where it is read using the dedicated function, decoded and put into action. It is also possible to include parameters in the code word, which are then used in the called functions. Except for some specific functions (e.g. the `stop` instruction to end the loop and terminate the *Particle Server*), the messages are "blocking", i.e. the 1T side halts at the point in the code where the instruction is sent and waits for a confirmation from the 5T Kasli (again, via TCP message) that the required action has been successfully performed (or failed, in which case an exception is raised) before continuing. In this way, the 5T side supplies particles (electrons, antiprotons) as instructed by the 1T system, which takes care of managing the overall procedure.

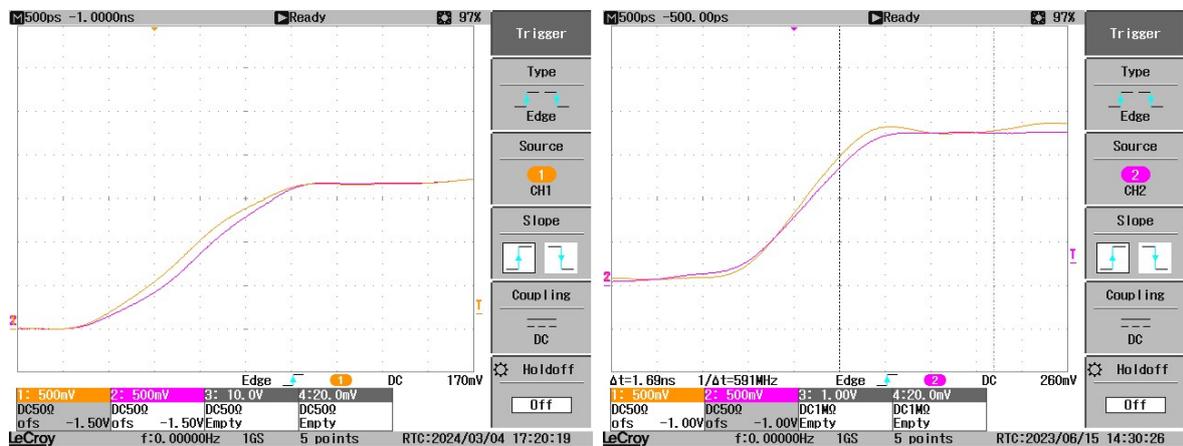


Figure 4.17: Oscilloscope images of synchronized TTL pulses on two Sinara digital I/O channels. One timing unit corresponds to 500 ps and one step on the voltage scale is 500 mV. Left: Channel 1 (orange) is a channel directly connected to the 5T master Kasli, while channel 2 (pink) is located in a satellite crate. Right: Channel 1 (orange) belongs to the 5T master setup and channel 2 (pink) is a channels of the 1T crate.

## 4.6 Results and reliable operation

The core task of the AERIALIST itself is the control of the electrodes involved in the capture and trapping of antiprotons as well as electrons and positrons and, during dedicated measurement periods, highly charged ions. These objectives have been well achieved during the data taking campaigns in 2021, 2022 and 2023, culminating in a record number of accumulated cold antiprotons and the development of a procedure for the formation of a pulsed beam of antihydrogen, which is reported on in Ch. 6 and Ch. 7. These achievements have only been possible thanks to the reliable operation of the AERIALIST and, more broadly, the CIRCUS control system in general. The same is true for the work performed with leptons in the experiment, which is detailed in Ch. 5.

Beyond these results, AEGIS has for the first time successfully achieved laser cooling of positronium atoms (see Sect. 7.5) with a dedicated laser system, whose synchronization is also only possible thanks to the new control system.

Utilizing the AERIALIST and its integration in TALOS, the entire procedure to produce

antihydrogen in the trap in a ballistic way for beam formation fits into less than 90 lines of Python code, including the preparation of all components of the apparatus and the detection of the annihilation products, through the use of the libraries of standardized functions (which have useful names and comments). This being one of the most complex procedures performed in AEGIS, which involves almost all aspects of the experiment, other tasks can be achieved with even less coding. In addition to saving significant amounts of time, this also renders the integration of new team members straightforward, a feature that has already been proven and exploited in several student measurement campaigns as well as by external collaboration members working on the development of sub-components of the experiment, such as a new detection module for the gravity measurement, who often only have a few days of test periods at CERN.

For the measurement campaigns in 2022 and 2023, the control system itself was online for approximately 65 % of the available time, reliably running the experiment and taking care of any exceptions. This corresponds to most of weekends and the nights of the period, during which the system was left to run the experiment autonomously, while the working days were generally dedicated to developments of procedures and operations on the apparatus. In 2023, TALOS had to be restarted only twice over the almost five months of operation due to unhandled errors. Any other exceptions were handled by the system in a safe way, i.e. applying the error handling routines detailed above and putting the entire experiment in a safe state when needed. The AERIALIST as such did not have to be rebooted on its own account at all but merely in case of power failures, which inhibit the clock signal and the network connection.

These observations show that CIRCUS, including the AERIALIST and TALOS, is well suited for the automated operation of an experiment for extended amounts of time without supervision. Furthermore, the AERIALIST can provide the precise timing synchronization and voltage supply that is needed in the context of antimatter and, more generally, quantum and atomic physics experiments. The system will therefore continue to be used for any experiments within AEGIS and can be adapted in straightforward ways for a multitude of different setups and procedures thanks to its modular, flexible nature.

## 4.7 Summary

A new, powerful experimental control system, CIRCUS, has been conceptualized, developed and commissioned, which is now running the entire AEGIS experiment with all its subsystems. It is founded on two main pillars: the TALOS slow-control unit, which is a LabVIEW™-based architecture and interconnects all computers used in the network, and the AERIALIST fast-control unit, which is the focus of this work and uses the Sinara hardware and an ARTIQ-based Python library structure to precisely synchronize the experiment's routines and provide the voltages required for the formation of the potentials on the electromagnetic trap electrodes. The two interact closely to coordinate all involved hardware and software parts of the experiment.

In the context of the AERIALIST, a variety of standard routines have been coded and included in the Python library and a framework involving git version control, easy access for new, individual procedures and programming efficiency has been put into place, which is used extensively by the entire collaboration. To supply the trap electrode voltages with the required accuracy of a few millivolts, a calibration procedure has been furthermore implemented and performed on each of the over 60 channels. Independent tests have validated the achieved accuracy and the robustness of the technique.

Thanks also to a stable error handling and the integration of self-optimization capabili-

ties, CIRCUS has been steering AEGIS reliably and without significant down-time over a period of several months at a time, occasionally autonomously, while at the same time providing precise routine synchronization on the nanosecond level. As such, it has been the basis for all scientific achievements of AEGIS in the past years, some of which are reported on in Ch. 5 through Ch. 7. It is designed to be flexible and modular, such that it can be furthermore employed in a variety of experiments outside of AEGIS.



# Chapter 5

## Electrons in AEGIS

*This chapter summarizes the work I have been involved in to commission the new AEGIS electron gun and include it in the CIRCUS control system. Furthermore, I have developed the electron procedures needed in the experiment. Aside from the plasma manipulation and antiproton cooling, which are described in the following chapter, these include any calibration, alignment and test techniques involving electrons. I have set up and conducted the experiments and analyzed the data presented here, guided by R. Caravita and B. Rienäcker.*

Electrons are needed in AEGIS to form mixed electron-antiproton plasmas in order to cool the antiprotons, but they are also used for detector calibration purposes and to perform various tests of the setup and electronics because they are easily available and controllable. Finally, they are additionally employed to develop plasma manipulation techniques, which are then applied to the mixed electron-antiproton plasma once the latter are available.

The objective of this part of the work has been the commissioning of the the source of electrons in AEGIS, i.e. the new electron gun, its incorporation in the newly implemented control system, the AERIALIST (see Ch. 4), and the development of all its required procedures within the upgrades experiment. Novel techniques involving the electrons for the calibration and alignment of experimental aspects have additionally been developed, with some of them now being used regularly. This chapter introduces the electron itself and its operation principles in the context of the AERIALIST. Subsequently, procedures introduced with electrons to test all electrodes of the electromagnetic traps, to align devices to the experimental axis, and to calibrate detector components are described.

### 5.1 The AEGIS electron gun

The source of electrons used in AEGIS is a so-called electron gun: simply put, a cathode which is heated to emit electrons (thermionic emission) and a metal extractor with a small hole in front of it which can be biased to form the emitted electrons into a focused beam. The cathodes used in AEGIS are barium oxide coated disc cathodes (Kimball Physics Inc. ES-015). In 2022, the electron gun setup was upgraded to hold two cathodes instead of one (one as a spare). Fig. 5.1 shows the new electron gun holder with both cathodes installed. The electrons are injected into the experiment through the small holes in the aluminium extractor frame. The new electron gun was installed, aligned, and commissioned for the ELENA antiproton run in fall of 2022, and both cathodes have been fully conditioned as well as tested. Since then, the new electron gun is routinely used for all experimental routines involving electrons.

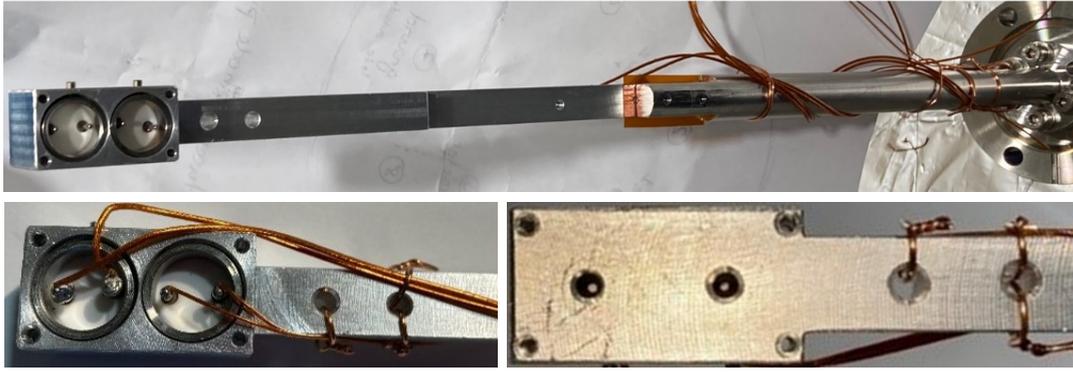


Figure 5.1: Photographs of the new AEGIS electron gun. Top: The electron gun holder on the flange used to attach it to the experiment entrance chamber. Bottom left: Back side of the double cathode head with copper wire connections. Bottom right: Front side of the double cathode head with two small holes in the extractor frame to let electrons pass through.

### 5.1.1 Using the electron gun

A schematic of the electronics circuit of the electron gun is shown in Fig. 5.2. For standard use, the cathode heating current from the corresponding power supply is ramped up to 1.15 A, ready to extract electrons. The energy of the extracted electrons depends on the setting of the bias voltage of the filament and increases with it. The extracted intensity of the electron beam, instead, is governed by the difference in the bias settings of the filament and the extractor: the higher the difference, the more electrons are extracted. As a general rule, a voltage difference between the two of 10 V is used, which was determined to generate electron emission reliably and efficiently in good amounts (tens of  $\mu\text{A}$ , i.e. large enough amounts to enable efficient antiproton cooling but not too many electrons to cause expansion losses of the mixed plasma, see Sect. 6.2.2) via a broad scan through the different bias settings. Furthermore, this value is also generally beneficial for the avoidance of magnetic bottle/magnetic mirror effects (see below) in the used voltage range. For example, to be able to use the electrons for an intercalibration of the detectors with positrons, the energies of the two types of particles entering the experiment need to be comparable. Since positrons can be injected with a minimum energy of the order of 110 eV due to the way they are accumulated and prepared (see Sect. 3.4.4 and [121]), the electron gun filament bias voltage is set to  $-110\text{ V}$  to bring the extracted electrons to the same energy, as they accelerate against the grounded trap. The extractor is then biased with  $-100\text{ V}$ . With this setting, an electron current of the order of 55 to 60  $\mu\text{A}$  is reliably produced. A scan through the extractor bias settings is shown on the left hand side of Fig. 5.3, with the filament voltage kept constantly at  $-110\text{ V}$  and the extracted current read by a Keithley 2100 6 1/2 digit multimeter. As expected, no current is extracted if the negative extractor bias is larger than the negative filament bias (or too similar in value to it), since the extractor essentially acts as a potential barrier for the electrons. This is the cause for the slow rise of the observed current with increasing extractor voltages; afterwards, the increase is approximately linear.

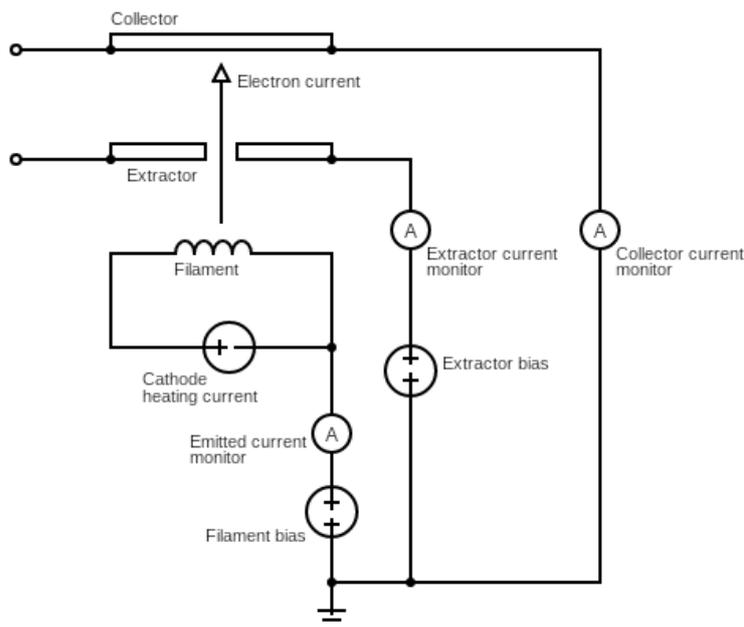


Figure 5.2: Schematic of the electric circuit of the electron gun as it is set up in AEGIS. The different components are discussed in the main text.

Varying the absolute values of the two bias voltages from above settings but keeping the difference constant leads to a change in the intensity of the extracted current. This feature can be observed on the right hand side of Fig. 5.3 and is related to the fact that the electron gun is not perfectly centered on the axis of the experiment, i.e. the magnetic axis: The electrons pass through the experiment on a trajectory that is not straight but slightly rotating around the central axis due to the magnetic force, causing small losses of those electrons that get back-reflected due to magnetic mirror effects [158, 159], depending on the ratio of axial and radial momentum. The higher the energy of the electrons is thanks to a higher filament bias, the higher is their axial velocity and thus the smaller is the probability to get reflected. The rotation of the electrons can also be observed on the image of the downstream MCP as a circular veil at high enough operation voltages. In addition to this relation, fluctuations in the extracted current are caused by the time the gun filament needs to heat up or cool down after extraction at a given filament temperature or after not being used. Overall, difference of up to  $\pm 5 \mu\text{A}$  are to be expected for the same settings when operating in different conditions. The right figure of Fig. 5.3 shows the current recorded by the digitized MCP front face at the downstream end of the experiment (MCP-FC detector, see Sect. 3.4.9) for different filament bias voltages, where the extractor bias is adapted to maintain a 10V difference between the two. For low negative voltages, very little current is observed, and an increase for higher bias settings is evident, which is not linear.

The error bars in both plots of Fig. 5.3 correspond only to the statistical errors on the mean values of the measurement points (between one and eight repetitions per data point). Furthermore, the shown data stem from a manual reading of the device signals and are subject to a resulting uncertainty of up to  $\pm 2 \mu\text{A}$ . The reading is performed at the end of a measurement, i.e. several seconds after the initiation of the extraction. This is relevant because the current needs a time of several microseconds to settle to a constant value on

the order of  $\mu\text{A}$  and even longer for increased stability.

In regular operation mode for antiproton cooling, the electron gun filament bias is kept between  $-40$  and  $-50$  V and the extractor voltage between  $-30$  and  $-40$  V, respectively, since the current arriving in the experiment is already sufficiently between  $40$  and  $50$   $\mu\text{A}$  for these settings, while the electron energy is still low enough to be easily manipulable by the low voltage electrodes (see Sect. 3.4.8).

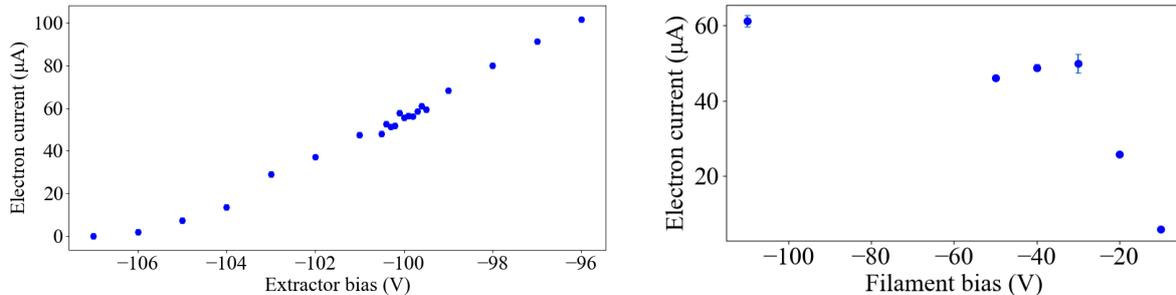


Figure 5.3: Systematics of the observed AEGIS electron gun current. Left: The current recorded by a multimeter at the downstream end of the experiment plotted as a function of the extractor bias voltage, for a constant filament bias of  $-110$  V. Right: The digitized MCP-FC front face signal as a function of the filament bias, with the extractor voltage adapted in each step to be 10 V below the filament bias.

The new electron gun has been integrated into the new control system of AEGIS and is fully programmatically controllable. While the filament is kept at a constant bias voltage by a power supply, the extractor bias is provided by one of the Fastino/amplifier channels of the Sinara control electronics (see Sect. 4.1.1) as needed. Thanks to a newly installed electric relay switched by a Sinara TTL channel on the 1 T side, the extractor bias can be programmed both from the 5 T and from the 1 T Kasli controllers: as a default, the 5 T channel controls the extractor bias, but the 1 T side can take control via the switching-on of the corresponding TTL line. This is needed to provide cooling electrons to both trapping regions and for tests of the electrodes and detectors on both sides of the experiment. The cathode heating current is managed by a dedicated LabVIEW™  $\mu\text{Service}$  (see Sect. 4.4.2) that has been developed to control the current and voltage of the connected power supply remotely. The same  $\mu\text{Service}$  also monitors the emitted current of the electron gun continuously via the Keithley 2100 6 1/2 digit multimeter ("Emitted current monitor" in Fig. 5.2) and sends this information to the data acquisition system upon request by the user (e.g. automatically at the correct time in the run routine).

The monitor of the extracted current is only connected for dedicated verification measurements and not usually used because it interferes with the passage of the particles through the apparatus, which is generally needed for experiments. The "collector" refers to the detector on the downstream end of the experiment that monitors the arriving charges. In standard configuration, this is the front face of the MCP, operated as a charge collector on its metallized front face and digitized by the NI 6133 oscilloscope with the  $10$  k $\Omega$  resistor (the MCP-FC detector, see Sect. 3.4.9), but it is equally possible to connect a multimeter, e.g. to verify the functionality.

In a first step, it is verified that the current recorded on the MCP front face digitizer at the end of the experiment is compatible with the reading of the emitted current monitor

in the front for different bias voltage settings. This is found to be the case within  $\pm 1 \mu\text{A}$  for the used configurations, with the digitizer signal typically slightly below the emitted current.

### 5.1.2 Re-conditioning of the electron gun

During the times when no antiprotons are available from ELENA and development work is ongoing in the experiment, there are typically extended periods of time in which the electron gun is not being used and is moreover not kept in vacuum conditions. In order to be able to use the gun efficiently again during the next run time, a re-conditioning procedure is needed which "wakes up" the electron gun and furthermore cleans it. The procedure which has proved to work effectively in AEgIS, with the electron gun inserted in the beam line, is detailed here.

1. The extractor bias voltage is set to a value that is sufficiently low to block the extraction current for the normal operation voltage of the cathode filament ( $-50 \text{ V}$ ), e.g.  $-100 \text{ V}$ .
2. The bias voltage of the cathode filament is set to  $-50 \text{ V}$ .
3. The cathode heating current is gradually ramped up to the final operating value (e.g. using the following steps:  $0.1 \text{ A}$  -  $0.5 \text{ A}$  -  $0.75 \text{ A}$  -  $0.9 \text{ A}$  -  $1.0 \text{ A}$  -  $1.05 \text{ A}$  -  $1.1 \text{ A}$  -  $1.15 \text{ A}$  (final value)). At each step, the vacuum of the experiment is observed to ensure that the pressure spikes are not going beyond a difference of an order of magnitude and the setting is left at each step for a few minutes to allow for the vacuum conditions to resettle. From around  $0.9 \text{ A}$ , significant vacuum spikes can be observed at each step.
4. The extractor bias voltage is set to  $-40 \text{ V}$  and then gradually increased (e.g. in steps of  $40 \text{ V}$ ) until an extracted current of the order of  $100 \mu\text{A}$  can be observed (this typically happens at an extractor voltage of around  $80 \text{ V}$ ).
5. The  $100 \mu\text{A}$  current is left on for one minute while carefully monitoring and then switched off by reducing the extractor bias voltage (e.g. to  $-100 \text{ V}$ ) and slowly ramping down the cathode heating current.

After re-conditioning, the amount of extracted current generally drifts slightly for the first 30 minutes, meaning that a stable operation requires a break before taking measurements. When the electron gun is not used for multiple months at a time (e.g. during the annual interruption of the CERN accelerator chain), it may become necessary to increase the cathode heating current up to  $1.24 \text{ A}$  (following the standard procedure for the activation of barium oxide filaments) to create free barium again. In this case, as a precaution, the difference of the extractor bias and the cathode filament voltage should remain no higher than a few tens of volts.

## 5.2 Tests of the new antihydrogen production trap

Following the installation of an entirely new antihydrogen production trap in the experiment in 2022, it was important to ensure the correct connection and functionality of all its electrodes, including the operation of high voltages in vacuum as well as the programmatic setup in the control system. As electrons are those particles in AEgIS that are most

easily available on demand after the commissioning of the new electron gun, they were the chosen candidate to perform the corresponding checks.

The straight-forward way of checking the connection of the electrodes is to block a stream of electrons with a given energy - depending on the exact settings of the electron gun filament and extractor bias - systematically by each electrode in turn. The procedure is conducted as follows: The stream of electrons from the electron gun is sent through the experiment continuously for five seconds and the extracted current monitored either with a multimeter or as the charge arriving on the MCP front face, digitized by the 6133 device. For the first second, all involved electrodes are kept at zero potential. A high enough voltage to block the passing electrons ( $-150\text{ V}$  were used to be on the safe side with the different settings) is then applied on one of the electrodes for another second, lowered to zero again, and left at zero for the last three seconds of the passing electron current.

An example of such a test on one of the electrodes of the new 1 T trap performed with the digitizer is shown in Fig. 5.4. If the procedure works as expected and the electrode is correctly connected and configured, the electron current ( $-125\text{ }\mu\text{A}$  in the shown example) can be observed on the multimeter during those times when the electrode is at  $0\text{ V}$  and disappears when ramping up the electrode potential. The inset plot shows a zoom to the range of time during which the voltage ramp (i.e. also the reduction in current) happens. The change is shown to happen within  $1\text{ ms}$ ; however, this is simply the acquisition rate set on the digitizer - the actual ramp happens much faster (tens of microseconds, for details see Ch. 4) within that time window and is not resolved.

This procedure is systematically repeated for all of the electrodes of the new trap (and additionally for those of the existing antiproton capture trap in the  $5\text{ T}$  region because it is newly connected to the control electronics as well) and their correct operation verified. Generally, it can be launched whenever needed to survey the whole functionality chain.

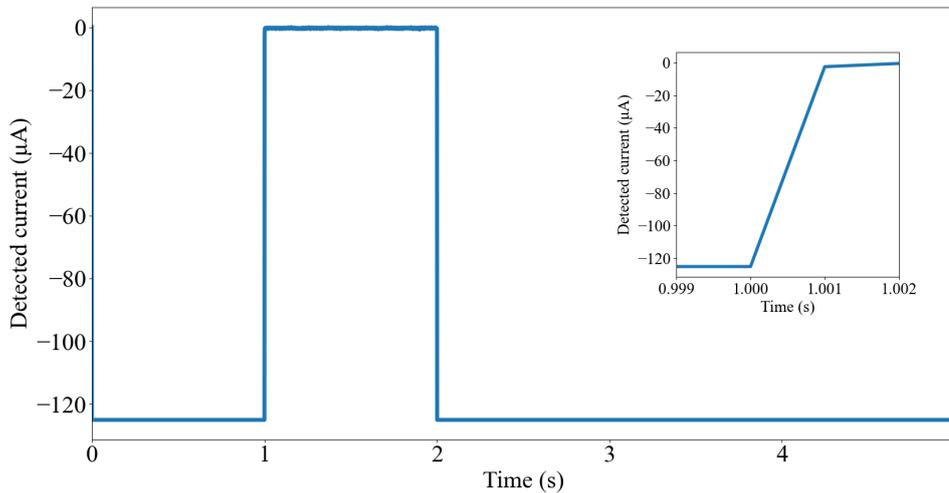


Figure 5.4: Successful test of connection and functionality of a trap electrode using the electron gun current. Plotted is the current observed by the digitized MCP front face versus time. The electron gun extracts continuously over the entire measurement time. First, all electrodes are left at  $0\text{ V}$  for one second. For the subsequent second, the voltage on the electrode under test is set to a high, negative value ( $-150\text{ V}$  in the shown case), before being ramped down to  $0\text{ V}$  again. The current arrives on the MCP when no negative voltage is applied; otherwise the current is blocked and does not arrive. The inset plot shows a zoom to the time range of the voltage ramp.

A further way of testing the entire setup and confirming its correctness is to individually load electrons into the different trapping regions of the experiment, leaving them in the traps for a few seconds (without manipulating them in any particular way), and then dumping them towards the MCP at the downstream end of the experiment: if the electrons can be observed at the correct time, there are no obstacles such as interfering stray potentials due to disconnected cables etc. and the electrodes behave as expected, at least on a broad scale. The procedures for the electron loading and dumping are explained in Sect. 5.3.

Both of the above procedures are regularly performed before commencing antiproton work after a break and in case of unexpected behaviour of any of the antimatter particles inside the experiment.

### 5.3 Experimental alignment procedures using electrons

In the course of the plasma preparation steps involving electrons, a new procedure has been developed to align the electron gun current on the axis of the experiment.

This procedure consists of the comparison of the exact positions of electrons arriving on the MCP at the end of the experimental apparatus: once directly from the electron gun and once after confining the electrons with the electrodes in one of the available trapping regions and then dumping them towards the MCP. The MCP itself can be operated either as a charge collector with its metallized front face (MCP-FC detector) or as an imaging device by enabling its amplifying voltage and acquiring the image of its phosphor screen (see Sect. 3.4.9 for details).

To begin with, electrons are loaded into one of the trap regions. The so-called C Trap (formed by electrodes C5 to C16, see Sect. 3.4.8) is used in the figures shown here. The incoming electrons have an energy of the order of 40 eV, as determined by the setting of the electron gun filament bias voltage. Fig. 5.5 outlines the reshaping operations of the electrode potentials used for the described procedures, i.e. the transition from an initial (cyan dashed line) to a new potential (blue solid line). The potentials are shown according to the voltage amplitude at each given relevant position along the axis of the experiment, from 1 m to 40 cm upstream of the defined center of the experiment (see Sect. 3.4.1). They take into account the lengths of the electrodes, with which the actually produced potential scales with respect to the programmed one, and the influence of the neighbouring electrodes, as obtained from a finite element calculation (see Sect. 3.4.8 and Appx. E). The locations and movements of the electrons are shown in green in the plots as well (only indicatively, not to scale). The transitions are performed using the defined operations of the AERIALIST control system (in the `Plasma Library`, see Ch. 4). As displayed on the top left of Fig. 5.5, to let the electrons enter into a trap, a positive potential well is formed by the "inlet" and "floor" electrodes (typically at 190 V), while a negative voltage is applied on the "outlet" electrode (of the order of  $-190$  V). The arriving electrons are reflected by the negative potential, collide inelastically with the onward stream of incoming electrons from the gun, losing some of their energy, and fall into the positive well. After several seconds of this loading procedure, the trap is closed by setting the same negative potential on the inlet electrode (top right of Fig. 5.5). The plasma is confined for a few seconds and then dumped towards the MCP by modifying the axial potentials as shown in the bottom row of Fig. 5.5: Firstly, the potential is reshaped to a well formed by negative voltages (left), including the floor electrodes, which

are also slightly negatively ( $-10$  V) biased, to allow for a smooth extraction. Secondly, the voltage on the outlet electrode is switched off to let the electrons escape downstream towards the MCP. Their arrival position is recorded via the light information of the fast phosphor screen on the back of the MCP by a CMOS camera image (see Sect. 3.4.9), marked with a cross on the shown image, and left in place on the camera monitor for the next acquisition.

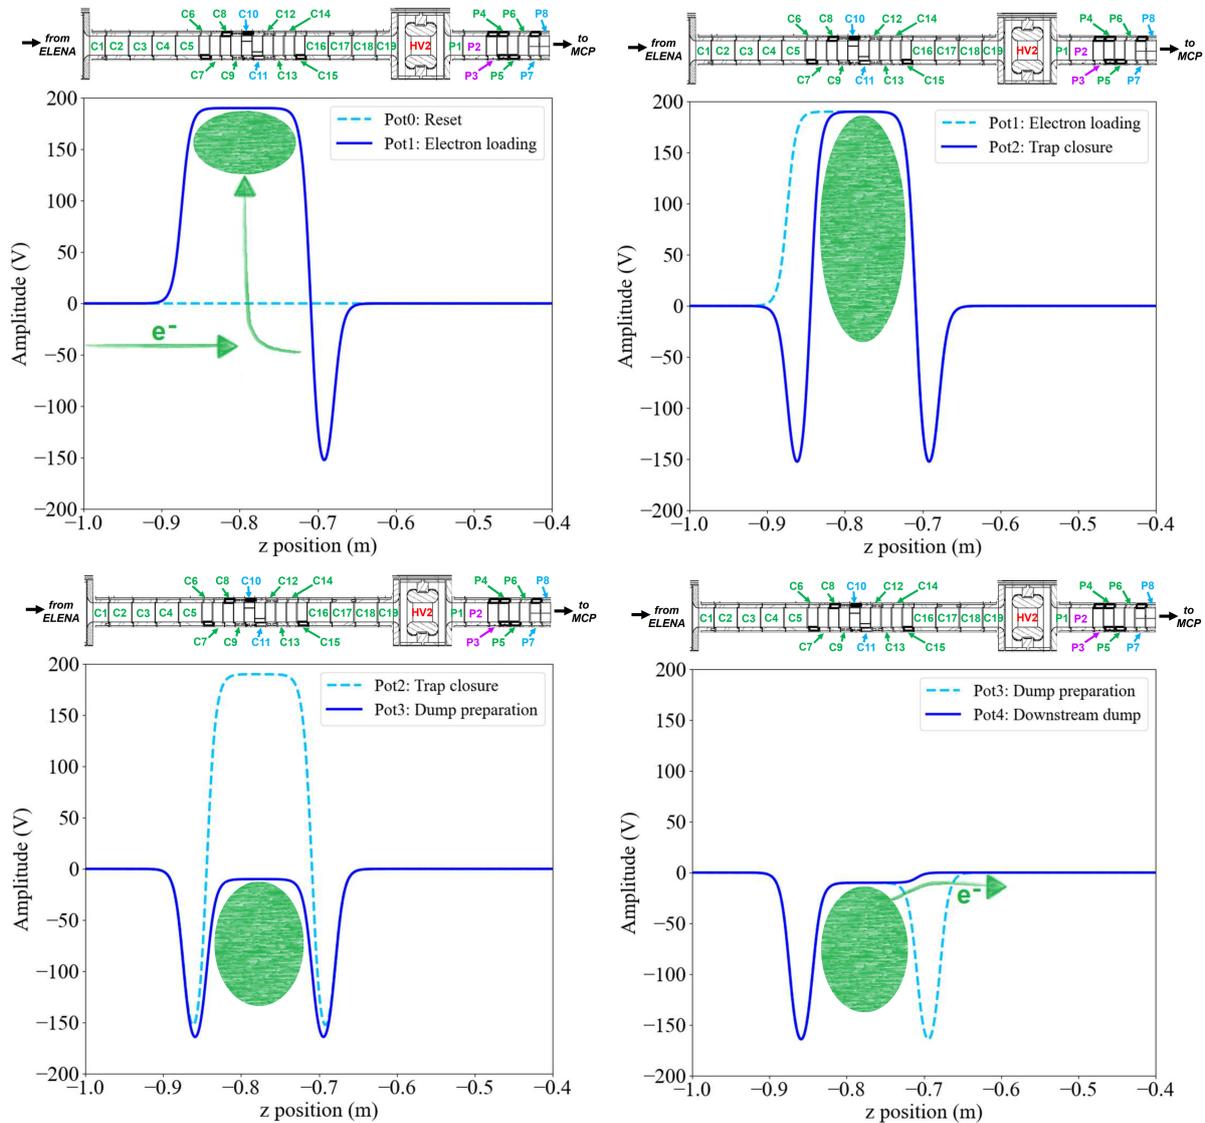


Figure 5.5: Potential reshaping operations to load electrons into the C Trap and dump them towards the MCP. Initial potentials are shown by the cyan dashed lines; the final potentials are the blue lines. The electron plasma in the final potential is schematically added in green in the plots; it is only meant to indicate its location (ellipse) and direction of movement (arrow) and is not drawn exactly to scale but to give a rough indication of the plasma expansion. Above each plot, a schematic of the involved range of trap electrodes is included, scaled to match the dimension of the x-axis, with the colors of the electrode labels having the same meaning as in Fig. 3.9. Top left: Electron loading into a positive well after reflection on a negative barrier. Top right: Closing the trap. Bottom left: Transition to a well formed by negative voltages. Bottom right: Opening of the outlet electrode to dump the electrons in the downstream direction.

---

The electron loading procedures depicted here are also the ones used to prepare electron plasmas for antiproton cooling, which is described in Ch. 6.

For the direct detection of the current on the MCP, whose metallized front face is used as a Faraday Cup charge collector here in the MCP-FC configuration, short bunches of  $1 \times 10^9$  electrons (electron gun pulse length  $\sim 10 \mu\text{s}$ ) are emitted from the gun and their position is recorded in the same way as for the plasma. According to the cross position of the electron plasma arrival, the position of the electron gun can be modified gradually, using rotation motors in all three dimensions, in between the bunches until the arrival position of the direct electron current matches that of the plasma from the trap. The left image of Fig. 5.6 shows the arriving direct current on the camera  $\mu\text{Service}$  together with its cross marker and the markers for the determined trap centers during the procedure.

The repetition of this procedure with the different possible trapping regions also presents an alignment check for the entire experimental axis. It can be observed that, according to the CMOS calibration explained in Sect. 5.4.1, the direct current and the centers of the C and P Traps are aligned to each other to within less than 1 mm, while the A Trap, which is further downstream, is misaligned by several mm. It is possible to adapt the angles of the new 1 T with the cryogenic actuators to improve the alignment.

Such a steering procedure is now performed routinely at the start of a new beam time.

A steering procedure using "electron images" is also used to align the arrival position of the positrons to the location of the positronium conversion target. The Ps target can be moved in and out of the trap axis on its holder with cryo motors via a dedicated  $\mu\text{Service}$ . Its position on that trajectory can be determined by the camera image, as described above: Electrons can be loaded into one of the trapping regions and subsequently dumped towards the MCP, passing the Ps target electrode. Their arrival on the MCP, operated nominally as an imaging detector using its amplification voltage and phosphor screen, then produces the image of their shape, with a square (or a shape depending on the level of insertion of the target, plus the thin holding structure) "shadow" remaining dark where the passing electrons are blocked by the target. For its fully inserted position, the center of this shadow can again be determined in the same way used for the electron current alignment and marked on the software monitoring the camera image. Left in place, the corresponding marker thus serves as the intended location for the positrons, which can be steered to arrive at this position. An image of this procedure on the camera  $\mu\text{Service}$  is shown on the right hand side of Fig. 5.6. It has proved very useful to determine the location of components in the beam line inside the cryostat, where the use of regular cameras is rendered very challenging by the cryogenic temperatures.

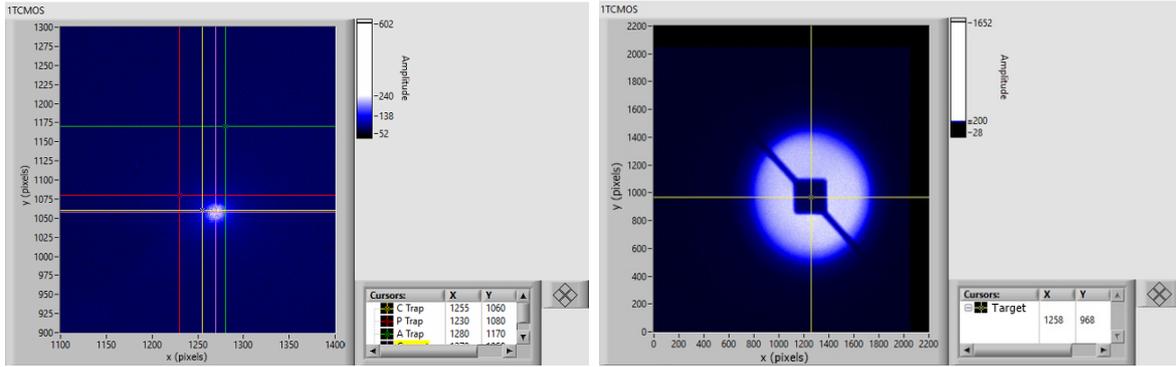


Figure 5.6: Images of the CMOS camera  $\mu$ Service during experimental alignment procedures using electron imaging. Left: Alignment of electron gun current (purple cross) and trap centers (yellow: C Trap, red: P Trap, green: A Trap). Right: Steering of the Ps target.

## 5.4 Detector calibration using electrons

Since the electrons are so readily available and comparatively well under control, they can be used to calibrate detector components, both by simply using them as a way to generate images on the CMOS camera and by using their signals on the detectors as references.

### 5.4.1 CMOS calibration

As explained in Sect. 5.3, the Ps target can be imaged by the passing electrons on the MCP CMOS camera when it is inserted on the experiment axis. Since the dimensions of the target are very precisely known ( $5 \times 5 \text{ mm}^2$  with its holder included), an image analysis of the produced shadow can also be used to perform the pixel to millimeter calibration of the camera itself, which is extremely useful to be able to determine radial dimensions of any cluster of particles arriving on the MCP, such as the radii of the arriving antiproton beam and plasmas dumped from the traps.

According to a visual analysis, the described conversion factor is determined as  $46.42 \pm 0.12 \text{ px/mm}$ . Using this conversion, the electron current arriving directly from the gun has a radial extension of the order of  $0.4 \times 0.4 \text{ mm}^2$ . This conversion is confirmed by another electron image: When filling up an entire trapping region (e.g. the "C Trap") with electrons and releasing them subsequently, the diameter of the plasma as imaged on the MCP/CMOS is found to be of the order of 30 mm, in very good agreement with the diameter of the trap electrodes.

### 5.4.2 MCP as an amplified Faraday Cup: an initial exploration

As discussed previously, the metallized front face of the MCP is frequently used as a Faraday Cup charge collector, without applying a bias voltage across it, i.e. without the gain element. This use of the MCP detector is referred to as "MCP-FC" here. For the electron current from the gun, this feature can be exploited to determine the number of arriving electrons. However, for other relevant particle types used in AEGIS, in particular antiprotons and positrons, the traditional Faraday Cup detection technique for the determination of the number of particles is not usually applicable due to the typically low available numbers of these antiparticles. The integration of amplifiers in the readout

chain, on the other hand, would have posed a threat to the electronics in case of excess charge for a too high amplification value, given the available amplifier choices: It would have been necessary to swap the used amplification value frequently for different measurements. Logistically, a safe operation in such a configuration would not have been efficient. Therefore, as an alternative, the calibration of the non-linear MCP gain amplification using electrons has been investigated: Having verified the compatibility of the digitizer reading with the emitted current, the digitized MCP front face signal can be calibrated using known the well-tunable electron current to obtain a conversion from MCP signal to the amount of deposited charge, which can then also be used to understand the signals generated by other particles in the experiment, whose intensities are not or not as precisely controllable.

The overall calibration procedure has been developed but the results reported on here are preliminary, and more systematic studies need to be carried out during 2024.

### Determination of the arriving electron numbers

First, it is important to gain an understanding of the amount of charge arriving on the detector for a given electron current per unit of time and whether the relation of number of detected electrons versus extraction time of the current is linear, as would be expected for an ideal source. For this purpose, the Faraday Cup technique alluded to previously is used. Since the bias of the extractor is controlled by the AERIALIST and connected to a Pulser channel (see Ch. 4), it is possible to precisely control the time length of the electron gun extraction in pulses, with a minimum pulse duration of 10 ns and a maximum of 30  $\mu$ s. A scan is thus performed over the electron current pulse length, leaving the filament bias at  $-110$  V and the extractor voltage at  $-115$  V, which is then pulsed for the desired duration with a positive 15 V to extract electrons with the default 10 V difference. A pulse rise and fall time of the Pulser of the order of 5 ns affects these measurements, especially for short pulse durations. From the observations described in Sect. 5.1.1, it is expected that the observed current should settle to around 60  $\mu$ A after an initial increase and fluctuations for long pulse lengths.

The digitized MCP front face signal is obtained for electron pulses over the entire range of possible Pulser settings, applying no bias voltage on the MCP. The measured voltage integral over the 10 k $\Omega$  is expected to be directly proportional to the deposited charge, measuring a negative potential peak from the current slowly flowing to ground from the charged up front face. Fig. 5.7 shows example distributions obtained with different pulse lengths, using a 2 ms acquisition window on the digitizer and a 0.4  $\mu$ s sampling rate. For a pulse length of 10 ns, no evident peak can be observed. For longer times, it starts to form and is already quite well defined at 50 ns, becoming clearer the longer the pulse gets. To compare the signals for different pulse durations, an integral over the signal in a commonly defined time window is used. The integration window is indicated by the vertical red lines in Fig. 5.7. For the analysis, the electronic noise background is subtracted, whose rate is determined from the counts within a time range before the peak signal itself. The vertical grey lines mark this window in the different distributions.

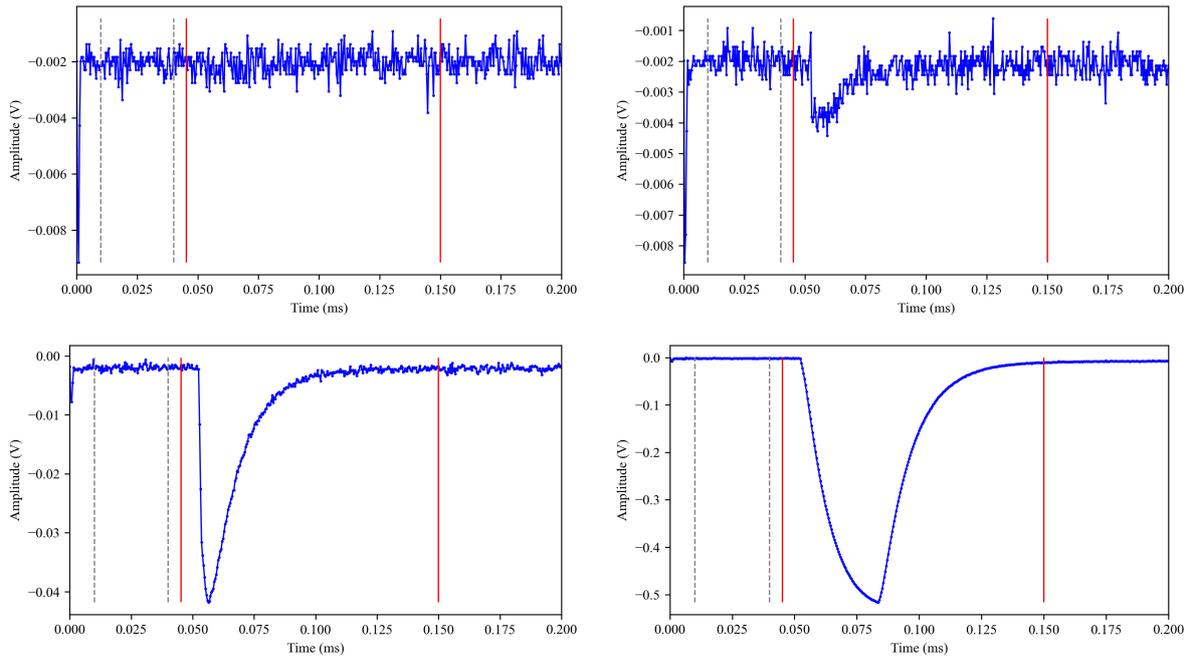


Figure 5.7: Distributions obtained by digitizing the MCP front face signal from an impinging electron current with different pulse lengths. The grey dashed lines indicate the range used to determine the background. The red lines mark the integration window. Top left: 10 ns. Top right: 50 ns. Bottom left: 1  $\mu$ s. Bottom right: 30  $\mu$ s.

The measurement with each pulse length is repeated three times. Fig. 5.8 shows the mean values and associated standard errors of the corresponding integrated signals for the different pulse durations. As the integral values themselves are of course negative (the negative spike from the deposition of the negatively charged electrons, reducing the potential), their negated values are plotted. Since the digitizer is used with an  $R = 10 \text{ k}\Omega$  resistor, it is possible to convert the read voltages  $U$  to currents  $I$  by applying Ohm's law, Eq. 5.1. From the integrated signal, one can thus obtain the total detected charge in Coulomb. By dividing through the elementary charge of an electron ( $\sim 1.6 \times 10^{-19} \text{ C}$ ), the number of detected electrons can be determined. The resulting values are given by the second (right hand side) vertical axis plotted in Fig. 5.8.

$$I = U/R \quad (5.1)$$

While the increase of the signal looks approximately linear for longer durations, the data is better fitted by a second order polynomial. Particularly large deviations are observed for short pulse lengths. This will become evident in the subsequent paragraph and can be attributed to the influence of the Pulser rise time. However, such short pulse lengths are used to intercalibrate the signal with other particles because it is safe to operate the MCP at nominal bias voltages without damaging it only for these low enough numbers of arriving charges, given the electron energies compatible with those of the positrons. In fact, positron pulse lengths of the order of 20 ns are used. Therefore, it is still essential to have this knowledge of the amount of current detected for these short pulse durations.

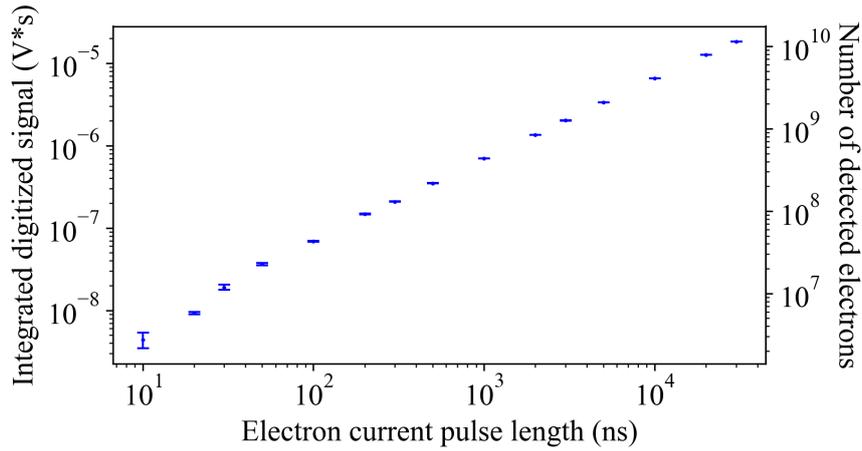


Figure 5.8: Integrated digitized MCP front face signals (negated values, left vertical axis) and corresponding detected numbers of electrons (right vertical axis) as a function of the electron gun current pulse length.

Via the division of each integrated charge value through the duration of the corresponding pulse, the detected electron current can be determined. The resulting values are plotted in Fig. 5.9. If the relationship in Fig. 5.8 really was a linear one, the current would be found to be a constant value. Instead, an increase in detected current can be observed with the first few increases in pulse length, which is expected due to the greater impact of the Pulser rise time: For shorter pulse lengths, the first and last few nanoseconds with lower emitted current have a stronger relative influence. For longer pulse durations, the observed values settle around  $60 \mu\text{A}$ , which is compatible with the continuous current produced after longer times (see Fig. 5.3), with saturation effects also visible. This finding means that it is not possible to use a longer pulse or direct current measurement to calculate the number of arriving electrons more generally. Instead, for intercalibration studies that use short-duration pulses, the number of detected electrons has to be determined from the shown measurements individually for each required pulse length (or, alternatively, via an interpolation).

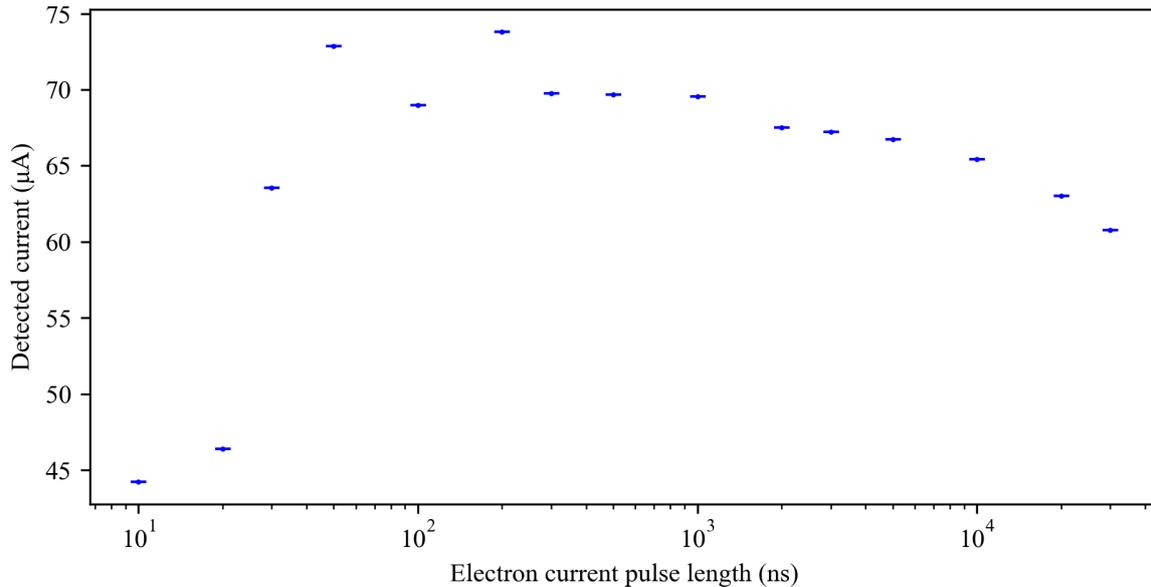


Figure 5.9: Observed current on the MCP front face as a function of the electron gun pulse length.

### Calibration of the MCP-FC using the gain signals

For regular operation, a bias voltage of the order of 1.4 kV is generally applied across the MCP (i.e. between front (lower positive voltage) and back (higher positive voltage) of the MCP) to obtain a good signal. This voltage difference between the two plates creates an electric field that accelerates the produced secondary electrons: In a strong enough field, they in turn liberate additional electrons from the MCP channel walls and charge multiplication, referred to as a cascade, is produced. This multiplication, i.e. the MCP gain, can be tuned by adapting the applied bias voltage.

From Fig. 5.7, it is evident that, for low numbers of arriving particles (i.e. very short electron pulse lengths), no negative peak from the electrons passing through the resistor to ground is observed in the digitized MCP front face signal. When applying a strong enough bias voltage, however, a positive peak corresponding to the charge amplification gain becomes visible, which is caused by the passage of electrons from ground (again, slowly through the resistor) to the positive front face (depleted of electrons additionally through the amplification process). Thus, another source of non-linearities for longer pulse durations is of course the partial balancing of the contributions and therefore of the positive and negative signals.

From the knowledge of the arriving current at different electron pulse lengths (the measurement above, Fig. 5.8 and Fig. 5.9), the relationship of detected charge versus current or, rather, number of arriving electrons, can be determined when performing a similar scan over the pulse lengths with such a voltage applied. Using the same integration procedure described above, the total signal can be extracted. Fig. 5.10 (left) shows an example signal obtained with a pulse length of 20 ns and an MCP voltage of 1.4 kV. The obtained integrated signals are shown on the right hand side of Fig. 5.10 as a function of the used short pulse lengths between 20 and 140 ns. The displayed data points are the mean values of two or three measurements each with the corresponding standard errors. Data with pulse durations of 80 and 100 ns have also been acquired but an error in the

DAQ has prevented the correct storage and corresponding analysis, the same feature that also limits the number of used repetitions to two for some of the shown points. It is noteworthy that only at these short pulse lengths, i.e. low particle numbers, it is safe to use the MCP with the applied bias; otherwise, there is a risk of damage to the front face.

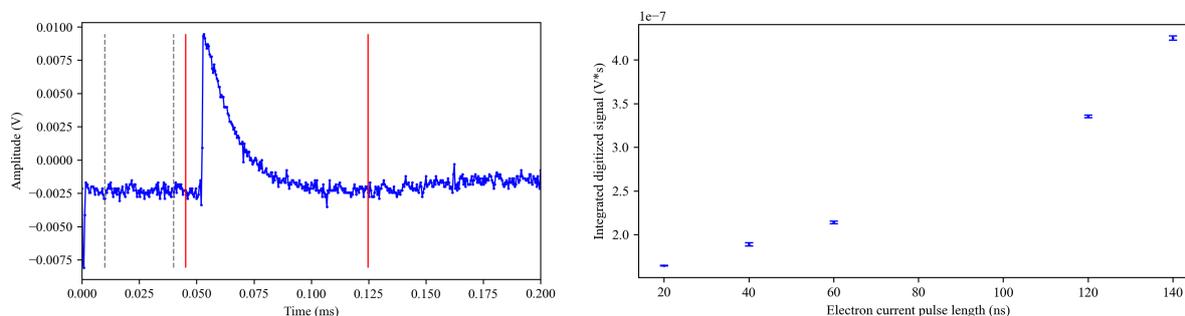


Figure 5.10: Observed MCP front face signals from an impinging pulse of electrons with an MCP front-to-back bias of 1.4 kV. Left: Example of a signal from a 20 ns pulse. The grey lines indicate the region used for background subtraction; the red lines mark the integration window used to obtain the signal strengths at different pulse lengths and MCP voltages. Right: Plot of the integrated signals as a function of a range of short electron pulse durations.

Fig. 5.11 shows the obtained signal with a 20 ns electron pulse as a function of the applied MCP bias voltage (mean values and standard errors of three measurements each): For voltages below 1.1 kV, no peak can be identified. From then on, the gain signal starts to develop until the clear peak shown on the left hand side of Fig. 5.10 for a voltage of 1.4 kV forms. Above scans of the electron current pulse length have been performed with the standard operation voltage of 1.4 kV; for an intercalibration at different bias voltages, such scans would have to be repeated, which is done within minutes.

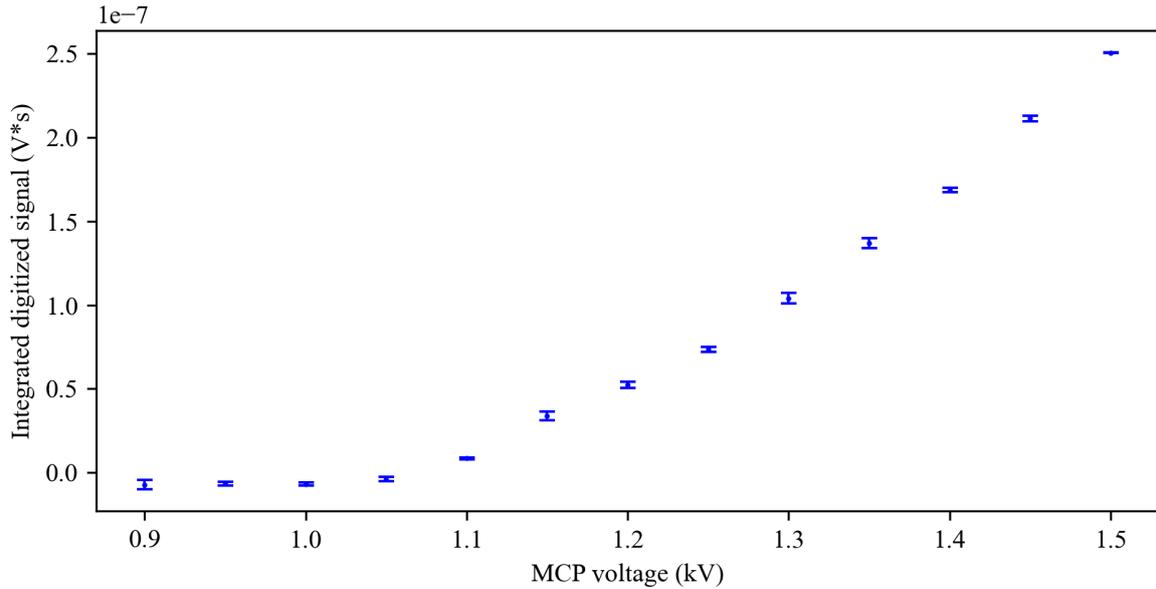


Figure 5.11: Integrated MCP front face signals as a function of the front-to-back MCP voltage for impinging 20 ns electron pulses.

The signals obtained with the applied bias voltage, i.e. with gain amplification, also depend on the radial extension of the arriving bunch or plasma of particles, as more channels get addressed and contribute to the cascade in an approximately linear fashion. As a secondary effect, if more channels of the MCP are covered by the stream of particles, a smaller saturation effect is produced in the individual channels. The saturation itself also depends slightly on the used bias voltage (which can be the topic of an additional, dedicated study). Therefore, the electron signals obtained here should be normalized to the covered area for the calibration, as should then those from the measurements whose (anti-)particle numbers are to be determined. This step makes the signals of different particle types actually comparable, since the direct electron current is extremely focused, for example with respect to the arriving positrons.

Fig. 5.12 shows an example analysis of the CMOS image of the arriving electron current. On the left hand side, an example image obtained with the CMOS camera from the phosphor screen of the MCP is displayed (normalized to a background value of 1, with the normalization factor obtained within the small red square drawn in the image), showing the arriving electron pulse for a pulse duration of 20 ns, with an MCP bias voltage of 1.4 kV. The cause of the shown distribution is the fact that the electrons do not enter the experiment perfectly on axis (simply because it has been impossible to align the electron gun perfectly in the manual way that is currently the default); instead, they come in at a small, not exactly known angle, which depends on their exact momentum, leading to a rotating trajectory that produces a main spot on the MCP with a slightly delayed tail. From this image, a value,  $S_{80\%}$ , is obtained, which corresponds to the integrated pixel intensity of all pixels with values of or above 80% of the determined peak intensity. This parameter is chosen to reduce the saturation effect contribution. The peak intensity is defined as the mean intensity of the twenty highest pixel values in the signal window (large red rectangle in the image), excluding outliers which are off by more than five standard deviations. For the presented electron gun pulse length scan, of the order of 300 pixels fulfill this criterion, slightly increasing in number with increasing pulse durations. This

analysis is performed for all available pulse length measurements (the same data used for Fig. 5.10). A plot of the resulting  $S_{80\%}$  value, normalized to the included MCP area (i.e. number of included pixels), is shown on the right hand side of Fig. 5.12.

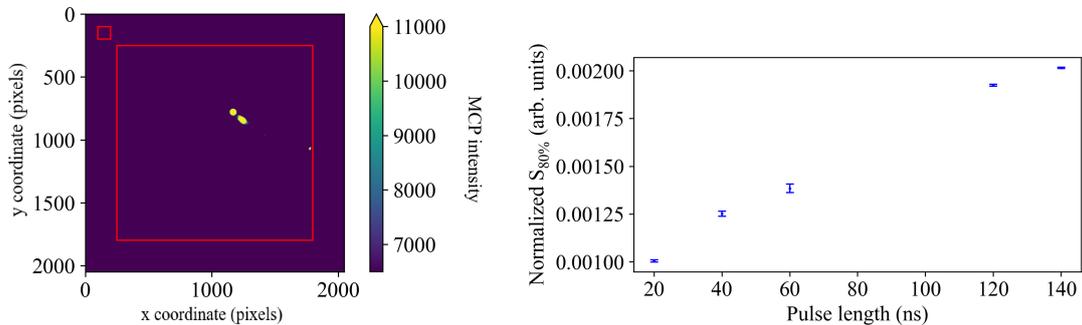


Figure 5.12: Analysis of the CMOS image of the MCP's phosphor screen stemming from the arriving electron current. Left: Example image taken of an arriving 20 ns electron pulse. The small red square in the top left corner indicates the window used for background subtraction. The larger red square is the signal region for the determination of  $S_{80\%}$  (see main text). Right: Plot of  $S_{80\%}$  as a function of the electron current pulse length.

Employing this preliminary calibration, the MCP can essentially be used as an amplified Faraday Cup. While the use as a non-amplified FC is common, the amplified usage requires a calibration with a known number of arriving particles. The well-controlled electrons available in AEGIS represent a candidate to perform such a calibration with. In this way, the door can be opened for MCP-FC measurements of (anti-)particles whose signal is too weak to determine their numbers without the amplification: The MCP gain significantly increases the sensitivity but eliminates the possibility for a direct counting of the incoming particles only.

### Determining absolute positron numbers from the MCP-FC calibration

The determined values for electrons can be compared to those of positrons to draw conclusions on the number of arriving positrons for a known number of detected electrons. Fig. 5.13 shows a comparable measurement for positrons, which, in AEGIS, have an unchangeable pulse length of the order of 20 ns. In this measurement, the positrons have an energy similar to that of the electrons ( $\sim 110$  eV), and an MCP bias voltage of 1.4 kV is used. On the left hand side, an example CMOS camera image of the phosphor screen is shown, taken upon the arrival of positrons on the MCP. The background and  $S_{80\%}$  signal are determined from it in the same way as for the electrons and the signal is normalized to the number of pixels included in  $S_{80\%}$  (in this case of the order of 5000 pixels). Compared to the image obtained during the electron run, the positron image seems brighter and furthermore shows a rim, originating from the reflection of the trap walls. This is due to the auto-adjustment of the camera settings according to the lower intensity of the positron bunch and is taken into account in the background subtraction.

The procedure is repeated ten times with the same settings, and the plot on the right hand side of Fig. 5.13 shows the obtained values as a function of the iteration number. The mean value determined from these runs is  $S_{80\%} = 1.85 \times 10^{-4} \pm 3 \times 10^{-6}$ . This can be compared to the number found for the electrons with these settings,  $S_{80\%} = 1.005 \times 10^{-3} \pm 6 \times 10^{-6}$ ,

extracting a factor of around 5.4 higher signal per area for the electrons. According to the Faraday Cup measurement shown in Fig. 5.8 and Fig. 5.9, a 20 ns electron pulse leaves around  $(5.8 \pm 0.2)$  million electrons impinging on the MCP. Division through the signal-per-area factor yields 1 to 1.1 million positrons. While turning this technique into an actual calibration certainly requires further systematic studies and much improved statistics, this value is in agreement with expectations from previous studies of the positron system, the source intensity at the time of the measurement, and trapping, cooling, pulsing and transport efficiencies [85].

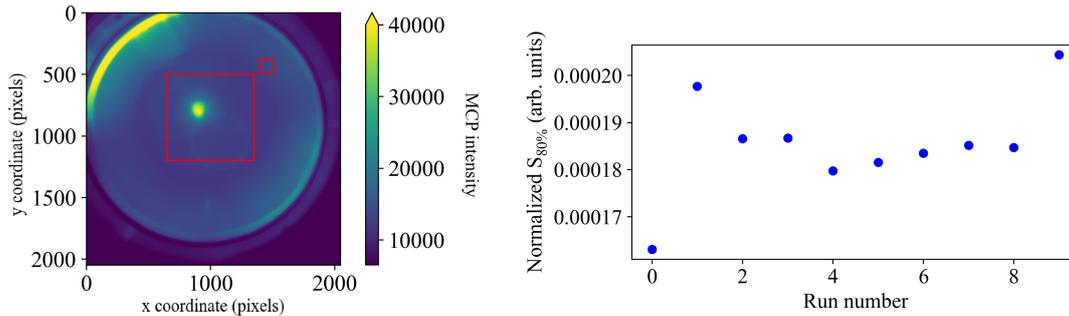


Figure 5.13: Analysis of the CMOS image of the MCP's phosphor screen stemming from arriving positrons. Left: Example image taken of an arriving positron bunch. The small red square indicates the window used for background subtraction. The larger red rectangle is the signal region for the determination of  $S_{80\%}$  (see main text). Right: Plot of  $S_{80\%}$  as a function of the run number.

One additional issue that should be kept in mind is that the resistance of MCPs increases strongly for decreasing temperatures, essentially reducing their gain for high count rates and, thus, the detection efficiency [160, 161]. The MCP in AEGIS is operated in a cryogenic environment, meaning that its performance is subject to such variations, which are currently the topic of systematic studies.

## 5.5 Summary

A new, double-head electron gun has been commissioned in the AEGIS apparatus. As its electrons form the basis of the plasma routines performed in the experiment, this  $e^-$  source has been incorporated in the new CIRCUS control system and carefully characterized, and a routine has been put into place to re-condition it when needed. All electron routines handled in the previous system have been converted and ported into CIRCUS.

For standard operation, the electron gun is operated with a filament bias voltage between  $-40$  and  $-50$  V and an extractor voltage between  $-30$  and  $-40$  V, providing a stable current of the order of 40 to 50  $\mu$ A.

With the  $e^-$  gun having been consolidated, new techniques have been developed and successfully implemented to test the functionality of the electrodes of the experiment's electromagnetic traps and to align them. Furthermore, a pixel-to-millimeter conversion of the downstream MCP's CMOS camera has been realized using "electron imaging", which is frequently used in the collaboration, and a novel procedure for the intercalibration of different particles' signals on the MCP front face itself, used as an amplified Faraday Cup, has been devised and preliminarily tested, for which further studies are currently ongoing.

## Chapter 6

# Efficient capture and accumulation of cold antiprotons in AEGIS

*This chapter presents the developments in AEGIS concerned with a record accumulation of cold antiprotons. They will be reported on in [162], a draft of which I have formulated as the first and corresponding author. I have designed the procedures needed on the hardware and software side to achieve this success, guided by the experience of R. Caravita. Together with R. Caravita, M. Volponi, J. Zieliński and T. Rauschendorfer, I have also been responsible for the involved data taking campaigns between 2021 and 2023. The data analysis for the antiproton accumulation and electron cooling systematics has been performed by me. The analysis of the antiproton capture data was led by T. Rauschendorfer with support from R. Caravita and me. A first publication regarding the record efficient antiproton capture from ELENA is also in the pipeline, with me as one of the main authors [163]; prior results from the optimization of the capture parameters have been published in [132], also with me as one of the main authors. The development and simulation of the degrader system were led by M. Doser and J. Zieliński.*

Antiprotons lie at the core of most research involving bound antimatter systems. For them to be useful for controlled experiments, they have to be relatively cold, i.e. in a first step lowly-energetic enough to be manipulable by voltages from standard electrodes (a few hundred volts to maximum tens of kilovolts) and eventually at sub-eV temperatures for precision studies. However, antiprotons can generally only be produced at high energies. For this reason, the objective of the work presented here has been the development and implementation of procedures, for the upgraded experimental setup of AEGIS and within the newly commissioned CIRCUS control system (see Ch. 4), for the capture of highly-energetic antiprotons from the CERN Antiproton Decelerator facility and their confinement and cooling in low-voltage potentials. For this purpose, various parameters of the experiment have been optimized, including the steering of the antiproton beam into the apparatus, the timing of and the voltages applied on the capture electrodes, the potentials used for the plasma confinement, and the conditions of the electron cooling. Furthermore, as the number of producible antihydrogen atoms in AEGIS scales with the number of available antiprotons, an additional goal has been the accumulation of as many cold antiprotons in the electromagnetic trap system as possible, stacking up multiple bunches arriving from the decelerators. This last step requires a solid understanding of the potentials as they are modified to move around and manipulate the trapped particles, which was not possible in such a straightforward way with the previous system. In general, to minimize plasma losses, trap configurations should be as symmetric and harmonic

as possible, and the reshaping of one potential configuration to another should happen relatively slowly, i.e. maximum on the order of milliseconds, to allow for an equilibration time and to avoid plasma heating and losses. Additionally, due to the very low efficiency in capturing the arriving highly energetic antiprotons from the Antiproton Decelerator itself, the accumulation of a large antiproton density in the trap has only become feasible after the commissioning of and connection to the lower-energy ELENA decelerator as well as the adaptation of the involved devices, including the energy degradation system and the electronics. For the accumulation, an optimization of the electron cooling and the plasma compression was mandatory in a prior step.

This chapter describes the procedures used for the efficient capture, cooling and accumulation of antiprotons in the upgraded AEGIS apparatus. For trapping antiprotons from the ELENA decelerator, record capture efficiencies of up to 70% are routinely achieved. The antiprotons are then sympathetically cooled with electrons in a low-voltage potential and compressed using the Rotating Wall technique [105–107]. Thanks to these developments, it has become possible to stack up many of the arriving ELENA bunches, accumulating over 100 million cold antiprotons in the small trap volume of less than  $100\text{ cm}^3$ . These achievements open the door to a large variety of fundamental research directly on the antiprotons and on large quantities of bound systems formed with them (e.g. antihydrogen, antiprotonic atoms), facilitating probes of basic symmetries involving antimatter, theories beyond the Standard Model of particle physics and Dark Matter hypotheses. For AEGIS in particular, the number of antihydrogen atoms produced from the employed charge exchange reaction depends linearly on the number of available cold antiprotons, making this result a crucial step toward increasing the flux of produced  $\bar{\text{H}}$  atoms for statistically significant measurements of the influence of gravity on antimatter.

## 6.1 Efficient capture of antiprotons

Fig. 6.1 shows a schematic of the electrode structure forming the 5 T trap used to capture antiprotons from ELENA and cool them sympathetically with electrons, a procedure already briefly outlined in Sect. 3.3.1. This structure is used to accumulate large numbers of cold antiprotons in a small volume of  $66\text{ cm}^3$ , confined by a set of seven low-voltage electrodes ("Short P Trap" between electrodes P3 to P9, see Sect. 3.4.8) located between the high-voltage electrodes HV2 and HV3.

The high-voltage electrodes (HV1-3) are highlighted in light red in Fig. 6.1 and can reach up to  $-14\text{ kV}$ . Generally, the voltage on HV3 is already ramped up prior to the arrival of the antiprotons (dark red block) and electrons are pre-loaded into a trap formed by the low-voltage electrodes. The reshaping procedures of the low-voltage potential are described in more detail in Sect. 6.1.1. The arriving ELENA bunch (with  $N_{\text{in}}$  as the number of antiprotons per bunch) then passes the degrader structure, inside which some of them are lost ("degrading losses",  $f_{\text{deg1}}$ ). The antiprotons that survive the moderation pass into the trapping region. Those whose energies have thereby been sufficiently reduced, are reflected by the potential of HV3, while the rest pass the structure and end up on the downstream MCP ("transmission losses",  $f_{\text{MCP}}$ ). After an optimized delay time (see Sect. 6.1.3) to maximize the number of trapped antiprotons, the same voltage used for HV3 is also ramped up on HV2 or HV1, employing a dedicated system based on a Behlke high-voltage nanosecond switch (dark red arrow as an example on HV1). Per default, HV1 is used as inlet but measurements have also been performed with HV2. Some of the antiprotons reflected by HV3 escape through the inlet before it closes ("reflected losses",  $f_{\text{deg2}}$ ), while the rest become trapped between the high-voltage electrodes

("trapped fraction",  $f_{\text{trap}}$ ) and eventually settle into the attractive low-voltage potential generally formed in the P Trap (also indicated in Fig. 6.1), where they start cooling through Coulomb collisions with the waiting electrons and are stored. During storage, a fraction of the antiprotons is lost, depending on the expansion of the combined plasma and resulting annihilations on the trap walls ( $f_{\text{store}}$ ). These losses can be mitigated by the use of dedicated plasma compression techniques (see Sect. 6.2.2). In principle, all loss contributions together with the trapped fraction of antiprotons should sum up to unity, as defined in Eq. 6.1.

$$1 = f_{\text{deg1}} + f_{\text{MCP}} + f_{\text{deg2}} + f_{\text{store}} + f_{\text{trap}} \quad (6.1)$$

Whenever needed, the voltage on the electrodes can be ramped down, releasing the confined particles ( $N_{\text{out}}$ ): Usually, the trap is opened in one of the two directions first to let the particles end up either on the downstream MCP for direct detection or on the degrader structure to detect annihilation products on the ESDA (see Sect. 3.4.9). The opening of the potentials formed by the high-voltage electrodes releases only those antiprotons that have not sufficiently cooled to remain in the low-voltage potential and are referred to as "hot". This procedure is therefore called "hot dump", and the preceding time of confinement in the high potential is referred to as "hot storage time". Equivalently, the ramping down of the low voltages to release all remaining particles is named "cold dump" and the low-voltage confinement "cold storage time".

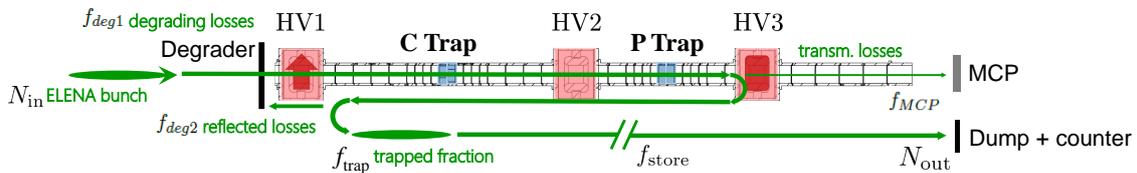


Figure 6.1: Schematic of the 5T trap structure used for the capture and cooling of antiprotons in AEGIS. The structure is formed by over thirty low-voltage electrodes and three high-voltage electrodes (HV1-3, marked in light red). The electrodes highlighted in blue are the ones employed for the Rotating Wall technique. The locations of the C Trap and the P Trap (between HV1 and HV2 or HV2 and HV3 respectively) are also indicated. Before entering the trap, the antiprotons arriving from ELENA pass the degrader structure. The MCP detector at the downstream end monitors the arriving charges.

### 6.1.1 Tail-less electron cooling plasma preparation

The reshaping procedures described here form the potentials given by the low-voltage electrodes from one configuration to another. This is done using the defined operations of the AERIALIST control system (in the Plasma Library, see Ch. 4). The software gives the option to perform these operations fast (i.e. as quickly as possible) or slowly. For the slow procedures, the voltages on all included electrodes are ramped up or down in steps, with the number of steps and the overall time of the operation as user-defined parameters. It is also possible to change potentials on one set of electrodes quickly and on another slowly at the same time and/or to and from different voltages. The slow operations are employed when it is important to minimize the disturbance of the trapped plasmas to limit

losses and plasma temperature increases. Generally, the faster the voltage ramp happens and the larger the steps are, the more disturbed the plasma becomes, as it has less time to become equilibrated in the new potential. Typical numbers of steps are around 100 and typical times range between hundreds of milliseconds and a few seconds for the operations described here. When dumping a large number of accumulated particles, up to several thousand steps in several minutes are also sometimes used, as this prevents saturation of the ESDA detector and allows to determine the particle energies. In general, all these values used for the slow potential reshaping operations yield a ramping speed of hundreds of millivolts to no more than a few volts every 10 ms. These gradients have proved to be sufficiently gentle while being compatible with the overall timing of the experiment; for significantly faster ramping speeds, increased particle losses are observed. Between the steps, typical delay times of 10 ns are given to the electronics.

The electrons employed for the  $\bar{p}$  sympathetic cooling are pre-loaded into the trap region prior to the arrival of the antiprotons. Typically the "Short P Trap" configuration between electrodes P3 and P9 (see Sect. 3.4.8) is used to cool the electrons. This range of electrodes is chosen because they all have the same dimensions, and a very symmetric, harmonic trap can thus be formed, minimizing disturbances of the confined plasma. Additionally, using only this range and not extending to electrode P12 has the advantage of being able to employ this same trap configuration also for the preparation of the antiproton launch procedure used for directional antihydrogen production, as described in Ch. 7. In principle, the same procedure outlined in Sect. 5.3 for the C Trap is also used here. However, to make the number of used cooling electrons and thus the plasma extension and cooling efficiency reproducible as well as to prevent the formation of tails far away from the plasma core, a two-step electron loading sequence is introduced [108]: Electrons are loaded into the C Trap in exactly the same way, but instead of dumping the entire trap content directly into one of the two directions, only a controlled fraction of the confined electrons is released downstream toward the P Trap. This fraction is determined by the value by which the outlet electrode of the C Trap is "lowered" (from a highly negative value toward zero<sup>1</sup>) for the release of a part of the space charge. The value is generally chosen to be between 30 and 120 V, leaving -160 to -70 V on the outlet<sup>2</sup>. The top left image in Fig. 6.2 visualizes this procedure for a lowering by 40 V, showing the realistic potentials. The location of the plasma (ellipse) and the movement of the particles (arrow) are indicatively shown in the plots as well (green for electrons), not to scale. After a fixed amount of time (several seconds), the outlet electrode is "closed" again, i.e. goes back to its full negative potential. Prior to this operation, the P Trap has already been shaped into the same configuration as the C Trap for electrons previously - that is a positive potential well, in this case at 150 V, and a negative barrier (-190 V). Therefore, the arriving electrons assemble in the positive potential of the P Trap. While the C Trap dumps the leftover electrons in the upstream direction and is eventually reset (top right and bottom left of Fig. 6.2), the P Trap is then formed into the cooling configuration by removing the negative barrier and setting the outlet electrodes to zero potential (bottom right). Antiprotons are subsequently captured and join the electrons in this trap, sketched by the orange arrows of movement. The yellow ellipse represents the mixed electron-antiproton plasma, not to scale.

<sup>1</sup>For readability reasons, when working with negative voltages, more negative values are from here on referred to as "higher", while values closer to zero are called "lower". Equivalently, "lowering" or "reducing" a voltage means bringing it closer to zero, and "raising" or "increasing" it brings it to more negative values.

<sup>2</sup>As explained in Sect. 3.4.8, the stated voltages do not correspond to those produced in reality due to differences in electrode lengths and the influence of the neighbouring ones.

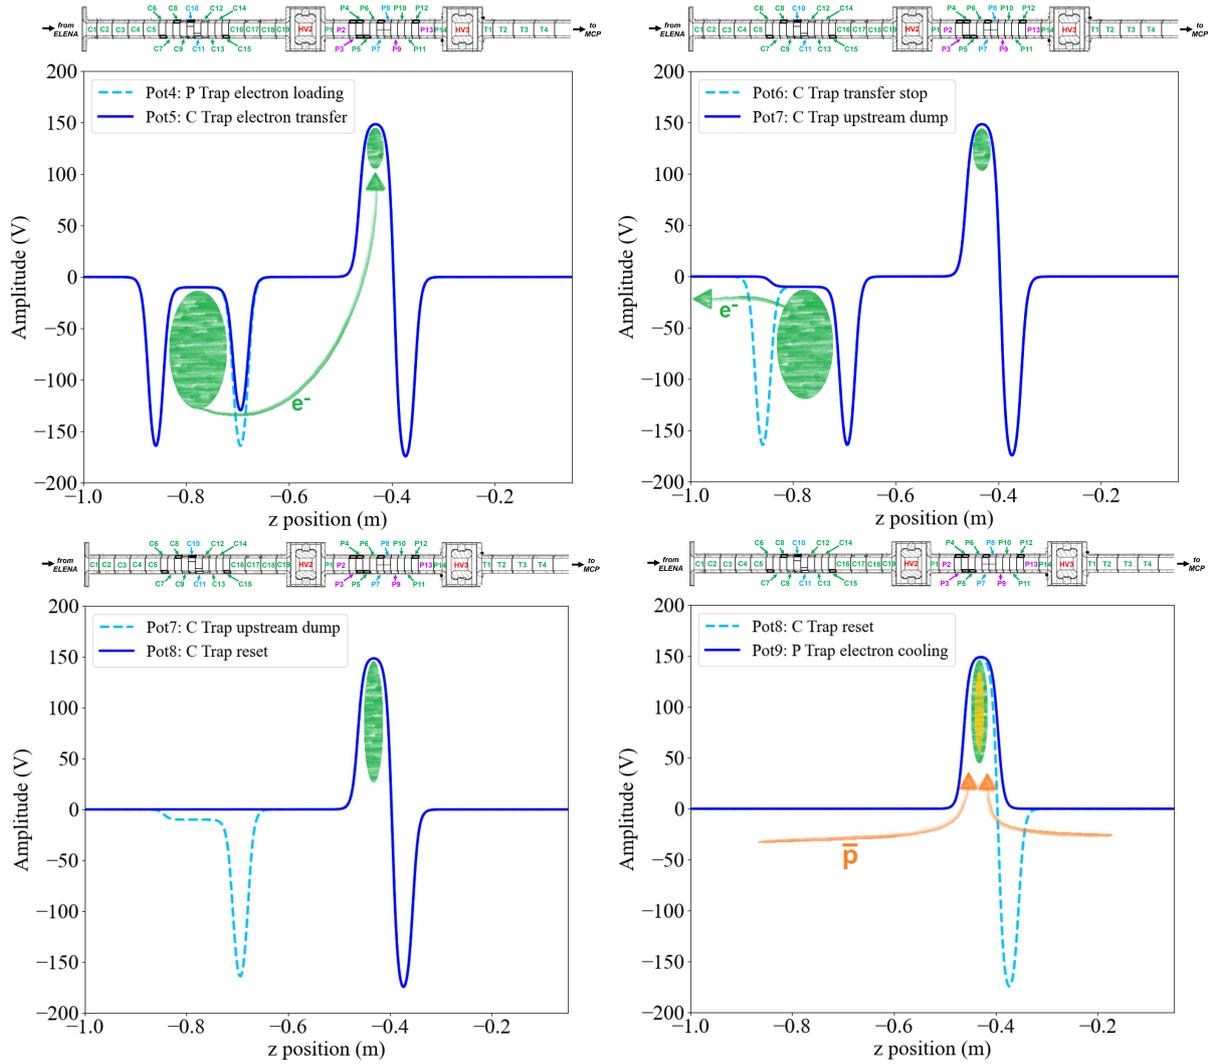


Figure 6.2: Low-potential reshaping operations to load electrons from the C Trap into the Short P Trap and prepare the trap for antiproton cooling. Initial potentials are plotted in cyan, dashed line; final potentials are shown by the blue solid line. Plasma locations (ellipses) and particle movements (arrows) are included as an indication, not to scale (green for electrons, orange for antiprotons, yellow for the mixed plasma). Above each plot, a schematic of the involved range of trap electrodes is included, scaled to match the dimension of the x-axis, with the colors of the electrode labels having the same meaning as in Fig. 3.9. Top left: Partial opening of the C Trap outlet electrode to transfer a fraction of space charge to the open P Trap. Top right: Upstream electron dump from the C Trap. Bottom left: C Trap reset. Bottom right: Transition of the P Trap from the electron loading configuration to the cooling trap.

While a positive trap configuration is useful to accumulate the mixed plasma of negatively charged particles and thus facilitate cooling, the plasma cannot be easily extracted from the trap in a controlled direction. Therefore, for any procedure that involves a plasma dump, the potential of the P Trap is beforehand reshaped into a negative well with strongly negative endcaps (e.g.  $-150$  V, one endcap each shown here) and a floor that is also slightly negatively biased ( $-10$  V shown here). This operation is depicted on the left hand side of Fig. 6.3, showing only the relevant  $z$  axis region for the P Trap. As is obvious from there, the applied  $-150$  V potential results in a barrier voltage of only

about  $-80$  V. This is because the voltage is only applied on one short (12.5 mm length) electrode each, with the neighbours kept at or close to 0 V, i.e. the produced potential is significantly lower (refer to Sect. 3.4.8 and Appx. E for details). From this trap, particles can be dumped into one of the two directions by simply bringing the corresponding electrode to zero potential, analogously to the procedure described for the pure electron plasma in Sect. 5.3. On the right hand side of Fig. 6.3, a downstream extraction toward the 1 T region and the MCP is shown as an example; an upstream extraction toward the degrader structure is equivalently possible.

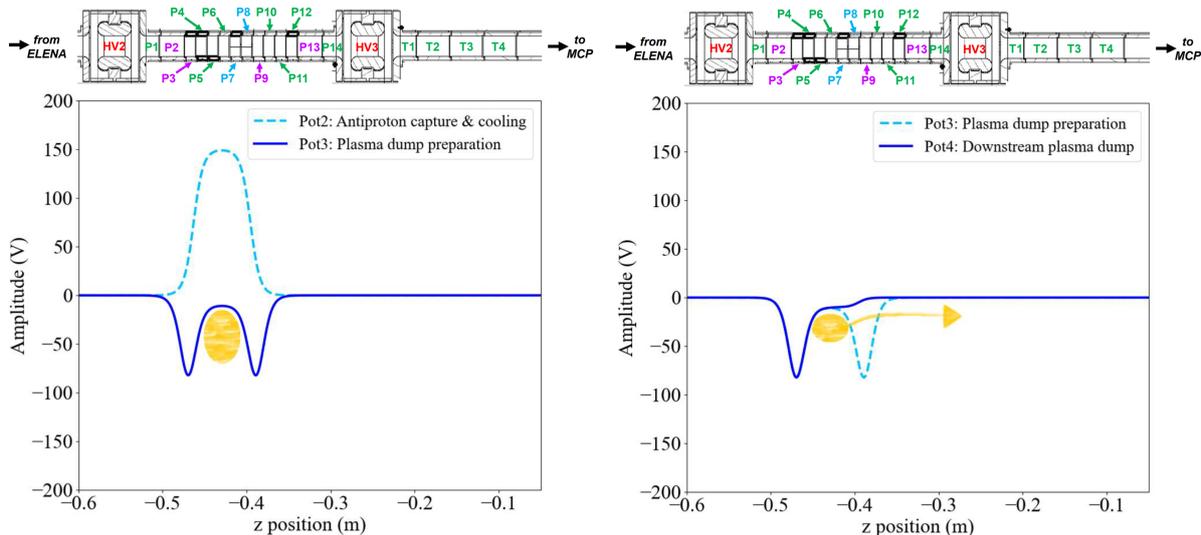


Figure 6.3: Low-potential reshaping operations to prepare the P Trap for a dump of the confined particles. Included are also the locations and movement of the mixed electron-antiproton plasma (yellow, not to scale). Above each plot, a schematic of the involved range of trap electrodes is included, scaled to match the dimension of the x-axis, with the colors of the electrode labels having the same meaning as in Fig. 3.9. Left: Transition from cooling trap to dump trap. Right: Downstream dump.

### 6.1.2 Beam steering

In order to maximize the number of available antiprotons from ELENA, it is possible to steer the parameters of the arriving beam to be optimal for the experiment setup. These parameters include the horizontal and vertical offsets as well as the corresponding angles, i.e. a four-dimensional optimization is needed. Furthermore, such optimizations have to happen every time the settings of ELENA are changed, which typically happens several times during the months-long measurement campaigns. In AEGIS, a procedure has been developed to perform this optimization scan automatically within the CIRCUS control system and the included ALPACA Bayesian optimizer [132]<sup>3</sup>. The parameter that is chosen to be maximized is the intensity of the signal obtained by the MCP detector (used in "MCP-FC" configuration with its metallized front face as a charge collector, see Sect. 3.4.9) when dumping the captured antiprotons in the downstream direction, which is proportional to the number of trapped particles. Fig. 6.4 shows the resulting intensity maps for the four parameters, with the measurement points indicated by the black dots and the color representing the MCP intensity distribution as predicted by the

<sup>3</sup>ALPACA's Bayesian optimizer uses *Scikit-optimize* with a Gaussian Processes surrogate model [164].

optimizer according to the acquired data points. Lighter colors indicate a higher MCP intensity, i.e. a larger number of captured antiprotons. To determine the optimal settings in an optimization series, a convergence criterion is defined as given in Eq. 6.2.  $\mu_{\text{best10}}$  and  $\sigma_{\text{best10}}$  are the mean values and corresponding standard deviation of the ten runs with the currently highest measured values of the parameter that is being maximized.  $\delta$  can be chosen freely; it has proven useful to choose values between 0.01 and 0.05 in the applications of AEGIS to reach the criterion efficiently (in fewer than 100 runs) but only for the actual maximum. The Bayesian optimizer takes decisions based on previous results and explores only areas around local maxima in detail instead of scanning the entire parameter space. It can be seen in Fig. 6.4 that each given parameter in one of the two directions finds its maximum for very similar settings when optimizing together with both the same parameter in the other direction and the other parameter in the same direction, indicating that the effects of offset and angle are independent.

$$\left| \frac{\sigma_{\text{best10}}}{\mu_{\text{best10}}} \right| < \delta \quad (6.2)$$

The ELENA parameters controlled by AEGIS are steered according to above optimized findings for every run. Details on the beam steering procedure are given in [132]. The implementation of this procedure in the CIRCUS/ALPACA setup both improves the performance of the optimization, as the entire parameter space is available without the need for human bias or pre-scans (or a measurement campaign over several weeks), and significantly increases the convergence speed compared to scanning over equidistantly chosen random values.

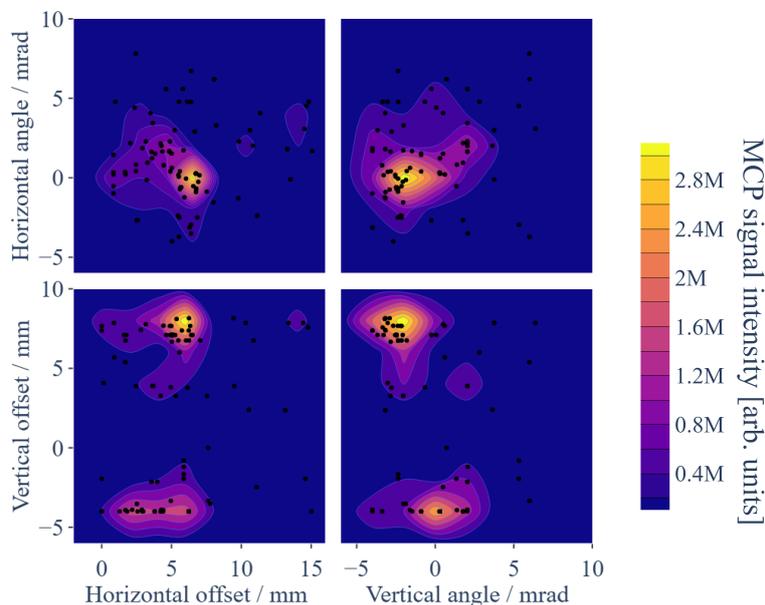


Figure 6.4: MCP intensity map corresponding to the distribution of the number of captured antiprotons obtained from the beam steering procedure. The optimization parameters are the vertical and horizontal offset of the ELENA beam entering the experiment as well as its entry angles. The black dots are the measurements points and the color distribution represents the MCP intensity as determined and predicted by the Bayesian optimizer based on the obtained data.

### 6.1.3 Optimization of the trap closing time

Utilizing the same convergence criterion as for the beam steering, it is also possible to optimize the trap closing time, i.e. the delay after which the high voltage on the inlet electrode (HV1 or HV2) is ramped up subsequently to the arrival of a trigger given by the ELENA decelerator for the injection of the antiprotons. Typical delay times are of the order of a few 100 nanoseconds to one microsecond. Also in this case, the aim is to maximize the number of particles caught in the trap, meaning here to leave the trap open long enough to let as many as possible enter but close it early enough to prevent a re-exit of the fastest antiprotons reflected on HV3 (i.e. minimizing also  $f_{\text{deg2}}$  in Fig. 6.1). In the measurements presented here, the observable is the (background-subtracted<sup>4</sup>) number of counts on SC34 of the ESDA detector (see Sect. 3.4.9) in a time window of a few tens of seconds after dumping the trap content onto the degrader structure in the upstream direction. Fig. 6.5 shows this number of counts, normalized to the maximum of the sequence, as a function of the closing time for three different optimization sequences: two with HV2 as inlet electrode, using the HV electrodes at  $-10$  kV (yellow) or  $-14$  kV (red), and one using HV1 as inlet, with a voltage of  $-12$  kV (blue). The found maximum points are indicated by the dashed lines, and the corresponding delay times are noted. The observed number of counts is assumed to be caused by the annihilations of escaping antiprotons and therefore proportional to the number of captured antiprotons. Very similar results to the ones shown here for SC34 are obtained using different ESDA units [163].

A first observation is the fact that, as would be assumed, for negative delay values the number of observed counts, i.e. the number of captured antiprotons, is consistently around zero. Clearly, when closing the inlet electrode before antiprotons can enter the trap, no capture is expected. The number of counts is maximized for longer delay times when utilizing lower voltages on the HV system because the average speed of the trappable fraction of the arriving antiprotons is lower. Independently, maximum capture is already achieved at shorter times when using HV1 instead of HV2. This is expected because HV1 is installed further upstream, meaning that the antiprotons from ELENA arrive at its location earlier. The typical bunch length provided by the ELENA decelerator is of the order of 100 to 200 ns, which corresponds broadly to the slope of the observed distribution for short closing times, however, clearly, the interaction with the degrader material distorts this distribution prior to their arrival. With the longer trap, using HV1 as inlet, a plateau forms in a closing time range with a length of around 800 ns, during which the number of captured particles changes only slightly. This plateau corresponds to those closing times which allow the large majority of antiprotons to enter the trap but not re-exit. Appropriately, for the fastest antiprotons captured by the 12 kV electrodes, a flight time back and forth between HV1 and HV3 (almost 0.8 m in one direction) of the order of 1  $\mu$ s is expected, i.e. in agreement with the rising slope time in combination with the plateau. Since using HV2 as inlet significantly shortens the trap to around 0.3 m, this plateau does not as prominently or at all form due to the shorter distance to be covered before escaping, especially for the higher barrier voltage, which can capture more highly energetic particles that are also faster to escape again before closing. For the fastest expected antiprotons at capture voltages of 10 kV and 14 kV, travel times back and forth of slightly above and slightly below 400 ns, respectively, are expected, again in rough agreement with the observed combined time of the rising slope and the peak (or

---

<sup>4</sup>Following the standard procedure described in the following sections, the first 7 s of the acquisition, prior to the antiproton capture, are used to determine the individual background rate of each run.

short plateau, if any). It further makes sense that the rising slope at shorter closing times is steeper than the falling slope for long times: The latter corresponds to the losses from re-exiting antiprotons, whose time spread is expected to have increased during the travel and reflection process due to their momentum spread.

The trap closing delay times that maximize the number of ESDA counts found in this way for the different settings are used as standard values in the antiproton capture procedure to optimize the number of captured antiparticles.

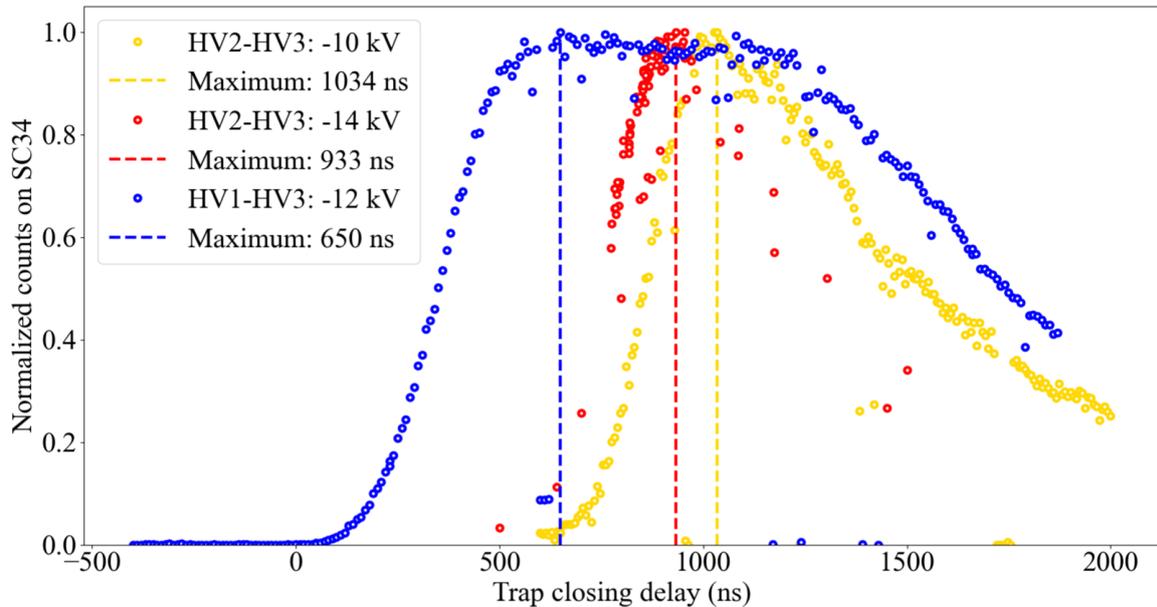


Figure 6.5: Plot of the number of counts on SC34 subsequent to a dump on the degrader structure as function of the 5 T trap closing time. The numbers are normalized to the maxima of the sequences. The maximum points are marked by dashed lines and the corresponding times are noted. Scans are performed in the longer HV configuration (HV1-HV3) with the electrodes operated at  $-12$  kV (blue) and the shorter one (HV2-HV3) at  $-10$  kV (yellow) and  $-14$  kV (red).

### 6.1.4 Degrader studies

The design of the AEgIS degrader structure (see Sect. 3.4.7) was driven by a GEANT4 Monte Carlo simulation of the degraded and trappable fractions of ELENA antiprotons after passage through the degrader foils, which incorporates also the magnetic field map of AEgIS (see Sect. 3.4.1). The simulation is reported on in [84] and will be detailed in [163]. Each simulated measurement directs  $1 \times 10^6$  antiprotons with a flat energy of 100 keV onto the center of the foils in a perpendicular direction. At the given energies, antiproton energy loss is expected to be primarily of electronic stopping nature, i.e. due to inelastic collisions with bound electrons in the degrader material, which scales with the ratio of the  $\bar{p}$  velocity to the velocity of light,  $\beta = v_{\bar{p}}/c$  [165, 166]. Fig. 6.6 displays the results of the simulation.

On the left hand side, a scan over the thickness of the main degrader Mylar foil is shown. Plotted is the fraction of antiprotons that passes through the foil and has a resulting axial momentum below a certain threshold, as a function of the thickness of Mylar used for the main degrader. The thresholds, 6 keV (green), 10 keV (red) and 14 keV (black),

correspond to different values of the reachable range of capture potentials applied on the HV electrodes. Clear maxima are observed for each of the different thresholds, which are found at slightly lower degrader thicknesses for increasing momentum thresholds. For thinner foils, the antiprotons' energies are less degraded by the foils (i.e. fewer lower-energy particles are observed); thicker foils lead to increased annihilation losses in all three cases. The highest fraction of resulting trappable antiprotons is found when allowing for the highest possible energies (i.e. using the highest safely feasible capture voltage, 14 kV) at a Mylar thickness of approximately 1400 nm, motivating the choice of the main degrader thickness in the experiment. With this, 80% of the antiprotons are expected to become trappable.

The study shown on the right hand side of Fig. 6.6 includes the 1400 nm Mylar main degrader together with Parylene degrader ladder foils of a varied thickness between 0 and 500 nm. Plotted is the resulting fraction of trappable antiprotons (i.e. particles with axial energies below the corresponding capture voltages) after passage through the entire degrader structure as a function of the used HV capture potential. For each Parylene thickness, the trappable fraction increased with the capture voltage, as is expected since more and more highly energetic antiprotons can be caught, until a plateau is reached, which is caused by the increased influence of  $\bar{p}$  annihilations in the foil and cannot be further compensated. As would also be intuitively assumed, the plateau is reached later with thinner foils. Which foil thickness is best suited to maximize the trappable  $\bar{p}$  fraction depends on the capture voltages; the higher the voltage, the thinner the optimal foils. From around 16 kV, the main degrader on its own, without any additional Parylene foil, is found to perform best. Between the voltage used as default now, 14 kV, and 16 kV, an additional 100 nm foil seems to have a beneficial influence. At 14 kV, 80% of antiprotons become trappable with the main degrader alone, while 90% could be reached with the added 100 nm Parylene. Even larger fractions beyond 95% become feasible with higher capture potentials closer to 20 kV when using the main degrader alone.

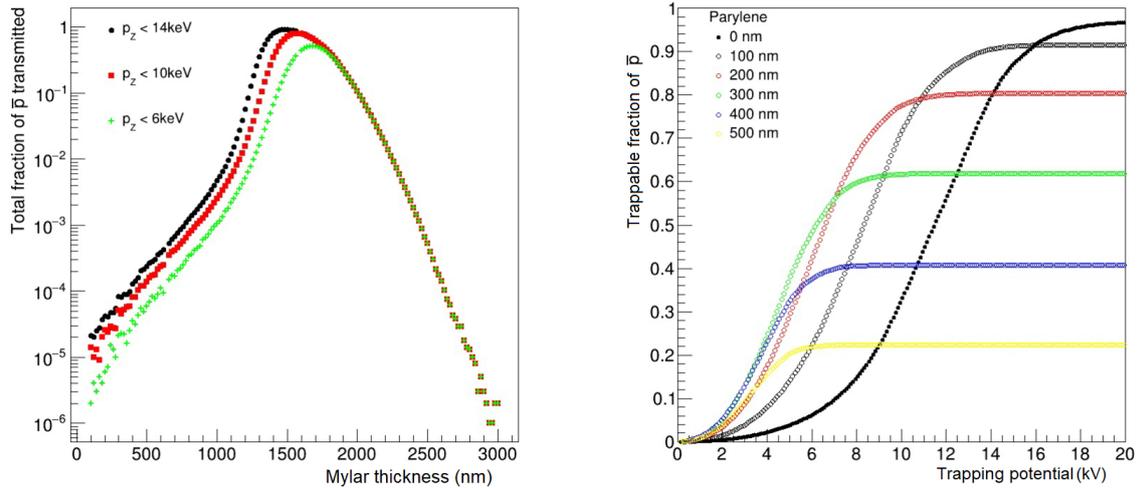


Figure 6.6: Simulation of the passage of antiprotons through the AEGIS degrader structure. Adapted from [84]. Left: Fraction of antiprotons transmitted by the main degrader foil with an axial momentum below a given threshold of 6 keV (green), 10 keV (red) or 14 keV (black), as a function of the used Mylar thickness. Right: Fraction of antiprotons with an axial momentum below the trapping threshold as a function of the capture potential applied on the HV electrodes, for the use of a 1400 nm Mylar main degrader together with a Parylene foil of different thicknesses in the degrader ladder.

Fig. 6.7 shows experimental studies of the degrader system installed in the apparatus in 2022. On the left hand side, the integrated, digitized signal of the MCP phosphor screen ("MCP-FC" detector, see Sect. 3.4.9) is plotted as a function of the high voltage applied on HV3. Only the main degrader is used in this study. The measurements are taken at the time of the arrival of the antiprotons from ELENA, i.e. the observed intensity corresponds to the number of antiprotons passing HV3 and not getting stopped by its potential, making their way to the MCP at the downstream end of the experiment. For the integration, the signal observed by the high-speed oscilloscope is summed up within  $2\mu\text{s}$  around the peak amplitude of the antiproton peak. The integrals plotted in Fig. 6.7 are the mean values (and standard errors) of five independent measurements, and the values are normalized to the highest intensity, which is, as expected, found with the lowest HV3 potential (3 kV shown here). As is also intuitively reasonable, the observed MCP intensity, i.e. the number of antiprotons passing the barrier voltage, decreases when increasing the voltage, since more and more highly energetic particles are blocked. In good agreement with the simulation shown on the right hand side of Fig. 6.6, for the thick degrader alone, the increase in the blocked  $\bar{p}$  fraction is relatively slow for the first few kV of the blocking potential increase but then becomes steep. When applying 14 kV on HV3, only around 20% of the original intensity are observed, implying that 80% of the antiprotons are blocked, i.e. become trappable. This value, as well, confirms the expectations from the simulation.

The experimental results shown on the right hand side of Fig. 6.7, on the other hand, exhibit significant deviations from the simulation. It should be kept in mind here that these measurements have failed, at least in part: As stated in Sect. 3.4.7, the thicknesses of the thin degrader foils cannot be determined with certainty, due to unclear starting conditions and their proneness to mechanical damage. The results are still presented here for completeness but, as discussed below, the decision has been made to discontinue the use

of the thin foils. Plotted on the right hand side of Fig. 6.7 is the number of (background-subtracted) counts observed on ESDA unit SC56<sup>5</sup> (see Sect. 3.4.9) subsequently to the release of one bunch of ELENA antiprotons captured in the 5 T high-voltage trap, as a function of the voltage applied on HV1 and HV3. The background rate of counts from cosmic rays and surrounding radioactive sources is determined for every run individually from the first 7 s of the acquired ESDA spectra. The summation of the signal is performed from the time of release (approximately 10 s after the start of the acquisition) until the end. For further details on ESDA spectra and the analysis methods, see Sect. 6.2 to Sect. 6.3. This scan over the capture voltage is performed for different configurations of the degrader setup: with the thick degrader ("TD") only, without the thick degrader and instead with a 100 nm or a 200 nm foil alone on the degrader ladder, and with the main degrader together with a 300 nm, 400 nm or 500 nm foil. Again, except for the thick degrader, these are only the foreseen foil thicknesses that may not correspond to the actual values. The number of counts on the ESDA is expected to be proportional to the number of annihilating particles, i.e. to the number of captured antiprotons. The left vertical axis shows the counts on SC56 normalized to the maximum value, which is found with the highest capture voltage (14.4 kV), using the thick degrader alone; the right vertical axis states the actual number of counts. As expected, the thin foils alone do not reduce the antiproton energy enough to allow for the capture of a significant fraction; the observed number of annihilation counts remains flat at zero with the 100 nm foil, while evidence of a small number of caught particles seems to form for very high capture voltages with the 200 nm foil. Furthermore, the overall capture efficiency generally increases with the applied capture voltage. Additionally, the main degrader seems to perform best on its own, especially at high capture voltages, as predicted by the simulation on the right hand side of Fig. 6.6. However, this feature develops already for lower voltages than expected. Furthermore, the simulated curve is not realistically reproduced in its shape for any of the degrader configurations, in particular the expected flattening out and resulting plateau for increasing voltages is not observed (at all, or at least for the most part, if one accepts the hint of a plateau for high voltages in some of the measurements as the onset of a flattening). Of course it is not expected that the plots on the right hand side of Fig. 6.6 and Fig. 6.7 agree very well in absolute numbers: The simulation gives only the estimated fraction of antiprotons that is trappable according to the energy distribution when emerging from the degraders (disregarding several factors, as explained above or in the next paragraph), not those that are actually expected to be trapped, as this entails all the loss contributions given in Fig. 6.1. Only a relative comparison between the different degrader configurations realistically makes sense.

Still, there are two main factors that render the comparison of the degrader simulations and the obtained data difficult: The first one is the fact that the main degrader is simulated as one cylinder of Mylar instead of a thick and a second thin foil directly behind, as is the case in the experiment. Similarly, the thin foils are modeled as single foils, while some of them are in reality formed by a stack. This can lead to scattering losses that are not correctly taken into account in the simulation. The distribution of antiproton energies and trajectories is also disregarded. More importantly, however, the thin Parylene foils have proven to not be very resistant (which is for example also why Mylar was chosen as the main degrader material): During the operation of the experiment, already small mechanical stress has caused them to rupture, producing holes in an irregular fashion and diminishing their degradation abilities. This feature has only become evident after an

---

<sup>5</sup>SC56 is shown here as an example because of its central location in the apparatus; the other ESDA units yield similar results.

opening of the vacuum chamber housing the foils at the end of the run period, and it is basically impossible to estimate at what point of the measurements the damage happened, i.e. for which measurements which foils were still present. For example, having found the degrader ladder foils in a rather battered state after opening, it is very realistically possible that only single foils of the 400 nm and 500 nm stacks survived and were present during the measurements. This could contribute to an explanation of the observations shown on the right hand side of Fig. 6.7, e.g. if the 500 nm stack consisted only of the 100 nm foil and the 400 nm stack of the 200 nm foil, higher antiproton capture efficiencies closer to that of the thick degrader alone would indeed be expected (and, in part, their yields would be flipped as found). Due to issues such as the latter, in combination with the unclear exact thicknesses of the foils to begin with, and the results shown in Fig. 6.6, the collaboration has decided in 2023 and 2024 to adopt the strategy of a further upgrade of the switch for the HV electrodes, enabling in the future capture potentials of up to 20 kV, instead of investing in improvements of the degrader ladder. At higher voltages, the main degrader on its own is expected to perform best to enable a maximum capture efficiency, and this feature is confirmed by the data for the highest tested capture voltages.

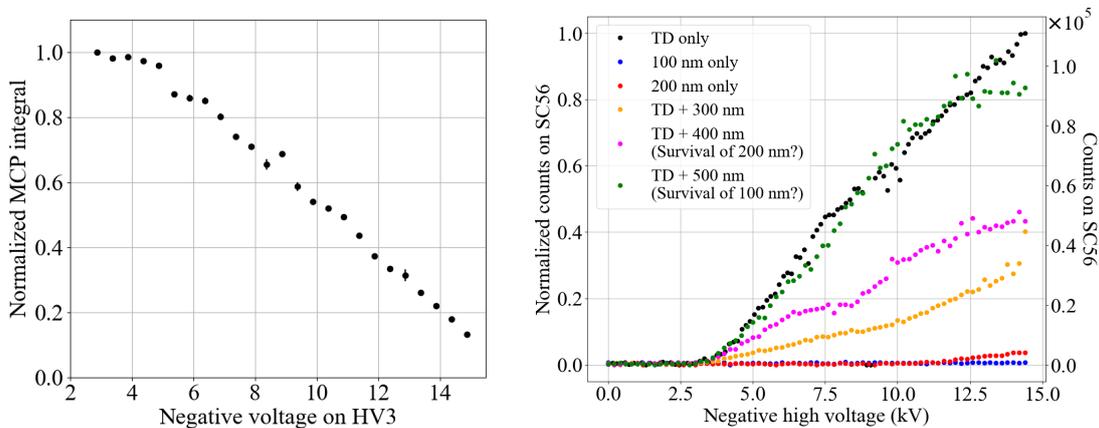


Figure 6.7: Experimental studies of the passage of antiprotons through the AEGIS degrader system. Left: The observed integrated MCP intensity at the time of arrival of the antiproton bunch as a function of the barrier voltage applied on HV3, using only the main degrader. Right: Number of counts observed on SC56 subsequently to the dump of captured particles from the trap, as a function of the capture voltage applied on HV1 and HV3, for different degrader configurations: with the thick main degrader only ("TD only"), with only the use of one of the thin Parylene foils at different thicknesses, and with a combination of both degraders. In the latter two cases, the used Parylene thickness is stated in the label. The configurations, especially those marked with a question mark, are furthermore discussed in the main text.

### 6.1.5 The antiproton capture efficiency

In order to determine the antiproton capture efficiency directly from the capture data, the number of counts on SC56 can be converted to the number of annihilating antiprotons using the conversion factors stated in Sect. 3.4.9. With the thick degrader alone and the default trapping voltage of 14 kV, counts of the order of  $1 \times 10^5$  are routinely achieved. Applying the intercalibration<sup>6</sup> to SC12 and the conversion from there to  $\bar{p}$  numbers, i.e.

<sup>6</sup>Due to a technical failure, SC12 itself did not acquire data during the displayed measurements.

factors  $f_{\text{SC56} \rightarrow \text{SC12}}$  and  $f_{\text{SC12} \rightarrow \bar{p}}$ , yields around 5.15 million captured antiprotons<sup>7</sup>. With the stated ELENA beam intensities typically of the order of  $7 \times 10^6$ , these numbers correspond to a capture efficiency between 70 and 75 %, in good agreement with values obtained from loss fraction measurements: Using Eq. 6.1, the trapped fraction of antiprotons  $f_{\text{trap}}$  can be calculated by determining the different loss contributions. In standard conditions (using only the main degrader and a capture voltage of 14 kV),  $f_{\text{MCP}}$  has been found to be of the order of 0.2, as shown on the left hand side of Fig. 6.7.  $f_{\text{store}}$  of course depends on the time for which the antiprotons are kept in the trap as well as on the plasma compression tuning (e.g. Rotating Wall) but does not amount to more than a few percent in usual conditions and for the short trapping time used here (a few seconds).  $f_{\text{deg1}}$  and  $f_{\text{deg2}}$  have been determined in dedicated measurements, whose detailed analysis is still ongoing, yielding preliminary first estimates of  $f_{\text{deg1}} \approx 0.01$  and  $f_{\text{deg2}} \approx 0.05$  [84]. These measurements rely on the separation in time of the two signals in the ESDA acquisition due to the antiproton travel times of the order of  $1 \mu\text{s}$  back and forth in the trap and on the normalization to the signal of the entire beam annihilating in the experiment (on the MCP, instead of on the degrader structure), i.e. relatively large uncertainties are expected. Given the preliminary results, however,  $f_{\text{trap}}$  would be calculated to be of the order of 64 to 74 %, confirming the direct measurement and its conversion.

With regard to previous standards before the commissioning of ELENA, antiproton capture efficiencies of the order of 70 % and beyond mark an improvement by approximately two orders of magnitude [68, 97]. A comparison with the other AD experiments furthermore shows that the efficiencies achieved by AEGIS with the upgraded setup largely exceed the rest, by at least a factor 2.5 [167].

## 6.2 Efficient antiproton cooling

Electron cooling is a well-established technique to cool negative ions sympathetically and, more recently and relevant here, antiproton beams (e.g. in ELENA) as well as those confined in electromagnetic traps, as it is more efficient than other mechanisms at the relevant energies (for example, stochastic cooling is more efficient at energies comparable to the initial ones in the AD, while adiabatic cooling works well for energies significantly below the low-voltage potential walls used here) [146, 147]. When mixing antiprotons of a given energy with electrons, the two species undergo collisions until they reach thermal equilibrium. In a strong magnetic field  $B$ , the resulting heating of the electrons (due to which they are for instance regularly exchanged in the ELENA apparatus) is counteracted by cooling through synchrotron radiation. The energy lost due to synchrotron radiation is proportional to  $B^2/m^4$ , with  $m$  as the particles' mass, so it is several orders of magnitude larger for the lighter electrons. For large enough electron densities (from  $\sim 1 \times 10^7 \text{ cm}^{-3}$ ) and a ratio of electron to antiproton numbers of  $1 \times 10^4$  or more, antiproton temperatures of the order of eV can be reached within  $\sim 1 \text{ s}$  [147].

In AEGIS, antiproton capture from ELENA is performed as described in Sect. 6.1, and the antiprotons are combined with electrons in the P Trap region of the electromagnetic trap stack as depicted in Fig. 6.2. Fig. 6.8 shows the number of counts on SC56 of the ESDA as a function of the time one captured bunch of ELENA antiprotons has passed in the P Trap. The time is given in seconds after the start of the detector acquisitions; the antiprotons arrive at the first visible thin peak (injection losses) and resulting slight increase in the

---

<sup>7</sup>It should be noted that this measurement happened before the installation of the PMT splitters, meaning that  $f_{\text{splitter}}$  is not applied.

---

count rate at beyond 200 s. The time range of the two plots differs because of a different arrival time of the antiprotons in the experiment but the overall length is kept the same. A binning of 100 ms is chosen for both plots for a good resolution of the individual peaks. The hot dump time is marked by a vertical red dotted line and the initialization of the cold dump by a yellow dotted line. Outside of the dump events, the observed counts are caused by unwanted annihilations of antiprotons which escape the trap and are lost, plus a constant background rate due to cosmic rays and radioactive sources in the environment of the experiment. The additional thin peaks which are evident in the distribution stem from AD injection or ELENA extraction events, resulting in a series of fast antiproton annihilations that are also detected in the AEGIS apparatus. These events are clearly tagged by software triggers such that the counts can be subtracted in the analysis. Only a fraction of them is visible here due to the fine binning. Further details on this are given in Sect. 6.3.

For the measurement shown at the top of Fig. 6.8, no loaded electrons are present in the trap with the antiprotons, while the plot on the bottom shows the count distribution of a standard run employing sympathetic cooling with the electrons loaded from the electron gun into the C Trap and then into the P Trap. It is found that, without electron cooling, all antiprotons are released in the hot dump, producing a strong peak in the number of counts, while no additional counts are observed in the cold dump. Instead, using the cooling technique, the count numbers in the hot dump on the right hand side of Fig. 6.8 are significantly reduced, while the cold dump events are clearly visible.

When dumping the particles "slowly", i.e. ramping down the trap wall voltages over many seconds, as done here, it is possible to observe the quick onset of counts from the cold dump with count rates decreasing with time, i.e. fewer particles remaining for ever lower trapping potentials. In thermal equilibrium, this procedure can be used to determine the particle temperatures, and a brief discussion on this can be found in Sect. 6.4.

In the context of the presented data, "cold" simply means cold enough to be confined by the low-voltage potentials.

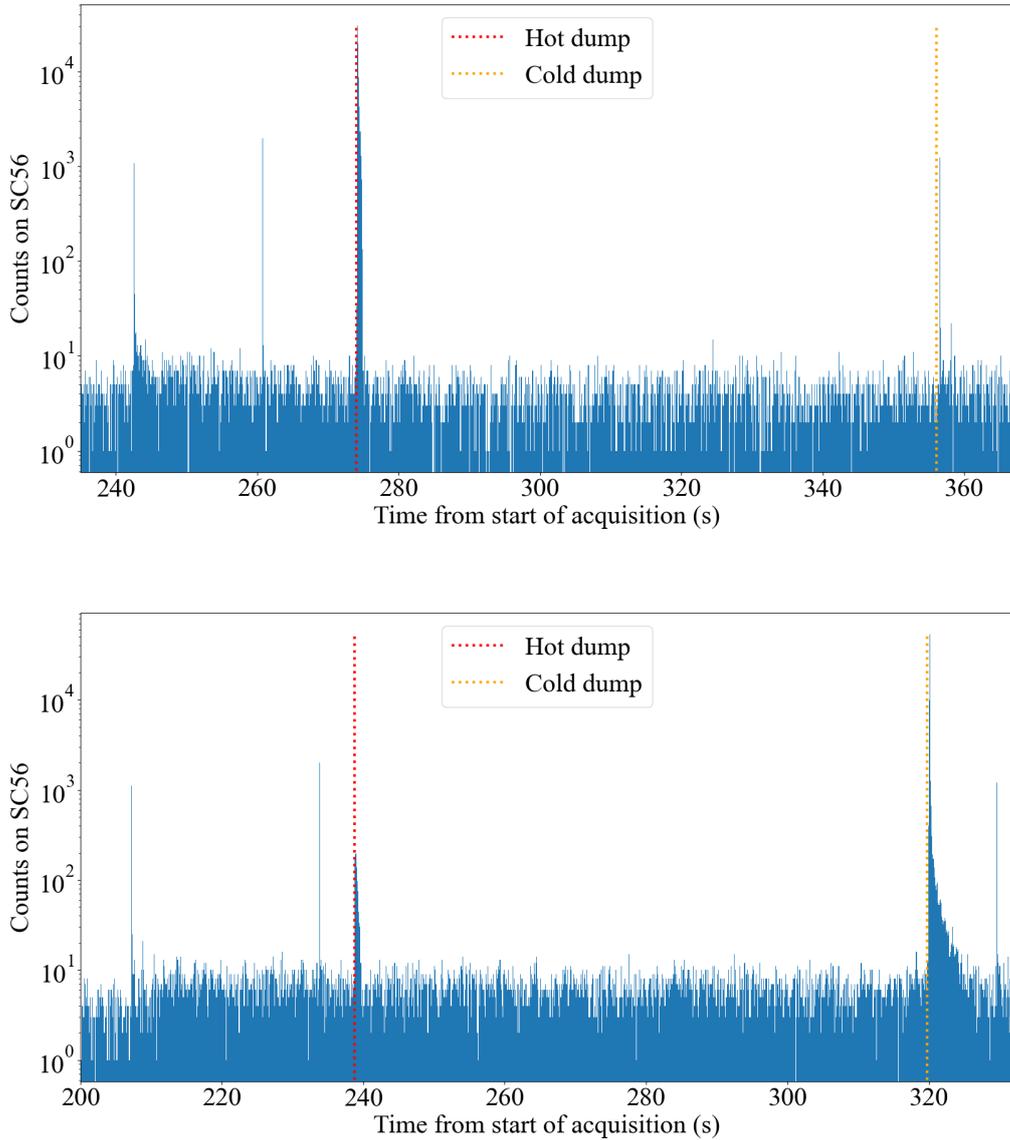


Figure 6.8: Comparison of the distribution of the number of counts on SC56 versus time when trapping one bunch of ELENA antiprotons in the P Trap for 80s without and with cooling electrons. The red dotted line indicates the starting time of the hot dump; the yellow dotted line shows the onset of the cold dump. The individual additional peaks are caused by antiproton annihilations from AD injection or ELENA extraction events, of which only a fraction is visible with the given binning of 100 ms. Top: No added cooling electrons. Bottom: Standard electron cooling.

### 6.2.1 Cooling efficiency

The cooling efficiency  $\eta$  is defined as the fraction of antiprotons that has cooled and remains in the trap until the cold dump. It is determined as the ratio of counts observed on a given ESDA unit in the cold dump ( $CD$ ) over the total number of counts in the cold and hot dump ( $HD$ ), as given in Eq. 6.3.

$$\eta = \frac{CD}{(CD + HD)} \quad (6.3)$$

Fig. 6.9 shows the so obtained cooling efficiency as a function of a scan over the hot storage time. The different colors represent different ESDA units. After an initial onset of around five seconds, during which the antiprotons settle into the electron trap and the combined plasma for cooling forms, cooling is achieved very quickly, on the order of a few seconds. The cooling efficiency therefore increases quickly with the time for which the particles are confined. This behaviour is observed on all ESDA units. However, those scintillator slabs located in the 1 T region of the experiment (SC1112, SC1314, SC1718, SC2324 shown here), further downstream and closer to the MCP, do not show close to zero counts in the hot dump for very short storage times. This feature is attributed to their miscounting of "hot" antiprotons that leak, i.e. spill over the downstream potential barrier due to heating upon opening the upstream barrier for the dump, and annihilate on the downstream MCP during the time assigned to the cold dump. Being much further away from the MCP, these annihilations are not registered in the 5 T region. Furthermore, the numbers of counts observed on SC12 and SC34 for longer storage times are significantly lower than those of the rest of the units, while the others are in good agreement. This development stems from a strong pile-up effect, i.e. saturation of the two units closest to the annihilation point after the cold dump.

For these reasons, to determine the overall cooling efficiency, the average values of SC56 and SC78 are used as a compromise, as they are not affected by either one of these issues as significantly thanks to their location and corresponding solid angle. The resulting efficiencies are plotted in Fig. 6.10 for the different hot storage times. The error bars correspond to the standard errors on the mean, corrected for low statistics. An exponential fit, despite being unphysical for negative efficiency values, fits well to the data when restricted to storage times between six and 60 seconds (i.e. the region between the third and last data point). As an eye guide, it is also shown in Fig. 6.10. From it, a cooling time constant of  $(2.2 \pm 0.1)$ s and a maximum reachable efficiency of  $(92.2 \pm 0.2)\%$  can be obtained.

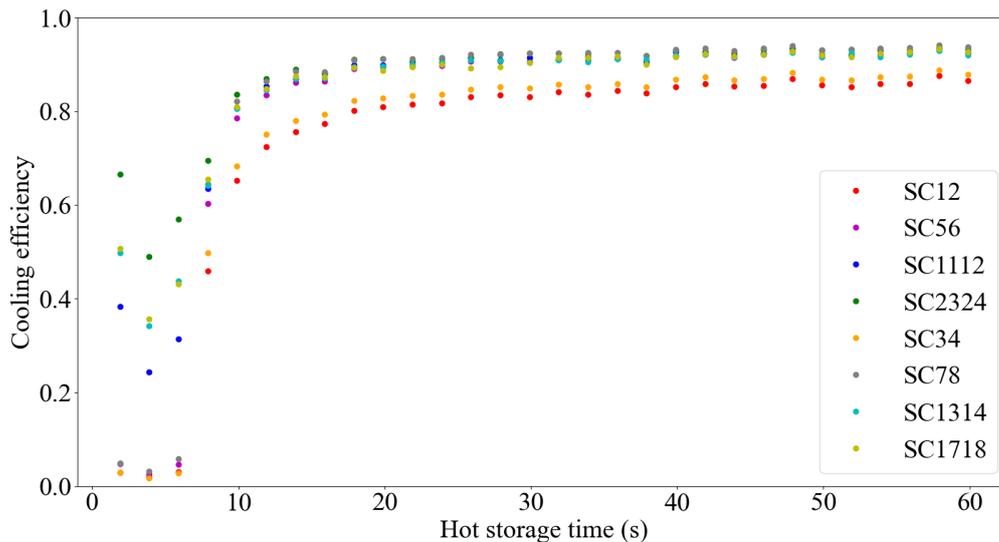


Figure 6.9: The achieved electron-antiproton cooling efficiency as a function of the hot storage time, as obtained on different ESDA units.

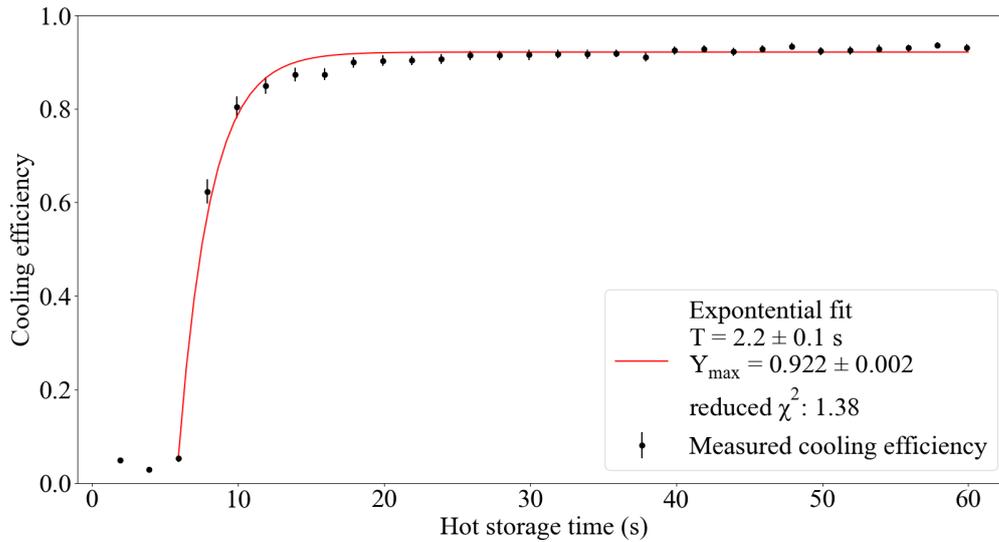


Figure 6.10: The achieved electron-antiproton cooling efficiency as a function of the hot storage time, as obtained from the average of SC56 and SC78. An exponential fit, confined to the range between the third and last point, is also shown as an eye guide.

In principle, cooling efficiencies closer to 100% can also be reached. This is possible by increasing the number of cooling electrons present in the trap. For the data presented in Fig. 6.9 and Fig. 6.10, the C Trap outlet was lowered by 45 V. When increasing this value and letting more electrons pass from the C Trap to the P Trap, the cooling efficiency improves. This can be observed in Fig. 6.11. Plotted is the cooling efficiency as a function of the amount of space charge transferred to the P Trap from the C Trap, as observed on the different ESDA units. Again, the units in the 1 T region do not go to an efficiency of zero for very few present electrons, as would be expected. Other than that, the efficiency generally improves with the amount of transferred electron plasma, excepting individual outliers. When transferring 75 V of space charge, efficiencies of the order of 98% are already reached.

However, to minimize losses from annihilations on the trap walls, it is generally favourable to keep the mixed plasma as compressed as possible and avoid expansion [108]. For too high numbers of cooling electrons, the radial extension of the electron plasma exceeds that of the arriving antiproton bunch. The antiprotons then adapt to the distribution of the electrons, their plasma radius increases as well, and losses are incurred. The right hand side plot of Fig. 6.12 shows the dependence of the radius of the electron plasma (i.e. also the mixed plasma) on the amount of space charge transferred from the C Trap to the P Trap (red points). For comparison, the radius of the arriving ELENA antiproton bunch is also shown (blue points). The plotted radii are determined as follows from the CMOS image obtained of the phosphor screen connected to the MCP at the downstream end of the experiment. The antiproton plasma image is taken upon the arrival of the ELENA bunch, of those particles that are not captured by the high voltage potentials and end up on the downstream MCP. The image of the mixed plasma (whose distribution is expected to be equivalent to that of the electron plasma) is obtained subsequently to the downstream dump of the confined plasma after a hot storage time of 40 s. As described in Sect. 3.4.9, a higher intensity observed on the camera at a given pixel corresponds to more charge, i.e. a larger number of particles, having arrived on the MCP in that location. Both

pictures show approximately circular plasma distributions, an example of which is shown on the left hand side of Fig. 6.12 for the mixed plasma with a transferred space charge of 95 V. The coordinates are given in mm according to the pixel-to-millimeter conversion described in Sect. 5.4.1. The color scale indicates the MCP intensity in a given location as observed by the camera pixels. Examples for other values and for the antiproton image are shown in Appx. I, with the dimensions given in pixel coordinates. The low-intensity spot close to the plasma center is due to a permanent fault in the MCP caused by an excessive amount of charge arriving in a too short time and is present independent of the image itself. The two-dimensional pixel intensity values of each image are used for analysis. First, a background intensity, obtained as the average of the pixel values in the small yellow square shown in the image ( $50 \times 50$  pixels), is subtracted from all pixel intensities to normalize the images. The two-dimensional pixel intensities are ordered in the x and y directions individually. Then, starting once from the lower end and once from the upper end each, the intensities are added up until 12% of the sum of all intensities are reached. The corresponding x and y values give the minimum and maximum included coordinates in both dimensions. 12% are chosen as one half of the percentile not included in the Full Width Half Maximum (FWHM, 76%), assuming a Gaussian distribution of intensities. The radii in both directions is subsequently determined as half the distance between the minimum and maximum included coordinates, and the average radius is calculated from the two. The red circle in the image shown on the left side of Fig. 6.12 is drawn with this average radius around the point defined by the midpoints between minimum and maximum in x and y direction. The average radius is the one plotted on the right hand side of Fig. 6.12 for different settings of transferred space charge.

As expected, while the latter remains approximately constant, the mixed plasma radius increases with the number of electrons present in the trap. With no or very few electrons, no significant fraction of the antiprotons is cooled, such that the radius of the mixed plasma cannot be determined. Ideally, the radii of the antiproton and electron plasma are compatible so that a complete overlap of the two is facilitated to enable full cooling but minimize the radial expansion within this boundary condition. Therefore, according to the plot on the right hand side of Fig. 6.12, a transferred space charge of 50 V is used as default. Here, a plasma radius of the order of 5 mm is observed without the application of any plasma compression techniques. However, when it is more important to cool arriving antiprotons efficiently than to minimize all individual losses, and when accumulating large numbers of antiprotons in the trap, which itself leads to a plasma extension due to space charge effects, higher values are used.

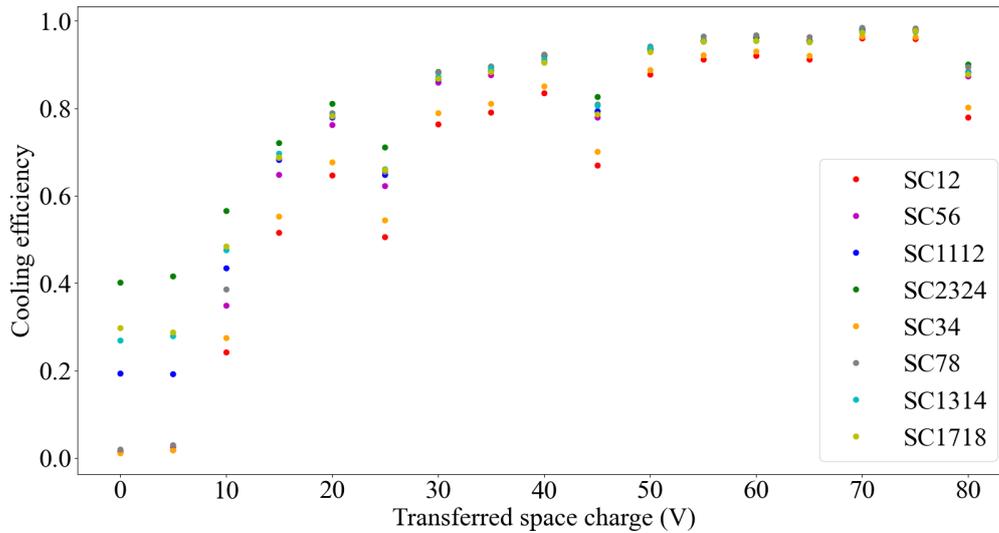


Figure 6.11: The achieved electron-antiproton cooling efficiency as a function of the electron space charge transferred from the C Trap to the P Trap, as obtained on different ESDA units.

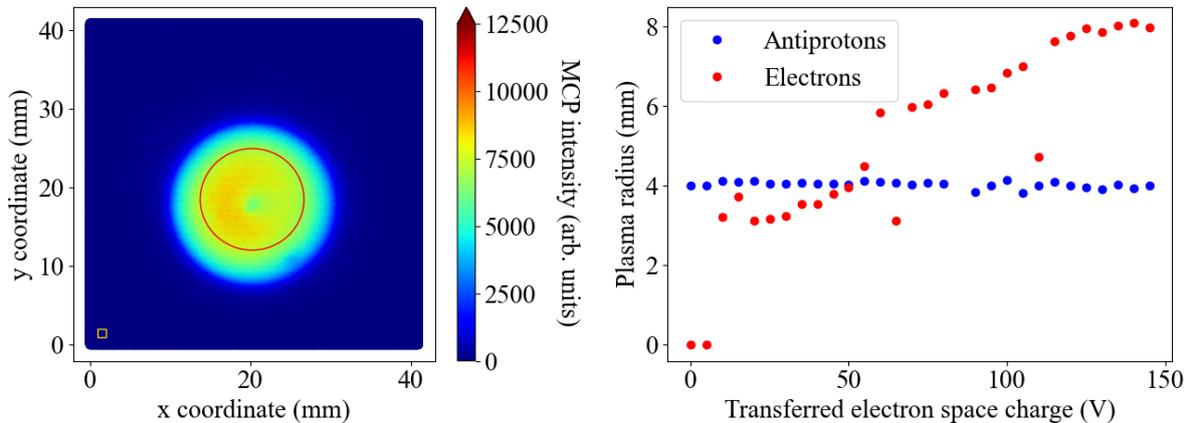


Figure 6.12: Analysis of the radii of plasmas of antiprotons and electrons for different amounts of electron space charge transferred from the C Trap to the P Trap. Left: Example CMOS image taken with the downstream MCP of the mixed electron-antiproton plasma upon the dump, with a transferred space charge value of 95 V. The pixel intensity is given by the color bar. The red circle indicates the pixel coordinates included by the average plasma radius as determined from the FWHM method described in the text. The yellow square borders the pixels used to obtain the average background intensity. Right: Plot of the determined antiproton (blue) and electron/mixed plasma (red) radii as a function of the transferred space charge.

## 6.2.2 Rotating Wall compression

While the introduction of electrons into the  $\bar{p}$  trap cools the antiprotons and improves their lifetime, after a certain time of storing the particles in the trap (over several hundred

seconds), heavy loss rates are still observed. It can be proven that these losses are of radial nature, i.e. can be prevented by actively compressing the mixed antiproton-electron plasma. One possibility to achieve compression is the use of the Rotating Wall (RW) technique (see Sect. 3.2.2). In AEGIS, this is facilitated by the application of  $90^\circ$  phase-shifted, sinusoidal frequencies on four-fold segmented electrodes, which form part of the low-voltage electrode stack, to create a rotating electric field perpendicular to the axis of symmetry of the contained plasma, producing a torque on the particles. For frequencies corresponding to rotation speeds faster than the natural plasma rotation, plasma compression can be achieved, essentially minimizing radial losses. The procedure in AEGIS has been studied in detail in [108], used also to obtain rough estimates for the starting points of the RW optimization in the data presented here. Typically, the Rotating Wall procedure is already switched on before the electrons enter the trap.

Fig. 6.13 compares  $\bar{p}$  annihilation events without and with the use of the RW technique versus time spent in the AEGIS antiproton trap, using a binning of 100 ms and the ESDA SC56 detector unit as an example. The binning is chosen to obtain a good resolution of the count increase. The increased losses from around 600 s present without RW (upper plot) are mitigated when actively compressing the plasma (bottom plot). The RW parameters used here are a frequency of 0.8 MHz and an amplitude of 0.05 V. These values are not optimized; optimization studies of RW parameters are presented below. In Fig. 6.13, the hot dump (red) and cold dump (yellow) times are indicated as well. The individual sharp peaks observed in between can be attributed to AD injection and ELENA extraction events, which cause antiproton annihilations that are observed in the experiment and are discussed in more detail in Sect. 6.3, as only a fraction of them is visible here due to the binning.

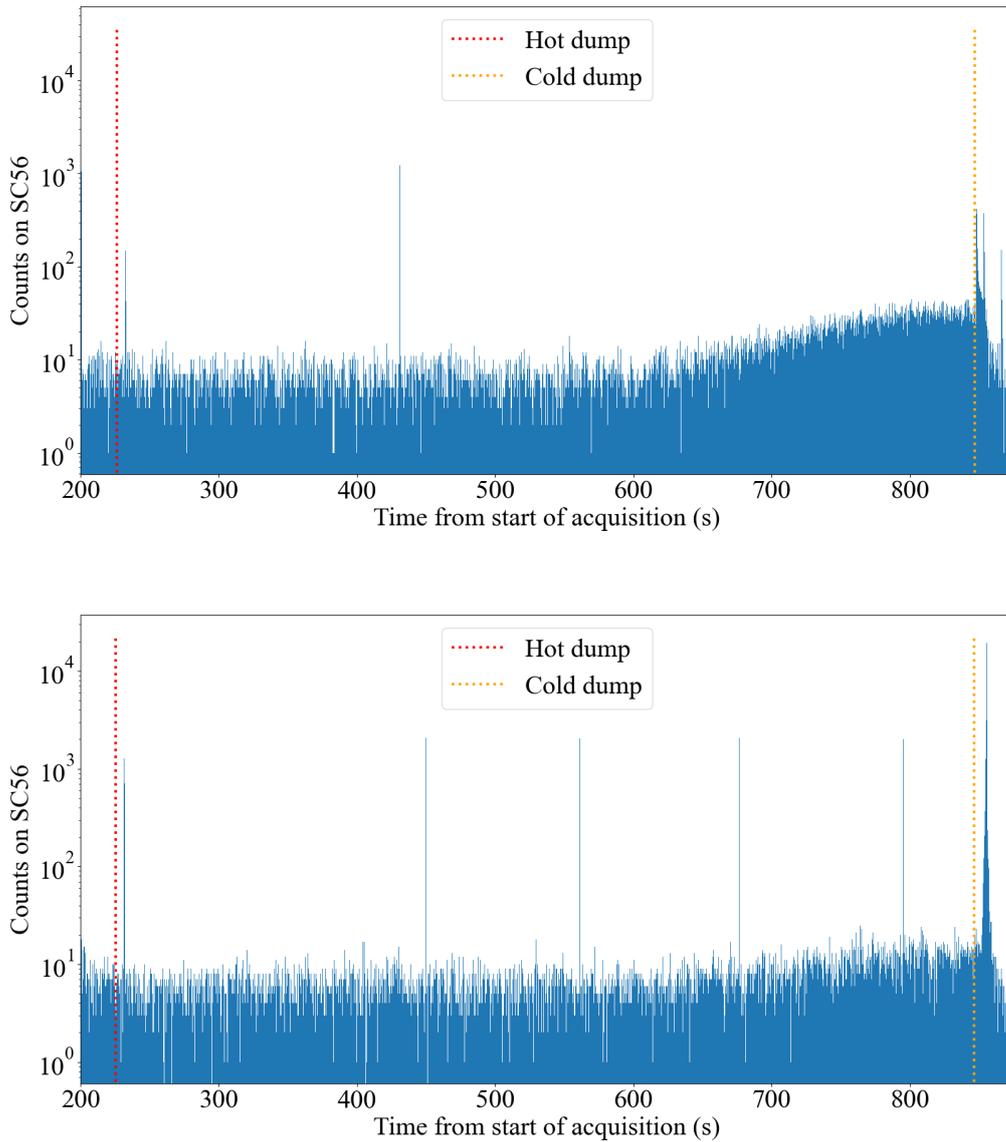


Figure 6.13: Comparison of the distribution of the number of counts on SC56 versus time when trapping one bunch of ELENA antiprotons in the P Trap for 600s without and with Rotating Wall compression. The red dotted line indicates the starting time of the hot dump; the yellow dotted line shows the onset of the cold dump. The individual peaks observed in between are caused by antiproton annihilations from AD injection or ELENA extraction events, of which only a fraction is visible with the given binning of 100 ms. Top: Rotating Wall off. Bottom: Standard Rotating Wall compression with non-optimized settings.

These parameters can be optimized to increase the number of antiprotons remaining in the trap after being stored by the low-voltage potential for a certain amount of time (i.e. the antiprotons in the cold dump), as shown in figure Fig. 6.14 for a storage time of 60s. Shown is the number of discriminated counts registered on SC1112 for the cold dump (after background subtraction), using different RW amplitudes, as a function of the employed RW frequency. To be able to compare different runs, the counts on SC1112 are normalized to the beam intensities of ELENA, i.e. the number of overall available

---

antiprotons: All values are divided by the ratio of the corresponding beam intensity to the maximum beam intensity found for a run in the presented scan.

Here, it is found that for the compression of a single ELENA stack in the trap, the number of remaining antiprotons is generally maximized when using higher Rotating Wall frequencies and lower amplitudes. This is in good agreement with the previous findings for the use of the RW technique in AEGIS, which are reported on in [108], including also further systematic studies. In this publication, the radius of the plasma as a function of the RW parameters is studied in detail, verifying that it is the working compression, i.e. reduction of the plasma radius, that prolongs the the lifetime of the antiprotons in the trap by preventing radial losses. This effect is also observed in the current data: Fig. 6.15 shows images obtained with the downstream MCP/CMOS detector after a dump of the confined plasma from the trap subsequently to a 60s cold storage with the use of two of the RW settings used in Fig. 6.14. For the plasma image on the left hand side, an RW amplitude of 0.9 V and a frequency of 0.2 MHz are used, i.e. settings that have yielded a low number of remaining antiproton counts in the measurement shown in Fig. 6.14. Instead, the plasma imaged on the right hand side has been compressed by the application of the RW technique with an amplitude of 0.1 V and a frequency of 1.8 MHz. In both images, a determination of the average background value is first performed, in the same way as explained for the image on the left hand side of Fig. 6.12, which is then subtracted from the pixel intensity values as a normalization. The yellow  $50 \times 50$  pixel square shown in the images indicates those pixels included in the background estimation. The color scale represents the intensity observed on the camera pixels, whose coordinates are converted to millimeters according to the calibration from Sect. 5.4.1. Apparently, the radial extension of the plasma core is significantly reduced on the right hand side compared to the left hand side, verifying that a reduction in the plasma radius corresponds to a larger fraction of antiprotons surviving in the trap. It should be noted that this effect is consistently observed despite Fig. 6.14 and Fig. 6.15 having been obtained in different measurements, once dumping the plasma toward the degrader structure to study the annihilation counts in the ESDA and once dumping toward the MCP to image the plasma.

Depending on the requirements of a given measurement and the corresponding particle numbers confined in the trap, the RW settings have to be individually optimized (see for example Sect. 6.3), and this study is only meant to show the general functionality of the technique.

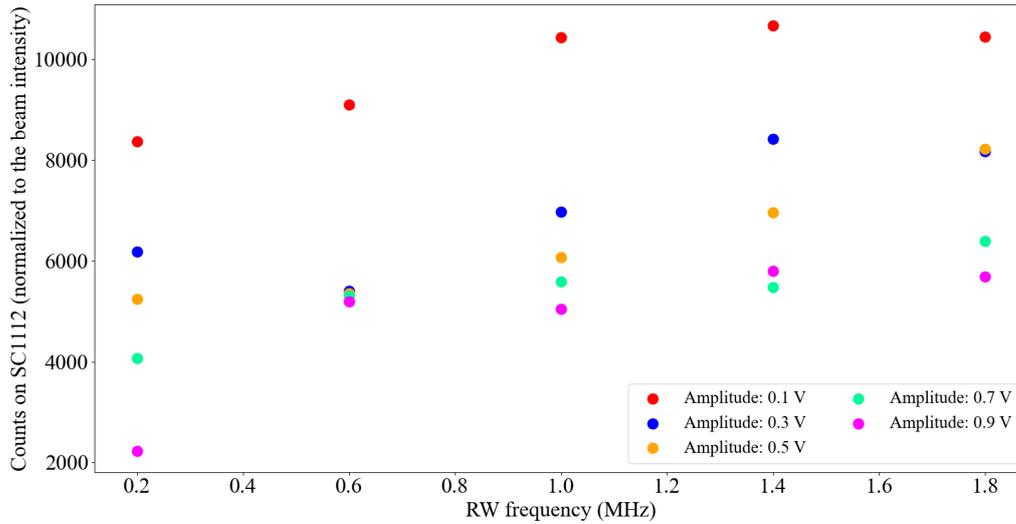


Figure 6.14: Measurement of the number of cold dump (60 s of cold storage) counts on an exemplary ESDA scintillation detector (SC1112) as a function of different Rotating Wall frequencies for various Rotating Wall amplitudes. The number of counts are normalized to the ratio of the ELENA beam intensity of a given run to the maximum observed intensity in this scan.

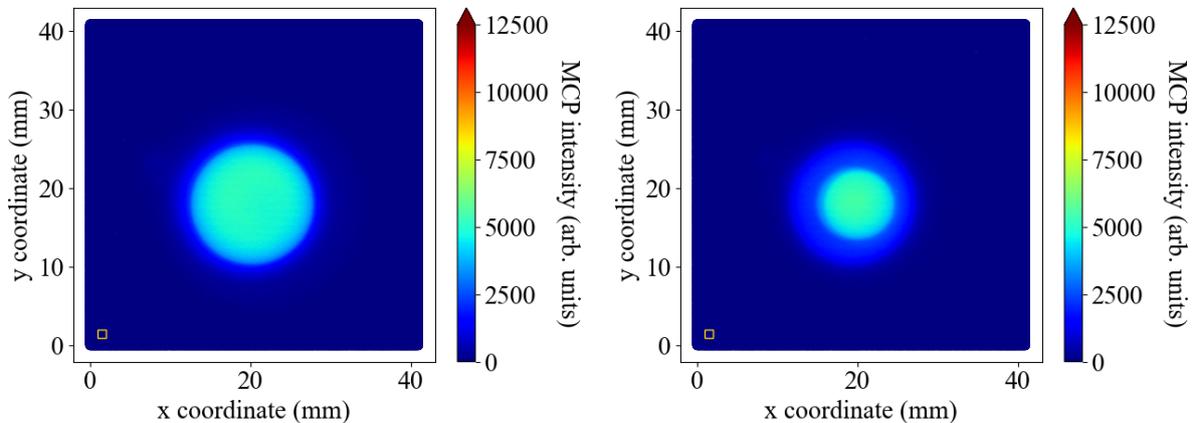


Figure 6.15: Example CMOS images taken on the downstream MCP detector of the electron-antiproton plasma dumped from the trap, subsequently to being subject to Rotating Wall compression in the trap, with two different settings, for a cold storage time of 60 s. The yellow square borders the pixels included in the determination of the background, which is subtracted from the pixel intensities. Left: RW amplitude of 0.9 V and frequency of 0.2 MHz. Right: RW amplitude of 0.1 V and frequency of 1.8 MHz.

### 6.3 Accumulation of cold antiprotons in the trap

With efficient electron cooling and the possibility to extend the antiproton lifetime inside the trap thanks to the Rotating Wall compression, it becomes feasible to repeat the proce-

cedure used to capture antiprotons from ELENA and keep them in the low-energy potential multiple times in a row without discarding of the already present particles. In this way, several bunches of antiprotons can be accumulated in the trap volume. This procedure is referred to as "stacking" of ELENA bunches and has been developed in CIRCUS with the following steps: antiproton capture – hot storage – hot dump – cold storage and compression – antiproton capture – ... – cold dump, where "..." includes the repetition of the previous steps for a chosen number of times.

For the application in AEGIS, the number of antihydrogen atoms produced through the used charge exchange reaction depends linearly on the number of available antiprotons, as is apparent from Eq. 3.14. The accumulation of cold antiprotons is therefore a crucial step toward the increase of the  $\bar{H}$  flux for the gravity measurement.

The accumulation limit is defined by space charge effects and losses of the antiprotons on the trap walls due to the radial expansion of the growing plasma. Therefore, with increasing numbers of stacked antiproton bunches (assuming relatively constant numbers of captured particles), the following evolution of trapped numbers of antiprotons is expected: For lower numbers, the increase is expected to be approximately linear until above limitations start to take effect and the increase flattens out toward a plateau, where space charge losses and newly added particles balance each other out. The ultimate limit depends on the quality of the plasma compression, i.e. optimization of the RW technique for a given trap configuration.

Fig. 6.16 is obtained from an example experiment, which captures, stores and accumulates eight antiproton bunches from ELENA in the trap. Shown is the distribution of the number of counts on SC1112, close to the center of the experiment, as a function of the time from the start of the acquisition. In these measurements, the typical hot storage time is 60 s and 120 V of space charge are transferred from the C Trap to the P Trap. The upper plot includes the entire acquisition time with a binning of 1 s, while the bottom plot shows a zoom to the range that is taken into account for the determination of the number of counts, using a 100 ms binning.

In regular intervals, the extraction of antiprotons from ELENA to the experiments and their injection into the AD for the new deceleration cycle cause a strong flux of annihilating antiprotons, which can be seen in the distribution as an increased number of counts in the ESDA. This is the recurring structure of two sharp peaks, the first one (ELENA extraction) slightly less intense than the second (AD injection), visible in the top plot. These events are observed at precisely known times, which are tagged in the analysis software, and if they happen during the signal window, i.e. the time range used for the summation of the counts, the additional counts can be subtracted. Generally, their number is found by summing the counts within a 200 ms window around the expected time of the peak, and the signal window size is reduced accordingly (which is relevant for background subtraction). It should be noted that for the accumulation of a very big number of ELENA bunches (from around 60), the tagging of these events in the software fails due to an integer overflow prompting the time counter to restart from zero. In these cases, the corresponding count numbers to be subtracted are obtained as an average of those ELENA extraction/AD injection events that can still be tagged, if the peaks are expected in the signal window.

The broader peak structures occurring in Fig. 6.16 in regular intervals between the described ELENA/AD events correspond to the hot dumps, i.e. removal of those antiprotons from the trap that have not cooled within the hot storage time by lowering the high-voltage potential. As expected their number increases with the number of accumulated antiprotons, since the trap operations to add new particles as well as their interactions partly

heat those antiprotons already present in the plasma, causing small additional losses here. Furthermore, the overall loss rate, independent of the peak structures, due to continuous antiproton losses also increases slightly with the number of particles present in the trap. This is because the higher particle density leads to an increased expansion of the plasma. This translates into an overall small increase in the number of counts.

Aside from the antiproton losses, cosmic ray interactions and sources of radioactivity present in the experimental area contribute to the continuous count rate. In fact, for a 10 cm wide scintillator slab, such as SC1112, a rate of 20 Hz would be expected from cosmic rays alone<sup>8</sup>, while the observed rate is approximately twice as large. The individual background rate that is subtracted from the absolute number of signal counts is obtained for every run individually. For its determination, the time range marked by the dotted red vertical lines in the top plot of Fig. 6.16 is used. It corresponds to the first microseconds (usually slightly below 10  $\mu$ s) of the acquisition, before the first ELENA extraction and subsequent antiproton capture. Typical observed background rates are of the order of 40 to 50 Hz, and the number of counts to be subtracted is calculated from it according to the individual signal window size of the run<sup>9</sup>. The background rate is indicated in both plots by the horizontal dashed pink line.

The solid yellow vertical lines shown in both plots of Fig. 6.16 indicate the used signal window that is employed to count the number of annihilations from the remaining cold antiprotons subsequent to the cold dump, whose potential operation is shown on right hand side of Fig. 6.3 and which refers to the opening of the low-potential trap in one direction to let the particles escape. In default operation, the trap content is dumped in the upstream direction onto the degrader structure and the speed of the lowering of the electrode voltages to release the particles depends on the exact measurement. Typically, as for the measurement shown here, the lowering happens slowly in 30 to 60 s. A lower limit on the ramping speed stems from the necessity to stay within one AD/ELENA cycle time so as not to miss the next round of available antiprotons; however, when loading a very large amount of particles into the trap, it is important to lower the potential in slow enough steps to not saturate the ESDA from the flux of escaping and annihilating particles. The bottom plot of Fig. 6.16 shows a zoom to the signal region of the same measurement and analysis depicted above. The two sharp peaks marked by the dotted cyan lines are the mentioned ELENA extraction and AD injection events, whose counts are subtracted. Within the signal region, there is a first peak of increased count numbers, which is caused by the final reshaping of the trap from the cooling trap to the dump trap, shown on the left hand side of Fig. 6.3. While the hot dump is included in every iteration of accumulating antiprotons, this reshaping is only performed once at the end of the sequence: To optimize cooling, the cooling trap is maintained during the accumulation, i.e. the potential configuration shown on the bottom right of Fig. 6.2 is kept the same while adding antiproton bunches to the trap. The loss of some of the contained, cold particles is expected during the reshaping, in particular because the available potential is reduced for the same present space charge, despite it being done slowly to minimize these losses. As this operation involves the entirety of the included trap electrodes and represents a strong change of the trap potential, it disturbs the plasma in general and causes spiked losses as opposed to a controlled opening of the potential on one side only. The second prominent

<sup>8</sup>The expected cosmic ray rate corresponds to a standard estimation of a little more than one particle per  $\text{cm}^2$  per minute.

<sup>9</sup>While the loss rate of the antiprotons themselves could also be referred to as background, the goal of the analysis is usually the determination of the number of cooled antiprotons left in the trap, and since those lost due to the expansion of the plasma in the trap after the hot dump are still expected to be cooled (as they are confined by the low-voltage potential), they may still be counted.

structure in the signal window is the peak originating from the actual cold dump. Its structure is discussed in Sect. 6.4 and it is defined by the distribution of the axial energies of the confined particles: Those with higher energies escape earlier on in the lowering process of the electrode voltages, while colder particles remain in the lower potentials, leading to the observed exponential shape of the peak (Maxwell-Boltzmann distribution in the case where the particles are in equilibrium). In the shown measurement, the signal window extends to the last detected count.

The number of counts observed on SC1112 in the signal window are summed up and the determined background rate is subtracted from the resulting value. To obtain an estimate of the number of cold antiprotons confined in the trap, the number of counts on the ESDA unit is converted according to the procedure outlined in Sect. 3.4.9.

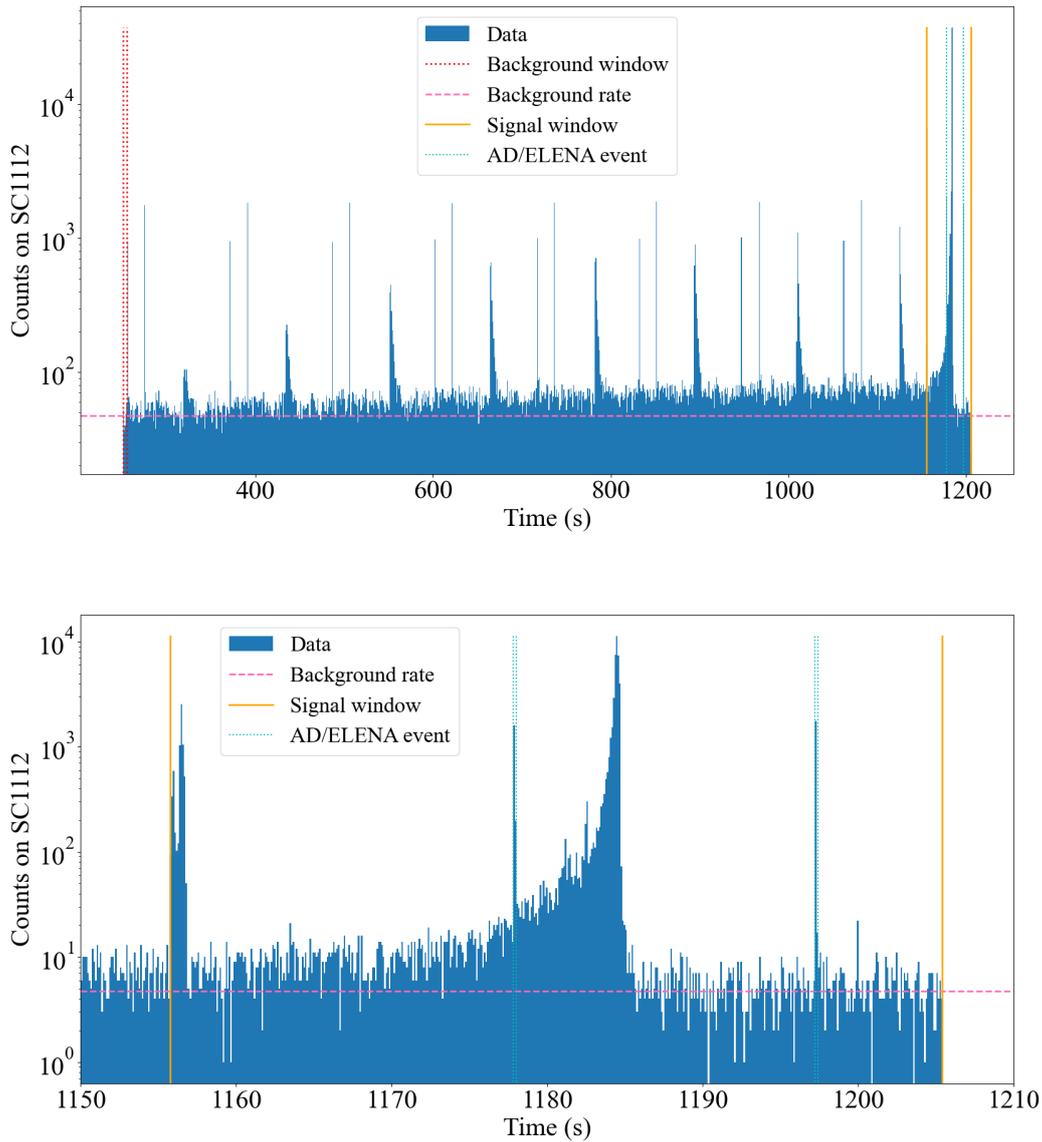


Figure 6.16: Distributions of the number of counts on SC1112 versus time for an example experiment accumulating antiprotons from eight ELENA bunches in the trap. The red dotted lines highlight the window used for the background estimation; the dashed pink line indicates the background count rate. The solid yellow lines mark the signal window. The dotted cyan lines indicate the last detection of the ELENA extraction and AD injection events, which happen in regular intervals that are tagged in the software. Top: The entire distribution from start to end of the acquisition. ELENA extraction and AD injection events appear in pairs of sharp peaks. In between, broader peaks correspond to the counts of the hot dump events. Bottom: Zoom to the signal window. The first sharp increase in the number of counts stems from the reshaping of the trap from the cooling to the dump configuration. It is followed by the broad peak of the slow cold dump. The ELENA extraction and AD injection spikes appearing in the signal window are subtracted from the overall counts in the analysis.

The ideal Rotating Wall settings do not only depend on the trap configuration but also on the number of particles confined in it, i.e. the plasma density (see Sect. 3.2.1 and

Sect. 3.2.2). For this reason, when accumulating several bunches of ELENA antiprotons in the trap, the RW parameters need to be re-optimized instead of using those settings that are found to be useful for single stacks. In principle, an ideal procedure would adapt the settings to the plasma density, or at least to the number of acquired ELENA bunches, in every iteration. A protocol to "sweep" the RW frequency for this purpose during the accumulation is in principle already developed and included in the AERIALIST code; however, due to time constraints during the limited duration of antiproton availability, an optimization has not yet been done for most of the points<sup>10</sup>. Therefore, for now, the settings obtained for a re-optimization with 45 accumulated ELENA bunches are used as the default and the full procedure will be implemented in the course of 2024. After a very broad scan, an overview of possible optimized settings is given in Fig. 6.17. Shown is the number of counts observed on SC1112 after the final cold dump of the accumulation procedure outlined above for three different RW amplitudes as a function of the RW frequency. The counts are normalized to the beam intensities in the same way as for the optimization with a single ELENA shot, which is described in Sect. 6.2.2, using the mean beam intensity of all ELENA bunches included in one accumulation as a factor for the given data point instead of the single intensity. Error bars are not included in the data points, but it can be seen from taking the same measurement multiple times that the deviation for the same setting is quite large. This is partly due to the fact that for such long accumulation runs, ELENA does not always operate perfectly stably over the several hours needed for the measurement and individual bunches are sometimes skipped; these are then added to the trap in the next iteration, however, continuous losses in the trap during the additional required time of course distort the results. Some outliers are also included, where for example the data acquisition system failed. These measurements are not performed to obtain absolute values but are instead supposed to show the overall trend, allowing to estimate useful parameters. These trends observed in Fig. 6.17 indicate that for lower RW amplitudes, the ideal RW frequency that maximizes the number of antiprotons left in the cold dump is higher (e.g. around 3.5 to 4 MHz for an amplitude of 0.15 V) than for higher RW amplitudes (around 1.5 MHz at 0.6 V). In agreement with [108], it makes sense that an upper limit for the RW frequency is observed: Likely due to space charge displacement (high on-axis particle density) and possibly additionally heating of the plasma from the excessive RW application, additional losses are observed for high RW frequencies. For the record accumulation of cold antiprotons in the trap, which is reported on here, an RW frequency of 3 MHz is chosen with an amplitude of 0.3 V.

---

<sup>10</sup>Since the capture of every single ELENA bunch takes of the order of two minutes (defined by the AD/ELENA cycle), the time for the accumulation scales with the number of bunches, and the time required to obtain a single data point exceeds several hours for long accumulations.

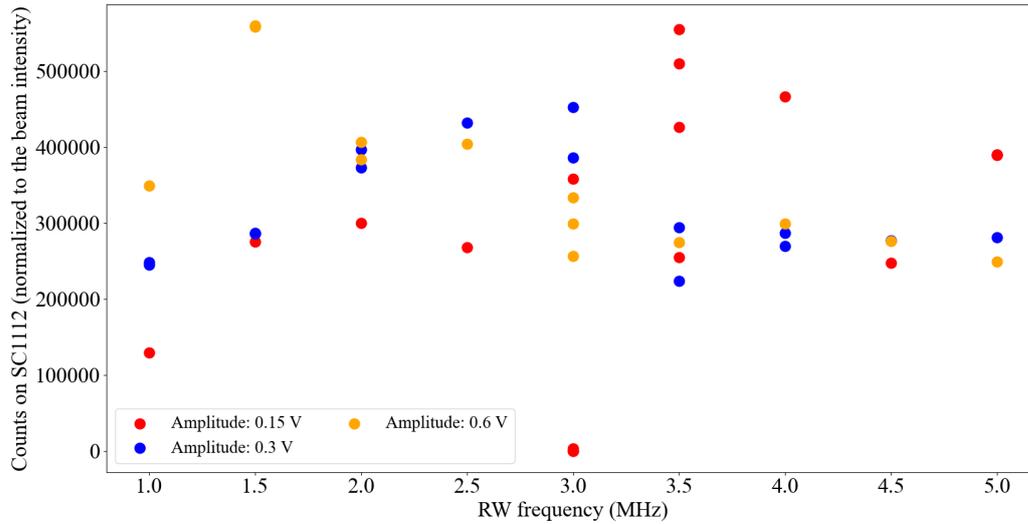


Figure 6.17: Measurement of the number of cold dump counts on ESDA unit SC1112 following an accumulation of 45 ELENA antiproton bunches in the trap as a function of the Rotating Wall frequency for three different Rotating Wall amplitudes. The number of counts are normalized to the ratio of the average ELENA beam intensity of a given accumulation to the maximum observed average intensity in this scan.

Using the optimized RW settings, a scan is performed over the number of ELENA bunches accumulated in the trap, using the described procedure. Fig. 6.18 shows the resulting number of cold dump counts on SC1112 as a function of the number of stacked ELENA bunches. An initial scan has been performed for the range between five and 70 ELENA bunches in steps of five. The data point at 60 bunches is taken three times, on different days, showing again the usual fluctuations in the produced numbers. Error bars taking these fluctuations into account are not shown. Additionally, two particularly long accumulations have been performed with 120 and 240 ELENA bunches, which are also plotted. One 240 bunch accumulation run takes over eight hours of continuous running and accumulation, if everything work ideally. Taking into account also individual empty bunches from ELENA and/or longer dead times and resulting retries, this time is much longer. Therefore, and because of the priority given to the development of the antihydrogen production procedure itself during the limited available antiproton time in 2023, such long accumulations have only been performed on single occasions and the limits of the procedure are currently being explored in 2024. The second vertical axis shown in Fig. 6.18 indicates the conversion of the number of counts on SC1112  $N_{\text{counts,SC1112}}$  to the number of trapped cold antiprotons  $N_{\bar{p}}$ . It is obtained using Eq. 6.4 with the conversion factors determined as described in Sect. 3.4.9. The relative error arising from the propagation of the uncertainties in the individually obtained factors included in the conversion and also noted in Sect. 3.4.9 is 3%. As an example, for the measurement having detected 500 million antiprotons, this corresponds to 15 million  $\bar{p}$ . The error bars plotted in Fig. 6.18 correspond to these uncertainties.

$$N_{\bar{p}} = N_{\text{counts,SC1112}} \cdot f_{\text{SC1112} \rightarrow \text{SC12}} \cdot f_{\text{SC12} \rightarrow \bar{p}} \cdot f_{\text{splitter}} \quad (6.4)$$

A relatively linear increase of the number of trapped antiprotons with the number of stacked ELENA bunches can be observed in Fig. 6.18 for the fine scan up to 70 bunches,

in good agreement with the number of antiprotons captured from ELENA per shot, as found in Sect. 6.1.5. The two data points for the long accumulations suggest a flattening of the curve and resulting slower increase in the numbers, which would be compatible with the onset of the plateau expected for the limitations of the plasma compression discussed above. An exponential fit has been applied to the dataset, which is included in Fig. 6.18 as a saturation curve eye guide; the saturation constant in terms of number of stacked ELENA bunches,  $N_{\text{sat}}$ , and the number of counts corresponding to the plateau,  $Y_{\text{max}}$ , are also stated. The highest number of cold antiprotons confined in the trap found in this way is slightly below 500 million. In the trap volume of approximately  $66 \text{ cm}^3$  between electrodes P3 and P9, this corresponds to a density of the order of 7.5 million  $\bar{p}$  per cubic centimeter and represents by far a record intensity source of controlled, cold antiprotons.

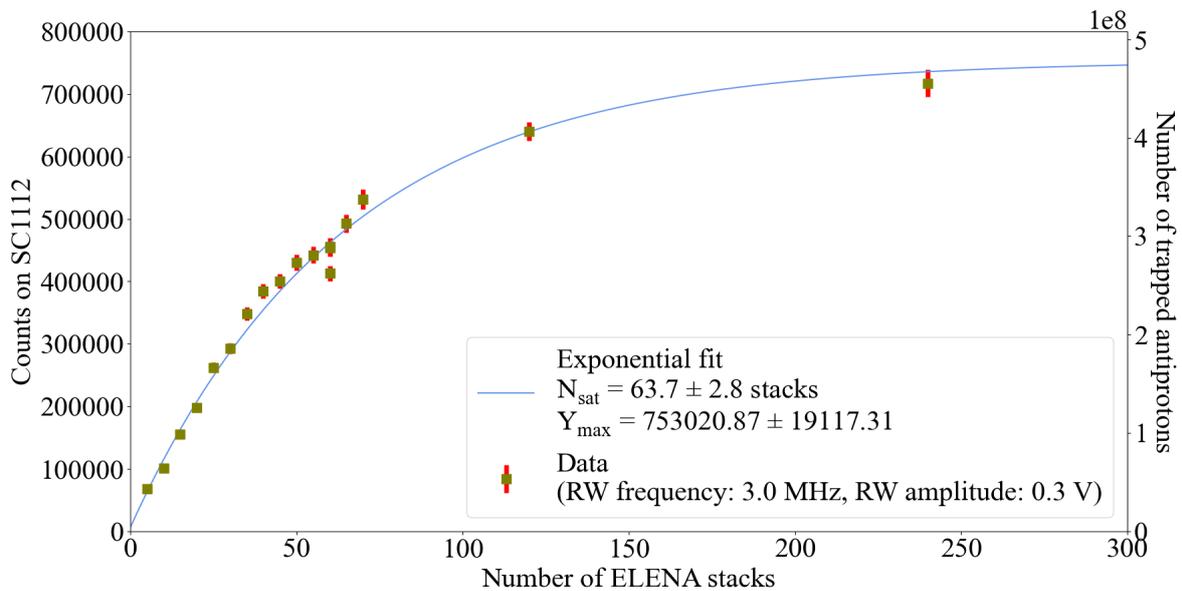


Figure 6.18: Number of observed cold dump counts on SC1112 (left vertical axis) and corresponding number of trapped cold antiprotons (right vertical axis) as a function of the number of ELENA bunches accumulated in the trap. The used RW settings are 3 MHz and 0.3 V, the hot storage time is 55 s, and 120 V of space charge are transferred from the C to the P Trap for electron cooling. As an eye guide, an exponential saturation curve fit applied to the dataset is shown as well (blue line).

## 6.4 Axial energy components of antiprotons in the trap

The temperature of a plasma is a measure of the average thermal kinetic energy of its constituents. The plasma temperature is only clearly defined in equilibrium conditions, when the energy distribution is Maxwell-Boltzmann-like, and it can be determined both in the radial and the axial direction (i.e. via the radial or axial energy components). It is generally possible to determine the axial energy component and, thus, the temperature of a plasma in equilibrium confined in a Penning-Malmberg-like potential configuration via a slow, controlled reduction of one of the confining endcap potentials, releasing hotter particles in every step. From the determination of the number of escaping particles in the high-energy tail in combination with the known ramping speed, the axial ensemble temperature can be inferred. This approach is referred to as the Parallel Energy Analyzer

technique [168]. Since the plasmas investigated here have likely not yet fully reached equilibrium, the details of this technique are not discussed here, as a systematic analysis will have to follow. It has already been performed in the past for antiprotons confined in AEGIS traps [169]. Essentially, following the derivation in [168] and [170], the axial plasma temperature  $T$  can be determined according to Eq. 6.5 from the slope of the number of escaping particles  $N$  per unit of reduction in the barrier voltage  $V_{\text{barrier}}$ , with around 5% accuracy (the factor of 1.05 being obtained as an average value for the second term of an asymptotic expansion).  $e$  is the fundamental charge and  $k_B$  is the Boltzmann constant.

$$T \simeq \frac{1.05}{\frac{d}{dV_{\text{barrier}}} \ln(N)} \cdot \frac{e}{k_B} \quad (6.5)$$

The minimum kinetic energy  $E_{\text{kin,min}}$  in the axial direction required for a particle to overcome the barrier is given by Eq. 6.6. Here,  $\phi_{\text{sc}}$  is the space charge potential produced by those particles that remain in the trap, which is of the order of eV. It is assumed to not be modified significantly for the first most highly energetic particles that escape. For this reason, these are the ones typically used for the analysis.

$$E_{\text{kin,min}} = (\phi_{\text{sc}} - V_{\text{barrier}})e \quad (6.6)$$

Using the discriminated ESDA data and assuming a proportionality of the number of detected counts to the number of escaping and annihilating antiprotons,  $\frac{d}{dV_{\text{barrier}}} \ln(N)$  can be obtained: From the knowledge of the initial barrier voltage and the ramping speed, i.e. the time over which  $V_{\text{barrier}}$  is lowered from the initial value to 0 V, one can determine the voltage difference per unit of time. In this way, an axis transformation from time to the energy domain can be performed by including those ESDA counts in a given energy bin that correspond to the time at which  $V_{\text{barrier}}$  is expected to be below the equivalent axial energy component of the escaping particles. The ramping time and number of steps can be precisely steered in the AERIALIST control system (see Ch. 4) and are therefore known. The upper plot of Fig. 6.19 depicts such a transformation, with the main plot showing the distribution of counts as a function of the particle energy and the inset plot the raw ESDA data (SC1112 in this case) in the time domain. Five stacked ELENA bunches of antiprotons are dumped from the Short P Trap. The initial  $V_{\text{barrier}}$  in this measurement is approximately  $-82.5$  V (55% of the  $-150$  V applied on barrier electrode P3) and it is ramped down to ground in 180 s (in 5000 ramping steps). As is also the default case, the trap floor is kept at  $-10$  V in this measurement (which are actually reached, since many electrodes form the floor), thereby reducing the effective initial  $V_{\text{barrier}}$  to  $-72.5$  V. It can be seen that during the first 125 or so seconds of the ramp no significant number of counts (above the background level, pink dashed line) is observed, i.e. the axial energy component of the vast majority of antiprotons is not close to 72.5 eV. From around 900 s, the highest-energy tail is observed, before the bulk of the plasma escapes, whose axial energy component is found to be below 10 eV by eye. It makes sense that the number of ESDA hits from escaping antiprotons is promptly reduced after the main peak and hardly any signal is observed in the last seconds of the ramp: The ramping includes also the last part from  $-10$  V to ground, however, the floor itself is kept at  $-10$  V, causing any remaining particles to already spill out around or slightly before this voltage is reached for  $V_{\text{barrier}}$ . In the energy domain, this feature is consistently observed as a lack of excess counts for negative energies.

As stated previously, the distribution in the upper plot of Fig. 6.19 seems to be composed of two components with different slopes on the high-energy ends and a transition in between: the main peak at lower energies with a steep slope and a more shallow slope at

higher energies. One interpretation is that the two parts correspond to different components of the confined plasma, one being made up of antiprotons that have already reached or are close to reaching equilibrium (main bulk) and the other representing those particles that would still require additional time for an equilibration (highest-energy tail). With respect to the timing of the measurement, this is a reasonable assumption since the plasma had very likely not fully reached equilibrium when the dump happened; it was given less than 1 s to re-equilibrate after the reshaping from the cooling to the dump trap configuration. The goal of these experiments had not, as such, been the determination of the plasma temperature but the development of the accumulation technique. Therefore, the obtained results should not be considered an actual temperature measurement, and more systematic measurements with longer delays are currently in the pipeline. Still, in particular because of the observation of the two components, an analysis of the axial energy components of the plasma particles is very interesting. It is very likely that the high-energy tail can only be as clearly distinguished from the background thanks to the large number of accumulated antiprotons and the resulting high count rates. In similar previous studies in AEGIS, such tails have not been observed [169].

It is possible to fit an exponential function to the high-energy slopes of both distribution components, yielding  $\frac{d}{dV_{\text{barrier}}} \ln(N)$  from Eq. 6.5 as the exponent parameter. However, due to their different nature, the two require a different binning of the data to allow for a good fit. The highest-energy tail requires a broader binning to provide sufficient statistics in every bin to perform the fit on while a fit to the peak of the bulk necessitates a finer binning to ensure a sufficient number of bins. For this reason, the distribution shown in the upper plot of Fig. 6.19, with a binning of 0.1 eV and 0.1 s, respectively, has only been fitted in the range of the highest-energy tail (orange line). For completeness, the shown fit converges on a temperature of  $\sim 42\,200$  K (reduced  $\chi^2 \sim 2.9$ ) when applying Eq. 6.5. The axial energy component of the plasma associated with this temperature would be around 3.6 eV. Results obtained in other runs and with different ESDA units are comparable.

The lower plot of Fig. 6.19 shows the same data as the upper plot, zoomed to the range of axial energies of the escaping particles between 0 eV and 20 eV (and the corresponding time range in the inset plot), i.e. to the main bulk of the distribution, with a finer binning of 0.01 eV and 0.01 s, respectively. The performed fit (represented by the orange line) to the high-energy tail of this part of the distribution yields a lower temperature of  $\sim 1200$  K (reduced  $\chi^2 \sim 1.5$ ), associated with an axial energy component of the order of 0.1 eV.

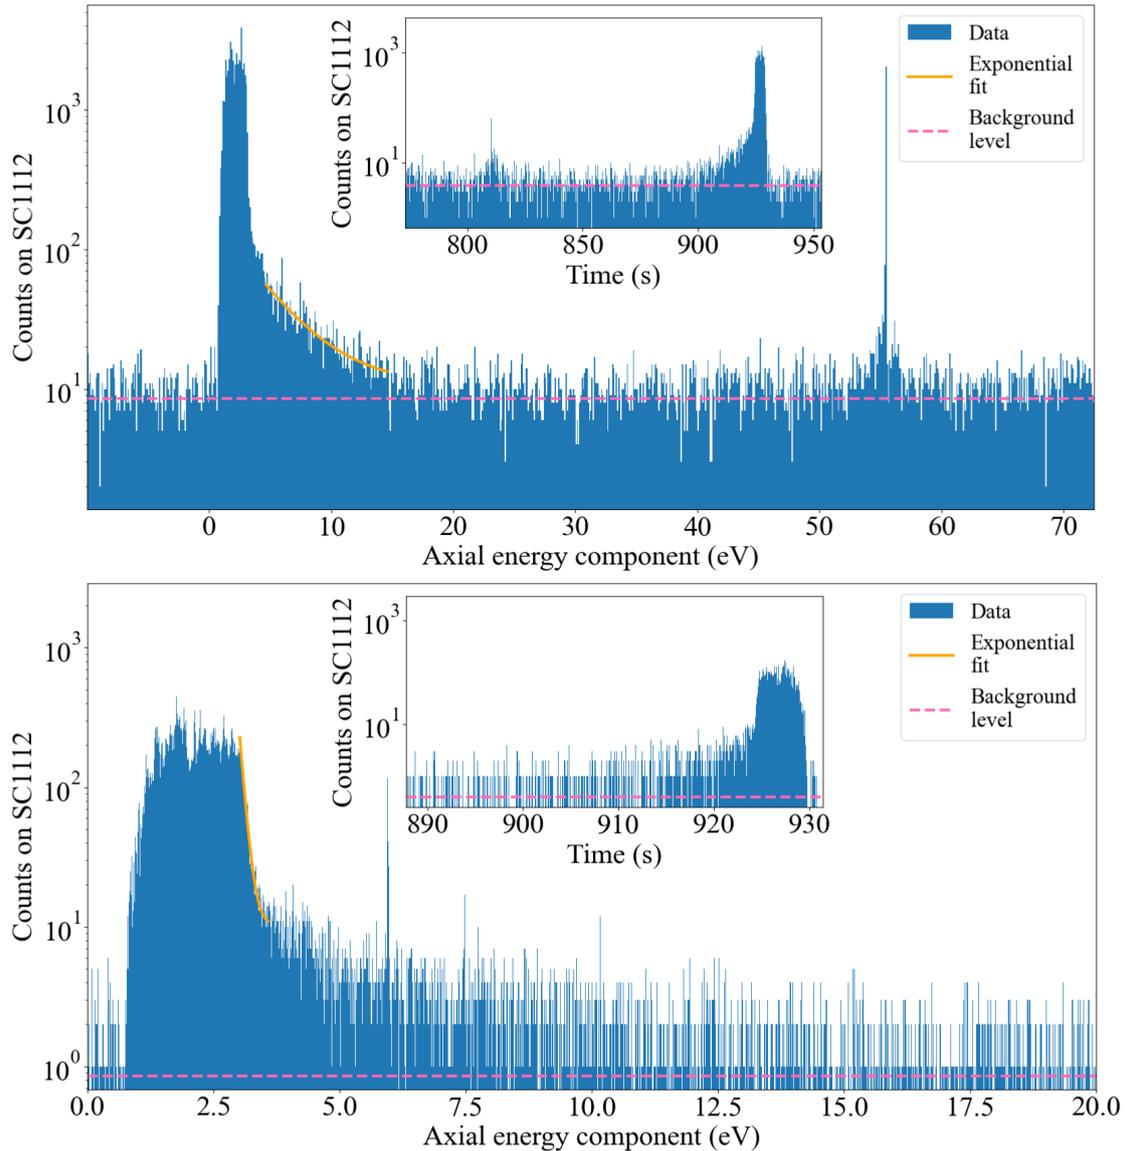


Figure 6.19: Distributions of counts registered on SC1112 in the energy (main plots) and time (insets) domains subsequent to a dump of five stacked ELENA bunches from the Short P Trap. The directly observed time is converted to the axial energy component of the escaping particles as explained in the main text. The orange lines correspond to exponential fits to the high-energy tails of parts of the distributions in the energy domain. The pink dashed line marks the average background level. Top: Plot of the number of observed counts during the entire dumping, with a binning of 0.1 eV and 0.1 s, respectively. Fitted is the highest-energy tail of the distribution, which likely stems from a plasma component that has not reached equilibrium. Bottom: Zoom to the energy range between 0 eV and 20 eV and the corresponding time range, with a binning of 0.01 eV and 0.01 s, respectively. Fitted is the high-energy tail of the main peak, which is interpreted as the part of the plasma that is close to equilibrium.

To conclude, it should be reiterated that, due to the lack of a plasma equilibration, the study performed here does not represent a temperature measurement of the plasma confined in the AEGIS trap system. However, if it is to be taken as a first indication of the plasma temperatures - being in good agreement with previous, more systematic measurements at similar equilibration time scales [169] -, it is apparent that an optimization

is, as expected, required to cool the antiprotons sufficiently for the produced antihydrogen to be at low enough axial temperatures for an efficient gravity measurement (see Sect. 3.5). Previous studies in AEGIS have shown that leaving several hundred seconds to the plasma for an equilibration can yield temperatures of a few hundred Kelvin to below 100 K [68, 169]. A longer re-equilibration time as well as an even more symmetric trap structure have already been implemented and are expected to result in lower axial temperatures. In general, avoiding any plasma disturbances and performing the necessary ones as slowly and gently as possible is beneficial; some optimizations of the plasma reshaping are also ongoing to improve this. To further reduce external noise, the outputs of the amplifier channels can also be decoupled (OptoMOS isolation, see Sect. 4.1.1). In any measurements done so far, this feature has not been exploited at all but it is one of the improvements foreseen when working toward a reduction of the antiproton temperatures for antihydrogen production.

## 6.5 Applications for large numbers of cold antiprotons

For AEGIS, the implementation of a source of large numbers of readily available antiprotons is a major step toward increasing the number of produced antihydrogen atoms, which scales linearly with the  $\bar{p}$  number (see Eq. 3.14). As discussed in Sect. 3.5, an increase in the  $\bar{H}$  number is crucial for the implementation of a statistically significant measurement of the influence of gravity on the antihydrogen atoms. Already the efficient capture and trapping of one ELENA shot, implemented in CIRCUS, yields an improvement factor of 6.5 (see Sect. 7.4.2 for a summary) for the expected  $\bar{H}$  numbers with respect to previous procedures in Phase I of the experiment. Clearly, the accumulation of several such  $\bar{p}$  bunches can significantly contribute to an even further increase in the  $\bar{H}$  numbers, making measurements feasible within a shorter time. It is, however, worth noting here that the increase in the  $\bar{H}$  number with that of the antiprotons is not expected to continue indefinitely for the proposed gravity measurement, for two main reasons: The increase in the  $\bar{p}$  density generally also increases the plasma rotation frequency and, with it, the resulting  $\bar{H}$  radial velocity, eventually requiring a forward-boost for the beam formation that is detrimental to a precision measurement. Additionally, an increase in space charge renders the axial velocity more disparate. Therefore, it remains to explore until which point exactly the increase in the number of available antiprotons is beneficial to the antihydrogen production.

In addition, the AEGIS collaboration is working on techniques for the controlled formation of antiprotonic atoms (or, more generally, antiprotonic bound systems) in vacuum by replacing one electron (or multiple electrons) of a regular atom with an antiproton. For this purpose, a source of anionic ions is being installed at the upstream end of the experiment. Cold antiprotons are to be confined in a trap structure together with these ions, which previously undergo photo-ionization and are then excited to Rydberg states. The two systems take part in a charge exchange reaction similar to the pulsed antihydrogen production technique routinely employed in AEGIS, forming antiprotonic species depending on the nature of the ions [171–173]. The simplest antiprotonic systems are protonium and antiprotonic helium, but more complex species are feasible as well. Antiprotonic atom formation is also possible starting from a source of cationic ions, which are co-trapped with cold antiprotons and undergo three-body interactions. This technique, as well, has been exploited for  $\bar{H}$  formation, however, with a relatively low cross section [72].

Already in 2023, techniques have been developed in AEGIS for the formation and manipulation of positively charged ions from the annihilations of trapped antiprotons on a

gas (e.g. Argon, Helium, Nitrogen) that is deliberately led into the experiment or, as in summer 2023, present in the apparatus due to a vacuum leak [85,174]. The nested trap procedure used is schematically shown in its simplest form in Fig. 6.20: Subsequently to the standard capture of antiprotons in the high-voltage potential (blue dashed line, between HV1 and HV3 in this case at  $-14$  kV), the entire set of low-voltage electrodes in the 5 Ttrap between the HV electrodes (i.e. from C2 to P14) is brought to a negative voltage (e.g.  $-190$  V shown here), producing the nested trap configuration plotted in red (solid line). In Fig. 6.20, the high voltage has additionally been reduced to  $-1$  kV in this step. The antiprotons undergo annihilations with the gas nucleons, producing, among others, positively charged fragments (e.g. ions, denoted " $A^+$ " here). When the HV electrodes are finally ramped down to ground (red dotted line), these fragments can end up captured in the upper part of the trap, while the antiprotons remaining in the lower part are released. For an analysis of the captured fragments, the ion trap can be reshaped to a positive potential well, analogously to the reshaping from the cooling to the dump trap configuration for antiprotons, and the fragments are finally ejected toward the downstream MCP detector in order to perform a time-of-flight measurement to identify their species and charge states, given a prior calibration with known particles (e.g. antiprotons). The data analysis and the preparation of a publication on the formation and spectroscopic study of highly-charged ions in this way are currently ongoing.

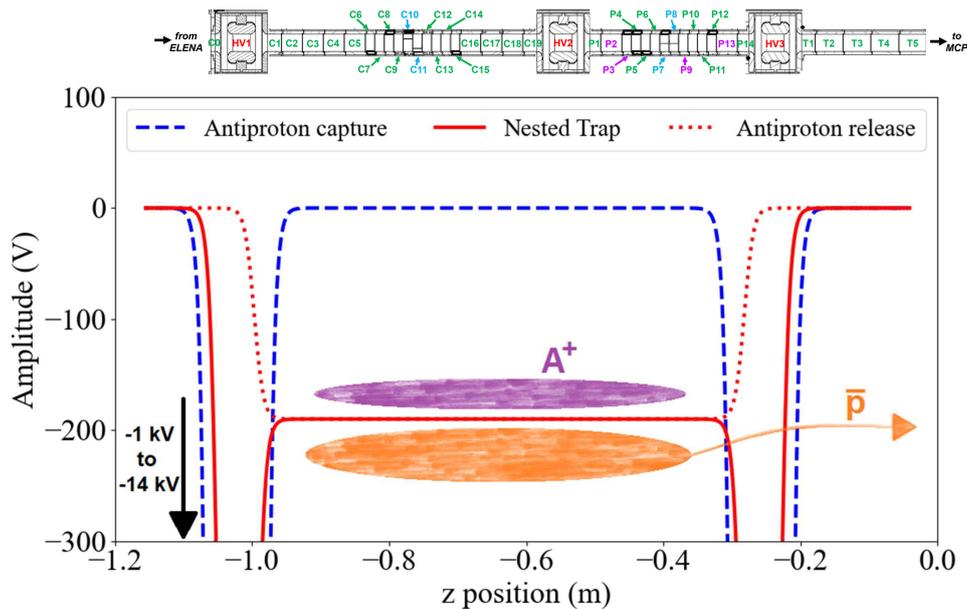


Figure 6.20: Nested trap potential used for the formation and capture of positively charged ions produced by the annihilation of trapped antiprotons on a gas. The antiprotons are first captured from ELENA in the standard high-voltage potential (blue dashed line) and interact with the surrounding gas. The thereby formed positively charged fragments (" $A^+$ ", purple ellipse) are captured in the upper part of a nested trap (red solid line) potential, which is subsequently shaped, while the antiprotons (orange) are released from the lower part when ramping the high-voltage potential to ground (red dotted line). Above the plot, a schematic of the involved range of trap electrodes is included, scaled to match the dimension of the x-axis, with the colors of the electrode labels having the same meaning as in Fig. 3.9.

$\bar{p}$ -induced nuclear fragmentation in traps represents a novel technique for the synthesis of highly-charged ions and radioisotopes that does not rely on the injection of a particle beam into a bulk of matter and subsequent deceleration, paving the way in particular for short-lived states and otherwise constrained species [174]. As such, above developments open up a playground for performing nuclear structure studies, while at the same time serving as the foundation for the formation of antiprotonic atoms by co-trapping negatively charged ions with antiprotons. Following the entire decay chain of such antiprotonic systems can yield insights into antimatter interactions as well as precision studies of fundamental symmetries. Possible scenarios for Dark Matter candidate searches, such as the hypothetical sexaquark ( $uuddss$ ), which is expected to be stable for a mass below approximately  $2\text{ GeV}/c^2$  and could have been abundantly produced in the early universe, are also feasible via the formation of antiprotonic atoms, the sexaquark for example being produced through the annihilations taking place in  $\bar{p}\text{-}^3\text{He}$  [173]. The low-energy environment of a formation in an antimatter trap yields favourable, cleaner conditions for a complete kinematical reconstruction of such a sexaquark compared to those of a high-energy experiment involving a large multiplicity of particles. In particular, for an annihilation into a sexaquark together with a  $\pi^-$  and two  $K^+$  mesons, the rest of the annihilation products are well detectable and the sexaquark could be reconstructed from the known energy and momentum of the initial  $\bar{p}\text{-}^3\text{He}$  system. Simulations and feasibility studies regarding such a measurement at CERN's Antiproton Decelerator are currently ongoing [173].

Of course, for the formation of antiprotonic systems and, generally, the interaction of antiprotons with atoms and ions in a trap, a large particle density is crucial. In this sense, the accumulation of cold antiprotons reported on here also marks an important step toward these goals.

Another application for a large number of accumulated antiprotons is the development of a portable antiproton trap, as currently pursued by the PUMA and BASE-STEP collaborations at CERN's AD [175,176]. Making antiprotons portable would allow to transport them to a calm experimental environment, free of the large magnetic field fluctuations present at the AD, to perform precision CPT tests [176], or to a location that offers experimental opportunities different from those in the AD. The latter include CERN's source of low-energy beams of radioactive nuclides, ISOLDE (Isotope Separator On Line DEvice [177]), where antiprotons are planned to be mixed with unstable atomic nuclei to perform studies of exotic phenomena in nuclear physics, in particular probing the nucleonic surface structure of the nuclei via the detection of the annihilation products [175]. To maximize the efficiency and maintain a reasonable particle density despite transport losses, starting from a large number of antiprotons in the trap is crucial; numbers of a few hundred million to one billion antiprotons are currently envisaged for portable traps [175], in line with the numbers achieved in AEGIS reported on here. Similarly, antiproton applications outside of fundamental research (e.g. biological/medical applications or material technologies) also become imaginable. As outlined in Sect. 3.5, the AEGIS collaboration itself has already succeeded in extracting the antiprotons from the trap backwards in the upstream direction, a fundamental step towards antiproton transportability.

Additionally, the confinement of a large number of cold composite antimatter systems could also, in itself, in the future provide windows for the observation of interesting effects of the behaviour of antimatter plasmas and, ultimately Bose-Einstein condensates. While such research likely still requires significant development, the production of any exotic atom or molecule that would contain more than one antiproton (e.g. the antihydrogen molecule), as well as any fundamental interaction process that requires multiple antiprotons in the initial state (e.g. antineutron production via the antimatter-equivalent

of the  $\beta^+$  decay, subsequent antideuteron formation, etc.) would benefit from such a cold, high-density  $\bar{p}$  source.

## 6.6 Summary

The procedure has been developed to routinely capture up to 70% of the antiprotons available from the newly commissioned and connected ELENA decelerator in the AEGIS electromagnetic trap structure. This achievement has required the optimization of the capture parameters for the ELENA antiprotons, which have an energy reduced by more than an order of magnitude compared to those previously supplied by the AD, and for the upgraded apparatus. These include the steering of the ELENA beam into the experiment, the determination of the best trap closing time for maximum capture efficiency, and the development of a degrader setup for an efficient energy degradation of the arriving antiprotons, as well as an upgrade of the high-voltage system to increase its reach for improved capture of the more highly energetic  $\bar{p}$  fraction. Furthermore, the involved routines needed for the capture on the hardware and software side have been translated to be used in the newly introduced CIRCUS control system, and the low-voltage potential operations of the AERIALIST fast-control unit have been studied in detail to provide symmetric and reproducible plasma confinement. The new control system has enabled a precise optimization of the parameters within a few hours, saving valuable time during the measurement campaigns. To determine the overall capture efficiency, both a direct study of the number of antiprotons being dumped from the trap, as obtained using an existing calibration of the ESDA detector system, and a separate investigation of the individual loss contributions have been performed, which are in good agreement.

As a standard procedure, to obtain 70% capture efficiency, the antiprotons are captured between high-voltage electrodes HV1 and HV3, with a closing time delay between 500 and 1000 ns on HV1, using the maximum 14 kV capture potential and only the main 1.4  $\mu\text{m}$  thick Mylar degrader, as the additional thin Parylene degrader foils have proved too mechanically unstable.

The captured antiprotons are then accumulated in the P Trap region of the electromagnetic trap stack, where a plasma of sympathetic cooling electrons has been previously prepared in a low-voltage (maximum  $\pm 200$  V) potential. After an optimization of the trap potentials and the amount of electrons used for cooling, a cooling efficiency of over 90% of the antiprotons to energies low enough to be confined by the low-voltage electrodes is achieved within a few seconds of combined confinement of the two species. To increase the mixed plasma's lifetime in the trap, it is being compressed by the use of the Rotating Wall technique, whose parameters have also been optimized, permitting a trapping of several hundred seconds without significant losses.

Finally, these developments together with a solid understanding of the involved procedures and potential reshaping operations in the traps have enabled the accumulation of multiple ELENA antiproton bunches in the trap. Using adapted electron cooling and compression parameters, several hundred bunches have been stacked up, yielding almost 500 million cold antiprotons in the 66  $\text{cm}^3$  trap volume.

In addition to achieving the goal of facilitating a significantly enhanced antihydrogen production in AEGIS - with the number of produced  $\bar{\text{H}}$  atoms being proportional to the available antiproton number -, this development also opens the door to techniques for the formation and study of antiprotonic atoms and highly-charged ions, the filling of portable antimatter traps, the production of more complex antimatter systems and the investigation of high-density antimatter interactions in general.

# Chapter 7

## Antihydrogen developments in AEGIS

*In this chapter, all developments I have been involved in towards an improvement of the antihydrogen production in AEGIS and the formation of a horizontal beam of the  $\bar{H}$  atoms are presented. As for the parts concerning the work with antiprotons, I have been the main person directly developing the potential routines to manipulate the plasmas in the needed ways with the AERIALIST control electronics and software. Together with M. Volponi and R. Caravita, I have also formed the core team running this part of the experiment during the data taking periods in 2021, 2022 and 2023. I have designed and implemented the parabolic potential procedure for the ballistic transfer of antiprotons required for the beam formation. The positronium formation work, as the second half of the antihydrogen production, was led by B. Rienäcker, with contributions from my side on the control system part and for the unification of the two parts. The procedures for the achievement of laser cooled positronium atoms, reported on in [116], of which I am a contributing author, have been realized on the basis of the ARTIQ/Python code library developed by me, and I have implemented the algorithms needed for the exact synchronization and control of the involved laser and positron systems, only possible thanks to the new control electronics built up by me and essential for the success of the approach. The analysis performed on the  $\bar{H}$  data with regard to the ESDA detector has also been done by me, based on the analysis framework developed by T. Rauschendorfer and guided by R. Caravita's experience.*

The main objective of the AEGIS collaboration is a precise measurement of the gravitational acceleration experienced by antimatter. The antimatter system of choice to perform this measurement with is antihydrogen in the form of a pulsed, horizontal beam of individual anti-atoms travelling in vacuum, outside of the magnetic field of the experiment, over a known distance through a deflectometer to end up on a time- and position-sensitive detector. Their vertical deflection due to the influence of gravity is measured, allowing to determine the experienced acceleration. Efficient antihydrogen production and the formation of a forward-boosted beam of the produced atoms are therefore at the core of the experimental procedures.

The general procedure employed in AEGIS to form antihydrogen is outlined in Sect. 3.3 and shown schematically in Fig. 3.4: Antiprotons are captured from CERN's ELENA decelerator, confined, sympathetically cooled and compressed in an electromagnetic trap, while positronium atoms are formed, through the implantation of positrons into a conversion target, and laser-excited to Rydberg states. The two components are combined to form antihydrogen via a charge exchange reaction. This last step has been previously performed with the positronium entering the trap where the comparatively stationary antiprotons were kept. A procedure has now been developed to instead accelerate the

antiprotons towards the Ps formation region in order to enable the formation of an anti-hydrogen beam.

The preparation of the first component, the antiprotons, has been upgraded significantly over the past few years and is described in detail in Ch. 6. Since the number of formed antihydrogen atoms scales linearly with the number of available cold antiprotons, an efficient  $\bar{\text{H}}$  production is only possible thanks to the much enhanced procedures and resulting abundant  $\bar{\text{p}}$  source outlined there.

With this development in place and a larger amount of positronium available at higher Rydberg levels thanks to the upgrades outlined in Sect. 3.3 and Sect. 3.4, newly expected numbers of produced  $\bar{\text{H}}$  atoms are opening the door to precision measurements. For this purpose, it has now become necessary to accelerate the anti-atoms in the desired forward direction towards the deflectometer. One way to achieve this is to accelerate the antiprotons in the direction of the positronium conversion target - and thereby in the direction of the deflectometer, which will be located further downstream - such that they combine in flight, via charge exchange, with the lighter positrons from the Ps atoms upon their creation and excitation, giving the resulting  $\bar{\text{H}}$  an immediate forward boost. To prevent a sudden expansion of the antiproton plasma during this acceleration, a procedure has been developed to guide the particles from the electromagnetic trap, from which they are launched, to the target along a parabolic potential formed by the set of electrodes located in between. This procedure allows to tune the average velocity of the produced antihydrogen atoms. It requires precise control of the involved low-voltage potentials and builds upon the structure of the newly commissioned CIRCUS/AERIALIST control system introduced in Ch. 4 as well as the thorough preparation of the antiprotons outlined in Ch. 6. For an improved efficiency, recycling routines for both the antiprotons and the cooling electrons have also been put into place, allowing to re-use those that have not been used for  $\bar{\text{H}}$  production.

The subsequent essential step is the synchronization of the antiproton routines to those of the positronium line such that antihydrogen is formed with maximum efficiency. Dedicated studies have been performed on the timing of the two, and they have eventually been combined into one profound routine. In the course of CERN's 2023 antiproton run campaign at the Antiproton Decelerator facility, almost 1000  $\bar{\text{H}}$  production runs have been performed with this technique, in three different modalities defined by the delay time given to the antiprotons prior to being launched, yielding expected  $\bar{\text{H}}$  production in the forward and backward directions and at rest. In addition, control runs without Ps laser excitation, i.e. without expected  $\bar{\text{H}}$  production, have been performed, allowing to gauge the differences in the resulting signals.

To draw conclusions on the  $\bar{\text{H}}$  signal, a dedicated analysis has been implemented for the data obtained with the AEGIS ESDA detector formed by scintillator/PMT units (see Sect. 3.4.9). The analysis includes an in-depth event selection and an interpretation of the observed signal timings, intensities and statistics. The result is a first preliminary analysis of the antihydrogen data obtained with AEGIS Phase II, in the on-axis production modality, using the upgraded apparatus and control system and the newly developed  $\bar{\text{p}}$  transfer method. Mainly due to issues encountered in the positronium system, the  $\bar{\text{H}}$  yield is extremely low and an analysis of the 2024 data will be needed to draw more statistically significant conclusions. An analysis of the combined statistics of the three obtained datasets does, however, yield a statistical significance of an excess signal compatible with antihydrogen production above the level of three standard deviations.

As an independent development, one-dimensional Doppler laser cooling of an ensemble of positronium atoms has been successfully achieved for the first time by the AEGIS collab-

oration in 2023 via a strong saturation of the  $1^3\text{S}-2^3\text{P}$  transition of ortho-positronium. Ps laser cooling represents a major result that opens the door to further fundamental research on positronium and has the potential to boost the  $\bar{\text{H}}$  production of AEGIS significantly thanks to the favourable charge exchange cross section. The Ps cloud is found to have cooled from  $(380 \pm 20)$  K to  $(170 \pm 20)$  K, as obtained from the observation of the velocity distributions from one-dimensional Doppler profiles. This achievement has required the development of a dedicated laser system and of a technique to transport Ps without the use of magnetic fields as well as the precise synchronization capabilities offered by the CIRCUS control system.

This chapter gives an overview over the mentioned developments in AEGIS to increase the antihydrogen production efficiency and form the produced anti-atoms into a pulsed, horizontal beam, which is needed for the measurement of the influence of gravity in the pursued scheme. In addition to the accumulation of cold antiprotons described in Ch. 6, these include the development of the antihydrogen procedures in the new control system CIRCUS, and with the newly installed 1 T trap, the successful forward-acceleration of the antiprotons prior to  $\bar{\text{H}}$  formation in a ballistic way, the recycling of the used particles and the implementation of laser cooling of positronium atoms. The influence of the different developments on the expected number of produced  $\bar{\text{H}}$  atoms is outlined as well, and the analysis of the ESDA signals obtained in 2023 is presented.

## 7.1 Ballistic transfer of antiprotons

In AEGIS Phase I, antihydrogen atoms were produced as an isotropic  $2\pi$  source, as the antiprotons were essentially at rest when the positronium atoms arrived to form  $\bar{\text{H}}$  [68]. When working towards the formation of a horizontal beam of these *neutral-charge* atoms, it is instead possible to accelerate the antiprotons prior to  $\bar{\text{H}}$  production in the desired travel direction. Since the added positron from the Ps has a lower mass than the antiproton by several orders of magnitude, its momentum has a negligible influence on the momentum of the final antihydrogen. With respect to the procedure shown on the right hand side of Fig. 3.4, the antiprotons are then not waiting in the 5 T trap for the positronium to enter but are launched from it towards the Ps target, where they are meant to be combined with the Ps atoms.

To transfer the antiproton bunch accumulated and cooled in the 5 T trap region of the experiment to the downstream end of the experiment in a ballistic way, without an abrupt expansion of the particle cloud, a parabolic potential is used. This allows for a slow acceleration toward the center of the flight path and a subsequent deceleration toward the Ps target, without blowing up the plasma, which would be the result of a sudden launch on a constant ground potential. A main advantage of the choice of this technique is the possibility to tune the velocity of the produced antihydrogen as needed by adapting the delay given to the antiprotons prior to being launched, i.e. choosing the velocity with which they meet the much lighter positronium (see the discussion in Sect. 7.4.3 for details). Additionally, by allowing the antiprotons to oscillate, i.e. swing back and forth in the parabolic potential, it becomes possible to recycle them and use them for more than one  $\bar{\text{H}}$  production run, improving the efficiency of the procedure.

### 7.1.1 Antiproton launch trap

The dump trap configuration shown in Fig. 6.3 formed by the Short P Trap between electrodes P3 and P9 is not the starting point of the launch for the ballistic transfer.

This is because, to enable a more efficient antihydrogen production, firstly, the electrons need to be removed from the mixed plasma before the antiprotons are launched, and secondly, a part of the P Trap is used to recycle electrons and antiprotons. This allows to reuse those particles that are not used in a production cycle for the subsequent ones. Furthermore, it is beneficial to confine the plasma in as small a space as possible so as to reduce the bunch separation during the launch and thus minimize the differences in the arrival time of the antiprotons in the target region. Therefore, prior to the launch, the dump trap is reshaped into a launch trap with P13 as downstream endcap electrode. The upstream-most electrode utilized in the configurations discussed here is either P2 (i.e. the Long P Trap is used) or P3. The reshaping steps to the launch trap, using P3, are shown in Fig. 7.1: First, all electrodes from P10 to P13 are brought to the same voltage as P9 (top left). The potential on all of them but P13 is then lowered to the slightly negative floor potential (top right), extending the plasma. Electrodes P3 to P9 are subsequently ramped up again to the same voltage as P13 (bottom left), slowly squeezing the confined plasma into the smaller space. It is to note that, because P13 is more than twice as long as P9, the same voltage programmed on P13 produces an approximately 30% stronger on-axis potential. On the other hand, by using several short electrodes together (P3-P9), the potential is further increased to even reach the programmed value ( $-150\text{ V}$  in this case). As a preparation for the parabolic launch, the voltage on P10 to P12, which form part of the parabola, is then raised from  $-10\text{ V}$  to  $-70\text{ V}$ .

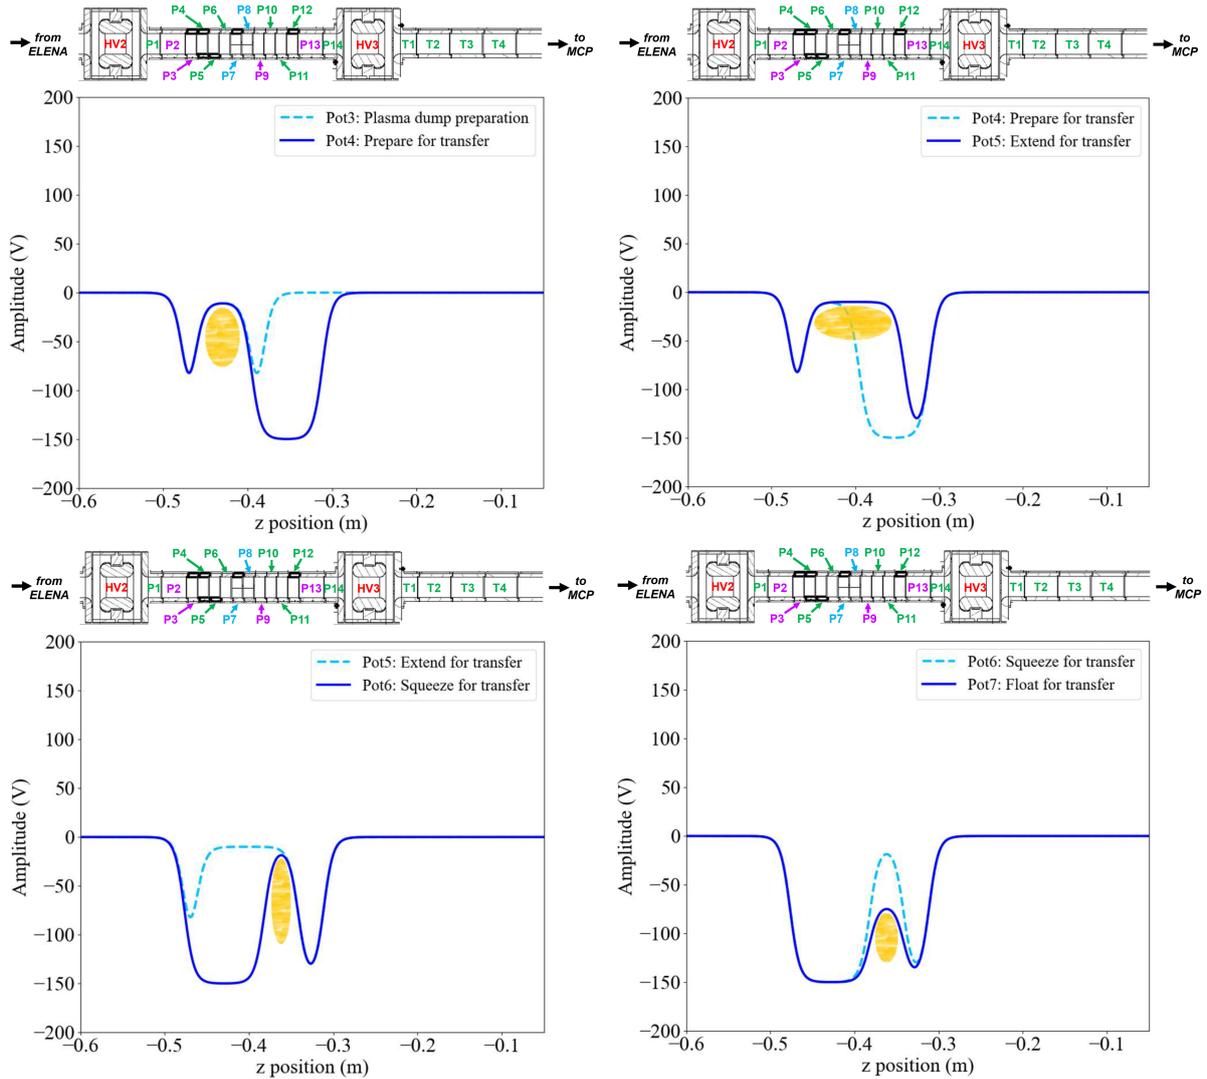


Figure 7.1: Low-voltage potential reshaping operations to transform the P Trap from the dump trap into the launch trap. Initial potentials are denoted by the cyan dashed line; final potentials are plotted in blue. The location of the mixed electron-antiproton plasma in each final potential is indicated by the yellow ellipse (not to scale). Above each plot, a schematic of the involved range of trap electrodes is included, scaled to match the dimension of the x-axis, with the colors of the electrode labels having the same meaning as in Fig. 3.9. Top left: Downstream extension of the outlet until P13. Top right: Extension of the trap floor until P12. Bottom left: Trap squeezing by extending the inlet in the downstream direction until P9. Bottom right: Raising of the floor voltage for the parabolic launch.

### 7.1.2 The parabolic potential

A carefully shaped parabolic potential is chosen for the transfer of the antiprotons towards the Ps target to be able to tune their axial arrival velocity and, thus, the velocity of the produced antihydrogen atoms, and to prevent a sudden expansion of the  $\bar{p}$  plasma during the launch, providing at the same time the possibility for a recycling of the used antiparticles.

Starting from the standard definition of a parabolic function, given in Eq. 7.1 for  $z$  as

a given horizontal position along the experimental axis and  $V(z)$  as the corresponding voltage along the parabolic potential, where the vertex of the parabola is defined by  $(b, c)$  and the opening is governed by the parameter  $a$ , the needed voltage at the center locations of all involved electrodes can be calculated.

$$V(z) = a(z - b)^2 + c \quad (7.1)$$

The known used parameters of the required parabola are:

- The locations, on the horizontal axis of the experiment, of the start and end of flight of the particles along the potential:  $z_{\text{launch}}$  as location of the center of the preparatory launch trap (i.e. the center of the central floor electrode of this trap) and  $z_{\text{end}}$  as the center location of the last electrode before the Ps target (i.e. the last electrode involved in the parabolic potential)
- The desired launch voltage to be applied on the central floor electrode of the launch trap,  $V_{\text{launch}}$
- The desired voltage of the last electrode involved in the parabolic potential,  $V_{\text{end}}$
- The desired voltage at the vertex of the parabolic potential,  $V_{\text{min}}$ .

Defining  $z_{\text{launch}}$  as the origin of the z-axis (i.e.  $z_{\text{launch}} = 0$ ) and shifting the locations of all electrodes accordingly, the parabola parameters can be determined as given in Eq. 7.2 through Eq. 7.5. The solution for  $b$  in Eq. 7.3 is valid if  $V_{\text{launch}} \neq V_{\text{end}}$ ; otherwise it simplifies to Eq. 7.4. In the former case, if  $V_{\text{launch}}$  and  $V_{\text{end}}$  are negative voltages, as is generally the case when working with negatively charged particles here, Eq. 7.3 produces a parabola for a positive sign in front of the square root; for positive voltages the minus sign has to be added. The opposite solution, respectively, yields a linear slope between  $z_{\text{launch}}$  and  $z_{\text{end}}$ .

$$c = V_{\text{min}} \quad (7.2)$$

$$b_{\neq} = \frac{\pm \sqrt{z_{\text{end}}^2 (V_{\text{launch}} - c)(V_{\text{end}} - c)} + z_{\text{end}}(V_{\text{launch}} - c)}{V_{\text{launch}} - V_{\text{end}}} \quad (7.3)$$

$$b_{=} = z_{\text{center}} = \frac{z_{\text{end}}}{2} \quad (7.4)$$

$$a = \frac{(V_{\text{launch}} - c)}{b^2} \quad (7.5)$$

The parabola used in standard conditions is formed between electrodes P10 and A1. For the current run period, the parabolic potential is optimized in such a way as to bring the required voltage on one of the electrodes near to its midpoint, T6, as close as possible to 0 V. This is needed because of a fault in the internal connection of the electrode that keeps it permanently at ground potential. This additional boundary condition has not inhibited the development and implementation of the approach significantly and is therefore accepted. Furthermore, a fix of the connection requires the opening of the most central region of the experiment, a procedure which is extremely delicate and time-consuming. However, a partial opening is planned for 2024, which will allow to access the electrode sufficiently to either replace its OVC cable or electrically connect it to the neighbouring electrode, T5, so as to be able to control the two together, essentially as one longer electrode, and give further flexibility to the procedure.

The voltages that are applied on each electrode involved in the parabolic potential are

plotted in Fig. 7.2 (red line). They are determined based on ideal assumptions, i.e. without taking into account the electrode lengths and the influence of the neighbouring potentials and by simply using the electrode center points as reference locations, calculated from their upstream position and their length that are noted in the `electrode_info.json` file. The red markers indicate the so given positions of the included electrodes, from P10 to A1. As seen in Fig. 3.9, the parabola spans over a large part of the electromagnetic traps, both in the 5 T and in the 1 T region of the experiment. The corresponding electrodes are therefore controlled by two different units of the AERIALIST and it is vital to be able to use the two parts together efficiently, as outlined in Sect. 4.5.2. Also plotted in Fig. 7.2 is the realistic potential (blue line) that is produced by the electrodes over the given range, obtained from the finite element calculation. As expected, the resulting potentials form a parabola that closely follows the desired values with small deviations according to the electrode lengths. The potential at the downstream end of the parabola deviates further due to the potentials applied beyond; they are briefly explained in Sect. 7.2.1.

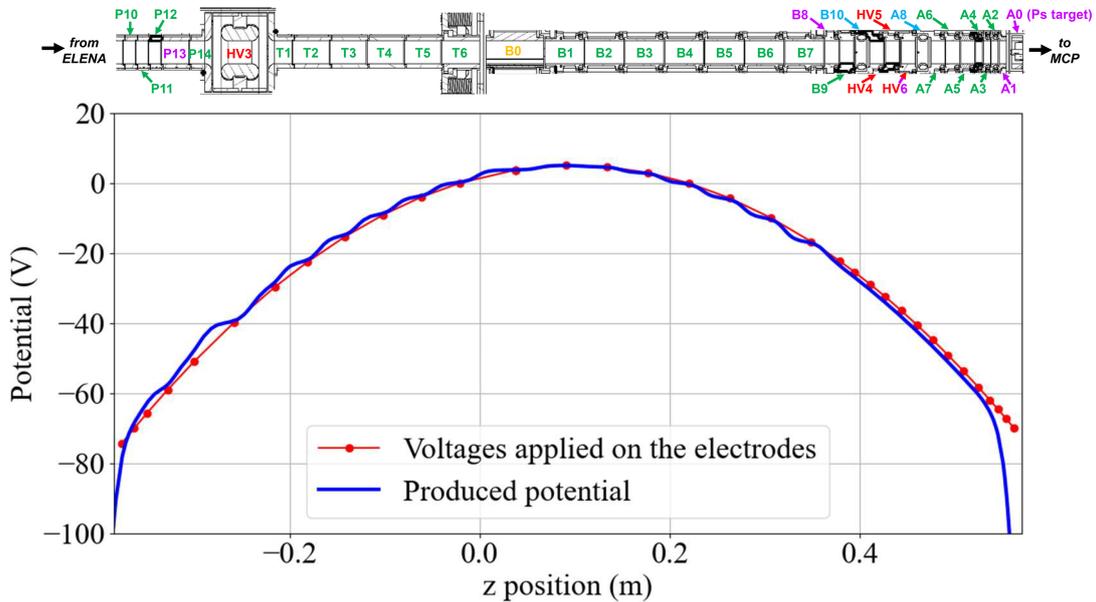


Figure 7.2: Representation of the parabolic potential formed by electrodes P10 to A1. The red markers indicate the center positions of the electrodes and the voltages applied on them. The blue line shows the resulting potential along the horizontal axis of the experiment, as calculated from a finite element COMSOL simulation. Above the plot, a schematic of the involved range of trap electrodes is included, scaled to match the dimension of the x-axis, with the colors of the electrode labels having the same meaning as in Fig. 3.9.

Before the launch, the particles are settled in the launch trap configuration described above. The parabola is then shaped on all involved electrodes except P13, as shown on the left hand side of Fig. 7.3. Only the most upstream part of the parabola is shown here as a representation. Since P10 - P12 are already at the launch potential and the voltages they are biased with for the parabola are within a few volts of this, the potential change in the area where the particles are held is very slight. At a precise time, it is subsequently possible to apply a positive pulse on P13 (from the Pulser, see Sect. 4.1.3, with a rise time of a few nanoseconds) that brings it from the endcap potential to the potential it needs to form part of the parabola, as depicted on the right hand side of Fig. 7.3. The

antiprotons are launched in the direction of the parabola. Usually, P13 is left open for a duration of the order of 10 to 15  $\mu\text{s}$  before the positive pulse is removed.

The antiprotons, once launched, thus make their way from the 5 T side to the 1 T, from a strongly negative potential (highest potential energy) over the slightly positive minimum (highest kinetic energy) towards the electrodes close to the Ps target, which are again more strongly negatively biased. The target electrode itself, A0, can also be biased. According to how the potential is tuned, the antiprotons "swing" in it, moving back and forth between P10 and A1/A0 or eventually passing over one of the barriers. On a broad scale, the evolution of the particle energy resembles a harmonic oscillator if the end potentials are similar.

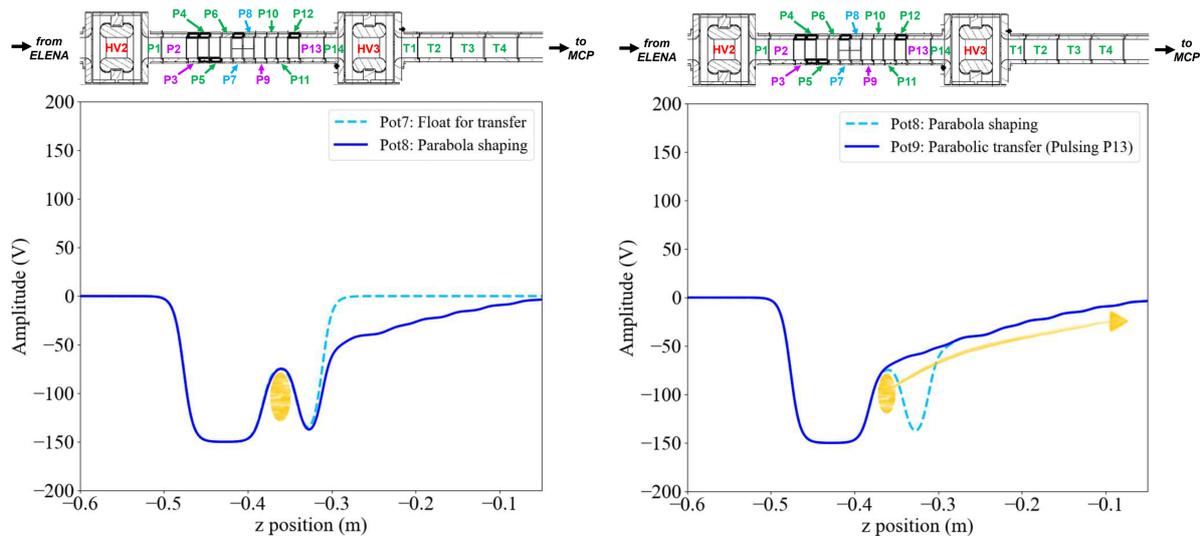


Figure 7.3: Low-voltage potential reshaping operations to form the parabolic potential and launch the particles towards the downstream side from the P Trap. Above each plot, a schematic of the involved range of trap electrodes is included, scaled to match the dimension of the x-axis, with the colors of the electrode labels having the same meaning as in Fig. 3.9. Left: Shaping of the parabola on all involved electrodes except P13. Right: Particle launch by applying a positive pulse on P13 (indicated by the yellow arrow).

### 7.1.3 Electron removal and recycling

The launch trap shown in Fig. 7.3 is still not the configuration actually used in the ballistic antiproton transfer procedure. Instead, a more flexible configuration has been designed, which allows to remove and/or recycle both particle types of the mixed electron-antiproton plasma residing in the trap. Essentially, two negative potential trap regions are formed instead of one: one between P9 and P13, i.e. the launch trap already introduced, and a second one between P2/P3 and P9, which is referred to as the recycling trap. All four of the mentioned endcap electrodes are connected to a channel of the Pulser. The potential configuration called "Ready" in Fig. 7.4 shows this trap, with the parabola already shaped and  $-180\text{ V}$  applied on the endcap electrodes (P2, P9, P13) instead of  $-150\text{ V}$ .

To send a pure antiproton plasma on its path towards the Ps target, the cooling electrons are removed just prior to the launch. As shown on the left hand side of Fig. 7.4, this can be done by applying a positive pulse of very short duration (less than 100 ns) on P9 and with a longer duration on P2 to open the barriers and let the electrons escape, while the majority of the heavier, slower antiprotons remains trapped. Detailed studies of this

procedure can be found in [108] and [112]; it has been adapted and implemented in the new AEGIS system. Once the barriers are closed again, the antiprotons are launched from the trap in the opposite direction towards the Ps target.

It is also possible to let the electrons escape from the launch trap but keep them accumulated in the recycling trap such that they can be reused to cool the next bunch of arriving antiprotons without loading new ones, or by topping them up from the electron gun and C Trap (see Sect. 3.4.8) to increase their density. As outlined on the right hand side of Fig. 7.4, this electron recycling is done in the same way as the electron removal with the difference being that P2 remains closed and blocks the stream of electrons to confine them together with the quickly closing P9.

Per default, in the antihydrogen production runs of 2023, electrons were planned to be removed from the trap and not recycled. However, due to a bug in the procedure, P3 was used as endcap of the recycling trap instead of P2 such that the opening of P2 did not let the electrons escape, leaving them in the recycling trap.

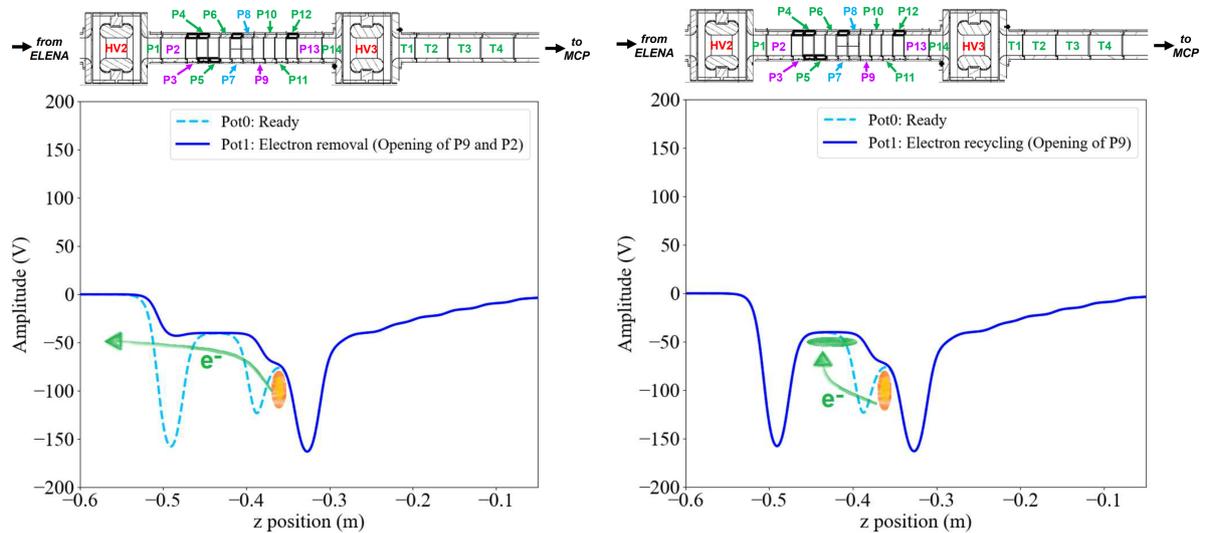


Figure 7.4: Low-voltage potential reshaping operations to remove the cooling electrons (green) from the launch trap, where the mixed plasma (yellow) of antiprotons (orange) and electrons is confined. Above each plot, a schematic of the involved range of trap electrodes is included, scaled to match the dimension of the x-axis, with the colors of the electrode labels having the same meaning as in Fig. 3.9. Left: Short pulses on P9 and P2 to let the electrons escape. Right: Short pulse on P9 to trap the electrons between P2 and P9 for recycling.

### 7.1.4 Antiproton launch and recycling

After the electron removal, the launch of the antiprotons from the full trap configuration is done as shown in the representation on the top left of Fig. 7.5: P13 is positively pulsed for 10 to 30  $\mu\text{s}$  to open to the potential needed for the parabola formation from P10. The antiprotons make their way towards the Ps target and either form antihydrogen, pass the last electrodes to annihilate on the downstream MCP, or are reflected back by the potential on A0. Generally, for antihydrogen production, a negative voltage is applied on A0, which is capable of reflecting back the majority of the antiprotons that do not form  $\bar{\text{H}}$ . This voltage is chosen to be of the order of -100 to -150 V. The reasoning for this decision is explained in Sect. 7.2.2.

It is possible to leave P13 open for long enough to recapture those antiprotons that have swung back from A0 along the parabolic potential. These antiprotons can then be launched again for the next production cycle but do not remain swinging between P13 and A0 to cause spurious losses and resulting background counts (which are otherwise observed, as pointed out in Sect. 7.2.2 and Sect. 7.3). Ideally, the recaptured antiprotons should be cooled while waiting for the next launch. For this reason, they should end up in the recycling trap, where the removed and/or new electrons have formed the next cooling plasma. This is why, in addition to P13, also P9 receives a positive pulse to let the antiprotons enter. The effect of this pulse is shown on the top right of Fig. 7.5. It happens  $5\ \mu\text{s}$  after the opening of P13 and is  $10\ \mu\text{s}$  shorter, such that P9 closes before P13 and the antiprotons are caught between P2 and P9 together with the electrons, as shown in the two figures in the bottom row of Fig. 7.5. Most of the electrons themselves do not escape in either direction through this pulse because still, the potential on P9 and the following electrodes is kept more negative than the floor of the recycling trap.

Fig. 7.6 shows the summed overall intensity observed on the downstream MCP/CMOS detector after dumping the re-captured particles, normalized to the maximum value, as a function of the time for which P13 is left open to allow the antiprotons to re-enter the trap. This intensity is expected to be proportional to the number of particles arriving on the MCP. Plotted is the data of two scans over this recapture time, once for a dump of the trap content towards the MCP itself (blue markers) and once for a dump towards the degrader structure at the entrance of the experiment (red markers). The scans are performed over recapture times between 10 and  $32\ \mu\text{s}$ , in steps of  $1\ \mu\text{s}$ , and a potential of  $-150\ \text{V}$  is applied on the target electrode to prevent the antiprotons from passing it. A maximum in the MCP intensity and, thus, the number of re-caught antiprotons, is observed for a recapture time between 20 and  $25\ \mu\text{s}$  for both scans. As discussed in Sect. 7.3.3, this time is compatible with the expected travel time of the antiprotons back and forth in the parabolic potential. It is therefore conclusive that for shorter opening times, fewer of them make it into the trap before closing, while for longer times, some of them exit the trap again while P13 is still open. This rough scan and analysis validate the general functionality of the antiproton recapture technique; a further fine-tuning of the procedure is ongoing.

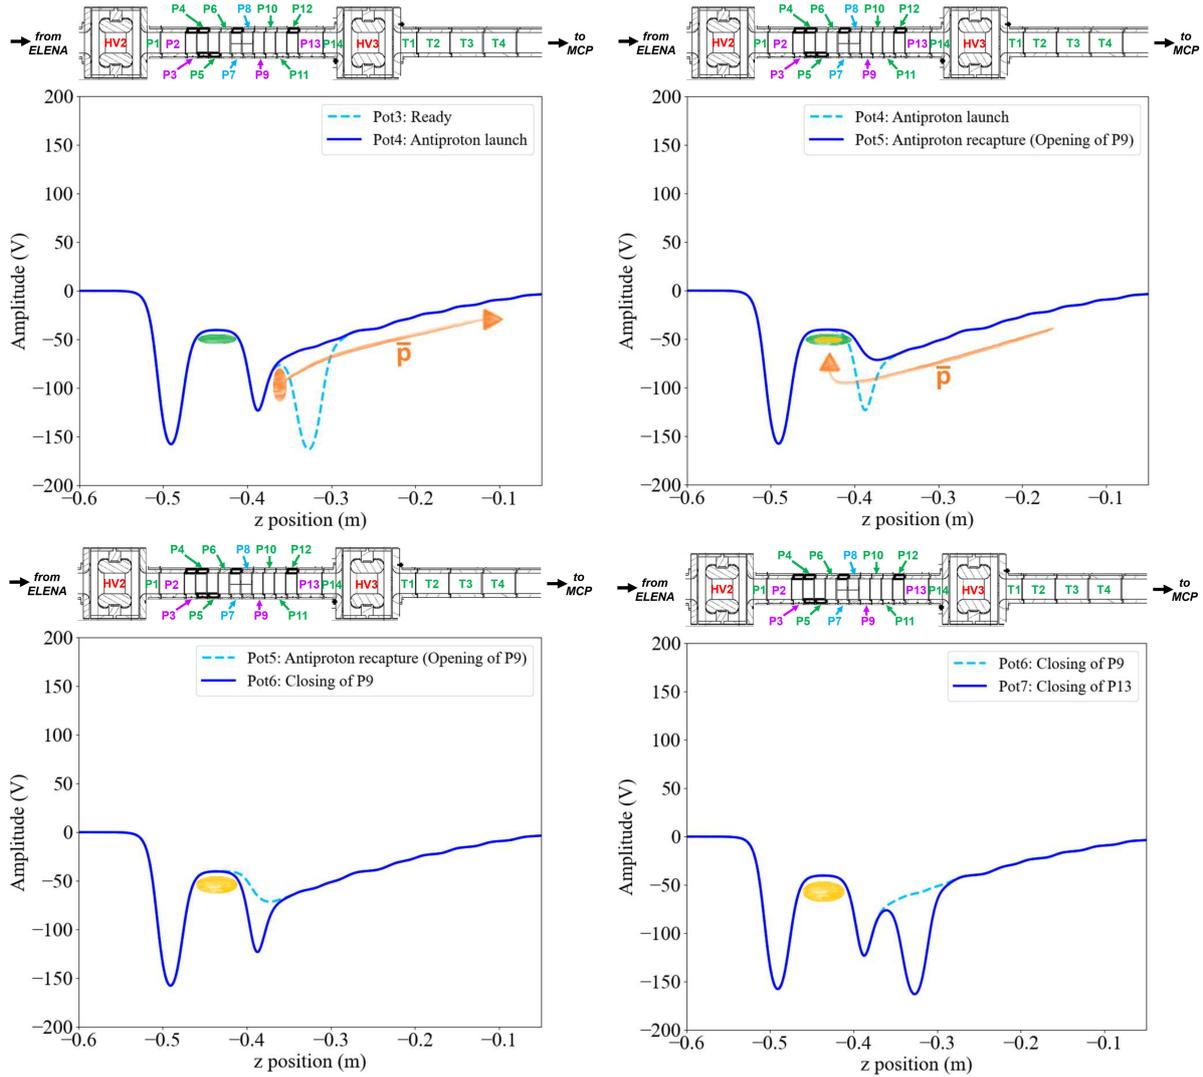


Figure 7.5: Low-voltage potential reshaping operations to launch and recapture antiprotons from/in the P Trap for antihydrogen production. Above each plot, a schematic of the involved range of trap electrodes is included, scaled to match the dimension of the x-axis, with the colors of the electrode labels having the same meaning as in Fig. 3.9. Top left:  $\bar{p}$  launch along the parabolic potential by an opening (positive pulse) of P13. Top right: Re-entering of  $\bar{p}$  into the recycling trap by opening P9. Bottom left: Closing of P9 to trap  $\bar{p}$  in the recycling trap. Bottom right: Subsequent closing of P13.

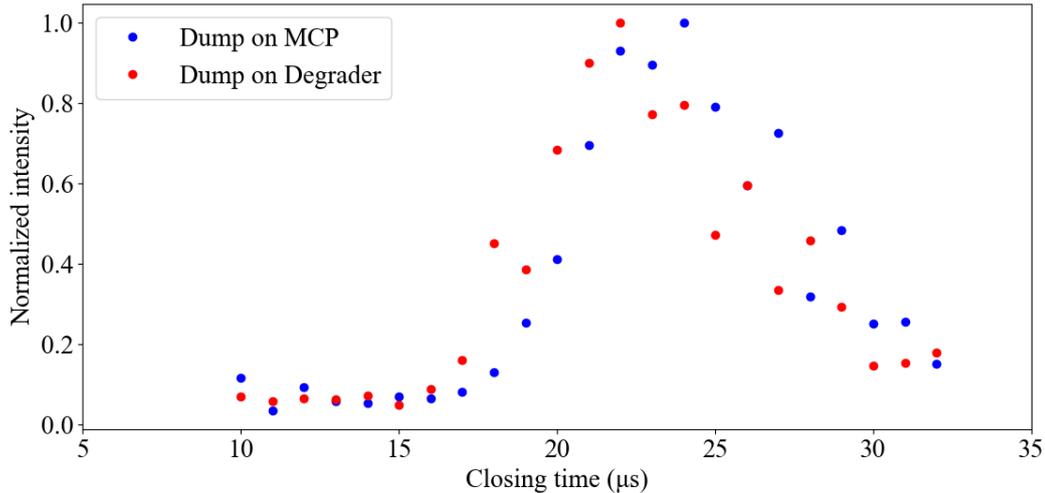


Figure 7.6: Normalized intensity observed on the MCP/CMOS detector when dumping antiprotons from the trap, subsequent to their recapture, as a function of the opening time of the barrier electrode. Two recapture time scans are plotted: one with a dump on the MCP (blue) and one with a dump on the degrader (blue). The intensity values are normalized to the maximum found in each scan.

In the antihydrogen production runs performed in 2023, the antiprotons were per default not recycled or accumulated, i.e. only one bunch of captured ELENA antiprotons was launched from the trap towards the  $\bar{\text{H}}$  formation in the Ps target area. This decision was taken as a precaution due to the already very limited available time for antihydrogen production and the novelty of the recycling and accumulation procedures. In 2024, to increase the  $\bar{\text{H}}$  yield, these developments are planned to be included.

To return from the recycling and launch trap configuration to the potential used for electron cooling and antiproton launch, and to include the recycled particles for the next round of antiproton preparations for  $\bar{\text{H}}$  production, the low-voltage reshaping operations shown in Appx. J are used.

## 7.2 Antihydrogen production

The production of antihydrogen in AEGIS relies on the unification of the preparation procedures of the two main components: the antiprotons, whose preparation and launch are discussed in the previous section and in Ch. 6, and positronium. The positronium procedures are outlined in Sect. 3.3.2. In CIRCUS, it has been possible to develop and test the two components entirely independently; in the end, the relevant time-critical parts of the antiproton routines have simply been inserted into the Ps part of the AERIALIST at the appropriate places and a combined script (preparation of detectors and trap configurations, closing procedure, etc.) has been built as a frame, allowing to be ready for  $\bar{\text{H}}$  production within one day without interfering with previous work on positronium and antiproton research.

It is vital to synchronize the two components in time and in particular tune the arrival of the positrons, with the subsequent Ps formation and (almost) immediate excitation, to that of the antiprotons.

The  $^{22}\text{Na}$  source used to provide positrons for Ps production has reached the end of its life in the experiment: Its activity is reduced to 200 MBq, strongly limiting the Ps yield. A new source, which is being installed in 2024, will have an activity that is close to an order of magnitude larger. Furthermore, the Ps conversion target used in 2023 had a lower than nominal yield: A positron to positronium conversion efficiency of only 1.5 to 2.5% could be reached; usually, 5 to 10% are routinely achieved. This further limits the possible amount of antihydrogen to be produced, since its number depends linearly on the number of excited positronium atoms, like it does on the number of available antiprotons, as given in Eq. 3.14. Possible reasons for the low yield include the clogging of the target pores through a long exposure to water/oxygen in air or to other pollutants, which were extensively present in the experiment during the work towards finding and fixing a leak in the vacuum apparatus. The Ps target is also being replaced in 2024.

### 7.2.1 Electric potentials at the antiproton launch

Fig. 7.7 shows the potentials formed by the low-voltage electrodes over the entire range from the P Trap to beyond the Ps target during the launch of the antiprotons (blue line, left vertical axis). P13 is open, the recycling trap is shown, and the parabola is shaped. A negative voltage of the order of  $-150\text{ V}$  is applied on A0, which corresponds to the kink before the two strong peaks to negative and positive potentials downstream of the target. These peaks are caused by the voltages applied on the movable grids installed as part of the 1 T trap (see Sect. 3.4.8). The grids are used to ionize atoms formed in excited states, such as the produced Rydberg antihydrogen, and their potentials are optimized to ionize the atoms with the expected principal quantum numbers with the highest efficiency, according to [178].

For reference, the magnetic flux density along the central axis of the experiment, as discussed in Sect. 3.4.1, is plotted in Fig. 7.7 as well (red line, right vertical axis). It can be observed that the transition between the 5 T and 1 T magnetic field regions happens at the locations of the T electrodes and electrode B0. The antiprotons released from the launch trap traverse the entire range of the parabolic potential, crossing this transition region.

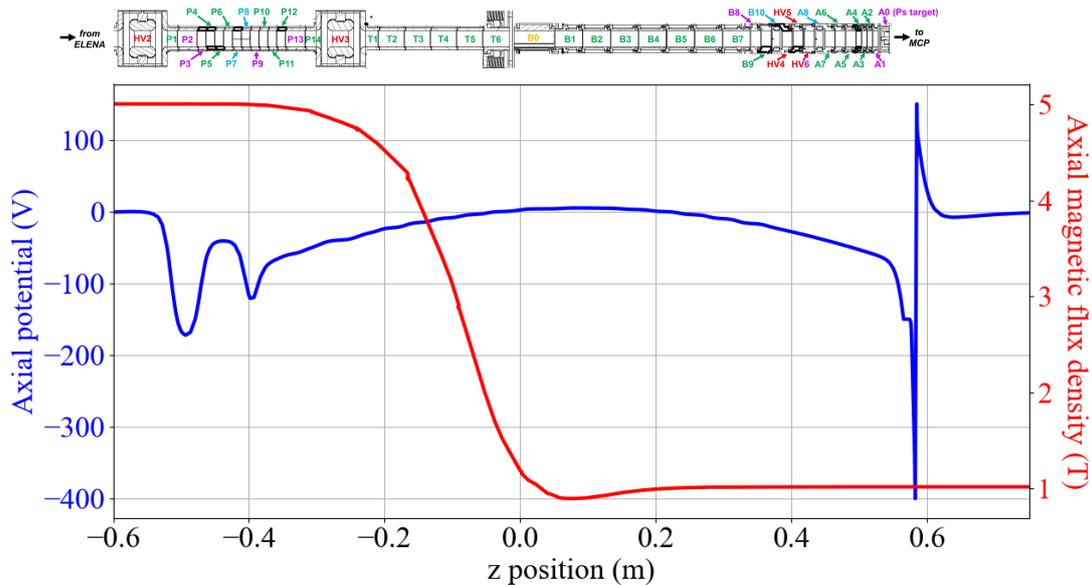


Figure 7.7: Low-voltage potentials along the experiment axis over the entire relevant range during the antiproton ballistic launch for antihydrogen production (blue line, left vertical axis). Above the plot, a schematic of the involved range of trap electrodes is included, scaled to match the dimension of the x-axis, with the colors of the electrode labels having the same meaning as in Fig. 3.9. The second vertical axis on the right hand side refers to the magnetic flux density along the central z-axis of the experiment, which is plotted by the red line for reference.

## 7.2.2 Timing considerations and the target electrode potential

It is possible to tune the antiproton launch instant to match the positronium formation timing: If the antiprotons arrive on the Ps target exactly at the same time as the positrons, antihydrogen is essentially produced at rest, as the Ps and  $\bar{p}$  combine at the target location at the point of maximum potential energy. If they arrive earlier than the positrons (with a shorter delay), they are reflected by the potential applied on the target electrode and pass through the Ps cloud while travelling back towards the 5 T trap. Instead, when using a longer  $\bar{p}$  delay, the antiprotons traverse the Ps on their way towards the target and antihydrogen is formed with a horizontal momentum towards the downstream end of the experiment. This configuration is the one that will eventually be used to form the  $\bar{H}$  beam.

During the first beam formation trials at the end of the antiproton campaign in 2023, the method employed to determine the necessary delay for the antiproton launch was the observation of the signal times on the ESDA units, i.e. the scintillator/PMT detector introduced in Sect. 3.4.9. Contrarily to the antiproton analysis presented in Ch. 6, here the digitized ESDA acquisition chain is used (see Sect. 3.4.9). For the purpose of understanding the origin of the different signal components observed on the PMTs, the combined  $\bar{H}$  production procedure is run in different ways. Thanks to the introduction of parameters in the AERIALIST run script, whose values can be passed individually for every run from the TALOS interface, it is possible to inhibit either the antiproton parts or the positron parts of the procedure, including a delay corresponding to the length of the respective procedures instead, but running the overall routine in exactly the same way, thereby avoiding the influence of systematic effects. The digitized PMT signals themselves

and their analysis is discussed in more detail in Sect. 7.3.

Furthermore, it is possible to bias the target electrode (A0) with any achievable value, thus either reflecting the arriving antiprotons back from the target region (for higher negative voltages than the corresponding particle energy) or letting them annihilate on the target and MCP (for lower voltages, letting the particles pass). By performing a scan over the voltage applied on A0, it is possible to "slice" the antiproton plasma at different potentials, determining its energy spread, which can be related to its original space charge. Since the voltage applied on A0 has a negligible influence on the more energetic positrons, in addition to runs with positrons or antiprotons alone, such a scan can yield insight into the nature of the different signal components obtained from the full procedure involving all particle types. Fig. 7.8 shows the spectrum measured by PMT 24, as an example, for a scan over the voltage applied on A0,  $V_{A0}$ , from  $-60$  to  $-130$  V in a set of runs with antiprotons only. The top plot shows the time range between 60 and 150  $\mu\text{s}$  from the start of the acquisition. Prior to this time, no interesting features are observed in the spectrum and the amplitude remains around the baseline; later on, the signal follows a similar behaviour to the shown range, with the peak structure smearing out further. The bottom plot shows a zoom to the time around the first distinctive structures, between 62 and 80  $\mu\text{s}$ . Two different kinds of signal components are found: The first one (in particular visible as the orange signal components for higher voltages, with the first occurrence at around 64  $\mu\text{s}$ ) corresponds to an accumulation of individual peaks, is independent of  $V_{A0}$  and happens in relatively regular intervals with a period of the order of 10  $\mu\text{s}$ , smearing out with time. The second one, instead, does depend heavily on  $V_{A0}$  and has a periodicity closer to 20  $\mu\text{s}$ . Its structure is marked by broader peak-like signals with a fast onset and a slower "decay". For its independence of  $V_{A0}$  and its periodicity, the first kind of signal is attributed to antiproton losses due to annihilations in the central region of the experiment, at the transition between the 5 T and 1 T magnetic fields, while swinging back and forth between the two. These kinds of signals are discussed in more detail in Sect. 7.3.3, as is the timing of the antiproton trajectory. The second kind of signal is expected to stem from those antiproton annihilations that occur when the swinging bunch arrives at the target and a fraction of it annihilates on it. Its estimated periodicity is therefore that of a full swinging period instead of half; no significant annihilation events are expected upon the antiproton arrival at the upstream end of the parabolic potential because, even if they do "spill over" P9, the larger potential on P2 (see Fig. 7.7) prevents a passage in the further upstream direction, where an annihilation for example on the degrader structure would be expected. Additionally, the expected intensity of this kind of signal decreases with higher negative voltages applied on A0, as more of the antiprotons are reflected and do not annihilate. These expectations are well validated by the observed second periodical structure with the first prominent peak at around 67.5  $\mu\text{s}$ . While many of the antiprotons pass lower potentials of the order of  $-60$  V and the annihilation peak is therefore very strong, its intensity decreases as expected for the higher potentials until all antiprotons seem to be reflected from  $-105$  V. This value is in agreement with the expected particle energies generated by the calculated actual voltages on the outermost electrodes involved in the parabola. The same is true for the following peaks, although it is less clearly visible in the way the data is plotted due to the smearing. Furthermore, the same behaviour is also observed on the other PMTs; PMT 24 is shown here as an example because the peak structures are very clearly visible thanks to its vicinity to the Ps target area. In Appx. K, the same data is depicted, this time in individual plots for each run with the different plots stacked underneath each other. The described evolution of the signals according to the applied  $V_{A0}$  is also clearly visible there. One thing that can

be mentioned here is the fact that by changing  $V_{A0}$ , the maximum axial arrival position of the plasma is also slightly adapted, influencing the number of impacts on the targets and, thus, the number of annihilations observed on the different detectors.

In order to reflect all antiprotons and prevent losses on the Ps target structure, in general antihydrogen production runs, the voltage on A0 is kept between -110 and -150 V. No significant difference is observed in the PMT spectra in this voltage range. For the future gravity measurement, a large enough reflection voltage on A0 also has the added benefit of "filtering out" those antiprotons that have not formed antihydrogen and letting only the electrically neutral  $\bar{H}$  atoms pass to the gravity module, allowing for cleaner data.

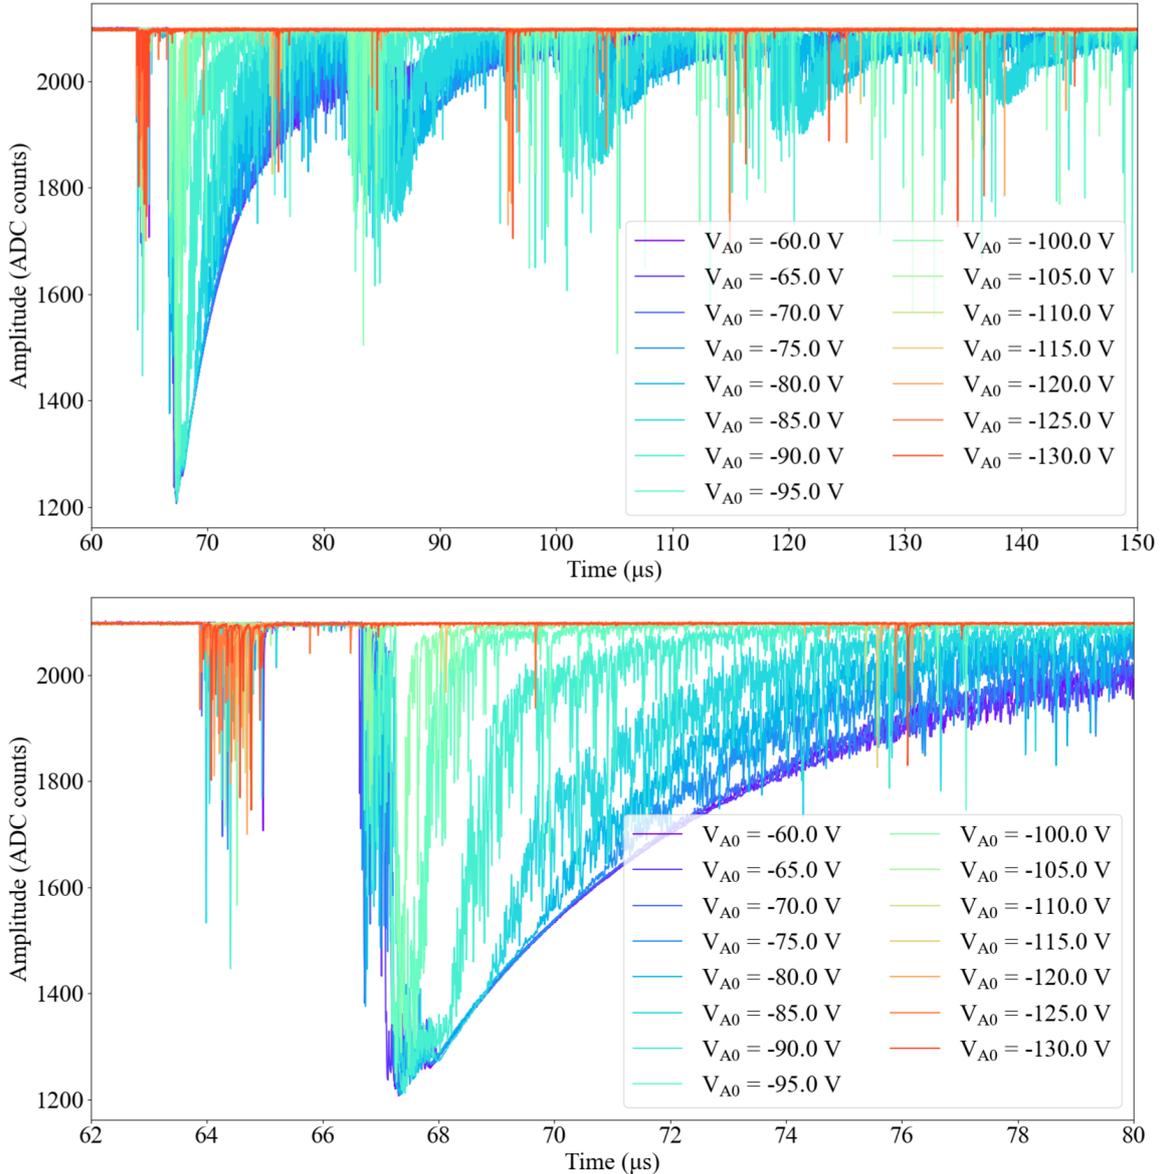


Figure 7.8: Signal distribution observed on the digitized PMT 24 as a function of time for a scan over the target electrode voltage in  $\bar{H}$  production runs without positrons. The time is given as the time from the start of the acquisition. Top: Range between 60 and 150  $\mu\text{s}$ . Bottom: Zoom to the region between 62 and 80  $\mu\text{s}$ .

Using the knowledge gained from the antiproton spectra, it is possible to tune the timing of the antiprotons such that their arrival on the Ps target (i.e. their annihilation

peak attributed to the losses on the target structure) coincides with that of the positron implantation in the target (or, for  $\bar{H}$  production with acceleration in the forward/backward direction, happens later/earlier respectively). For this purpose, knowledge of the positron timing is required. From extensive Ps experiments, it is already known that the positrons arrive in the target area reliably at around  $70.7\ \mu\text{s}$  on the PMT time axis, within a few tens of nanoseconds. Additionally, dedicated  $\bar{H}$  production runs without antiprotons, i.e. with positrons only, have also been carried out to verify the observation of this behaviour on the PMTs. The digitized signal of PMT 11 for an example of such a run is plotted in Fig. 7.9, red line. None of the structures observed in the  $\bar{p}$  runs are visible; instead, a sharp peak is found at  $70.7\ \mu\text{s}$  after the start of the acquisition. This peak corresponds to the annihilation of the majority of the positrons on the target and is discussed again in more detail in Sect. 7.3.2. For comparison, an individual antiproton run is also shown in Fig. 7.9 (blue line), for an A0 voltage of  $-90\ \text{V}$ . The transition losses (distributed low-intensity peaks) as well as the target annihilation peak (high-intensity peak with fast onset and slower tail) are clearly discernible, matching those in Fig. 7.8 in time, while no signal is observed around the positron peak time. The third spectrum included in Fig. 7.9 (orange line) is obtained from a run with both positrons and antiprotons included in the  $\bar{H}$  procedure in the standard way, with  $V_{A0}$  set to  $-90\ \text{V}$  as well. Since the maximum of the antiproton peak on the target had been determined to occur at around  $67.37\ \mu\text{s}$ , a delay of  $3.33\ \mu\text{s}$  is introduced here in the software for the launch of the antiprotons, with the goal of synchronizing the antiprotons to the positron arrival on the target. This procedure shows the desired effect: The signals attributed to the antiprotons occur with the applied delay, meaning that the antiproton annihilations on the target structure coincide with the positron peak (and the corresponding peaks are therefore not distinguishable from each other, as they overlay). Without the delay, both peaks are individually visible in the spectrum. PMT 11 is chosen here as an example thanks to its distance from the target area, rendering the individual peaks less broad and more easily separable.

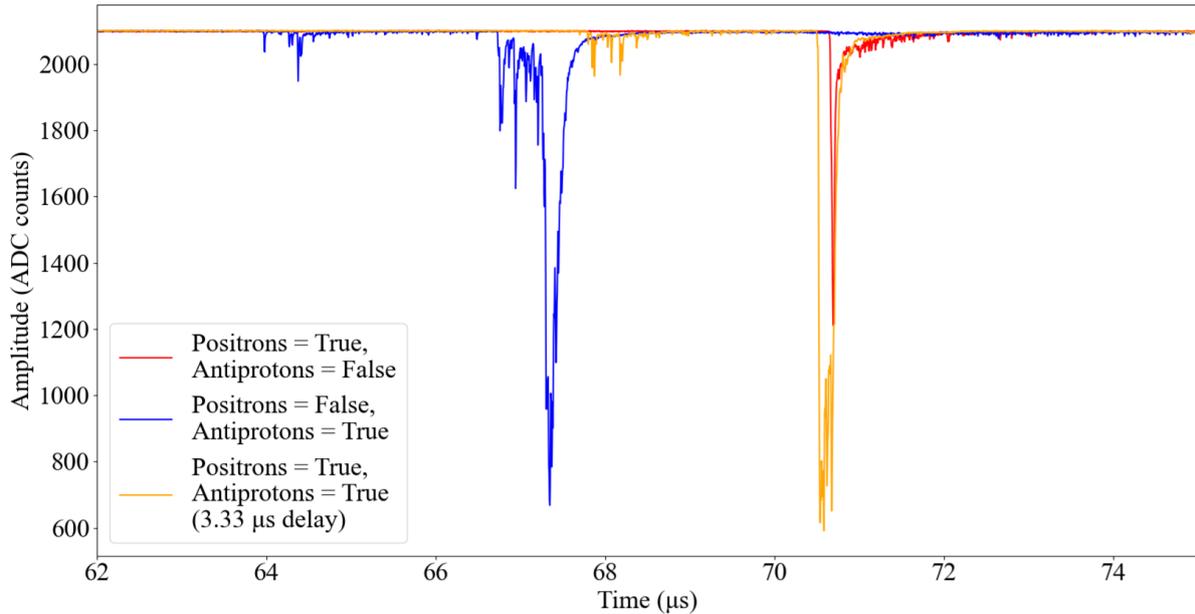


Figure 7.9: Signal distribution observed on the digitized PMT 11 as a function of time for  $\bar{\text{H}}$  production runs with different settings. The time is given as the time from the start of the acquisition, in the relevant range between 62 and 75  $\mu\text{s}$ . Included are two runs without any delays introduced in the software, with positrons only (red) and with antiprotons only (blue), while the other component is disabled respectively, and one run with both positrons and antiprotons (orange), where a delay of 3.33  $\mu\text{s}$  is applied to the antiproton launch.

### 7.3 ESDA analysis

For the analysis of the antihydrogen signal, two detector systems are employed: the downstream MCP and the ESDA. The ESDA analysis is presented here; details on the MCP analysis can be found in [156]. Since neither of the detector systems can distinguish between annihilation events from free antiprotons and from those within a formed antihydrogen atom, the analysis relies on a comparison between runs with and without antihydrogen production. As described in Sect. 3.4.6,  $\bar{\text{H}}$  formation can be inhibited without a modification of the overall procedure or setup by the use of an IR filter that prevents the required Ps excitation. In this way, data with and without  $\bar{\text{H}}$  formation is obtained alternately.

Like for the determination of the necessary antiproton launch delay, also here the digitized ESDA acquisition chain (see Sect. 3.4.9) is employed to analyze the signals. This means that the half of the split PMT signal that is not used for the discriminated counts is tenfold attenuated and then digitized. For this reason, the final amplitude values obtained in the data are multiplied by a factor of 20 for them to correspond to the full energy that is deposited (factor two because of the splitter and factor ten due to the attenuation). The CAEN V1720 ADC used for the digitization has a 12-bit resolution and an input range of 2 V, allowing to convert the measured ADC value to a voltage in steps of the order of 0.49 mV.

An example of a raw distribution of the amplitudes (in ADC counts and rebased to the lowest value) versus time of PMT 15 is shown in Fig. 7.10, obtained for an antihydrogen formation run. The plot at the top shows the entire acquisition range, while the bottom

plot is a zoom to the region around the positron peak, which is identified according to the findings in Sect. 7.2.2. The goal of the analysis is the identification of peaks in the spectra that correspond to antiproton annihilation event candidates. The first step is therefore the tagging of all peaks in the distribution. Subsequently, those peaks that are compatible with  $\bar{p}$  annihilations are used to perform the comparison analysis. This analysis is performed for all ESDA units located in the 1 T region of the experiment (see Sect. 3.4.9), and the detected events on all are cumulated for the final distribution.

### 7.3.1 Peak tagging

For a clean peak analysis, a baseline level is defined for each PMT in each run. In order to be useful for the peak discrimination, this baseline has to follow the trend of the spectrum itself: For example due to electronic noise or pile-up, there are parts in the spectrum where the ADC level is elevated, independent of individual peaks. To still be able to identify peaks above, the baseline is made to follow the evolution of the spectrum itself, meaning that for every recorded amplitude value  $A_i$ , the baseline  $A_{\text{base},i}$  it is compared to is defined as given in Eq. 7.6, depending on the previous baseline  $A_{\text{base},i-1}$  with 99.5 % and on the amplitude itself with 0.5 %. The starting baseline is determined as the mean value of the amplitudes in the first 10  $\mu\text{s}$  of the spectrum. The resulting adaptive baseline is plotted in red in Fig. 7.10.

$$A_{\text{base},i} = (1.0 - 0.005) \cdot A_{\text{base},i-1} + 0.005 \cdot A_i \quad (7.6)$$

The baseline is used to take a decision on whether or not a recorded amplitude is identified as a peak, with the standard deviation of the first 10  $\mu\text{s}$  of the amplitudes,  $A_{\text{rms}}$ , serving as determination factor: For each amplitude, two thresholds are determined according to Eq. 7.7 and Eq. 7.8. To be defined as the onset of a new peak, an amplitude has to exceed the upper threshold  $T_{\text{up},i}$ , i.e. exceed the baseline by at least eight standard deviations, to avoid double-counting. In this case, the corresponding time and amplitude are appended to a list of the peaks on the given PMT for that run. If the following amplitudes are again higher, the peak attributes are modified accordingly. Instead, once a following amplitude is reduced to beyond the lower threshold  $T_{\text{low},i}$ , the peak is defined to be complete and the search for the next peak begins from the following recorded amplitude.  $T_{\text{low},i}$  is required with a difference of only four standard deviations due to hysteresis effects, i.e. the electronics level lagging slightly behind the actually produced amplitudes for high values. Because of an electronic ringing in the temporal vicinity of the positron peak, which is observed on several PMTs,  $T_{\text{up},i}$  is required to exceed the baseline by twelve standard deviations instead of eight for amplitude recording times below 74.5  $\mu\text{s}$  (indicated by the (\*) in Eq. 7.7). Both thresholds are shown in Fig. 7.10 ( $T_{\text{up},i}$  in yellow,  $T_{\text{low},i}$  in green); their adaptive character as well as the jump in the lower threshold after the electronic ringing can be observed. The amplitudes that are found to be peak maxima are marked by cyan dotted vertical lines.

$$T_{\text{up},i} = A_{\text{base},i} - 8^{(*)} \cdot A_{\text{rms}} \quad (7.7)$$

$$T_{\text{low},i} = A_{\text{base},i} - 4 \cdot A_{\text{rms}} \quad (7.8)$$

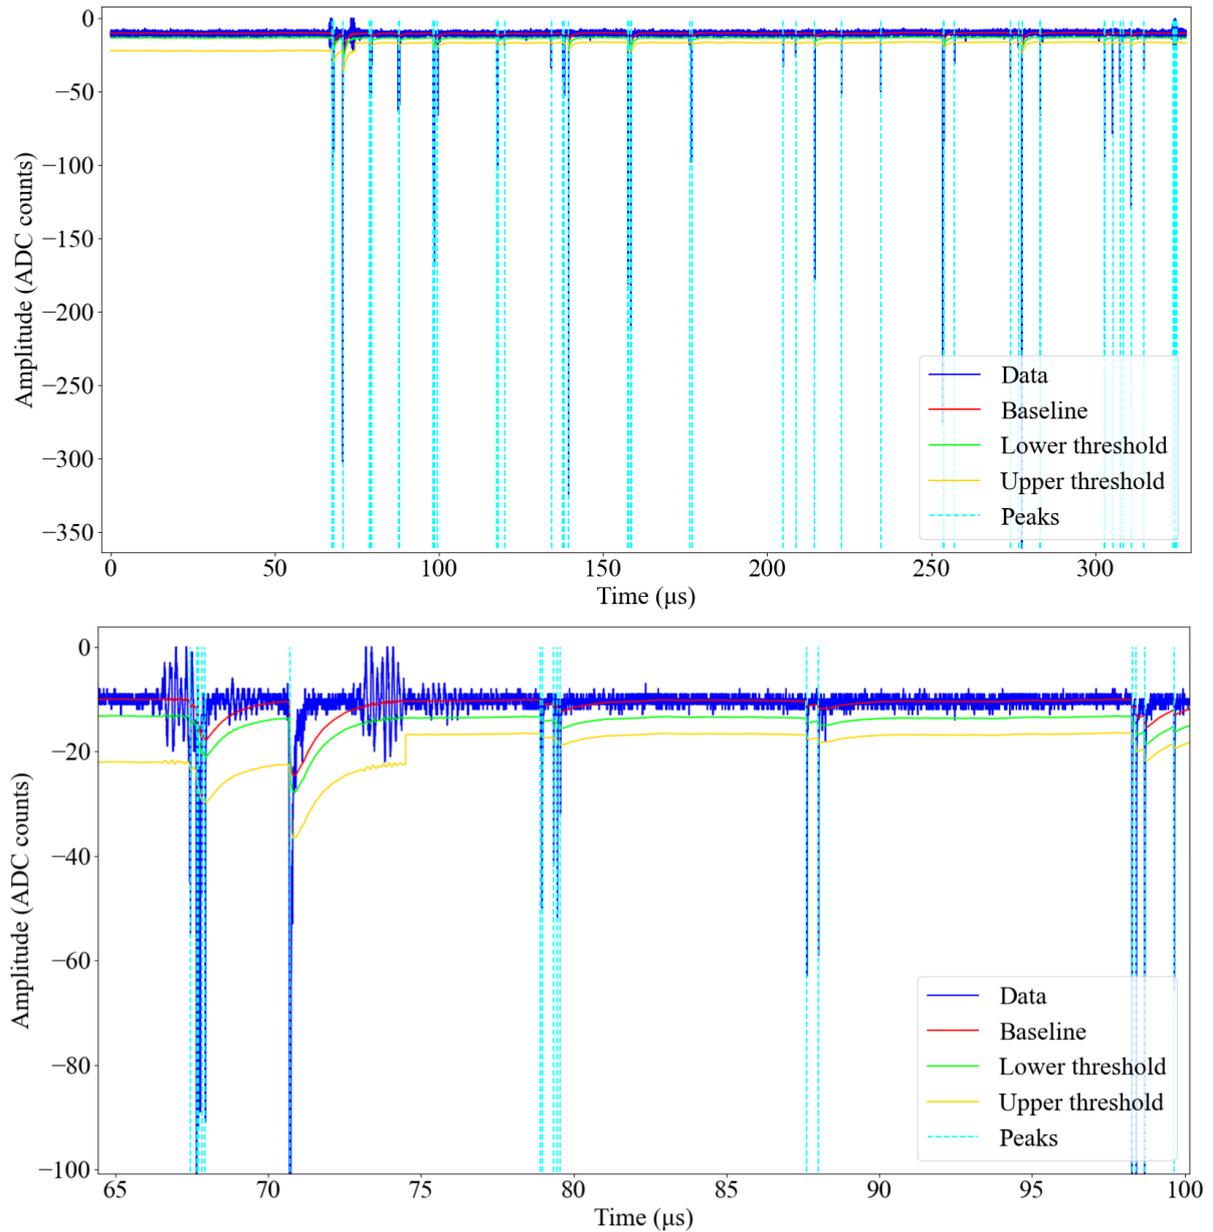


Figure 7.10: Distribution of digitized amplitudes versus time of an example ESDA PMT for an antihydrogen production run. The spectrum is rebased to the lowest obtained ADC value. The measured amplitudes are plotted in dark blue. Also shown are the baseline (red) as well as the upper (yellow) and lower (green) thresholds used for the peak search. The identified peaks are tagged by dotted cyan lines. Top: Full range. Bottom: Zoom to the region around the positron peak, for visibility.

It can be noted that, due to the baseline being adapted to the previous measured signal intensities, a slight reduction in the event detection efficiency in the first couple of microseconds of the spectrum following the positron peak might be expected. This is because, as will become clear in the following Sect. 7.3.2, an observed peak is required to exceed a certain amplitude above the baseline to pass the event selection, as antiproton annihilation events usually cause larger amplitudes than the expected background. The details of the cut are also discussed there. For the selected cut amplitude (350 mV), the required excess above the baseline would be of the order of 36 ADC counts. Since, as

shown in Fig. 7.10, the baseline increase from the positron peak can be of the order of 10 ADC counts, i.e. almost 100 mV in voltage amplitude, some of the relevant peaks may not be included in the selection. A comparison of the amplitudes of the selected and discarded events, which is also presented in Sect. 7.3.2, indicates an increase in the number of selected events of the order of 20% in this range. However, not all of these will have likely occurred in the relevant time of a couple of microseconds after the positron peak, and the influence of the baseline decreases steeply until around 3  $\mu$ s after the positron peak. In addition, some of the hypothetically added are expected to originate from the gamma ray signals, which need to be excluded. The actual reduction in the  $\bar{p}$  annihilation detection efficiency is therefore likely smaller. Furthermore, it can be argued that the importance of performing a clean analysis, excluding background events as strictly as possible, outweighs an improvement in the statistics. An optimization of the involved parameters and an adaptation of the baseline algorithm may improve the overall efficiency in the future by yielding a better compromise.

### 7.3.2 Event discrimination and selection

As explained above, the antihydrogen formation measurements are performed with and without the transmission of the IR laser into the experiment, i.e. with and without expected antihydrogen formation. Here, the label "IR off" refers to those measurements in which the IR laser is filtered out and not used to excite Ps for  $\bar{H}$  production. This run modality is also referred to as the control run modality. Furthermore, as also discussed previously, attempts have been made to produce antihydrogen in three different modalities: with an acceleration in the downstream direction towards the MCP for a longer delay between positron insertion and antiproton launch, with an acceleration back in the upstream direction when antiprotons arrive earlier and are reflected back, and with an arrival time of the antiprotons tuned to that of the positrons to not accelerate the  $\bar{H}$  in either direction. For the latter modality, a delay of 3.33  $\mu$ s is chosen for the antiprotons, and the other two delays are reduced or increased by 500 ns, to 2.83  $\mu$ s and 3.83  $\mu$ s respectively. These delay times provide the labels for the corresponding measurements. The "IR off" measurements are performed with a delay of 3.33  $\mu$ s. To prevent an influence of systematic effects, the run modalities have been alternated in the following cycle during measurements: "IR off,  $\bar{p}$  delay = 3.33  $\mu$ s" – "IR on,  $\bar{p}$  delay = 2.83  $\mu$ s" – "IR on,  $\bar{p}$  delay = 3.33  $\mu$ s" – "IR on,  $\bar{p}$  delay = 3.83  $\mu$ s" – "IR off,  $\bar{p}$  delay = 3.33  $\mu$ s" – ... . In all runs, the potential of the target electrode is kept at around  $-150$  V.

To exclude dark count events and afterpulses registered on the PMTs, a coincidence is required for the two PMTs connected to the same scintillator: An identified peak has to occur on the second PMT within a time window of 50 ns around the time of the identified peak of the first PMT. This time window has been optimized in the past for the used scintillator slabs [68, 179]. If such a peak is found, the mean values of the two peak times and amplitudes are formed and stored as information of the combined event. Otherwise, the peak is disregarded.

As described in [68] and [135], the antiproton events can be distinguished from the background of late positron annihilations: The photons resulting from the  $e^+$  annihilations either enter the PMTs directly to produce a signal, in which case requiring a coincidence on the two PMTs of the same scintillator slab can mitigate the contribution, or produce free, excited electrons (through Compton effect and photoelectric effect) in the surrounding material. These electrons, when undergoing de-excitation, in turn induce a scintillation signal in the ESDA, whose amplitude, however, does not exceed 250 mV according to

simulation results [135], meaning that a threshold amplitude cut can be employed to minimize the resulting background.

During the first few hundred nanoseconds after the positron injection for Ps formation, however, the ESDA is blinded by an over-saturation caused by of the order of 90% of the positrons annihilating in the converter target: While individual contributions are relatively small, the overall energy deposited in the very short time directly after injection exceeds the amplitudes of the antiproton annihilation signals. This is the strong, sharp peak visible in Fig. 7.10 at around 70.7  $\mu\text{s}$ , in agreement with the findings from the timing studies in Sect. 7.2.2. Since the positron injection can be precisely tuned in time, the onset of this "positron peak" occurs reliably within a few tens of nanoseconds around the expected time and can therefore be well tagged in the analysis. For simplification, the time range up to 500 ns after Ps formation is generally excluded from the  $\bar{H}$  signal search region.

The expected time of antihydrogen formation is within a range of the order of a few microseconds from the time of the Ps excitation laser, given by the lifetime of the Ps and the geometry [68]. The distribution of the peaks found in the range from 0.5 to 20.5  $\mu\text{s}$  for all four run modalities described above, accumulated over all ESDA units and all runs (approximately 300 runs per modality), is shown in Fig. 7.11 with and without the discussed event discrimination. Those peaks attributed to the Ps formation are removed. In the upper left plot, neither is the PMT coincidence required nor is a cut on the peak amplitude applied. Clear antiproton annihilation peaks are observed between 7.5 and 10  $\mu\text{s}$  and around 17.5  $\mu\text{s}$  after Ps formation. The periodicity of these peaks and the difference between the peaks found for the different formation formalities are discussed in Sect. 7.3.3 and Sect. 7.3.4. However, it is already worth noting that the onsets, maxima and tails observed in the antiproton annihilation peak accumulations for the different run modes behave as expected: While "IR off,  $\bar{p}$  delay = 3.33  $\mu\text{s}$ " and "IR on,  $\bar{p}$  delay = 3.33  $\mu\text{s}$ " have the same timing and their distributions in time are compatible, the ones for the modified delays occur approximately 500 ns in advance or later, respectively. Additionally, a few peaks very shortly after Ps formation remain, and a prominent structure of events within a few microseconds around 3  $\mu\text{s}$  after Ps formation can be discerned, whose timing is independent of the formation mode. These peaks can be attributed to PMT dark counts and afterpulses, which is evidenced by their disappearance when requiring the PMT coincidence described above, as shown in the top right plot of Fig. 7.11.

For comparison, the bottom left plot shows the distribution without the coincidence requirement but with an amplitude cut of 350 mV: All peaks with a lower amplitude are disregarded. This cut, as well, removes the dark count/afterpulse peaks efficiently and at the same time has a stronger influence on the remaining peaks at very short times after Ps formation. Fig. 7.12 shows the number of remaining candidate events in a signal region defined between 0.5 and 5.5  $\mu\text{s}$  after Ps formation for different threshold amplitudes when applying the cut. The time before 0.5  $\mu\text{s}$  is excluded from the signal region due to the ESDA blinding from the positron implantation in the target; the time after 5.5  $\mu\text{s}$  is not used because the contributions from the transition losses of the antiprotons "swinging" back from the target to the 5 T region become a significant background that cannot be excluded. Coincidence filtering is also included here. The pink line indicates the used threshold at 350 mV. It is apparent that from this value to higher thresholds, the three formation modalities which include the use of the IR laser have a higher event count than the one without Ps excitation. 350 mV is chosen as a more conservative threshold value than the expected energy deposition values from gamma rays determined in [135] due to the higher errors expected in the ADC quantization following the installation of the

attenuators in the PMT digitization chain.

Finally, the plot on the bottom right of Fig. 7.11 shows the remaining event distribution with the application of both the coincidence requirement and the amplitude threshold cut. Both the afterpulse and the positron annihilation contributions are removed but the antiproton annihilations are still clearly observed.

It is worth explaining one feature of Fig. 7.11 that is due to the way the analysis is performed: Except for those peaks that are explicitly removed by the event filtering, the absolute number of counts found in the different structures of those plots on the right hand side (with the PMT coincidence filtering active) is approximately half of that found on the left (without coincidence filtering) - both in the upper row (without the amplitude cut) and in the bottom row (with the amplitude threshold applied). This is intuitively expected and further confirms the coincidence filtering approach: When including the data of all involved PMTs, the events occurring in a given scintillator are essentially counted twice - once by each of the connected PMTs. As outlined previously, the coincidence requirement combines the double events into single events with the average amplitude. In this way, the apparent reduction by a factor two in the statistics of the peak structures attributed to antiproton annihilations is simply this merging of events. While the targeted structures that likely stem from PMT dark counts and afterpulses is effectively removed, there is, thus, no significant reduction in the actual statistics from the coincidence filtering.

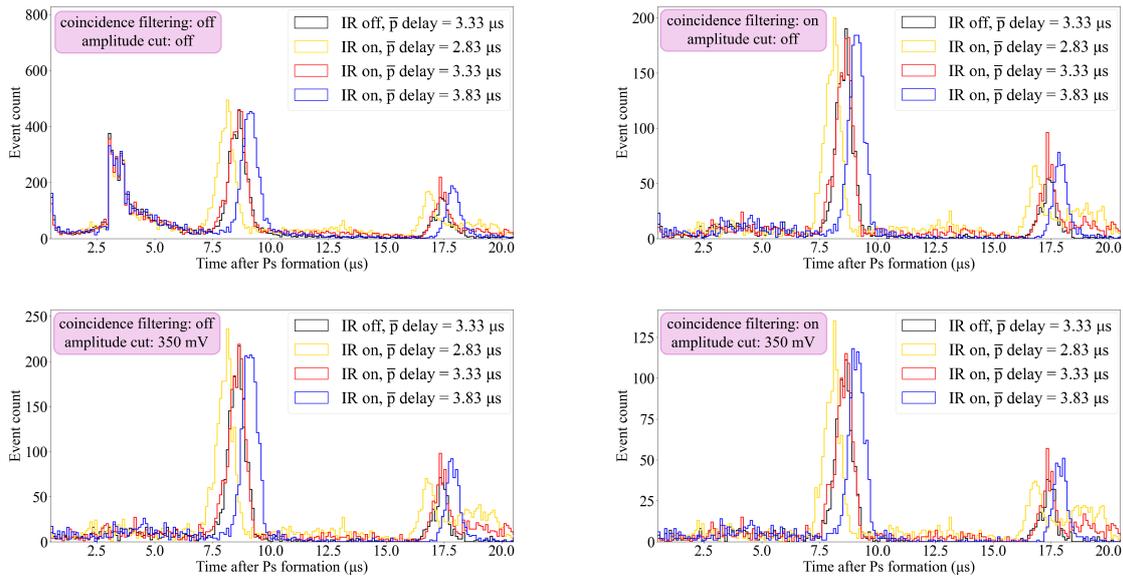


Figure 7.11: Distributions of the number of identified peak events in the ESDA as a function of time after the expected Ps formation in the range between 0.5 and 20.5  $\mu\text{s}$ . Shown are the events in all PMTs in different run modalities: With the use of the IR Ps excitation laser for  $\bar{\text{H}}$  formation for three different delay times between positron insertion and antiproton launch (yellow, red, blue), and without the IR laser for the intermediate delay time (black). Top left: Without PMT coincidence filtering and peak amplitude cut. Top right: With PMT coincidence filtering but without peak amplitude cut. Bottom left: Without PMT coincidence filtering but with peak amplitude cut (350 mV). Bottom right: With PMT coincidence filtering and peak amplitude cut (350 mV).

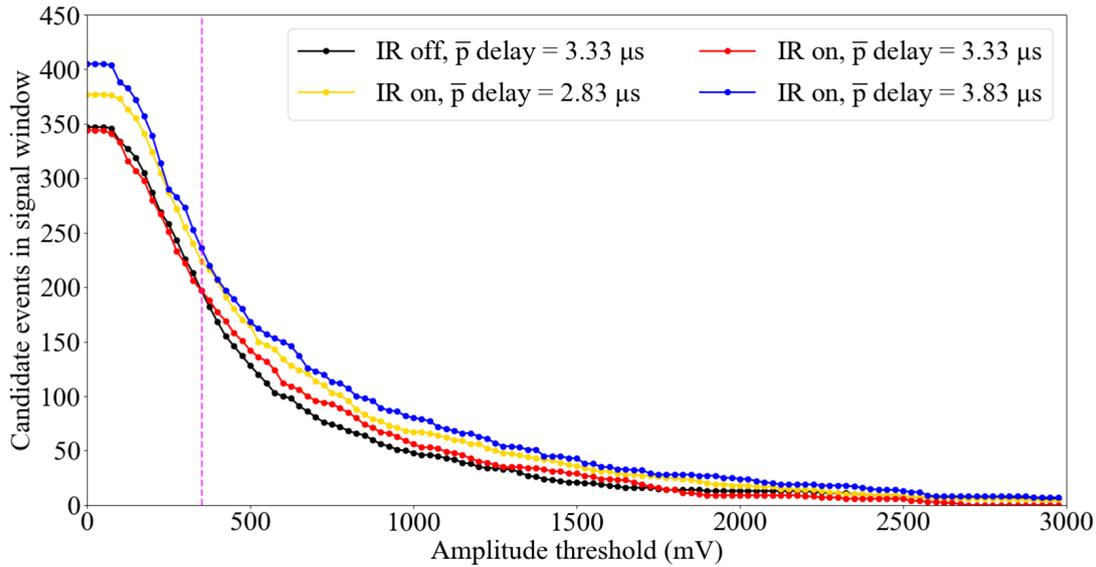


Figure 7.12: Number of remaining event peaks in the signal window between 0.5 and 5.5  $\mu\text{s}$  after the expected Ps formation for different amplitude threshold cuts, with PMT coincidence filtering applied. The results for the four different run modalities are plotted for a range of cut thresholds (circular markers); the connecting line is included as an eye guide. The vertical pink line indicates the 350 mV cut applied as default.

Prior to the detailed analysis, some runs are filtered out and not included in the datasets, for example if the DAQ has not saved the data correctly, if no positron peak is observed on the PMTs or if the laser signal (monitored by a separate oscilloscope) does not correspond to the expectations. For many issues, such as empty  $\bar{p}$  shots from ELENA or dysfunctional detector components, CIRCUS takes care of discarding non-usable runs and repeating their acquisition instead.

Overall, 1201 successful antihydrogen campaign runs have been performed in 2023 that are taken into account here (306 with the IR laser off, 296 with a  $\bar{p}$  delay of 2.83  $\mu\text{s}$ , 301 with a  $\bar{p}$  delay of 3.33  $\mu\text{s}$  and 298 with a  $\bar{p}$  delay of 3.83  $\mu\text{s}$ ), having yielded in total slightly below 200,000 events that have passed the described selection.

Fig. 7.13 shows the number of selected events as a function of the run number, with the used formation mode for each run indicated by the color of the marker. A clear jump in the event count is observed from run number 404640 onwards, with most runs previously having yielded less than 100 events, while several hundred counted events are routinely produced afterwards. This jump can be attributed to a restearing of the ELENA parameters exactly prior to this run, leading to a much higher flux of antiprotons entering the experiment and then swinging in the parabolic potential.

The  $\bar{H}$  runs shown here happened over a span of eleven days in November 2023, mostly during the nights and weekends thanks to the CIRCUS automation. Those runs in between that are not included in Fig. 7.13 were not used for antihydrogen production but instead focused on other developments, such as the accumulation of large numbers of antiprotons in the trap (see Ch. 6) and studies of the positron procedures to try to increase the Ps yield.

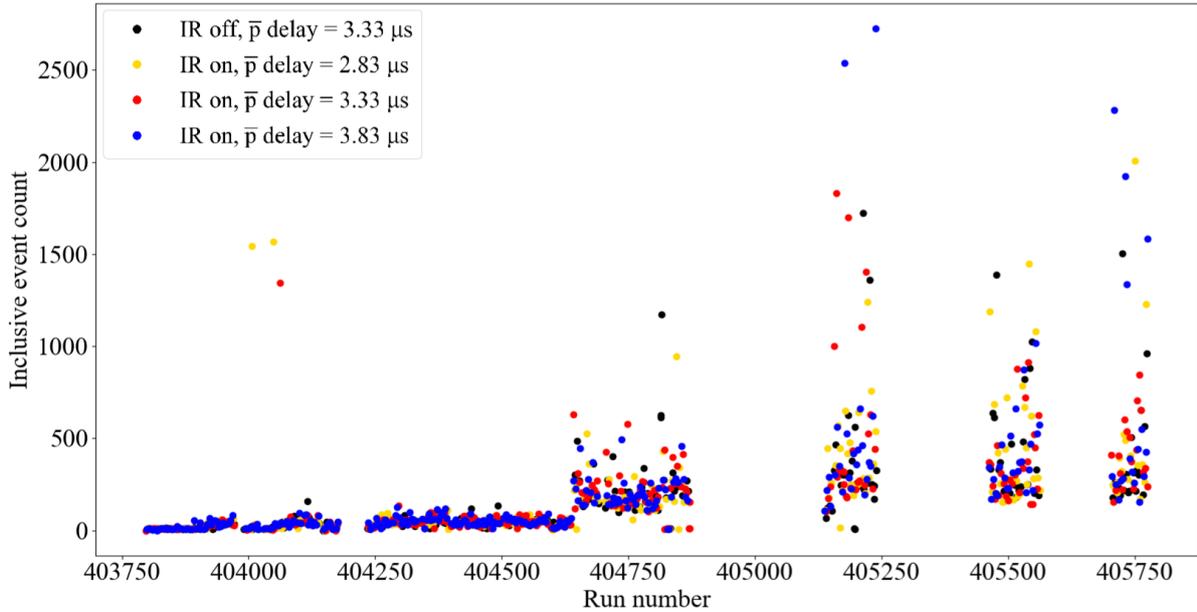


Figure 7.13: Distribution of the number of selected events per run as a function of the run number during the antihydrogen production campaign in 2023 for the four different run modalities.

### 7.3.3 Antiproton oscillations

In ideal conditions, the evolution of the potential felt by the antiprotons on their path towards the target follows that of a harmonic oscillator. The difference of the potential  $\Delta V(\Delta x)$  between two locations along the experimental axis  $\Delta x$  is therefore theoretically given by Eq. 7.9, where  $m$  is the antiproton mass and  $\omega$  is the angular frequency, which is related to the travel time  $\Delta t$  of a particle between two locations as given in Eq. 7.11. Using the knowledge of the voltages applied on the electrodes and the symmetry of the parabolic potential that is used, the potential difference  $\Delta V(\Delta x) = V_{\text{start}} - V_{\text{mid}} = V_{\text{end}} - V_{\text{mid}}$  for the antiprotons between the start (or end) of their journey, at  $x_{\text{start}}$ , i.e. the center location of electrode P11, (or  $x_{\text{end}}$ , i.e. the center location of electrode A1), and the mid-point of the trajectory  $x_{\text{mid}}$  can be determined. Plugging in the values programmed on the respective electrodes,  $V_{\text{end}} = -70 \text{ V}$  and  $V_{\text{mid}} = 5.038 \text{ V}$ , and their known locations,  $\Delta x = x_{\text{end}} - x_{\text{start}}/2 \approx 0.925 \text{ m}$ , allows to solve the harmonic oscillator equation for  $\omega$ , as shown in Eq. 7.10, and thus determine the expected time of flight of the antiprotons from Eq. 7.11<sup>1</sup>.

$$\Delta V(\Delta x) = \frac{1}{2} m \omega^2 (\Delta x)^2 \quad (7.9)$$

$$\omega = \frac{1}{\Delta x} \sqrt{\frac{2\Delta V}{m}} \quad (7.10)$$

$$\Delta t = \frac{\pi}{\omega} \quad (7.11)$$

<sup>1</sup>The potentials  $V_{\text{end}}$  and  $V_{\text{mid}}$  felt by the antiprotons due to the voltages applied on the electrodes are stated here in units of electronvolt to facilitate an intuitive understanding. However, to calculate the expected time of flight  $\Delta t$  in SI units, the resulting potential difference is converted to Joule by dividing by the elementary charge.

The expected time of flight of the antiprotons in the parabolic potential determined in this way is approximately  $12\ \mu\text{s}$ . However, as clearly shown in Fig. 7.2, while the potential actually produced at the mid-point of the parabola is in good agreement with the voltage applied on the electrodes, that at the end points deviates significantly and is rather of the order of  $-100\ \text{V}$  than the programmed  $-70\ \text{V}$ . Plugging this potential into the harmonic oscillator calculations instead yields a significantly shorter flight time of approximately  $10\ \mu\text{s}$ .

While these calculations are merely performed to obtain an estimated time of flight to be expected and several diverging effects (e.g. damping due to air resistance, imperfections in the potential) are not taken into account, the comparison of the measured distributions in time of the antiproton annihilation events to these expectations is noteworthy.

"Swinging" back and forth in the parabolic potential between the launch trap and the target region, the antiprotons are expected to primarily annihilate in the central region of the experiment: Right at the transition between the 5 T and 1 T regions, where the magnetic field strength decreases from both sides (see Sect. 3.4.1), the plasma is largest due to magnetic bottle effects, and radial losses are more likely. This behaviour also implies that the effect is expected to be stronger for those passages of the antiprotons from the 5 T region to the 1 T side than vice versa.

Fig. 7.14 depicts the distribution of observed event-selected ESDA peaks versus their time after expected Ps formation, for the four different run modes. The top plot shows the range between  $-10$  and  $150\ \mu\text{s}$  in logarithmic scale, including the sharp positron peak at around  $0\ \mu\text{s}$  (independent of the formation mode). Also included here is a first strong accumulation of peaks before the positrons arrive. In the bottom plot, only the peaks after the positron peak are taken into account and a linear scale is used. In both diagrams, a periodic pattern of peak accumulations is clearly visible: These are the annihilation events caused by the antiprotons during their "swinging" in the parabolic potential. The first of these happen already before the positron peak, i.e. while the antiprotons are on their way from the launch trap to the Ps target for the first time. From here on, a stronger and a less strong peak alternate in the distribution, with approximately  $20\ \mu\text{s}$  periodicity, quite well compatible with the antiproton annihilations expected between the magnetic field regions (stronger for the downstream direction from 5 T to 1 T and less strong for the upstream direction) for the calculated travel time of  $10\ \mu\text{s}$  in one direction. Due to the fewer antiprotons that are subsequently left and the resulting smaller plasma expansion, it also makes sense that the annihilations become fewer with the number of swings. Furthermore, the annihilation events within one swing differ in time further and further due to the continued divergence of the  $\bar{p}$  bunch, broadening the peak structures.

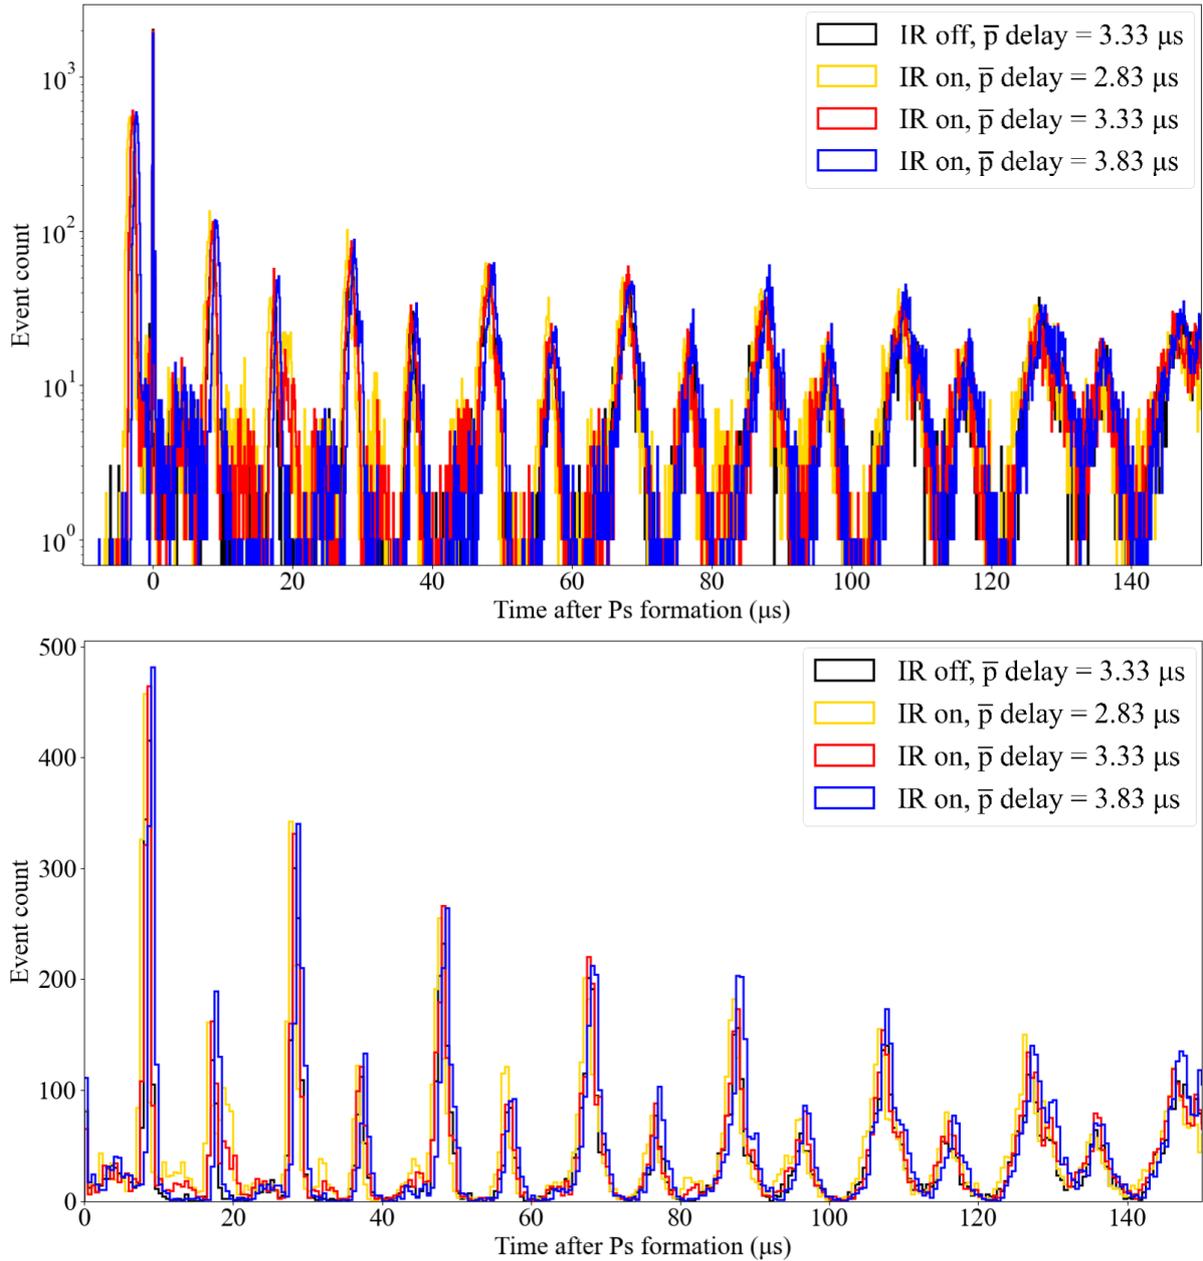


Figure 7.14: Distributions of the number of identified peak events in the ESDA as a function of time until 150  $\mu\text{s}$  after the expected Ps formation. Shown are the events in all PMTs in different run modalities: With the use of the IR Ps excitation laser for  $\bar{H}$  formation for three different delay times between positron insertion and antiproton launch (yellow, red, blue), and without the IR laser for the intermediate delay time (black). Top: Logarithmic vertical axis; the horizontal axis starts at  $-10 \mu\text{s}$ . Events attributed to the positron peak are included. Bottom: Linear vertical axis; the horizontal axis starts at 0  $\mu\text{s}$ . Events attributed to the positron peak are excluded.

Fig. 7.15 depicts a zoom to the region between  $-10$  and  $20.5 \mu\text{s}$  of the top plot of Fig. 7.14, i.e. showing the first three annihilation event peaks of the swinging antiproton bunches as well as the positron peak between the first two. As expected, these peaks are shifted slightly depending on the timing of the used  $\bar{H}$  formation mode (at rest, forward- or backward-boosted), while the timing of the positron peak remains constant. Here, a feature that is already adumbrated in Fig. 7.8 and Fig. 7.14 becomes clearly

visible: The annihilation events originating from the passage of antiprotons from one magnetic field region to the other do not appear as equidistant peaks in the spectrum. In fact, slightly longer and slightly shorter distances alternate, with a difference of a few microseconds (yielding absolute times above and below  $10\ \mu\text{s}$ , respectively). These findings are compatible with the antiprotons travelling for a shorter time in the 5 T field region than in the 1 T part during their trajectory along the parabolic potential. A comparison between the plot of the magnetic field map in the experiment, Fig. 3.6, which shows the transition region at an axial position of  $z \approx 0\text{ m}$ , and Fig. 7.2 or Fig. 7.7 confirms this observation: The antiproton trajectory after launch is significantly shorter for the first part in the 5 T field before reaching the field transition area, where the annihilations are expected to occur, and continuing in the 1 T field past the parabola vertex towards the Ps target. For comparison, the midpoints between the maximum bins (chosen for simplicity, not accuracy) of the first two bunch annihilation peaks are marked in the distribution in Fig. 7.15 by the dashed lines in the colors corresponding to the formation modes with IR laser excitation. It is again interesting to note that the (albeit roughly estimated) observed time differences between these midpoints confirm the consistency of the time-wise accurate steering of the antiproton launch and corresponding travel, since the antiproton arrival in the magnetic field transition area differs by approximately 500 ns for the different modalities, as would be expected from the applied delays. Therefore, the here developed procedure seems to be a suitable, stable technique for the time steering of the antihydrogen production from this point of view.

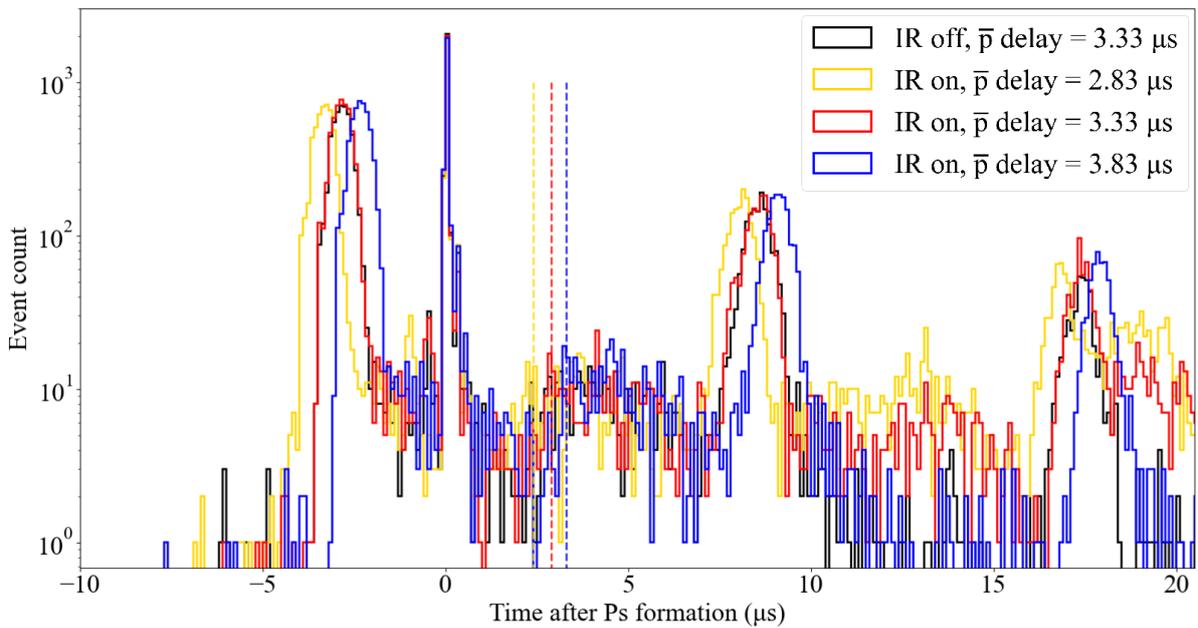


Figure 7.15: Distributions of the number of identified peak events in the ESDA as a function of time between  $-10$  and  $20\ \mu\text{s}$  after the expected Ps formation. Shown are the events in all PMTs in different run modalities: With the use of the IR Ps excitation laser for  $\bar{\text{H}}$  formation for three different delay times between positron insertion and antiproton launch (yellow, red, blue), and without the IR laser for the intermediate delay time (black). The dashed lines indicate the midpoints between the maximum bins of the first two bunch annihilation peaks in the three run modalities with the IR laser.

### 7.3.4 The antihydrogen signal

The candidate event peaks remaining in the signal window between  $0.5$  and  $5.5 \mu\text{s}$  from Ps formation after applying the coincidence filtering and the amplitude threshold cut, summarized in Fig. 7.12, can be plotted as a function of their time of occurrence. In Fig. 7.16, they are grouped by a moving window<sup>2</sup> with a length of  $1.5 \mu\text{s}$  starting from the time of each bin that is plotted to favour the separation of signal and background. The step from one bin to the next is  $1 \text{ ns}$ . The three distributions of the events obtained with the IR laser peak at different times after the Ps formation. The bins with the highest number of counted events are indicated by the dashed vertical lines for each. They are found at  $2.2 \mu\text{s}$  ( $2.83 \mu\text{s}$   $\bar{p}$  delay mode),  $2.7 \mu\text{s}$  ( $3.33 \mu\text{s}$   $\bar{p}$  delay mode) and  $3.2 \mu\text{s}$  ( $3.83 \mu\text{s}$   $\bar{p}$  delay mode), respectively. The differences of these values are again in good agreement with the expected  $500 \text{ ns}$  difference in the delays of the antiproton signals.

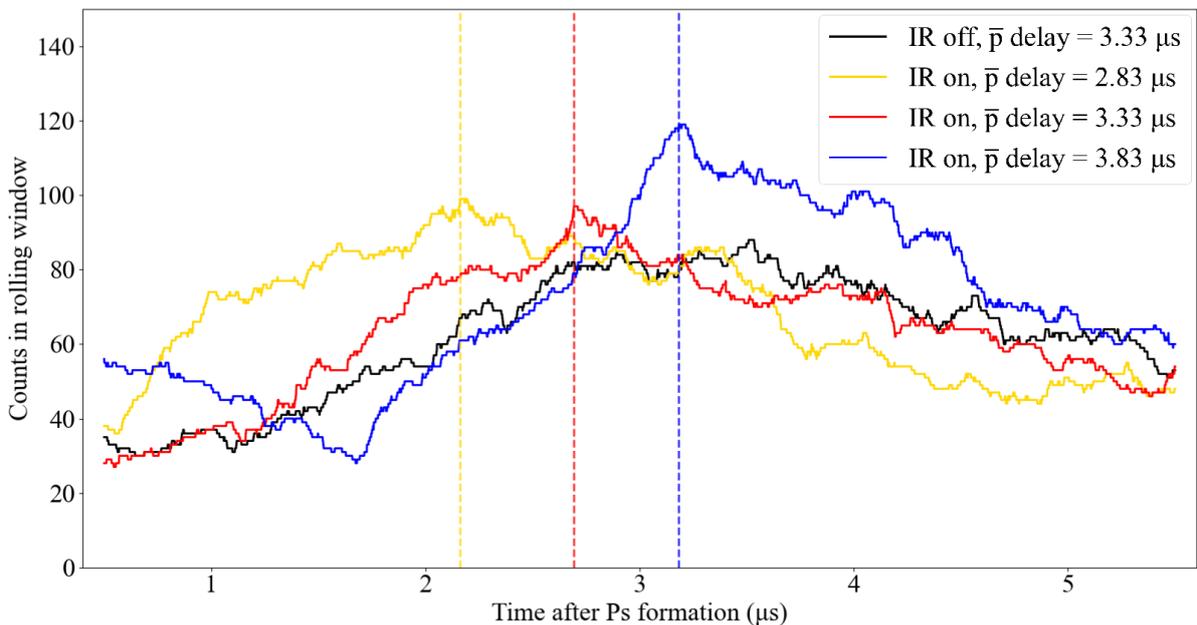


Figure 7.16: Antihydrogen candidate events in the signal window as a function of time after Ps formation. Plotted is the content of a rolling window with a  $1.5 \mu\text{s}$  length and a  $1 \text{ ns}$  step size, for the four different  $\bar{H}$  production modes. The signal window extends from  $0.5$  to  $5.5 \mu\text{s}$  after the expected Ps peak. For the three distributions obtained with the use of the IR laser, the dashed vertical lines indicate the respective bins with the highest event count.

A "null hypothesis" can be formulated that corresponds to the assumption that no excess is found for the number of events in the case of the use of the IR laser with respect to the control runs without the laser, i.e. no antihydrogen has been formed. The significance  $S$ , in standard deviations, of the deviation from this assumption can be

<sup>2</sup>For the moving window method, the counts found with time stamps included in a given window length are added in every step to give the final value for that step, with the time increase in every step chosen as a constant: in every step, the earliest time included in the given window, i.e. its starting time, is increased by the step time, such that counts having occurred earlier are dropped, as is its end time, such that counts with later time stamps are additionally included. The values obtained in every step are then plotted. Here, the time used as the horizontal axis value for a given data point is the respective starting time of the window.

determined according to Eq. 7.12 to disprove the null hypothesis [180].  $N_{\text{IR}}$  and  $N_{\text{off}}$  refer to the numbers of counts in a selected bin of Fig. 7.16,  $N_{\text{IR}}$  for the datasets obtained with the use of the IR laser in the three modalities and  $N_{\text{off}}$  for the case of the disabled IR laser with a  $\bar{p}$  delay of 3.33  $\mu\text{s}$ .  $\alpha$  corresponds to the ratio of the number of runs taken in the two cases (or, equivalently, the measurement acquisition time of both). As outlined in Sect. 7.3.2, since approximately the same number of runs have been performed in each modality and the timing of each run is the same,  $\alpha \approx 1$  can be assumed. As the measurements were performed by alternating all four run modalities, systematic effects, such as the ELENA intensity, are expected to average out and not influence the obtained numbers.

$$S = \sqrt{2 \left( N_{\text{IR}} \cdot \ln \left[ \frac{1 + \alpha}{\alpha} \cdot \frac{N_{\text{IR}}}{N_{\text{IR}} + N_{\text{off}}} \right] + N_{\text{off}} \cdot \ln \left[ (1 + \alpha) \cdot \frac{N_{\text{off}}}{N_{\text{IR}} + N_{\text{off}}} \right] \right)} \quad (7.12)$$

For the determination of  $S$ , the maxima of the three distributions of the datasets with the IR laser in Fig. 7.16, marked by the vertical dashed lines, are used. The corresponding bin content is compared to that of the "IR off" case in the same bin. Table 7.1 summarizes the selected event numbers and the results obtained for  $S$  in the different cases.

Overall, the statistical analysis is severely limited by the low numbers, and the analysis to be performed in 2024 with the expected higher numbers should yield more significant conclusions. The maximum significance observed here is of 2.8 standard deviations, found for the case of the largest antiproton delay. i.e. where the produced antihydrogen should be boosted in the forward direction. The combined statistics, i.e. the summation of the event numbers in the selected bins over all three datasets both with and without the use of the IR laser yield an overall significance of 3.7 standard deviations. This combined value motivates a statistical confidence in the data that matches an evidence of an excess signal corresponding to antihydrogen production.

Dividing the excess number of counts observed with the different run modalities by the number of performed runs yields a value of 0.05 to 0.13 produced antihydrogen atoms per run, depending on the modality, and 0.10  $\bar{\text{H}}$  atoms per run for the combined analysis. These numbers are discussed in the context of the  $\bar{\text{H}}$  formation enhancement expected from different improvements in the experiment in Sect. 7.4.

Production mode ( $\bar{p}$ delay)	$N_{\text{IR}}$	$N_{\text{off}}$	$S$
2.83 $\mu\text{s}$	99	67	2.5
3.33 $\mu\text{s}$	97	82	1.1
3.83 $\mu\text{s}$	119	80	2.8
Combined	315	229	3.7

Table 7.1: Summary of the selected event numbers of the antihydrogen production runs ( $N_{\text{IR}}$  with the IR laser on,  $N_{\text{off}}$  with the IR laser not used) in the different formation modes (i.e. different time delays of the antiproton launch) and the resulting significance  $S$  in number of standard deviations.

It should be mentioned that, for a more convincing comparison, control runs without the IR laser should have also been performed for the cases in which the antiproton launch is delayed by 2.83  $\mu\text{s}$  and 3.83  $\mu\text{s}$ . Due to a lack of available antiproton beam time, these runs have been omitted in 2023 and the data obtained with the use of the excitation laser in these modalities are compared to the control run data acquired with a 3.33  $\mu\text{s}$  delay as

well.

The  $\bar{\text{H}}$  analysis using the MCP detector system, which is currently being developed in AEgIS and a first glimpse of whose outcome for the 2023 dataset is presented in [156], does not yield any conclusive results as of yet. It is currently being fine-tuned, and first observations seem to indicate strongest traces of a signal excess in the formation modality with the largest  $\bar{\text{p}}$  launch delay, in agreement with the ESDA results presented here. However, due to low statistics, also in that analysis, it will be more interesting to see the 2024 results.

## 7.4 Expected $\bar{\text{H}}$ production improvement

The upgrades on the AEgIS apparatus outlined in Ch. 3 and in [82] together with the resulting possible collinear antihydrogen formation with high-efficiency antiproton accumulation from ELENA and 24-hour operation also thanks to the new CIRCUS control system, as well as an option for colder positronium, are expected to significantly increase the  $\bar{\text{H}}$  production rate and overall produced numbers. The influence of the different developments is summarized here.

### 7.4.1 The Ps Rydberg level reach

One of the most important upgrades is the implementation of an on-axis antihydrogen production (see Sect. 3.3.4), preventing asymmetry losses and in particular allowing for higher excited states of the positronium atoms by drastically reducing Stark ionization [110,111,114], the latter also being possible thanks to the upgraded laser system and lower velocity distribution of the Ps from the upgraded conversion target (see Sect. 3.4).

To elaborate on the second point, the magnitude  $E$  of the electric field  $\vec{E}$  experienced by the excited Ps atoms traveling in the magnetic field  $\vec{B}$  with velocity  $\vec{v}$  is given by

$$E = \left| \vec{v} \times \vec{B} \right| = vB \sin(\theta), \quad (7.13)$$

where  $\theta$  is the angle between the Ps travel direction and the magnetic field lines (equivalent to the trap axis in this case). At the same time, the minimum electric field inducing significant ionization probabilities due to the Stark effect, i.e. the ionization limit,  $E_{\text{ion}}$ , can be expressed as Eq. 7.14, depending on the principal quantum number  $n$  of the Ps atoms [111, 114]. Here,  $e$  denotes the elementary charge,  $\epsilon_0$  is the vacuum permittivity constant and  $a_0$  is the Bohr radius.

$$E_{\text{ion}}(n) = \frac{e}{16\pi\epsilon_0 a_0^2} \frac{1}{9n^4} \quad (7.14)$$

Eq. 7.13 and Eq. 7.14 can be combined for the determination of the maximum Rydberg level  $n_{\text{Ps,max}}$  allowed by the ionization limit from Eq. 7.15.

$$n_{\text{Ps,max}} = \sqrt[4]{\frac{e}{16\pi\epsilon_0 a_0^2} \frac{1}{9vB \sin(\theta)}} \quad (7.15)$$

Assuming the typical Ps velocities of  $10^5 \text{ ms}^{-1}$  [113] and plugging in the 1 T magnetic field in the vicinity of the Ps target (see Sect. 3.4.1), an angle  $\theta = 90^\circ$  of the incoming Ps atoms yields  $n_{\text{Ps,max}} = 19$ . Conservatively, to ensure that the fraction of ionized Ps atoms is kept below , the Rydberg excitation had been limited to  $n_{\text{Ps}} = 17$  in AEgIS Phase I for

the first antihydrogen production runs [68].

Using the geometry of the experimental procedure for the collinear  $\bar{\text{H}}$  production ( $\bar{\text{p}}$  plasma with a 5 mm radius at a distance of 35 mm from the Ps target), a maximum angle  $\theta = 8^\circ$  can be assumed for the Ps trajectory towards the antiprotons, allowing for  $n_{\text{Ps,max}} = 32$  according to Eq. 7.15. For an interaction with only the innermost couple of millimeters of the  $\bar{\text{p}}$  plasma, angles of the order of  $\theta = 3^\circ$  are realistic, yielding  $n_{\text{Ps,max}} = 40$ . Increasing the fraction of antiprotons in this plasma core region is possible with an optimized Rotating Wall technique. The corresponding increase in the  $\bar{\text{H}}$  production cross section according to Eq. 3.11 or Eq. 7.17 is a factor 12.6 or 30.7, respectively for  $n_{\text{Ps,max}} = 32$  and  $n_{\text{Ps,max}} = 40$ .  $n_{\text{Ps}} = 40$  is the desired reachable state with the upgraded EKSPLA laser system (see Sect. 3.4.6). As obvious from Eq. 7.15, slower Ps velocities and weaker magnetic fields have a favourable effect on  $n_{\text{Ps,max}}$ , motivating the developments on Ps laser cooling, a first implementation of which has been successfully achieved by the AEGIS collaboration for the first time in 2023 (see Sect. 7.5), and a reduction of the magnetic field strength in the  $\bar{\text{H}}$  production region, which is currently being implemented and tested.

Fig. 7.17 shows a map of the charge exchange cross section for  $\bar{\text{H}}$  production,  $\sigma_{\text{Ps}^*+\bar{\text{p}}\rightarrow\bar{\text{H}}^*+\text{e}^-}$  (here expressed as  $\sigma$ ), as a function of both the Ps axial velocity and the excited state of the Ps. The plotted values for the cross section are obtained from a Classical Trajectory Monte Carlo (CTMC) simulation [77]. At first order,  $\sigma_{\text{Ps}^*+\bar{\text{p}}\rightarrow\bar{\text{H}}^*+\text{e}^-}$  can be expressed analytically (in units of meter) as given in Eq. 7.17, with  $k_v$  defined in Eq. 7.16 as the ratio between the Ps center of mass velocity  $v_{\text{Ps}}^{\text{cm}}$  and the velocity of the positron in the classical circular orbit,  $1/2n_{\text{Ps}}$ .

$$k_v = \frac{v_{\text{Ps}}^{\text{cm}}}{2n_{\text{Ps}}} \quad (7.16)$$

$$\sigma_{\text{Ps}^*+\bar{\text{p}}\rightarrow\bar{\text{H}}^*+\text{e}^-} = \frac{(1.2 + 0.15k_v^{-1.88}) \cdot 10^{-19}}{1 + (k_v/1.83)^{18}} n_{\text{Ps}}^4 \quad (7.17)$$

The maximum usable excited state of the Ps according to Eq. 7.15 defined by the Stark ionization limit is also plotted in Fig. 7.17 as a function of the Ps velocity, assuming Ps entering the trap at two different angles of  $90^\circ$  and at  $10^\circ$  (red lines). The velocity distribution and used  $n_{\text{Ps}}$  of AEGIS Phase I are indicated by the white cloud. The analysis of the positronium velocity distribution with the new setup is currently ongoing. Furthermore, a simulation investigating the influence of the ballistic  $\bar{\text{H}}$  formation scheme on the number of the produced atoms is also being implemented.

Generally, the expectation from AEGIS Phase II is a shift of the properties of the Ps for  $\bar{\text{H}}$  production to lower velocity distributions (thanks to the improved Ps target configuration and prospects of Ps laser cooling) and higher Rydberg levels (thanks to the collinear setup and upgraded laser system), i.e. unlocking the upper left region of Fig. 7.17, opening the door to orders of magnitude larger cross sections.

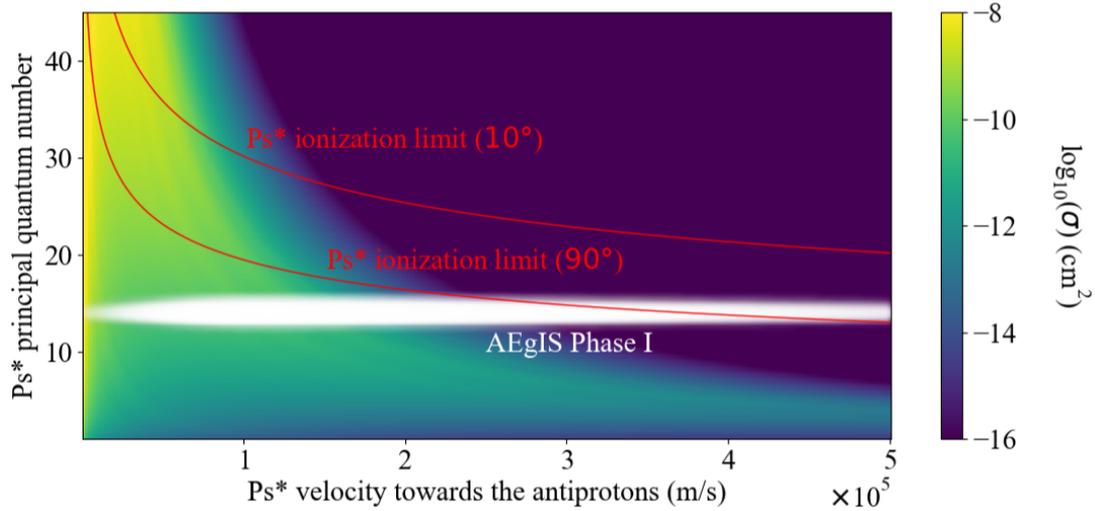


Figure 7.17: Map of the cross section  $\sigma_{\text{Ps}^*+\bar{p}\rightarrow\bar{\text{H}}^*+e^-}$ , here named  $\sigma$ , for charge exchange antihydrogen production as a function of the Ps axial velocity and excited state. The color range indicates the standard logarithm of the cross section value. Also plotted is the maximum Ps principal quantum number available from the ionization limit as a function of the Ps velocity, for two different trajectory angles of the Ps atoms with respect to the magnetic axis (red lines). The white cloud depicts the Ps velocity and excitation distribution obtained in AEGIS Phase I.

## 7.4.2 Summary of the expected $\bar{\text{H}}$ number increase

Table 7.2 summarizes the influence of the upgrades for AEGIS Phase II on the number of antihydrogen atoms produced per formation cycle, one cycle being determined by the AD/ELENA timing of the order of 110s. The values listed for Phase I are those measured [68], while those expected for Phase II are obtained from simulations and simple calculations outlined here above or in the corresponding sections of the main text. According to the combined improvement factors, Phase II is expected to boost the  $\bar{\text{H}}$  production by a factor between 1000 and 3000, increasing the formed number of atoms per cycle from 0.05 to several tens or even above 100.

For an evaluation of the influence of the increased  $\bar{\text{H}}$  numbers on the possible sensitivity of the gravity measurement (see Sect. 3.5), it is important to take into account the changed  $\bar{p}$  provision modality from ELENA as well (explained in Sect. 3.1.4): Instead of in 8-hour shifts, antiprotons are now available 24 hours, seven days a week. This invokes another improvement factor up to 6 (because of the six AD experiments, but in reality somewhat lower since not all experiments were requesting antiprotons during all of their allocated time in the past) when determining the final number of antihydrogen atoms available to the measurement over an extended amount of time. This is a logical consideration, given that the statistics for the gravity measurement are to be accumulated over several months. The last column in Table 7.2 itemizes the performance of the different components in the  $\bar{\text{H}}$  production run in 2023. While the available number of antiprotons has been greatly increased thanks to the developments described here (ELENA connection, CIRCUS control system, novel degrader setup, new trap structure, implementation of improved antiproton capture and trapping) and the reachable excited states of the Ps have been measurably higher already, the under-performance of the Ps chain overall largely limits the achieved

antihydrogen production, keeping the expected  $\bar{H}$  numbers only slightly above those observed in Phase I. Incidentally, the number of expected  $\bar{H}$  atoms per production cycle in this configuration is in good agreement with the numbers found from the experimental ESDA analysis reported on in Sect. 7.3.4 for the 2023 run (0.05 - 0.13  $\bar{H}$  atoms per cycle, depending on the modality), despite the statistical significance being limited.

Upgrade effect	Phase I (meas.)	Phase II (exp.)	Improvement factor (exp.)	Phase II (2023)
$N_{\bar{p}}$	$2.5\% \cdot 3 \times 10^7$	$70\% \cdot 7 \times 10^6$ (**)	6.5 (**)	$70\% \cdot 7 \times 10^6$ ( $\times 6.5$ )
$e^+$ source	900 MBq	1850 MBq	2	200 MBq ( $\times 0.2$ )
$e^+ \rightarrow Ps$ efficiency	5 - 10 %	28 %	3.7	1.5 - 2.5 % ( $\times 0.3$ )
Ps excitation	15 %	30 %	2	15 % ( $\times 1$ )
$n_{Ps}$ reach	17	$\geq 32$	12.6 (*)	$\leq 24$ ( $\times 4$ )
$\bar{H}$ atoms/cycle	0.05	60 - 150 (*)	1200 - 2900 (*)	0.078 (exp.)

Table 7.2: Summary of the effects of the most relevant upgrades from AEGIS Phase I to Phase II on the antihydrogen production. These include: 1. the available number of cold antiprotons  $N_{\bar{p}}$  thanks to the ELENA commissioning and the upgraded control system, degrader setup and trap region, 2. the replaced  $e^+$  source, 3. the new Ps target yielding a more efficient  $e^+ \rightarrow Ps$  conversion, 4. the new excitation laser covering a larger fraction of the Ps atoms, and 5. the higher available Ps Rydberg states,  $n_{Ps}$ , thanks to the collinear production scheme, lower Ps velocities from the improved conversion target, and upgraded laser setup. The upgrades themselves are outlined in the corresponding sections of the main text. The second column shows the values of the different items measured in Phase I [68], while the third column depicts the expected values for Phase II with the incorporated upgrades. In the fourth column, the improvement factor gained from every item is calculated. If a range is stated for the values, the mean value is used for the determination of the improvement factor. The last column shows the values obtained in the transition phase in 2023, whose results are presented in this work, with the resulting expected change factors (improvement or degradation with respect to Phase I) stated below in parentheses. In the final line, the number of antihydrogen atoms produced per production cycle ( $\approx 110$  s) is given for the different setups, measured in Phase I and expected for Phase II according to the combination of the improvement factors. (\*): The reachable  $n_{Ps} = 32$  and resulting improvement factor quotes the values for an interaction of the Ps atoms with the outermost antiprotons in the plasma (lower bound); the values using  $n_{Ps} = 40$  for an interaction with the plasma core are also mentioned (upper bound). (\*\*): The improvement factor 6.5 for  $N_{\bar{p}}$  is valid when using only one shot of ELENA antiprotons directly for the  $\bar{H}$  formation. If the  $\bar{p}$  accumulation and recapture techniques presented in Ch. 6 are used, even more antiprotons can clearly be available.

### 7.4.3 Considerations of the $\bar{\text{H}}$ axial velocity

As obvious from Eq. 3.16 and described in the corresponding Sect. 3.5, slower axial  $\bar{\text{H}}$  velocities  $v_{\parallel}$  and resulting longer travel times  $\Delta t$  in the Moiré deflectometer of the gravity module cause larger vertical deflections due to the influence of gravity and are hence beneficial to the achievable precision in the determination of  $g$ . On the other hand, as also outlined there, to form the antihydrogen into a useful forward-boosted beam towards the gravimeter, the ratio of the axial to the transversal velocity, i.e. to the thermal velocity  $v_{\text{th}}$  in one dimension, Eq. 7.18, should be significant (e.g. an order of magnitude). In Eq. 7.18,  $k_{\text{B}}$  is the Boltzmann constant,  $T$  is the  $\bar{\text{H}}$  temperature and  $m$  its mass.

$$v_{\text{th}} = \sqrt{\frac{k_{\text{B}}T}{m}} \quad (7.18)$$

Thanks to the much smaller mass of the positron/positronium ( $m_{\bar{\text{p}}}/m_{\text{e}^+} \approx 1836$ ), the momentum transferred from this system during the charge exchange reaction is negligible and the  $\bar{\text{H}}$  atoms are assumed to be formed with the temperature of the arriving antiprotons. Previous typical  $\bar{\text{p}}$  temperatures are of a few hundred Kelvin, with values below 100 K having been achieved as well [68, 169]. From Eq. 7.18, resulting thermal velocities between  $600 \text{ m s}^{-1}$  (at  $T = 50 \text{ K}$ ) and  $2000 \text{ m s}^{-1}$  (at  $T = 500 \text{ K}$ ) are obtained.

Expanding on the harmonic oscillator considerations from Sect. 7.3.3, the velocity  $v$  of the antiprotons at a certain time  $t$  during their trajectory in the parabolic potential can be determined according to Eq. 7.19, with  $\Delta x$  and  $\omega$  as defined there, plugging in the determined flight time  $\Delta t \approx 10 \mu\text{s}$ .

$$v(t) = \frac{\Delta x}{2} \omega \sin(\omega t) \quad (7.19)$$

For the attempts at the forward-boosted  $\bar{\text{H}}$  beam reported on here, the antiprotons were launched with a delay of 500 ns with respect to their arrival synchronized to the Ps atoms, meaning that  $\bar{\text{H}}$  is formed at  $t = \Delta t - 500 \text{ ns}$ . Evaluating Eq. 7.19 accordingly yields axial velocities of the antiprotons and, thus, the antihydrogen of the order of  $2.3 \times 10^4 \text{ m s}^{-1}$ , i.e. at least an order of magnitude larger than the expected thermal velocities and allowing for a significant forward boost to increase the number of  $\bar{\text{H}}$  atoms arriving in the gravity module. At the same time, it is worth mentioning that these axial antiproton velocities are still an order of magnitude slower than those of the positronium atoms, which arrive at around  $1 \times 10^5 \text{ m s}^{-1}$ , meaning that they should have a negligible influence on the  $\bar{\text{H}}$  production cross section, for whose calculations the antiprotons have been assumed to be at rest here so far.

Lower antiproton temperatures would clearly allow for slower antihydrogen (i.e. larger vertical deflections and better precision on the gravity measurement). For example, a  $\bar{\text{p}}$  launch delay of only 150 ns would cause theoretical axial velocities of the order of  $6 \times 10^3 \text{ m s}^{-1}$ , still an order of magnitude larger than the thermal velocity at 50 K. Accepting a less forward-boosted beam, delays of the order of 40 to 50 ns can yield axial velocities between 1800 and  $2000 \text{ m s}^{-1}$ , i.e. a boost factor  $\beta \approx 3$  at 50 K; 25 to 30 ns delays give  $\beta \approx 2$ . A comparison of achievable numbers of produced  $\bar{\text{H}}$  atoms and the measurement precision at different temperatures is given in Sect. 3.5.1.

The AEGIS collaboration is investing in the development of additional cooling techniques for antiprotons, including sympathetic cooling on laser-cooled anions and resistive cooling [68, 103, 109, 148], opening the door to sub-Kelvin temperatures that allow for a significant boost while maintaining a high sensitivity on the gravity measurement.

## 7.5 Laser cooling of positronium

The second component needed to form antihydrogen, in addition to antiprotons, are positrons, which, as described in Ch. 3, are supplied in AEGIS from previously created and excited positronium (Ps) atoms. The formation and characterization of the positronium have therefore been a subject of extensive studies in AEGIS, having yielded several fundamental results as well [113, 123, 124, 129–131].

Recently, AEGIS has demonstrated the first successful laser cooling of Ps atoms [116]. By significantly narrowing down the velocity distribution of a formed Ps cloud, this development opens the door to a range of precision experiments, such as Ps spectroscopy for QED tests and positron/electron mass ratio measurements, as well as the formation of positronium Bose-Einstein condensates, allowing to study stimulated annihilation in the gamma radiation range [181–183]. For the purposes of AEGIS, a reduction of the Ps temperature can furthermore significantly enhance the cross section of the charge exchange reaction Eq. 3.10 for antihydrogen formation [77]. In this way, as discussed in Sect. 7.4.1, Ps laser cooling has the potential to contribute to an increase in the produced number of antihydrogen atoms in AEGIS, a development that is crucial to the main objective of the collaboration: a precise measurement of the influence of gravity on antimatter.

### 7.5.1 Setup for broadband laser cooling of positronium

Generally, laser cooling of an ensemble of atomic systems exploits the Doppler effect to reduce their mean kinetic energy [184–186]. When irradiated by a light source whose frequency is tuned to one of its electronic transitions, the atoms with mass  $m$  absorb the incoming photons, experience a velocity change  $\Delta v$  according to the momentum of the photon ( $\Delta v = \hbar k/m$  with  $k$  as the photon's wave number<sup>3</sup> and  $\hbar$  as the reduced Planck constant [187]), and enter an excited state. For de-excitation, a photon is emitted in a random direction by each atom, again changing the atom's momentum accordingly. If the light arrives continuously from the same direction, while the re-emission happens randomly, a repetition of the process leads to a net change in the momentum. For those atoms moving in the direction of the incoming light, this means a momentum reduction; for those moving away, an increase. Thanks to the Doppler effect, it is possible to favour the absorption of photons by those atoms moving towards the light, if its frequency is slightly detuned to below the frequency of the targeted transition ("red-detuned"): The frequency  $\omega$  experienced by the atoms is slightly shifted with respect to the rest frame frequency  $\omega_0$  according to their velocity  $v$  relative to the light source, as given in Eq. 7.20, toward higher frequencies for those atoms moving towards it, causing an increased absorption (and emission), as the experienced frequency is thus closer to resonance [187]. The absorption probability thus becomes dependent on the velocity of the atoms. As a result, the overall mean momentum of the ensemble and, thus, its temperature is reduced.

$$\omega = \omega_0 + kv \tag{7.20}$$

AEGIS uses a specifically designed broadband, frequency-tunable, long-pulse laser setup (the "Alex" alexandrite 243 nm laser system, see Sect. 3.4.6) [86] to strongly saturate the  $1^3S$ - $2^3P$  transition of ortho-positronium<sup>4</sup> for cooling. This laser is referred to as the "cooling laser" in this section and irradiates the Ps atoms in one dimension from both

<sup>3</sup>The wave number  $k$  is the magnitude of the photon's wave vector,  $k = 2\pi/\lambda$  with  $\lambda$  as the photon's wavelength.

<sup>4</sup>Due to its very short lifetime of 125 ps, laser-cooling of para-positronium is not currently feasible.

sides, slightly red-detuned with respect to the transition resonance, for 70 ns immediately after their formation<sup>5</sup>. The long pulse length allows for multiple (up to eleven) cooling cycles, as the spontaneous emission lifetime of the  $1^3\text{S}-2^3\text{P}$  transition is only 3.19 ns and one cooling cycle accordingly takes 6.38 ns on average to address all Ps atoms [188]. The Ps in AEgIS is formed by positron implantation into a converter target (see Sect. 3.4), with the positrons having to be transported to the target over a significant distance to a vacuum test chamber ("Positron test chamber" in Fig. 3.5). To avoid a reduction of the cooling efficiency and an increase of the annihilation rate in the magnetic field ("magnetic quenching")<sup>6</sup> generally present for  $e^+$  transport, a new procedure to direct the positrons with an electrostatic field only has been developed [189,190].

Since the velocity distribution of the Ps ensemble is not a direct experimental observable, instead, the changes in the time distribution of Ps annihilations induced by the cooling laser are employed to quantify the cooling via a comparison of the Doppler profiles with and without cooling: Spectral lines observed from atoms in an ensemble are deformed depending on their velocity distribution due to the different Doppler shifts of the emitting or absorbing atoms, a broader distribution (i.e. higher ensemble temperature) leading to a broader line profile ("Doppler broadening"). The aim is to obtain the resulting so-called Doppler profile, which is the line profile of a given transition, i.e. the distribution of the measured emission or absorption intensity as a function of the light detuning around the transition resonance. For this purpose, the atoms can be velocity-selectively manipulated. In AEgIS, ground state Ps is excited to the  $3^3\text{P}$  level by a 205 nm pulse of the "EKSPLA UV" laser (see Sect. 3.4.6). Only that part of the Ps population is excited for whom the wavelength of this laser, Doppler-shifted according to the atoms' individual velocity towards the laser, is close to the  $1^3\text{S}-3^3\text{P}$  resonance. In this way, the velocity distribution of the Ps atoms is probed, which is why the 205 nm laser is also called "probing laser" in this section. Those atoms excited to the  $3^3\text{P}$  level are subsequently photoionized by the direct 1064 nm laser pulse, referred to as "photoionization laser". This technique reduces the overall Ps population by a value proportional to the number of excitable atoms, i.e. atoms within the spectral bandwidth of the probing laser according to their velocity. A scan over the probing laser wavelength around this resonance can thus yield the corresponding Doppler profile if the reduction in the number of Ps atoms is quantified. For this reason, measurements are performed in different configurations: without the use of any lasers, with only the cooling laser, with only the probing and photoionization lasers, and with all lasers.

The used cooling, probing and photoionization transition frequencies are summarized in Fig. 7.18 a) together with the annihilation lifetimes of Ps in the different levels. The three lasers have to be extremely well synchronized (nanosecond precision) with each other and with the timing of the positrons/positronium atoms to facilitate efficient cooling in multiple cycles and to be able to infer the relative Ps quantities from the convoluted data. This has been achieved thanks to the novel control system, CIRCUS, and the fast-control unit, AERIALIST, developed in AEgIS, which are introduced in Ch. 4.

To quantify the number of annihilating Ps atoms, so-called single-shot positron annihilation lifetime spectroscopy (SSPALS) spectra [191] are used. They show the distribution in time of gamma radiation resulting from Ps annihilations, with the number of annihilations being proportional to the number of atoms remaining in an ensemble.  $t = 0$  is

<sup>5</sup>The pulse duration of the "Alex" laser can in principle exceed 100 ns but is interrupted at 70 ns for this measurement using the included switch, see Sect. 3.4.6.

<sup>6</sup>The increased Ps annihilation rate in external fields stems from singlet-triplet state mixing in the excited state manifold.

defined as the time of positron implantation into the Ps target. An example is shown in Fig. 7.19. In AEGIS, SSPALS spectra are obtained from an oscilloscope<sup>7</sup> reading out the signal of a PMT<sup>8</sup> connected to a  $25 \times 25 \times 25 \text{ mm}^3$   $\text{PbWO}_4$  scintillator, which is placed above the Ps converter. It produces fluorescence according to the intensity of arriving gamma radiation, both from positrons entering and directly annihilating in the converter target before having formed positronium (strong peak around  $t = 0$  in Fig. 7.19) and from Ps annihilations after formation (later contributions to the spectra in Fig. 7.19). Differences of the spectra obtained with and without the use of the different lasers are discussed in Sect. 7.5.2. Schematics of the used setup for the Ps cooling measurements inside the vacuum test chamber are shown in Fig. 7.18 b), including the Ps conversion, the lasers, and the detector. Details can be found in [131] and [116]. The cooling laser (blue,  $\lambda_{243} = 243 \text{ nm}$ ) is synchronized to the positron implantation, i.e. to the Ps formation, in time and space (with the laser being ramped up already slightly before to ensure its intensity is maximal at  $t = 0$ ), while the probing (purple,  $\lambda_{205} = 205 \text{ nm}$ ) and photoionization (dark red,  $\lambda_{1064} = 1064 \text{ nm}$ ) laser impact is located at a distance of 7 mm and delayed by 75 ns to correct for the travel of the Ps atoms (peak velocity component) during the 70 ns cooling pulse and to allow for a 5 ns de-excitation time of the excited atoms to the ground state prior to the probing.

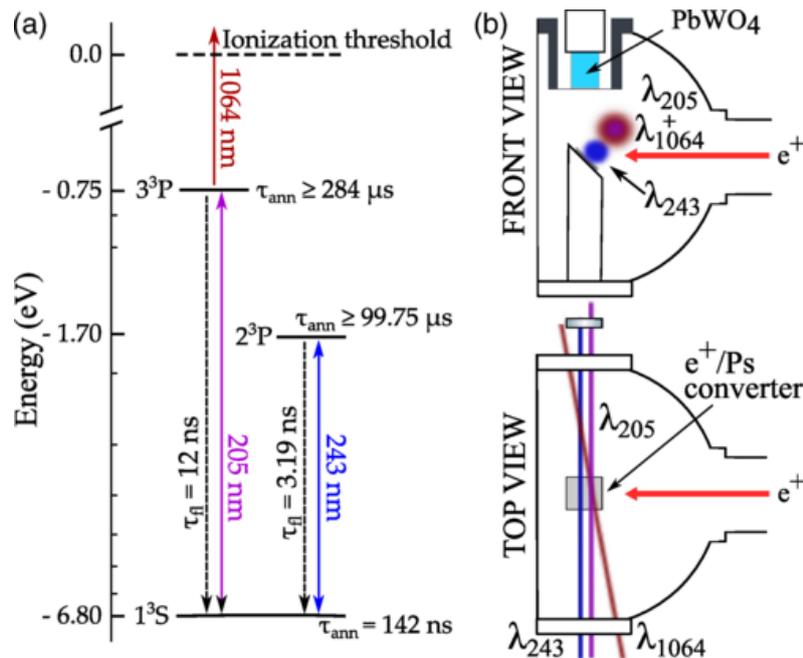


Figure 7.18: Schematics of the experimental setup used for positronium laser cooling in AEGIS. a) Transition diagram of the used Ps energy levels. b) Front and top views of the experiment chamber with the positron implantation (red arrow) into the Ps target (gray rectangle), the entering laser beams, and the detection module. The laser beams are shown in the same colors as the corresponding transitions in a): the cooling laser (blue,  $\lambda_{243} = 243 \text{ nm}$ ), the probing laser (purple,  $\lambda_{205} = 205 \text{ nm}$ ), and the photoionization laser (dark red,  $\lambda_{1064} = 1064 \text{ nm}$ ).

<sup>7</sup>The used oscilloscope is a Teledyne LeCroy model HDO4104A.

<sup>8</sup>The installed PMT is a Hamamatsu model R11265U-200.

### 7.5.2 Experimental results of positronium laser cooling

SSPALS spectra, i.e. spectra of Ps annihilation events as a function of time, are obtained for different configurations with and without the interaction of the lasers, as shown in Fig. 7.19, to extract the parameter  $S$  according to Eq. 7.21. Here,  $f_{\text{off}}$  (without the use of any lasers) and  $f_{\text{on}}$  (with the use of at least one of the lasers,  $S$  is further distinguished by the choice of the used lasers) correspond to the integrated SSPALS spectra from 150 to 400 ns (gray band in Fig. 7.19), averaged over multiple acquisitions. This integration window is chosen because direct  $e^+$  annihilations upon implantation into the target prior to Ps formation dominate at earlier times (strong, partly shown peak around  $t = 0$ ), and the signal is too low to be distinguishable from the noise level later on.  $S$  thus quantifies the relative increase or decrease in the Ps annihilation signal caused by the laser interaction and is generally stated in %. To deconvolute the contribution of the cooling laser from the combined effect of all three lasers, individual measurements are necessary. The long, exponential tail in Fig. 7.19 corresponds to the lifetime of Ps in vacuum (142 ns for the  $1^3\text{S}$  ground state, black dotted line). The cooling laser excites some of the atoms to the  $2^3\text{P}$  state (green dash-dotted line), extending the average annihilation lifetime by several orders of magnitude and causing an increased rate of annihilation events at later times. The use of the probing and photoionization lasers only (red dashed line), instead, velocity-selectively excites a fraction of the Ps atoms and subsequently photoionizes them, causing a relative increase of the detected annihilations directly after the probing and photoionization laser pulses (bump at around 90 ns) from photo-detached positrons hitting the converter target followed by a subsequent reduction of the Ps annihilations in the tail corresponding to the reduced numbers of ground state Ps. The spectrum obtained with the use of all three lasers shows the combined effects (blue line).

$$S = \frac{f_{\text{on}} - f_{\text{off}}}{f_{\text{off}}} \quad (7.21)$$

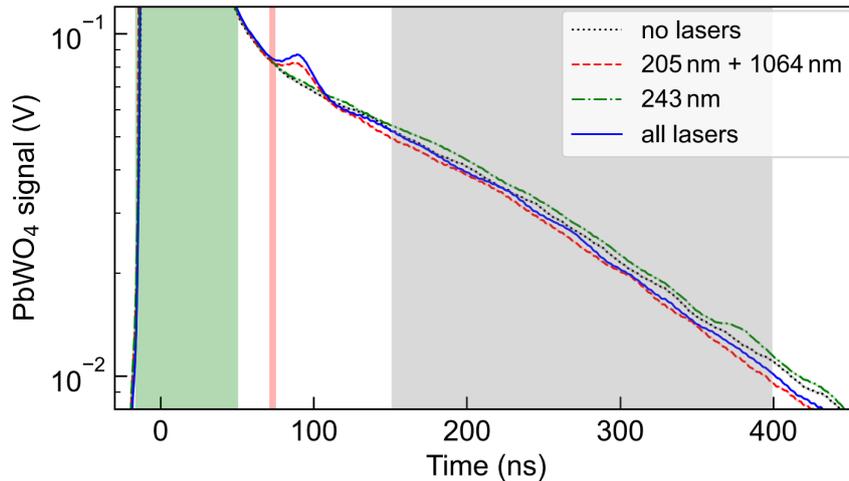


Figure 7.19: Positronium SSPALS spectra obtained for different laser configurations: without the use of any lasers (black dotted line), with only the cooling laser (green dash-dotted line), with only the probing and photoionization lasers (red dashed line), and with all lasers (blue solid line). Each spectrum is the average of 90 individual curves taken in independent measurements with the given laser configuration. The resulting statistical errors are smaller than the curve widths.  $t = 0$  is defined as the instant of positron implantation into the Ps target. The green band shows the active time of the cooling laser and the vertical red line indicates the firing of the probing and photoionization lasers. The integration of the spectra is performed in the time window of the gray band.

The parameter  $S$  is determined for the different laser configurations and, ultimately,  $S_{\text{cool}}$  is defined as the difference in  $S$  between all lasers being present ( $S_{243+205+1064}$ ) and the exclusive use of the cooling laser ( $S_{243}$ ), as given in Eq. 7.22, to quantify the normalized number of Ps atoms getting excited and photoionized by the corresponding lasers after interacting with the cooling laser.

$$S_{\text{cool}} = S_{243+205+1064} - S_{243} \quad (7.22)$$

To infer the Ps velocity distributions with and without laser cooling<sup>9</sup>, a scan over the probing laser detuning, a so-called Doppler scan, is performed, while the wavelength of the cooling laser is kept at 243.061 nm, corresponding to a detuning of 200 GHz below the resonance wavelength of 243.024 nm.  $S_{\text{cool}}$  is determined as a function of the probing laser detuning. The upper plot of Fig. 7.20 shows the resulting Doppler profile (blue line), obtained from a moving average with a square window of  $w_{\text{width}} = 350$  GHz width over the individual  $S_{\text{cool}}$  values<sup>10</sup>, compared to the same measurement without letting the Ps interact with the cooling laser beforehand,  $S_{205+1064}$  (red dotted line). Accordingly, the distributions are the result of a convolution of the actual physical distribution of the Ps ensemble with the probing laser bandwidth and the precision defined by the

<sup>9</sup>Preceding the combined measurement presented here, Doppler scans have been performed separately for the detuning of each laser system to deepen the understanding of the interactions of the individual lasers on the positronium. The corresponding results are outlined in [116] and are compatible with the data shown here.

<sup>10</sup>For the moving average window method, the mean value of a constant number of consecutive data points (given by the width of the window) is calculated in every step, with every step meaning here that the single lower-most data point is removed from the analysis while the next one following is newly included.

averaging window. The distribution after cooling is narrower thanks to a larger fraction of Ps atoms having close to zero velocity, i.e. being addressed by a zero detuning of the probing laser. The widths of the two distributions, as defined by the corresponding standard deviation values, are determined from a Gaussian fit ( $(330 \pm 2)$  GHz without the cooling laser and  $(269 \pm 1)$  GHz with cooling). Then, the corresponding Ps root mean square velocities<sup>11</sup>  $v_{\text{rms}}$  are yielded by a deconvolution<sup>12</sup> of the probing laser bandwidth,  $\sigma_{205} = (179 \pm 9)$  GHz, and the standard deviation of the rolling average method with the used window,  $w_{\text{width}}/\sqrt{12}$ . They result in  $(5.4 \pm 0.2) \times 10^4 \text{ m s}^{-1}$  for the non-cooled Ps and  $(3.7 \pm 0.2) \times 10^4 \text{ m s}^{-1}$  with the use of the cooling laser. According to Eq. 7.23, the respective one-dimensional temperatures  $T$  of the Ps ensemble can be calculated from the one-dimensional thermal velocity  $v_{\text{th}}$ , with  $k_{\text{B}}$  as the Boltzmann constant<sup>13</sup> [192]. Using the Ps rms velocities as  $v_{\text{th}}$ , the temperatures result in  $(380 \pm 20)$  K for non-cooled Ps and  $(170 \pm 20)$  K for cooled Ps, representing a reduction of the Ps temperature of  $(210 \pm 30)$  K through the described cooling technique, which is in agreement with theoretical predictions for the Doppler limit, i.e. with the minimum temperature achievable with the simple approach used here<sup>14</sup>, which is a result of the balancing of the cooling and a heating through spontaneous emission [116, 188, 190].

$$T = \frac{mv_{\text{th}}^2}{k_{\text{B}}} \quad (7.23)$$

A second Doppler scan has been performed to determine the maximum fraction of Ps atoms that can be cooled to velocities below  $3.7 \times 10^4 \text{ m s}^{-1}$  by leaving the probing laser on resonance and detuning the cooling laser. The result is the profile shown in the lower plot of Fig. 7.20, obtained with the same procedure described above, with a rolling average window with a 200 GHz width. Without the cooling laser,  $S_{205+1064}$  (red dashed line) is found to be  $(8 \pm 0.2)\%$ .  $S_{\text{cool}}$  is given by the blue line. For a given detuning, the difference between the two curves represents the fraction of Ps atoms cooled by the AEgIS technique. The maximum relative increase is found to be  $(58 \pm 9)\%$  when the cooling laser is detuned by 350 GHz below resonance.

<sup>11</sup>The root mean square velocity of an ensemble is defined as the square root of the average of the square of the velocities.

<sup>12</sup>The term deconvolution here refers to the subtraction of the squared quantities from the square of the obtained distribution widths  $f_{\text{width}}$  before taking the square root and multiplying by the probing laser wavelength  $\lambda_{205}$ :  $V_{\text{rms}} = \lambda_{205} \cdot \sqrt{f_{\text{width}}^2 - \sigma_{205}^2 - w_{\text{width}}^2/12}$ .

<sup>13</sup>Eq. 7.23 can be obtained by equating the classical expression for the kinetic energy,  $E_{\text{kin}} = 1/2mv^2$ , and the one for the one-dimensional average translational thermal kinetic energy,  $E_{\text{kin}} = 1/2k_{\text{B}}T$ .

<sup>14</sup>It is possible to cool below the Doppler limit using more sophisticated techniques, such as Sisyphus cooling, which utilize laser beams with opposite polarization going in opposite directions, producing coherent photon scattering between the beams [193].

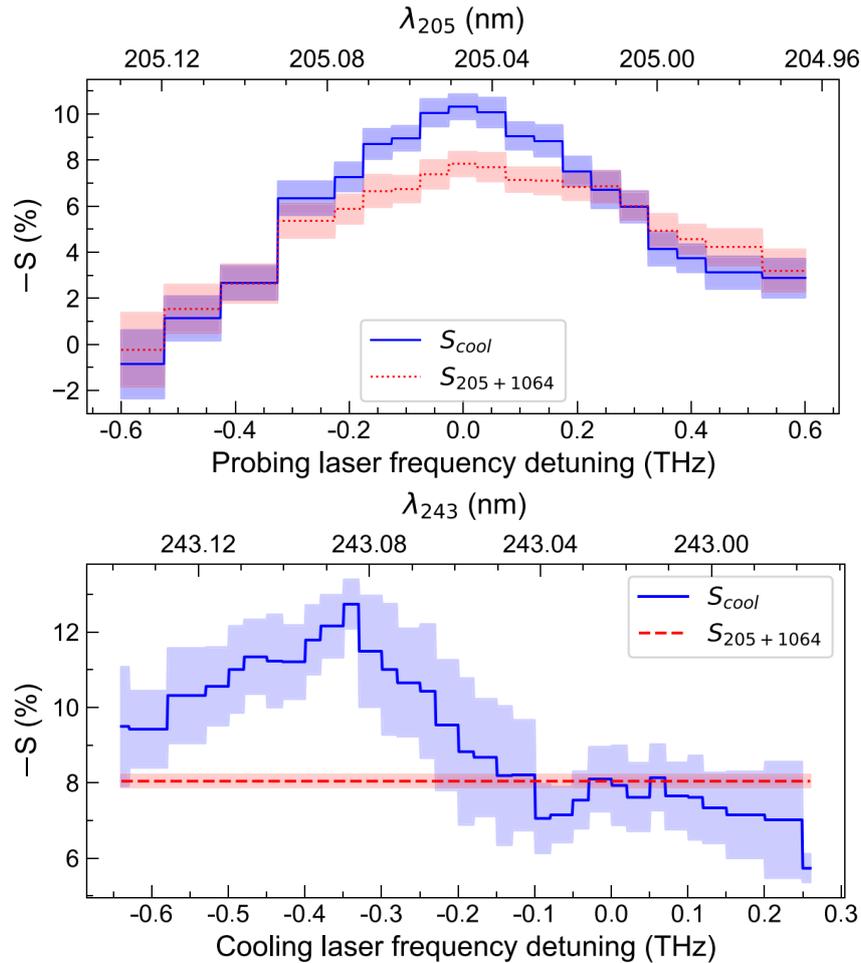


Figure 7.20: One-dimensional Doppler profiles of the Ps cloud. The blue line is the profile for the full cooling procedure,  $S_{cool}$  as defined in the main text; the red line is obtained with only the probing and photoionization lasers,  $S_{205+1064}$ . The semi-transparent bands indicate the statistical uncertainties (one standard deviation). Top: Scan over the probing laser detuning for a fixed cooling laser frequency (200 GHz below resonance). Bottom: Scan over the cooling laser detuning with the probing laser at resonance.

## 7.6 Summary

A procedure has been developed to produce antihydrogen in a pulsed, horizontal beam instead of the isotropic source obtained in previous  $\bar{\text{H}}$  productions in AEGIS. This development is needed to render a measurement of the effect of gravity on antimatter feasible using the collaboration's deflectometer approach and pulsed  $\bar{\text{H}}$  production through charge exchange of cold antiprotons and highly excited positronium. Essentially, instead of combining stationary antiprotons with Ps entering the trap, the antiprotons are accelerated from the trap in which they are captured and cooled in the direction of the positronium-positronium conversion target, precisely synchronized to Ps formation and excitation. For the purpose of  $\bar{\text{p}}$  acceleration, procedures for the reshaping of the potentials generated by the low-voltage electrodes of the electromagnetic trap structure, which confine the electron and antiproton plasmas, have been carefully devised in the newly implemented AERIALIST framework. The mixed plasma is kept in a "launch trap" configuration, whose two endcap electrodes have the option of a fast pulse to change their voltages.

This feature is used on the upstream endcap to remove the cooling electrons prior to the acceleration of the antiprotons. The potential beyond the launch trap in the downstream direction is formed into a parabola, which, upon pulsing the trap's upstream endcap, allows a passage of the antiprotons from the trap toward the Ps target with a slow acceleration, preventing a sudden plasma expansion, and enables a tuning of the antiproton velocity at the instance of antihydrogen formation. The pulsing on the endcaps can also be employed to let the antiprotons re-enter the trap for recycling.

The procedures detailed in Ch. 6 for the preparation of cold antiprotons in the AEgIS trap as well as their acceleration to the Ps target have been combined in a joint procedure with the Ps formation and excitation routines in the new control system, CIRCUS. For an efficient antihydrogen production, a synchronization of the procedures is essential - due to the short lifetime of the Ps and the requirement of a spacial overlap of the two species -, as is a minimization of losses. The synchronization is enabled by the CIRCUS/AERIALIST control system and verified by detailed studies of the annihilation signals observed in the ESDA detector at different times in individual and combined runs of positronium and antiprotons, and the involved trap potentials have been carefully optimized.

Using these procedures, first antihydrogen production runs in the upgraded apparatus, AEgIS Phase II, have been performed. In these runs, the delay given to the antiprotons prior to launch is varied with the aim of producing antihydrogen in different modalities: with an upstream acceleration, at rest, and accelerated in the downstream direction (as required for the beam formation). Control runs have also been taken, during which the firing of the excitation laser is inhibited such that positronium is not excited and no  $\bar{\text{H}}$  formation is expected. An analysis method for the detection of the antihydrogen annihilations in the ESDA signal has been implemented, involving an event selection and the comparison of the signals in the different modalities. This analysis verifies the antiproton acceleration procedure, as it finds clear traces of the antiprotons "oscillating" back and forth along the parabolic potential with the expected timing, and shows hints of excess signals at the estimated times for those runs with an expected  $\bar{\text{H}}$  production. The statistical significance observed from a combination of the datasets in the different run modalities is found to be above the evidence level of three standard deviations. However, the overall statistics are still low and future runs will yield more meaningful results.

On the positronium side, first-ever laser cooling of a fraction of a Ps ensemble has been achieved by strongly saturating the  $1^3\text{S}-2^3\text{P}$  transition with a broadband, long-pulsed alexandrite laser. To probe and quantify the cooling, Ps atoms are velocity-selectively excited to the  $3^3\text{P}$  level, exploiting the Doppler effect, and then photoionized by the "EKSPLA" laser system, synchronized to the cooling thanks to CIRCUS, reducing the subsequently observed Ps annihilations according to the number of atoms removed from the ensemble. A scan over the probing laser detuning yields insight into the thermal velocity distribution of the Ps population and, thus, its temperature. The cooling is found to reduce the Ps ensemble's temperature from  $(380 \pm 20)$  K to  $(170 \pm 20)$  K and increase the fraction of Ps with a velocity below  $3.7 \times 10^4$   $\text{m s}^{-1}$  relatively by  $(58 \pm 9)\%$ .

Given the collectivity of the upgrades performed in AEgIS, including most prominently the efficient capture and accumulation of cold antiprotons presented in Ch. 6 and more efficient positronium formation on the axis of the experiment with reduced Ps velocities and higher excitation levels, an increase in the producible number of antihydrogen atoms by a factor 1000 or more compared to previous production has become realistic, yielding the possibility for several tens to over 100 anti-atoms per production cycle of a few minutes. The acceleration of the antiprotons along a parabolic potential prior to  $\bar{\text{H}}$  formation also allows for a tuning of the velocity of the resulting antihydrogen atoms, which can be

used to maximize the precision achievable by the gravity measurement by a balancing of the number of available  $\bar{\text{H}}$  atoms and their velocity.

# Chapter 8

## Discussion and outlook

The direct study of antimatter systems represents a window into unknown physics, a possibility to shed light on some of the mysteries of the origin and evolution of the universe and the question how it has turned out to be so dominated by regular matter. The production of antimatter and its exploration have become feasible thanks to a large range of technological advancements over the past decades, following the first discovery of antiparticles in the 1930s.

The developments and studies performed in this work have been driven by the requirements of the AEGIS experiment at CERN's unique Antiproton Decelerator facility for the production of low-energy antiprotons, which lie at the core of most bound antimatter systems. In its main line of research, AEGIS produces antihydrogen atoms in a pulsed scheme through a charge exchange reaction between cold, trapped antiprotons and positronium laser-excited to Rydberg states. The main objective is a precise measurement of the influence of gravity on the anti-atoms via the determination of their vertical deflection when having passed through a deflectometer. AEGIS has undergone several significant upgrades in the past few years to improve its production of antihydrogen and render the experiment more versatile. Some of these upgrades form part of the work presented here, including work on the commissioning of a new electron gun and a new electromagnetic trap and the incorporation of multiple components into the newly developed control system.

Naturally, an experiment designed for the formation of antimatter requires very specific conditions, including high vacuum, strong confinement and precise timing, and relies on the combination of techniques from different fields as well as a multitude of independent subsystems. The **first main achievement** of this work is the development and implementation of the CIRCUS framework with the AERIALIST fast control system, which is capable of managing the entire experiment steadily and in an easily adaptive fashion, essentially without human supervision, and provides nanosecond synchronization to the experimental procedures thanks to Sinara hardware and an extensive library structure built from the ARTIQ software base. This system has become a core pillar of AEGIS and has enabled the introduction of improved efficient experimental routines, notably major improvements to the steps involved in antihydrogen production as well as the successful implementation of a procedure to laser-cool an ensemble of positronium atoms, a milestone in antimatter physics. CIRCUS has run the experiment reliably over the recent data taking campaigns since 2021, even autonomously for extended periods of time, and has proved to be ideally suited for the needs of an antimatter/atomic physics experiment, forming a solid base for future undertakings.

Following the remodelling of the electron gun and the electron plasma procedures in CIRCUS, as part of this work, new calibration and alignment techniques have been developed

that rely on the electron system, some of which have become part of the standard processes of AEGIS. Generally, these electron procedures are one of the most integral parts of the experiment, as they form the solid starting point for the development of more intricate routines involving antimatter.

As is the case for other antimatter experiments as well, antiprotons play a crucial role for AEGIS - in combination with positronium for the  $\bar{H}$  production and for the formation of most other bound antimatter systems. The **second main achievement** presented here is the implementation of a source of controlled, readily available, cold antiprotons. Through the development of procedures to routinely capture of the order of 70% of antiprotons available from CERN's ELENA decelerator, apply efficient sympathetic  $\bar{p}$  cooling and compress them using the Rotating Wall technique, it has become possible to accumulate and store large numbers. A density of around 7.5 million cold  $\bar{p}$  per cubic centimeter has been successfully achieved in one of the electromagnetic traps of AEGIS: Several hundred million cooled antiprotons have been confined together. This value represents a record accumulation of cold antiprotons by orders of magnitude and creates a playground for a variety of experiments involving antimatter.

For the gravity measurement of AEGIS, the antiprotons have to be accelerated in the forward direction such that the produced antihydrogen atoms are boosted towards the deflectometer. As the **third main achievement** of this work, the procedures involved in this technique have been elaborately developed and implemented via the formation of a parabolic potential along the axis of the experiment, and the  $\bar{p}$  routines have been combined with the positronium components for antihydrogen production and formation of a horizontal beam. For the  $\bar{H}$  production, the arrival of the antiprotons in the Ps formation region has been studied and, within the AERIALIST system, precisely timed to the excitation of the positronium such as to synchronize the two processes for improved efficiency.

For the antihydrogen production run of 2023, the analysis of the data acquired with a scintillator/PMT detector system has been analyzed as part of this work, involving an intricate event selection.  $\bar{H}$  formation has been attempted in four different modes, depending on the timing of the antiprotons and the enabling or disabling of the IR laser used for the Ps excitation. Oscillations of the antiprotons swinging in the parabolic potential have been clearly observed in all four modalities with the expected timing, verifying the functionality of the  $\bar{p}$  boost. Indications of a successful  $\bar{H}$  production have also been observed as an excess signal in the expected time range for those runs with the IR laser enabled. However, mainly due to heavy impediments on the positronium line, the statistics of the run are too low to be significant. Given the entirety of the upgrades in AEGIS, the number of formed antihydrogen atoms per production cycle is expected to increase by at least a factor 1000 compared to previous runs, giving of the order of 100 atoms every few minutes. Such an improvement renders the statistics required for a relative 1% precise antihydrogen gravity measurement feasible within a few months, a realistic time scale of available  $\bar{p}$  beam time. It will be very interesting to see the influence of the different components on the  $\bar{H}$  yield in the full run of 2024.

This work has contributed to the AEGIS experiment with a variety of original developments and results. With all the upgrades in place, starting from the powerful CIRCUS control system and a potent source of cold antiprotons and ranging to laser cooling of positronium and the development of the antihydrogen beam formation techniques, AEGIS has great potential for upcoming measurements, for precise tests of the gravitational interaction of antimatter thanks to a strongly boosted  $\bar{H}$  production and for studies of exotic antimatter systems beyond.

# List of Figures

3.1	Schematic of the CERN accelerator (and decelerator) complex. . . . .	16
3.2	Schematic of the AD hall at CERN. . . . .	20
3.3	Drawing of the operation principle of a Penning trap to confine a plasma. .	21
3.4	Schematics of antihydrogen production with the AEGIS experiment. . . . .	27
3.5	Simplified schematic setup of the AEGIS experiment. . . . .	28
3.6	Magnetic flux density in the AEGIS experiment. . . . .	29
3.7	Photographs of the Ps converter target and its holder, mounted in the center of the last trap electrode. . . . .	31
3.8	Photographs of the two degrader units installed in the AEGIS apparatus. .	35
3.9	Schematic of the Penning-Malmberg trap regions in the AEGIS apparatus. .	36
3.10	Potentials produced by the application of a voltage of $-150$ V on a single electrode along the horizontal axis of the experiment. . . . .	38
3.11	Images of the Penning-Malmberg trap installed in the 1 T apparatus of the AEGIS experiment. . . . .	39
3.12	Plot of the number of counts observed on SC12 with and without a splitter installed on the corresponding PMTs, for a set of measurements releasing captured antiprotons from the P Trap, plotted against the number of counts on SC56. . . . .	42
3.13	Experimental and simulated ESDA amplitude spectra used for the determination of useful threshold cuts. . . . .	43
3.14	Schematic functionality of a Moiré deflectometer, as designed for the gravity measurement of AEGIS. . . . .	44
3.15	Results of light and antiproton measurements in a Moiré deflectometer prototype. . . . .	45
3.16	Schematic illustration of the preliminary design of the AEGIS gravity module. .	47
3.17	A prototype high-resolution detector for reconstruction of antiproton annihilations. . . . .	50
4.1	Fully equipped electronics crate of the AEGIS experiment composed of Sinara hardware. . . . .	53
4.2	Rough verification measurements of amplifier characteristics. . . . .	55
4.3	Simple electronics circuit of a resistive voltage divider. . . . .	56
4.4	Oscilloscope views of the external clock input signals to the Sinara Kaslis without and with installed voltage dividers. . . . .	57
4.5	Oscilloscope image of the external clock input signals to the Sinara Kaslis without and with attenuators included in the circuit. . . . .	58
4.6	Schematic of the ARTIQ/Python library structure developed in AEGIS. . .	65
4.7	run method of the AEGIS Class. . . . .	66

4.8	Experimental scripts in the ARTIQ environment defining the routine to produce a well potential trap formed by three electrodes upon the arrival of an external trigger signal. . . . .	68
4.9	Calibration fit for the trap electrode amplifier channels. . . . .	70
4.10	Distribution of extracted calibration constants for the amplifier channels used for the AEGIS trap electrode voltages. . . . .	71
4.11	Voltage accuracy of one example Sinara amplifier board before and after calibration. . . . .	73
4.12	Systematic checks of the output stability of the calibrated amplifier channels. . . . .	74
4.13	Visualization of the constituents of the CIRCUS control system (AERIAL-IST and TALOS) as well as its relationship to experimental subsystems. . . . .	75
4.14	Schematic of the TALOS structure. . . . .	76
4.15	Screenshot of the Electron Gun $\mu$ Service running in the CIRCUS control system. . . . .	78
4.16	Synchronous voltage ramp from 0 to 20 V on three Sinara Fastino/amplifier channels subsequent to a common trigger signal. . . . .	80
4.17	Oscilloscope images of synchronized TTL pulses on two Sinara digital I/O channels. . . . .	83
5.1	Photographs of the new AEGIS electron gun. . . . .	88
5.2	Schematic of the electric circuit of the electron gun as it is set up in AEGIS. . . . .	89
5.3	Systematics of the observed AEGIS electron gun current. . . . .	90
5.4	Successful test of connection and functionality of a trap electrode using the electron gun current. . . . .	92
5.5	Potential reshaping operations to load electrons into the C Trap and dump them towards the MCP. . . . .	94
5.6	Images of the CMOS camera $\mu$ Service during experimental alignment procedures using electron imaging. . . . .	96
5.7	Distributions obtained by digitizing the MCP front face signal from an impinging electron current with different pulse lengths. . . . .	98
5.8	Integrated digitized MCP front face signals and corresponding detected numbers of electrons as a function of the electron gun current pulse length. . . . .	99
5.9	Observed current on the MCP front face as a function of the electron gun pulse length. . . . .	100
5.10	Observed MCP front face signals from an impinging pulse of electrons with an MCP front-to-back bias of 1.4 kV. . . . .	101
5.11	Integrated MCP front face signals as a function of the front-to-back MCP voltage for impinging 20 ns electron pulses. . . . .	102
5.12	Analysis of the CMOS image of the MCP's phosphor screen stemming from the arriving electron current. . . . .	103
5.13	Analysis of the CMOS image of the MCP's phosphor screen stemming from arriving positrons. . . . .	104
6.1	Schematic of the 5 T trap structure used for the capture and cooling of antiprotons in AEGIS. . . . .	107
6.2	Low-potential reshaping operations to load electrons from the C Trap into the Short P Trap and prepare the trap for antiproton cooling. . . . .	109
6.3	Low-potential reshaping operations to prepare the P Trap for a dump of the confined particles. . . . .	110

6.4	MCP intensity map corresponding to the distribution of the number of captured antiprotons obtained from the beam steering procedure. . . . .	111
6.5	Plot of the number of counts on SC34 subsequent to a dump on the degrader structure as function of the 5 T trap closing time. . . . .	113
6.6	Simulation of the passage of antiprotons through the AEGIS degrader structure. . . . .	115
6.7	Experimental studies of the passage of antiprotons through the AEGIS degrader system. . . . .	117
6.8	Comparison of the distribution of the number of counts on SC56 versus time when trapping one bunch of ELENA antiprotons in the P Trap for 80s without and with cooling electrons. . . . .	120
6.9	The achieved electron-antiproton cooling efficiency as a function of the hot storage time, as obtained on different ESDA units. . . . .	121
6.10	The achieved electron-antiproton cooling efficiency as a function of the hot storage time, as obtained from the average of SC56 and SC78. . . . .	122
6.11	The achieved electron-antiproton cooling efficiency as a function of the electron space charge transferred from the C Trap to the P Trap, as obtained on different ESDA units. . . . .	124
6.12	Analysis of the radii of plasmas of antiprotons and electrons for different amounts of electron space charge transferred from the C Trap to the P Trap. . . . .	124
6.13	Comparison of the distribution of the number of counts on SC56 versus time when trapping one bunch of ELENA antiprotons in the P Trap for 60s without and with Rotating Wall compression. . . . .	126
6.14	Measurement of the number of cold dump (60s of cold storage) counts on an exemplary ESDA scintillation detector (SC1112) as a function of different Rotating Wall frequencies for various Rotating Wall amplitudes. . . . .	128
6.15	Example CMOS images taken on the downstream MCP detector of the electron-antiproton plasma dumped from the trap, subsequently to being subject to Rotating Wall compression in the trap, with two different settings, for a cold storage time of 60s. . . . .	128
6.16	Distributions of the number of counts on SC1112 versus time for an example experiment accumulating antiprotons from eight ELENA bunches in the trap. . . . .	132
6.17	Measurement of the number of cold dump counts on ESDA unit SC1112 following an accumulation of 45 ELENA antiproton bunches in the trap as a function of the Rotating Wall frequency for three different Rotating Wall amplitudes. . . . .	134
6.18	Number of observed cold dump counts on SC1112 and corresponding number of trapped cold antiprotons as a function of the number of ELENA bunches accumulated in the trap. . . . .	135
6.19	Distributions of counts registered on SC1112 in the energy and time domains subsequent to a dump of five stacked ELENA bunches from the Short P Trap. . . . .	138
6.20	Nested trap potential used for the formation and capture of positively charged ions produced by the annihilation of trapped antiprotons on a gas. . . . .	140
7.1	Low-voltage potential reshaping operations to transform the P Trap from the dump trap into the launch trap. . . . .	147
7.2	Representation of the parabolic potential formed by electrodes P10 to A1. . . . .	149

7.3	Low-voltage potential reshaping operations to form the parabolic potential and launch the particles towards the downstream side from the P Trap. . .	150
7.4	Low-voltage potential reshaping operations to remove the cooling electrons from the launch trap. . . . .	151
7.5	Low-voltage potential reshaping operations to launch and recapture antiprotons from/in the P Trap for antihydrogen production. . . . .	153
7.6	Normalized intensity observed on the MCP/CMOS detector when dumping antiprotons from the trap, subsequent to their recapture, as a function of the opening time of the barrier electrode. . . . .	154
7.7	Low-voltage potentials along the experiment axis over the entire relevant range during the antiproton ballistic launch for antihydrogen production. .	156
7.8	Signal distribution observed on the digitized PMT 24 as a function of time for a scan over the target electrode voltage in $\bar{\text{H}}$ production runs without positrons. . . . .	158
7.9	Signal distribution observed on the digitized PMT 11 as a function of time for $\bar{\text{H}}$ production runs with different settings. . . . .	160
7.10	Distribution of digitized amplitudes versus time of an example ESDA PMT for an antihydrogen production run. . . . .	162
7.11	Distributions of the number of identified peak events in the ESDA as a function of time after the expected Ps formation in the range between 0.5 and 20.5 $\mu\text{s}$ . . . . .	165
7.12	Number of remaining event peaks in the signal window between 0.5 and 5.5 $\mu\text{s}$ after the expected Ps formation for different amplitude threshold cuts, with PMT coincidence filtering applied. . . . .	166
7.13	Distribution of the number of selected events per run as a function of the run number during the antihydrogen production campaign in 2023 for the four different run modalities. . . . .	167
7.14	Distributions of the number of identified peak events in the ESDA as a function of time until 150 $\mu\text{s}$ after the expected Ps formation. . . . .	169
7.15	Distributions of the number of identified peak events in the ESDA as a function of time between -10 and 20.5 $\mu\text{s}$ after the expected Ps formation. .	170
7.16	Antihydrogen candidate events in the signal window as a function of time after Ps formation. . . . .	171
7.17	Map of the cross section $\sigma_{\text{Ps}^* + \bar{\text{p}} \rightarrow \bar{\text{H}}^* + \text{e}^-}$ , here named $\sigma$ , for charge exchange antihydrogen production as a function of the Ps axial velocity and excited state. . . . .	175
7.18	Schematics of the experimental setup used for positronium laser cooling in AEGIS. . . . .	180
7.19	Positronium SSPALS spectra obtained with different laser configurations. .	182
7.20	One-dimensional Doppler profiles of the Ps cloud. . . . .	184
B.1	Feynman diagrams of processes leading to the production of positrons. . .	210
F.1	Internal setup of a Sinara Kasli controller. . . . .	216
F.2	JSON file used for gateway image compilation, corresponding to the setup of a typical Kasli used in AEGIS. . . . .	217
G.1	Schematic of the electronic circuit of one channel of the custom-designed amplifier board. . . . .	218
G.2	Photograph of an oscilloscope showing the oscillations of the amplifier signal.	219

---

G.3	Photograph of one amplifier board with the added capacitors. . . . .	219
G.4	Output voltage ramp from 0 to 10 V on an example Fastino channel, including a significant overshoot, as recorded by an oscilloscope. . . . .	219
G.5	Output voltage ramp from 0 to 50 V on three example channels of Fastino and custom amplifier with no observed overshoot, as recorded by an oscilloscope. . . . .	220
H.1	Calibration results of all Sinara amplifier channels of the AEGIS apparatus. . . . .	222
I.1	CMOS images taken with the downstream MCP of the mixed electron-antiproton plasma for different amounts of electron space charge transferred from the C Trap to the P Trap. . . . .	224
I.2	Example CMOS image taken with the downstream MCP of the incoming ELENA antiproton bunch. . . . .	225
J.1	Low-voltage potential reshaping operations to return from the recycling and launch trap to the dump trap configuration. . . . .	227
K.1	Signal distributions observed on the digitized PMT 24 as a function of time for a scan over the target electrode voltage in $\bar{H}$ production runs without positrons. . . . .	228
K.2	Signal distributions observed on the digitized PMT 24 as a function of time for a scan over the target electrode voltage in $\bar{H}$ production runs without positrons. . . . .	229

# Bibliography

- [1] Mary K. Gaillard, Paul D. Grannis, and Frank J. Sciulli, “The standard model of particle physics.” *Rev. Mod. Phys.* 71, S96, 1999.
- [2] Carl D. Anderson, “The Positive Electron.” *Phys. Rev.* 43, 491, 1933.
- [3] Paul Adrien Maurice Dirac, “The quantum theory of the electron.” *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, Vol. 117, No. 778, pp. 610–624,, 1928.
- [4] Peter W. Higgs, “Broken Symmetries and the Masses of Gauge Bosons.” *Phys. Rev. Lett.* 13, pp. 508-509, October 1964.
- [5] F. Englert, R. Brout, “Broken Symmetry and the Mass of Gauge Vector Mesons.” *Phys. Rev. Lett.* 13, pp. 321-323, October 1964.
- [6] The CMS Collaboration, “Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC.” *Phys. Lett. B* 716, 30-61, 2012.
- [7] The ATLAS Collaboration, “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC.” *Phys.Lett. B* 716, 1-29, 2012.
- [8] Edvige Corbelli, Paolo Salucci, “The extended rotation curve and the dark matter halo of M33.” *Monthly Notices of the Royal Astronomical Society* 311(2):441-447, 2000.
- [9] P. A. R. Ade et al. (Planck Collaboration), “Planck 2015 results: XIII. Cosmological parameters.” *A& A* 594, A13, 2016.
- [10] Steven Weinberg, “Implications of dynamical symmetry breaking.” *Phys. Rev. D* 13, 974, 1976.
- [11] Eldad Gildener, “Gauge-symmetry hierarchies.” *Phys. Rev. D* 14, 1667, 1976.
- [12] H.P. Nilles, “Supersymmetry, supergravity and particle physics.” *Phys. Rept.*, vol. 110, pp. 1–162, 1984.
- [13] Christoph Englert et al., “Exploring the Higgs portal.” *Phys.Lett.B* 703, pp. 298-305, 2011.
- [14] D. N. Spergel, “The dark side of cosmology: Dark matter and dark energy.” *Science*, vol. 347, pp. 1100–1102, 2015.
- [15] Chamberlain, Owen and Segrè, Emilio and Wiegand, Clyde and Ypsilantis, Thomas, “Observation of Antiprotons.” *Phys. Rev.* 100, 947, 1955.

- 
- [16] M. Hori et al. (ASACUSA Collaboration), “Buffer-gas cooling of antiprotonic helium to 1.5 to 1.7 K, and antiproton-to-electron mass ratio.” *Science* 354; 6312, pp. 610–614, 2016.
- [17] Borchert, M.J., Devlin, J.A., Erlewein, S.R. et al.), “A 16-parts-per-trillion measurement of the antiproton-to-proton charge–mass ratio.” *Nature* 601, 53–57, 2022.
- [18] Nagahama, H., Smorra, C., Sellner, S. et al. (BASE Collaboration), “Sixfold improved single particle measurement of the magnetic moment of the antiproton.” *Nat Commun* 8, 14084, 2017.
- [19] S. Sellner et al. (BASE Collaboration), “Improved limit on the directly measured antiproton lifetime.” *New J. Phys.* 19 083023, 2017.
- [20] Stephen H. Geer and Dallas C. Kennedy, “A New Limit on the Antiproton Lifetime.” *ApJ* 532 648, 2000.
- [21] S. Geer, J. Marriner, M. Martens et al. (APEX Collaboration), “New Limit on CPT Violation.” *Phys. Rev. Lett.* 84, 590, 2000.
- [22] G. Lüders, “On the Equivalence of Invariance under Time Reversal and under Particle-Antiparticle Conjugation for Relativistic Field Theories.” *Kong.Dan.Vid.Sel.Mat.Fys.Med.* 28N5 5, 1-17, 1954.
- [23] G. Lüders, “Proof of the TCP theorem.” *Annals Phys.* 2 1-15, 1957.
- [24] L. Canetti, M. Drewes, and M. Shaposhnikov, “Matter and antimatter in the universe.” *New J. Phys.* 14.095012, 2012.
- [25] Tanabashi, M. et al. (Particle Data Group), “Review of Particle Physics.” *Phys.Rev.D* 98, 3, 030001, 2018.
- [26] Latham Boyle, Kieran Finn, Neil Turok, “CPT-Symmetric Universe.” *Phys. Rev. Lett.* 121, 251301, 2018.
- [27] H. Alfvén, “Antimatter and the Development of the Metagalaxy.” *Rev. Mod. Phys.* 37, 652, 1965.
- [28] Kowitt, Mark E., “Gravitational repulsion and Dirac antimatter.” *Int J Theor Phys* 35, 605–631, 1996.
- [29] M. Villata, “CPT symmetry and antimatter gravity in general relativity.” *EPL* 94 20001, 2011.
- [30] M.J.T.F. Cabbolet, “Elementary Process Theory: a formal axiomatic system with a potential application as a foundational framework for physics supporting gravitational repulsion of matter and antimatter.” *Annalen der Physik* 522(10):699 - 738, 2010.
- [31] M. M. Nieto and T. Goldman, “The arguments against "antigravity" and the gravitational acceleration of antimatter.” *Phys. Rep.* 205, pp. 221–281, 1991.
- [32] Lehnert, B., “Problems of matter-antimatter boundary layers.” *Astrophys Space Sci* 46, 61–71, 1977.

- [33] M. M. Nieto and T. Goldman, “The arguments against “antigravity” and the gravitational acceleration of antimatter.” *Phys. Rep.* 205, pp.221, 1991.
- [34] A.D. Sakharov, “Violation of CP Invariance, C asymmetry, and baryon asymmetry of the universe.” *Pisma Zh. Eksp. Teor. Fiz.* 5, 32-35, 1967.
- [35] J. H. Christenson, J. W. Cronin, V. L. Fitch, and R. Turlay, “Evidence for the  $2\pi$  Decay of the  $K_2^0$  Meson.” *Phys. Rev. Lett.* 13, 138, 1964.
- [36] Y. Amhis et al. (Heavy Flavor Averaging Group Collaboration), “Averages of  $b$ -hadron,  $c$ -hadron, and  $\tau$ -lepton properties as of 2021.” *Phys. Rev. D* 107, 052008, 2023.
- [37] G. Gabrielse, A. Khabbaz, D. S. Hall, C. Heimann, H. Kalinowsky, and W. Jhe, “Precision Mass Spectroscopy of the Antiproton and Proton Using Simultaneously Trapped Particles.” *Phys. Rev. Lett.* 82, 3198, 1999.
- [38] M. Ahmadi et al., “Observation of the 1S–2S transition in trapped antihydrogen.” *Nature* Vol. 541 506-512, 2017.
- [39] P.A. Zyla et al. (Particle Data Group), “Review of Particle Physics.” *Prog. Theor. Exp. Phys.* 2020, 083C01, 2020.
- [40] Kostelecký, V. Alan and Russell, Neil, “Data tables for Lorentz and  $CPT$  violation.” *Rev. Mod. Phys.* 83, 11, 2011, last updated 2021.
- [41] Sebastian Bruggisser et al., “CP-violation for electroweak baryogenesis from dynamical CKM matrix.” *JCAP*11 034, 2017.
- [42] T. Cohen, D.E. Morrissey, A. Pierce, “Electroweak baryogenesis and Higgs signatures.” *Phys. Rev. D* 86, 013009, 2012.
- [43] D.E. Morrissey and M.J. Ramsey-Musolf, “Electroweak baryogenesis.” *New J. Phys.* 14 125003, 2012.
- [44] Maxim Pospelov, Adam Ritz, “Electric dipole moments as probes of new physics.” *Annals Phys.* 318: 119-169, 2005.
- [45] Yohei Ema, Ting Gao, and Maxim Pospelov, “Standard Model Prediction for Paramagnetic Electric Dipole Moments.” *Phys. Rev. Lett.* 129, 231801, 2022.
- [46] T.E. Chupp, P. Fierlinger, M.J. Ramsey-Musolf, and J.T. Singh, “Electric dipole moments of atoms, molecules, nuclei, and particles.” *Rev. Mod. Phys.* 91, 015001, 2019.
- [47] Tanya S. Roussy et al., “An improved bound on the electrons electric dipole moment.” *Science* 381, 6653, 2023.
- [48] C. Abel et al. (nEDM Collaboration), “Measurement of the Permanent Electric Dipole Moment of the Neutron.” *Phys. Rev. Lett.* 124, 081803, 2020.
- [49] Galilei, Galileo, “Discorsi e dimostrazioni matematiche intorno a due nuove scienze attenenti alla meccanica ed i movimenti locali.” *Gli Elsevirii*, 1638.

- [50] Newton, Isaac, “Philosophiae naturalis principia mathematica.” J. Societatis Regiae ac Typis J. Streater, 1687.
- [51] Einstein, Albert, “Über das Relativitätsprinzip und die aus demselben gezogenen Folgerungen.” Jahrbuch der Radioaktivität 4, pp. 411-462, 1907.
- [52] Roland v. Eötvös, Desiderius Pekár, Eugen Fekete, “Beiträge zum Gesetze der Proportionalität von Trägheit und Gravität.” Annalen der Physik, vol. 373, Issue 9, pp. 11-66, 1922.
- [53] C. M. Will, “The confrontation between general relativity and experiment.” Living Rev. Relativ. 17, 4, 2014.
- [54] T. A. Wagner, S. Schlamminger, J. H. Gundlach, E. G. Adelberger, “Torsion-balance tests of the weak equivalence principle.” Class. Quantum Grav. 29 184002, 2012.
- [55] Pierre Touboul et al. (MICROSCOPE Collaboration), “MICROSCOPE Mission: Final results of the test of the Equivalence Principle.” Phys. Rev. Lett. 129, 121102, 2022.
- [56] G. Gabrielse, A. Khabbaz, D. S. Hall, C. Heimann, H. Kalinowsky, and W. Jhe, “Precision Mass Spectroscopy of the Antiproton and Proton Using Simultaneously Trapped Particles.” Phys. Rev. Lett. 82, 3198, 1999.
- [57] Richard J. Hughes and Michael H. Holzschteier, “Constraints on the gravitational properties of antiprotons and positrons from cyclotron-frequency measurements.” Phys. Rev. Lett. 66, 854, 1991.
- [58] G. Chardin and G. Manfredi, “Gravity, antimatter and the Dirac–Milne universe.” Hyperfine Interact. 239, 45, 2018.
- [59] F. Witteborn and W. Fairbank, “Experiments to determine the Force of Gravity on Single Electrons and Positrons.” Nature, 220:436-440, 1968.
- [60] Nelson Jarmie, “A measurement of the gravitational acceleration of the antiproton: An experimental overview.” Nucl.Instrum.Meth. B24/25, pp. 437-441, 1987.
- [61] Anderson, E.K., Baker, C.J., Bertsche, W. et al., “Observation of the effect of gravity on the motion of antimatter.” Nature 621, 716–722, 2023.
- [62] Martin Deutsch, “Evidence for the Formation of Positronium in Gases.” Phys. Rev. 82, 455, 1951.
- [63] Y. Kataoka, S. Asai, and T. Kobayashi, “First test of  $O(\alpha^2)$  correction of the orthopositronium decay rate.” Phys. Lett. B, 32:219-223, 2009.
- [64] Ahmadi, M., Alves, B.X.R., Baker, C.J. et al., “Characterization of the 1S–2S transition in antihydrogen.” Nature 557, 71–75, 2018.
- [65] Ahmadi, M., Alves, B.X.R., Baker, C.J. et al., “Observation of the 1S–2P Lyman- $\alpha$  transition in antihydrogen.” Nature 561, 211–215, 2018.
- [66] The ALPHA Collaboration, “Investigation of the fine structure of antihydrogen.” Nature 578, 375–380, 2020.

- [67] Ahmadi, M., Baquero-Ruiz, M., Bertsche, W. et al., “An improved limit on the charge of antihydrogen from stochastic acceleration.” *Nature* 529, 373–376, 2016.
- [68] Amsler, C., Antonello, M., Belov, A. et al. (AEGIS collaboration), “Pulsed production of antihydrogen.” *Commun. Phys.* 4, 19, 2021.
- [69] Adrich, P., Blumer, P., Caratsch, G. et al., “Production of antihydrogen atoms by 6 keV antiprotons through a positronium cloud.” *Eur. Phys. J. C* 83, 1004, 2023.
- [70] G.Baur et al., “Production of antihydrogen.” *Physics Letters B*, 368 (3), 251-258, 1996.
- [71] G. Blanford et al., “Observation of Atomic Antihydrogen.” *Phys. Rev. Lett.*, 80 (14), 3037-3040, 1998.
- [72] Amoretti, M., Amsler, C., Bonomi, G. et al., “Production and detection of cold antihydrogen atoms.” *Nature* 419, 456–459, 2002.
- [73] G. Gabrielse et al. (ATRAP Collaboration), “Background-Free Observation of Cold Antihydrogen with Field-Ionization Analysis of Its States.” *Phys. Rev. Lett.* 89, 213401, 2002.
- [74] Y. Enomoto et al., “Synthesis of Cold Antihydrogen in a Cusp Trap.” *Phys. Rev. Lett.* 105, 243401, 2010.
- [75] G. B. Andersen et al., “Production of antihydrogen at reduced magnetic field for anti-atom trapping.” *J. Phys. B: At. Mol. Opt. Phys.* 41 011001, 2008.
- [76] C. H. Storry et al., “First laser-controlled antihydrogen production.” *Phys. Rev. Lett.* 93, 263401, 2004.
- [77] D. Krasnický, R. Caravita, C. Canali, and G. Testera, “Cross section for Rydberg antihydrogen production via charge exchange between Rydberg positroniums and antiprotons in a magnetic field.” *Phys. Rev. A* 94, 022714, 2016.
- [78] N. Zurlo et al., “Production of slow protonium in vacuum.” *Hyperfine Interact.* 172:97–105, 2006.
- [79] L. Venturelli et al. (ATHENA collaboration), “Protonium production in ATHENA.” *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 261 (1–2): 40–43, 2007.
- [80] M. Iwasaki et al., “Discovery of antiproton trapping by long-lived metastable states in liquid helium.” *Phys. Rev. Lett.* 67; 1246, 1991.
- [81] T. Yamazaki et al., “Formation of long-lived gas-phase antiprotonic helium atoms and quenching by H<sub>2</sub>.” *Nature*, volume 361, pp. 238–240, 1993.
- [82] S. Huck et al. (AEGIS Collaboration), “Toward a pulsed antihydrogen beam for WEP tests in AEGIS.” *EPJ Web of Conferences* 282, 01005, 2023.
- [83] Doser, M. et al. (AEGIS Collaboration), “AEGIS / AD-6 annual report 2022.” CERN-SPSC-2022-004, SPSC-SR-303, 2022.

- [84] Caravita, R. et al. (AEgIS Collaboration), “AEgIS / AD-6 annual report 2023.” CERN-SPSC-2023-003, SPSC-SR-321, 2023.
- [85] Caravita, R. et al. (AEgIS Collaboration), “AEgIS / AD-6 annual report 2024.” CERN-SPSC-2024-002, SPSC-SR-339, 2024.
- [86] N. Gusakova, A. Camper, R. Caravita, L. Penasa, L. T. Glöggler, T. Wolz, V. Krumins, F. P. Gustafsson, S. Huck, M. Volponi, B. Rienaecker et al., “An alexandrite laser system for positronium laser cooling.” *Opt. Laser Technol.* 182, 112097, 2025.
- [87] F. Guatieri, M. Berghold, D. Orsucci et al., “Real-time antiproton annihilation vertexing with sub-micron resolution.” Under review, <https://doi.org/10.48550/arXiv.2406.16044>, 2025.
- [88] M. Doser et al., “Exploring the WEP with a pulsed cold beam of antihydrogen.” *Class. Quantum Grav.* 29 184009, 2012.
- [89] Alanzeau, C. et al., “Extra Low ENergy Antiproton (ELENA) ring and its Transfer Lines: Design Report.” CERN-2014-002, 10.5170/CERN-2014-002, 2014.
- [90] L. Ponce et al., “ELENA - From Commissioning to Operation.” JACoW IPAC, 2391-2394, 2022.
- [91] Lopienska, Ewa, “The CERN accelerator complex layout in 2022.” CERN-GRAPHICS-2022-001, <https://cds.cern.ch/record/2800984>, 2022.
- [92] Einstein, Albert, “Ist die Trägheit eines Körpers von seinem Energiegehalt abhängig?” *Annalen der Physik* 18:639, 1905.
- [93] D. Möhl, “Production of low-energy antiprotons.” *Hyperfine Interactions* 109, pp. 33-41, 1997.
- [94] D. Boimond, M. Frauchiger, T. Kurtyka, M. Lubrano di Scampamorte, R. Maccaferri, S. Maury, Nikitina L., and J. C. Schnuriger, “Consolidation of the 400 kA magnetic horn for AAC antiproton production.” CERN/PS 94-02 (AR), 1994.
- [95] Liouville, Joseph, “Sur la Theorie de la Variation des constantes arbitraires.” *Journal de mathématiques pures et appliquées* 3: 342–349, 1838.
- [96] J. W. Gibbs, “On the Fundamental Formula of Statistical Mechanics, with Applications to Astronomy and Thermodynamics.” *Proceedings of the American Association for the Advancement of Science*, 33, 57–58, 1884.
- [97] W. Oelert et al. (ELENA Collaboration), “The ELENA Project at CERN.” *Acta Phys. Pol. B* 46, 181, 2015.
- [98] N. Kuroda, H. A. Torii et al., “Confinement of a Large Number of Antiprotons and Production of an Ultraslow Antiproton Beam.” *Phys. Rev. Lett.* 94, 023401, 2005.
- [99] Martina Knoop, Niels Madsen, Richard C. Thompson, “Trapped Charged Particles.” World Scientific Publishing Europe Ltd., 2016.
- [100] F. Penning, “Die Glimmentladung bei niedrigem Druck zwischen koaxialen Zylindern in einem axialen Magnetfeld.” *Physica* 3(9), 873-894, 1936.

- [101] H. Dehmelt, "Radiofrequency spectroscopy of stored ions I: Storage." *Advances in Atomic and Molecular Physics*, Vol. 3, pp. 53-72. New York, Academic Press, 1968.
- [102] Daniel H. E. Dubin and T. M. O'Neil, "Trapped nonneutral plasmas, liquids, and crystals (the thermal equilibrium states)." *Rev. Mod. Phys.* 71, 87, 1999.
- [103] Brown, Lowell S. and Gabrielse, Gerald, "Geonium theory: Physics of a single electron or ion in a Penning trap." *American Physical Society, Rev. Mod. Phys.*, 58, 1, pp. 233-311, 1986.
- [104] W. Paul et al., "Ein Ionenkäfig." *Forschungsberichte des Wirtschafts- und Verkehrsministeriums Nordrhein-Westfalen*, Westdeutscher Verlag Köln-Opladen. 415, 1-67, 1958.
- [105] R. G. Greaves and C. M. Surko, "Inward Transport and Compression of a Positron Plasma by a Rotating Electric Field." *Phys. Rev. Lett.* 85, 1883, 2000.
- [106] R. G. Greaves and C. M. Surko, "Radial compression and inward transport of positron plasmas using a rotating electric field." *Phys. Plasmas*, 8(5):1879-1885, 2001.
- [107] R. G. Greaves and C. M. Surko, "Emerging science and technology of antimatter plasmas and trap-based beams." *Phys. Plasmas*, 5(11):2333, 2004.
- [108] Aghion, S., Amsler, C., Bonomi, G. et al., "Compression of a mixed antiproton and electron non-neutral plasma to high densities." *Eur. Phys. J. D* 72, 76, 2018.
- [109] G. Drobychev et al. (AEGIS collaboration), "Proposal for the AEGIS experiment at the CERN antiproton decelerator (Antimatter Experiment: Gravity, Interferometry, Spectroscopy)." CERN-SPSC-2007-017; SPSC-P-334, 2007.
- [110] E. A. Hessels, D. M. Homan, and M. J. Cavagnero, "Two-stage Rydberg charge exchange: An efficient method for production of antihydrogen." *Phys. Rev. A* 57, 1668, 1998.
- [111] F. Castelli, I. Boscolo, S. Cialdi, M. G. Giammarchi, and D. Comparat, "Efficient positronium laser excitation for antihydrogen production in a magnetic field." *Phys. Rev. A* 78, 052512, 2008.
- [112] Ruggero Caravita, "Towards measuring gravity on neutral antimatter." PhD thesis, 2017.
- [113] M. Antonello et al. (AEGIS Collaboration), "Rydberg positronium velocity and self-ionization studies in a 1T magnetic field and cryogenic environment." *Phys. Rev. A* 102.013101, 2020.
- [114] T. F. Gallagher, "Rydberg atoms." *Rep. Prog. Phys.* 51 143, 1988.
- [115] Kellerbauer, A. et al. (AEGIS collaboration), "Proposed antimatter gravity measurement with an antihydrogen beam." *Nucl. Instrum. Meth. B* volume 266, pp. 351-356, 2008.
- [116] L. T. Glöggler, N. Gusakova, B. Rienäcker, A. Camper, R. Caravita, S. Huck, M. Volponi, T. Wolz, L. Penasa, V. Krumins, F. P. Gustafsson, D. Comparat et al.,

- “Positronium laser cooling via the  $1^3S-2^3P$  transition with a broadband laser pulse.” *Phys. Rev. Lett.*, 2024.
- [117] Dudarev, A. et al., “Design of a Superconducting Magnet System for the AEGIS Experiment at CERN.” *IEEE Trans. Appl. Supercond.* 21, 1721-1724, 2011.
- [118] J.M. Lafferty, “Foundations of Vacuum Science and Technology.” John Wiley & Sons, Inc. (US), 1998.
- [119] Vincent Baglin, “Cryopumping and Vacuum Systems.” Proceedings of the CERN Accelerator School course on Vacuum for Particle Accelerators, Glumslöv, (Sweden), 2017.
- [120] P. J. Förster, P. Mascher, A. P. Knights, and P. G. Coleman, “Implantation profile of  $^{22}\text{Na}$  continuous energy spectrum positrons in silicon.” *J. Appl. Phys.*, 101:043702, 2007.
- [121] Benjamin Rienäcker, “Creation and manipulation of positronium for efficient anti-hydrogen production at AEGIS.” PhD thesis, 2021.
- [122] S. Aghion et al., “Positron bunching and electrostatic transport system for the production and emission of dense positronium clouds in vacuum.” *Nucl. Instrum. Methods Phys. Res., B* 362, 86-92, 2015.
- [123] S. Mariazzi, P. Bettotti, and R. S. Brusa, “Positronium cooling and emission in vacuum from nanochannels at cryogenic temperature.” *Phys. Rev. Lett.* 104.243401, 2010.
- [124] S. Mariazzi, P. Bettotti et al., “High positronium yield and emission into the vacuum from oxidized tunable nanochannels in silicon.” *Phys. Rev. B* 81.235418, 2010.
- [125] P. J. Schultz and K. G. Lynn, “Interaction of positron beams with surfaces, thin films and interfaces.” *Rev. Mod. Phys.*, 60(3):701-779, 1988.
- [126] E. Lazzarini, “A comprehensive model for positronium formation and related phenomena of its inhibition and enhancement.” *Radiat. Phys. Chem.*, 28:49-54, 1986.
- [127] S. Van Petegem, C. Dauwe et al., “Diffusion length of positrons and positronium investigated using a positron beam with longitudinal geometry.” *Phys. Rev. B*, 70:115410, 2004.
- [128] Y. Nagashima, Y. Morinaka et al., “Origins of positronium emitted from  $\text{SiO}_2$ .” *Phys. Rev. B*, 58:12676, 1998.
- [129] S. Mariazzi et al., “High-yield thermalized positronium at room temperature emitted by morphologically tuned nanochanneled silicon targets.” *J.Phys. B: At. Mol. Opt. Phys.* 54.085004, 2021.
- [130] S. Cialdi et al., “Efficient two-step Positronium laser excitation to Rydberg levels.” *Nucl. Instrum. Methods Phys. Res. B*, 269(13):1527 – 1533, 2011.
- [131] S. Aghion et al. (AEGIS Collaboration), “Laser excitation of the  $n = 3$  level of positronium for antihydrogen production.” *Phys. Rev. A* 94, 012507, 2016.

- [132] M. Volponi, J. Zieliński, T. Rauschendorfer, S. Huck, R. Caravita et al., “TALOS (Total Automation of Labview Operations for Science): a framework for autonomous control systems for complex experiments.” *Rev. Sci. Instrum.* 95, 085116, 2024.
- [133] K. L. Brown and G. W. Tautfest, “Faraday-Cup Monitors for High-Energy Electron Beams.” *Rev. Sci. Instrum.* 27, 696–702, 1956.
- [134] N. Zurlo, G. Bonomi, A. Fontana, A. Rotondi, “New Calibration of the PMT-Scintillators System.” AEGIS internal note: 2017/1, 2017.
- [135] N. Zurlo et al., “Calibration and Equalisation of Plastic Scintillator Detectors for Antiproton Annihilation Identification Over Positron/Positronium Background.” *Acta Phys. Pol. B* 51, 213-223, 2020.
- [136] Markus K. Oberthaler et al., “Inertial sensing with classical atomic beams.” *Phys. Rev. A* 54, 3165, 1996.
- [137] Aghion, S., Ahlén, O., Amsler, C. et al., “A moiré deflectometer for antimatter.” *Nat Commun* 5, 4538, 2014.
- [138] Philippe H. M. Bräunig, “Atom Optical Tools for Antimatter Experiments.” PhD Thesis, 2014.
- [139] Andrea Demetrio, “Feasibility of a gravity measurement on antimatter using a Talbot-Lau interferometer.” PhD Thesis, 2018.
- [140] M. Kimura, M. et al., “Development of nuclear emulsions with 1  $\mu\text{m}$  spatial resolution for the AEGIS experiment.” *Nucl. Instr. Meth. A* 732, 2013.
- [141] Tim H. Wolz, “Controlled manipulation of atoms in Rydberg quantum states for application in experiments with antihydrogen.” PhD Thesis, 2021.
- [142] J. Rodewald, N. Dorre, et al., “Isotope-selective high-order interferometry with large organic molecules in free-fall.” *New Journal of Physics*, 033016, 2018.
- [143] V. Krumins et al., “OTIMA interferometry for an antihydrogen gravity measurement.” Work in progress, 2024.
- [144] Klaus Hornberger, John E. Sipe, and Markus Arndt, “Theory of decoherence in a matter wave Talbot-Lau interferometer.” *Phys. Rev. A* 70, 053608, 2004.
- [145] B.M. Latacz et al. (BASE Collaboration), “Orders of Magnitude Improved Cyclotron-Mode Cooling for Nondestructive Spin Quantum Transition Spectroscopy with Single Trapped Antiprotons.” *Phys. Rev. Lett.* 133, 053201, 2024.
- [146] Gabrielse, G. et al., “Cooling and slowing of trapped antiprotons below 100 meV.” *Phys. Rev. Lett.* 63.1360, 1989.
- [147] Rolston, S.L., Gabrielse, G., “Cooling antiprotons in an ion trap.” *Hyperfine Interact* 44, 233–245, 1989.
- [148] G. Cerchiari, P. Yzombard, A. Kellerbauer, “Laser-assisted evaporative cooling of anions.” *Phys. Rev. Lett.* 123, 103201, 2019.

- [149] P. Yzombard, M. Hamamda, S. Gerber, M. Doser and D. Comparat, “Laser Cooling of Molecular Anions.” *Phys. Rev. Lett.* 114, 213001, 2015.
- [150] J. Storey et al., “Particle tracking at 4K: the Fast Annihilation Cryogenic Tracking (FACT) detector for the AEgIS antimatter gravity experiment.” *Nucl. Instrum. Methods Phys. Res., Sect. A*, 732:437–441, 2013.
- [151] Michael Berghold, Vassily Vadimovitch Burwitz, Lucian Mathes, Christoph Hugen-schmidt, Francesco Guatieri, “Imaging low-energy positron beams in real-time with unprecedented resolution.” *Sci Rep* 13, 18526, 2023.
- [152] L. Glögger, R. Caravita et al., “High-resolution MCP-TimePix3 imaging/timing detector for antimatter physics.” *Meas. Sci. Technol.* 33, 115105, 2022.
- [153] Volponi, M., Huck, S., Caravita, R. et al., “CIRCUS: an autonomous control system for antimatter, atomic and quantum physics experiments.” *EPJ Quantum Technol.* 11, 10, 2024.
- [154] Grzegorz Kasprowicz et al., “ARTIQ and Sinara: Open Software and Hardware Stacks for Quantum Physics.” *Optical Society of America, OSA Quantum 2.0 Conference, QTu8B.14*, 2020.
- [155] Bourdeauducq, Sébastien et al., “ARTIQ 1.0.” *Zenodo*, <https://doi.org/10.5281/zenodo.51303>, 2016.
- [156] M. Volponi, “Towards the formation of a forward-boosted beam of neutral antihydrogen for gravitational studies.” PhD thesis, Università di Trento, Department of Physics, 2024.
- [157] A.J. Peters and L. Janyst, “Exabyte Scale Storage at CERN.” *J. Phys.: Conf. Ser.* 331 052015, 2011.
- [158] Wayne M. Saslow, “Electricity, Magnetism, and Light. Chapter 10: How Electric Currents Interact with Magnetic Fields.” *Academic Press, San Diego*, ISBN: 978-0-12-619455-5, 2002.
- [159] M. Vogel, “Magnetic Bottles as Implemented in Penning Traps.” *Particle Confinement in Penning Traps*, pp. 319-334, 2018.
- [160] K.-U. Kuehnel, C.D. Schröter and J. Ullrich, “Operating MCP detectors at cryogenic temperatures.” *Conf.Proc.C 0806233 TUPC055*, 2008.
- [161] E. K. Anderson et al., “Improving detection efficiency in a cryogenic environment - implications for DESIREE.” *Journal of Physics: Conference Series* 635 022039, 2015.
- [162] S. Huck, T. Rauschendorfer, J. Zieliński, M. Volponi, R. Caravita et al., “New perspectives from record-high numbers of trapped cold antiprotons.” *Work in progress*, 2024.
- [163] T. Rauschendorfer, S. Huck, J. Zieliński, M. Volponi, R. Caravita et al., “Highly efficient antiproton capture from the ELENA decelerator.” *Work in progress*, 2024.
- [164] F. Pedregosa, G. Varoquaux et al., “Scikit-learn: Machine Learning in Python.” *JMLR* 12, 2825–2830, 2011.

- [165] J. Lindhard, “On the Properties of a Gas of Charged Particles.” Kgl. Danske Videnskab. Selskab, Mat.-Fys. Medd. 28, No. 8, 1954.
- [166] J. Lindhard, M. Scharff, and H.E. Schiøtt, “Range concepts and heavy ion ranges (Notes on atomic collisions, II).” Kgl. Danske Videnskab. Selskab, Mat.-Fys. Medd. 33, No. 14, 1963.
- [167] C. Amsler et al. (ASACUSA Collaboration), “Injection and capture of antiprotons in a Penning–Malmberg trap using a drift tube accelerator and degrader foil.” Nucl. Instr. and Meth. A 1065, 169529, 2024.
- [168] D.L. Eggleston et al., “Parallel energy analyzer for pure electron plasma devices.” Phys. Fluids B 4, 3432–3439, 1992.
- [169] Ingmari Christa Tietje, “Characterisation and Control of Antiproton-Electron Plasmas for Pulsed Antihydrogen Production.” PhD Thesis, Technische Universität Berlin, 2021.
- [170] Bret Robert Beck, “Measurement of the Magnetic and Temperature Dependence of the Electron-Electron Anisotropic Temperature Relaxation Rate.” PhD Thesis, University of California, San Diego, 1990.
- [171] M. Doser, “Antiprotonic bound systems.” Prog. Part. Nucl. Phys. 125, 103964, 2022.
- [172] Sebastian Gerber, Michael Doser and Daniel Comparat, “Pulsed production of cold protonium in Penning traps.” Phys. Rev. A 100, 063418, 2019.
- [173] M. Doser, G. Farrar and G. Kornakov, “Searching for a dark matter particle with anti-protonic atoms.” Eur. Phys. J. C 83, 1149, 2023.
- [174] G. Kornakov, G. Cerchiari, J. Zieliński, L. Lappo, G. Sadowski, and M. Doser, “Synthesis of cold and trappable fully stripped highly charged ions via antiproton-induced nuclear fragmentation in traps.” Phys. Rev. C 107, 034314, 2023.
- [175] T. Aumann, W. Bartmann, O. Boine-Frankenheim et al., “PUMA, antiProton unstable matter annihilation.” Eur. Phys. J. A 58, 88, 2022.
- [176] C. Smorra, F. Abbass, D. Schweitzer et al., “BASE-STEP: A transportable antiproton reservoir for fundamental interaction studies.” Rev. Sci. Instrum. 94, 113201, 2023.
- [177] M.J.G. Borge, B. Jonson, “ISOLDE past, present and future.” Phys. G: Nucl. Part. Phys. 44, 044011, 2017.
- [178] Kolbinger, B., Amsler, C., Cuendis, S.A. et al., “Measurement of the principal quantum number distribution in a beam of antihydrogen atoms.” Eur. Phys. J. D 75, 91, 2021.
- [179] M. Fanì et al. (AEGIS collaboration), “Developments for pulsed antihydrogen production towards direct gravitational measurement on antimatter.” Phys. Scr. 95 114001, 2020.
- [180] Ti-pei Li and Yu-qian Ma, “Analysis methods for results in gamma-ray astronomy.” Astrophysical Journal, Vol. 272, p. 317-324, 1983.

- [181] Cassidy, D.B., “Experimental progress in positronium laser physics.” *Eur. Phys. J. D* 72, 53, 2018.
- [182] Adkins, G. S., Cassidy, D. B., Pérez-Ríos, J., “Precision spectroscopy of positronium: Testing bound-state QED theory and the search for physics beyond the Standard Model.” *Phys.Rept.* 975, 1-61, 2022.
- [183] H.K. Avetissian, A.K. Avetissian, G.F. Mkrtchian, “Self-Amplified Gamma-Ray Laser on Positronium Atoms from a Bose-Einstein Condensate.” *Phys. Rev. Lett.* 113, 023904, 2014.
- [184] C. Doppler, “Über das farbige Licht der Doppelsterne und einiger anderer Gestirne des Himmels.” *Abhandlungen der k. böhm. Gesellschaft der Wissenschaften*, Bd. 2, 465-482, 1842.
- [185] T.W. Hänsch, A.L. Schawlow, “Cooling of gases by laser radiation.” *Opt. Commun.* 13, 68-69, 1975.
- [186] W. Neuhauser, M. Hohenstatt, P. Toschek, and H. Dehmelt, “Optical-Sideband Cooling of Visible Atom Cloud Confined in Parabolic Well.” *Phys. Rev. Lett.* 41, 233, 1978.
- [187] C. J. Foot, “Atomic Physics.” 1st ed. Oxford University Press, 2005.
- [188] D. B. Cassidy, H. W. K. Tom, and A. P. Mills, “Fundamental physics with cold positronium.” *AIP Conf. Proc.* 1037, 66, 2008.
- [189] K. P. Ziocck et al., “Optical saturation of the  $1^3S$ - $2^3P$  transition in positronium.” *J. Phys. B: At. Mol. Opt. Phys.* 23 329, 1990.
- [190] C. Zimmer, P. Yzombard, A. Camper, and D. Comparat, “Positronium laser cooling in a magnetic field.” *Phys. Rev. A* 104.023106, 2021.
- [191] D. B. Cassidy, S. H. M. Deng, H. K. M. Tanaka, and A. P. Mills Jr, “Single shot positron annihilation lifetime spectroscopy.” *Appl. Phys. Lett.* 88, 194105, 2006.
- [192] H. J. Metcalf and P. van der Straten, “Laser Cooling and Trapping.” 1st ed. Springer, 1999.
- [193] J. Dalibard and C. Cohen-Tannoudji, “Laser cooling below the Doppler limit by polarization gradients: simple theoretical models.” *J. Opt. Soc. Am. B.* 6.11, 1989.

# List of publications

In the course of the work presented here, I have significantly contributed to the following list of publications within the AEGIS collaboration. Additional publications are in the process of being finalized/published, which are referred to in the declarations at the beginning of the corresponding chapter.

- "An alexandrite laser system for positronium laser cooling", *Opt.Laser Tech.* 182 (2025) 112097
- "TALOS (Total Automation of LabVIEW Operations for Science): A framework for autonomous control systems for complex experiments", *Rev.Sci.Instrum.* 95 (2024) 8, 085116
- "CIRCUS: an autonomous control system for antimatter, atomic and quantum physics experiments", *EPJ Quant. Technol.* 11 (2024) 1, 10
- "Positronium laser cooling via the  $1^3S$ - $2^3P$  transition with a broadband laser pulse", *Phys.Rev.Lett.* 132 (2024) 8, 083402
- "Toward a pulsed antihydrogen beam for WEP tests in AEGIS", *EPJ Web Conf.* 282 (2023) 01005

# Appendix A

## List of acronyms

<b>AD</b>	Antiproton Decelerator
<b>ADC</b>	Analog to Digital Converter
<b>AEgIS</b>	Antimatter Experiment: Gravity, Interferometry, Spectroscopy
<b>AERIALIST</b>	Antimatter Experiment (or AEgIS) Realtime Integration of Artiq Libraries and Sinara Technology
<b>ALPACA</b>	All Python Analysis Code of AEgIS
<b>ALPHA</b>	Antihydrogen Laser PHysics Apparatus
<b>ARTIQ</b>	Advanced Real-Time Infrastructure for Quantum physics
<b>ASACUSA</b>	Atomic Spectroscopy And Collisions Using Slow Antiprotons
<b>ATHENA</b>	AnTiHydrogEN Apparatus
<b>ATRAP</b>	Antihydrogen TRAP experiment
<b>BASE</b>	Baryon Antibaryon Symmetry Experiment
<b>BASE-STEP</b>	Baryon Antibaryon Symmetry Experiment - Symmetry Tests in Experiments with Portable antiprotons
<b>BSM</b>	Beyond the Standard Model
<b>CERN</b>	Conseil Européen pour la Recherche Nucléaire: European Organization for Nuclear Research/ European Laboratory for Particle Physics
<b>CIRCUS</b>	Computer Interface for Reliably Controlling, in an Unsupervised manner, Scientific experiments
<b>CMB</b>	Cosmic Microwave Background
<b>CMOS</b>	Complementary Metal-Oxide Semiconductor
<b>CPT</b>	Charge, Parity, and Time reversal
<b>CPU</b>	Central Processing Unit
<b>CTMC</b>	Classical Trajectory Monte Carlo (simulation)
<b>DAC</b>	Digital to Analog Converter
<b>DAQ</b>	Data AcQuisition system
<b>DC</b>	Direct Current
<b>DDS</b>	Direct Digital Synthesis
<b>DIO</b>	Digital Input/Output
<b>DMA</b>	Direct Memory Access
<b>DRTIO</b>	Distributed Real-Time Input/Output
<b>EDM</b>	Electric Dipole Moment
<b>EEM</b>	Eurocard Extension Module
<b>ELENA</b>	Extra Low ENergy Antiproton ring
<b>EP</b>	Equivalence Principle

<b>ESDA</b>	External Scintillating Detector Array
<b>FACT</b>	Fast Annihilation Cryogenic Tracking detector
<b>FC</b>	Faraday Cup
<b>FHG</b>	Fourth Harmonic Generation
<b>FIFO</b>	First In-First Out
<b>FPGA</b>	Field Programmable Gate Array
<b>FWHM</b>	Full Width (at) Half Maximum
<b>GBAR</b>	Gravitational Behaviour of Antimatter at Rest experiment
<b>GUI</b>	Graphical User Interface
<b>HV</b>	High-Voltage
<b>IDC</b>	Insulation-Displacement Contact
<b>IR</b>	InfraRed
<b>ISOLDE</b>	Isotope Separator On Line DEvice
<b>JSON</b>	JavaScript Object Notation
<b>JTAG</b>	Joint Test Action Group
<b>KEK</b>	High Energy Accelerator Research Organization, Japanese: Kō Enerugi Kasokuki Kenkyū Kikō
<b>LEAR</b>	Low Energy Antiproton Ring
<b>LHC</b>	Large Hadron Collider
<b>LINAC</b>	LINear ACcelerator
<b>MCP</b>	Micro-Channel Plate
<b>MCX</b>	Micro CoaXial connector
<b>MOS</b>	Metal-Oxide-Semiconductor
<b>OPA</b>	Optical Parametric Amplification
<b>OPG</b>	Optical Parametric Generation
<b>OTIMA</b>	optical TIme-domain ionizing MAtter wave
<b>OVC</b>	Outer Vacuum Chamber
<b>PCB</b>	Printed Circuit Board
<b>PMT</b>	Photomultiplier Tube
<b>PS</b>	Proton Synchrotron
<b>PSB</b>	Proton Synchrotron Booster
<b>PUMA</b>	antiProton Unstable Matter Annihilation
<b>QED</b>	Quantum ElectroDynamics
<b>QDC</b>	Charge to Digital Converter
<b>QCD</b>	Quantum ChromoDynamics
<b>RC</b>	Resistor-Capacitor
<b>RF</b>	Radio Frequency
<b>rms</b>	Root Mean Square value
<b>RPC</b>	Remote Procedure Call
<b>RTIO</b>	Real-Time Input/Output
<b>RW</b>	Rotating Wall
<b>SCSI</b>	Small Computer System Interface
<b>SDRAM</b>	Synchronous Dynamic Random Access Memory
<b>SEP</b>	Strong Equivalence Principle
<b>SFG</b>	Sum Frequency Generation
<b>SFP</b>	Small Form Factor pluggable
<b>SHG</b>	Second Harmonic Generation
<b>SM</b>	Standard Model (of particle physics)
<b>SMA</b>	SubMiniature version A (connector)

<b>SPS</b>	Super Proton Synchrotron
<b>SSPALS</b>	Single-Shot Positron Annihilation Lifetime Spectroscopy
<b>TALOS</b>	Total Automation of LabVIEW™ Operations for Science
<b>THG</b>	Third Harmonic Generation
<b>TCP</b>	Transmission Control Protocol
<b>TTL</b>	Transistor–Transistor Logic
<b>UHV</b>	Ultra-High Vacuum
<b>UFF</b>	Universality of Free Fall
<b>UV</b>	Ultra-violet
<b>VBG</b>	Volume Bragg Grating
<b>VI</b>	(LabVIEW™) Virtual Instrument
<b>WEP</b>	Weak Equivalence Principle

# Appendix B

## Examples of positron-producing processes

Fig. B.1 depicts Feynman diagrams of two naturally occurring processes that produce positrons. During the  $\beta^+$  decay shown on the left, one of the up-quarks of a proton (uud) undergoes weak interaction via a  $W^+$  boson and converts into a down-quark. In this way, the proton "decays" into a neutron (udd) and a positron and an electron neutrino are produced from the intermediate  $W^+$ . The inverse process, the conversion of a neutron into a proton with the creation of an electron and an anti-electron neutrino is referred to as  $\beta^-$  decay. The process on the right side shows the pair production of an electron and a positron from an energetic photon with a "collision partner" (denoted by the large circle).

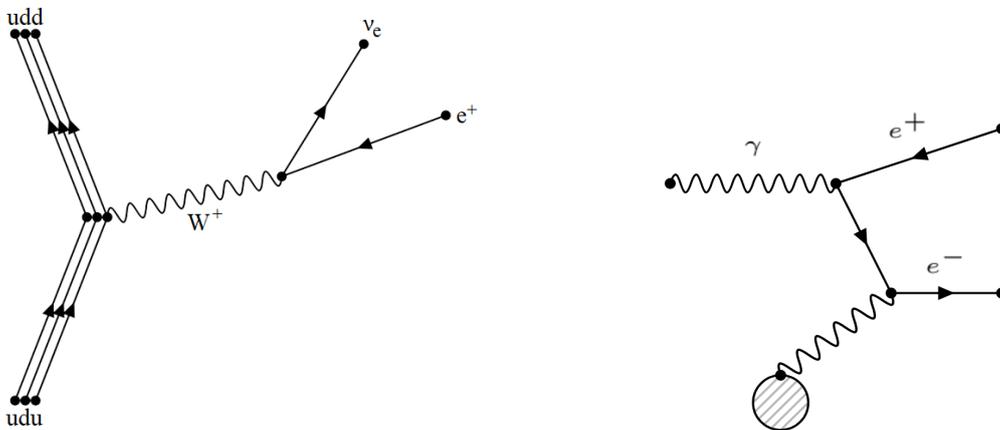


Figure B.1: Feynman diagrams of processes leading to the production of a positron. Left:  $\beta^+$  decay. Right:  $e^+e^-$  pair production.

# Appendix C

## The minimum proton energy for proton pair production

Starting from the basic relativistic energy momentum relation in Eq. C.1, one can employ four-momentum conservation laws to determine the minimum total energy of an incoming proton required to facilitate proton pair production in a static target,  $p+p \rightarrow p+p+p+\bar{p}$ <sup>1</sup>. Here,  $E_{\text{tot}}$  denotes the total energy of a system,  $p_{\text{tot}}$  the magnitude of its momentum,  $m$  the rest mass, and  $c$  the speed of light.

$$E_{\text{tot}}^2 = p_{\text{tot}}^2 c^2 + m^2 c^4 \quad (\text{C.1})$$

The situation for the two involved protons before the interaction is the following. The incoming proton has a total energy  $E_0$  and a momentum vector  $\vec{p}_0$  in three dimensions, while the target proton is considered to be at rest so that its total energy reduces to its rest mass,  $m_p c^2$ . For convenience, it is useful to analyze the system in the center of mass frame, i.e.  $p_{\text{tot}} = 0$ , with the center of mass energy,  $E_{\text{cm}}$ , related to the total energy and momentum of the system as given in Eq. C.2.

$$E_{\text{cm}}^2 = E_{\text{tot}}^2 - p_{\text{tot}}^2 c^2 \quad (\text{C.2})$$

Choosing a coordinate system such that the incoming proton is travelling in x-direction, one can utilize relation Eq. C.2 to calculate  $E_{\text{cm}}^2$  of the system before the interaction takes place, as shown in Eq. C.3.

$$E_{\text{cm}}^2 = \left( \begin{pmatrix} E_0 \\ p_0 c \\ 0 \\ 0 \end{pmatrix} + \begin{pmatrix} m_p c^2 \\ 0 \\ 0 \\ 0 \end{pmatrix} \right)^2 = \begin{pmatrix} E_0 + m_p c^2 \\ p_0 c \\ 0 \\ 0 \end{pmatrix}^2 = E_0^2 + m_p^2 c^4 + 2E_0 m_p c^2 - p_0^2 c^2 \quad (\text{C.3})$$

Since the center of mass energy of the system after the pair production is made up of the masses of the four protons and  $E_{\text{cm}}$  is a conserved quantity, applying relation Eq. C.1 solved for the squared rest mass yields Eq. C.4, relating the center of mass energy before and after the interaction.

$$2m_p^2 c^4 + 2E_0 m_p c^2 = 16m_p^2 c^4 \quad (\text{C.4})$$

Above equation can now be solved for  $E_0$  to determine the necessary minimum total energy of the incoming proton for the pair production process to be possible. Equivalently, by

---

<sup>1</sup>The reaction is equivalently possible within the field of a target nucleus neutron, whose mass differs from that of a proton by merely 0.1%.

subtracting the rest mass of the proton, one can determine the minimum required kinetic energy:

$$E_{0,\min} = 7m_p c^2 \approx 6.6 \text{ GeV} \quad (\text{C.5})$$

$$E_{\text{kin},\min} = E_{0,\min} - m_p c^2 = 6m_p c^2 \approx 5.6 \text{ GeV} \quad (\text{C.6})$$

# Appendix D

## The $E \times B$ drift in Penning-Malmberg traps

Due to the interplay of the electric field produced by a charged plasma confined in a Penning-Malmberg trap,  $E_{\text{plasma}}$ , and the radially confining magnetic field  $B$ , the particles spin around the trap axis with a rotation velocity  $v_{\text{rot}} = E_{\text{plasma,rad}}/B$  [99]. Accordingly, the plasma's rotation frequency  $f_{E_{\text{plasma}} \times B}$  is given by Eq. D.1, where  $r$  denotes the plasma radius.

$$f_{E_{\text{plasma}} \times B}(r) = \frac{v_{\text{rot}}}{2\pi r} = \frac{E_{\text{plasma,rad}}}{2\pi B r} \quad (\text{D.1})$$

Naturally, the electric field  $E_{\text{plasma}}$  satisfies Gauss' Law and can therefore be related to the charge density  $\rho(r)$  as shown in Eq. D.2, with  $q$  as the charge of the trapped particles and  $\epsilon_0$  as the vacuum permittivity constant. In cylindrical coordinates, the radial component can be expressed as given in Eq. D.3.

$$\nabla \cdot E_{\text{plasma}} = \frac{q\rho(r)}{\epsilon_0} \quad (\text{D.2})$$

$$\frac{1}{r} \frac{\partial}{\partial r} (r E_{\text{plasma,rad}}) = \frac{q\rho(r)}{\epsilon_0} \quad (\text{D.3})$$

Eq. D.3 can be multiplied by  $r$  and integrated to yield Eq. D.4. Assuming the plasma to be in thermal equilibrium allows to see it as a rigidly rotating object without shear forces, i.e. with a uniform radial density:  $n(r) = n$  [102].

$$E_{\text{plasma,rad}} = \frac{q}{r\epsilon_0} \int_0^r \rho r' dr' = \frac{q\rho r}{2\epsilon_0} \quad (\text{D.4})$$

Inserting Eq. D.4 into Eq. D.1 finally yields the relation for the  $E_{\text{plasma}} \times B$  drift frequency given below in Eq. D.5, demonstrating the linear dependence on the plasma density.

$$f_{E_{\text{plasma}} \times B} = \frac{q\rho}{4\pi\epsilon_0 B} \quad (\text{D.5})$$

# Appendix E

## List of the electrodes of the AEGIS trap system

Table E.1 lists the (high-voltage and low-voltage) electrodes which form the system of electromagnetic traps in the cryogenic apparatus of the AEGIS experiment. In addition to their reference names (which are a mostly random choice), their locations (i.e. axial distances from the upstream end of the C0 electrode) and lengths are also specified. This table contains the same information that is included in the `electrode_info.json` file of the AERIALIST control system and which serves as the basis for the determination of produced potential configurations along the trap axis.

HV electrodes			C electrodes			P electrodes		
Electrode name	Location (mm)	Length (mm)	Electrode name	Location (mm)	Length (mm)	Electrode name	Location (mm)	Length (mm)
HV1	23.5	40.0	C0	0.0	13.5	P1	533.5	23.0
HV2	483.5	40.0	C1	73.5	26.5	P2	557.5	29.0
HV3	786.5	40.0	C2	101.0	29.0	P3	587.5	12.5
HV4	1483.75	40.0	C3	131.0	29.0	P4	601.0	12.5
HV5	1500.75	40.0	C4	161.0	29.0	P5	614.5	12.5
HV6	1517.75	40.0	C5	191.0	29.0	P6	628.0	12.5
			C6	221.0	12.5	P7	641.5	12.5
			C7	234.5	12.5	P8	655.0	12.5
			C8	248.0	12.5	P9	668.5	12.5
			C9	261.5	12.5	P10	682.0	12.5
			C10	275.0	12.5	P11	695.5	12.5
			C11	288.5	12.5	P12	709.0	12.5
			C12	302.0	12.5	P13	722.5	29.0
			C13	315.5	12.5	P14	752.5	24.0
			C14	329.0	12.5			
			C15	342.5	12.5			
			C16	356.0	29.0			
			C17	386.0	29.0			
			C18	416.0	29.0			
			C19	446.0	27.5			
T electrodes			B electrodes			A electrodes		
Electrode name	Location (mm)	Length (mm)	Electrode name	Location (mm)	Length (mm)	Electrode name	Location (mm)	Length (mm)
T1	836.5	26.5	B0	1071.8	62.0	A8	1533.75	15.0
T2	864.0	39.0	B1	1134.75	42.0	A7	1549.75	15.0
T3	904.0	39.0	B2	1177.75	42.0	A6	1565.75	15.0
T4	944.0	39.0	B3	1220.75	42.0	A5	1581.75	15.0
T5	984.0	39.0	B4	1263.75	42.0	A4	1597.75	7.5
T6	1024.0	41.0	B5	1306.75	42.0	A3	1606.25	7.5
			B6	1349.75	42.0	A2	1614.75	7.5
			B7	1392.75	42.0	A1	1623.25	7.5
			B8	1435.75	15.0	A0	1636.75	12.0
			B9	1451.75	15.0			
			B10	1467.75	15.0			

Table E.1: List of the electrodes included in the AEGIS trap system. Noted are their names, their locations and their lengths. The locations refer to the position of their upstream end on the horizontal axis of the experiment with respect to that of electrode C0.

# Appendix F

## Setup of the Sinara Kasli

Fig. F.1 shows a visualization of the PCB inside one of the Sinara Kasli boards including the connectors of the front panel and the internal connections for the Eurocard extension modules. The typical connections used for the extension modules of the AEgIS core control electronics are marked in the schematic (shown for the "5TC1" crate). Depending on the number of DIO units connected per Kasli and additional extension modules (such as Urukuls), further extension connectors are used.

The JSON file describing the shown setup is sketched in Fig. F.2. The gateway image file compiled from it is flashed into the Kasli board to configure its logic gates.

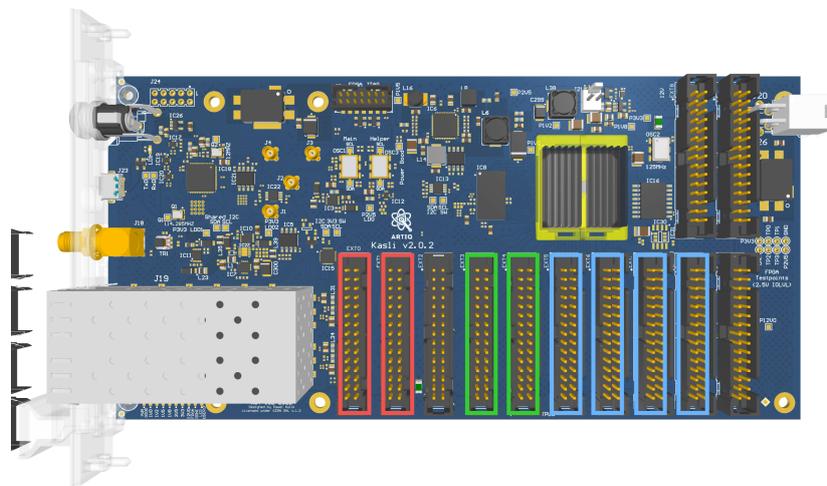


Figure F.1: Internal setup of a Sinara Kasli controller. In AEgIS, the extension modules are connected via the connectors ("EXT") marked as follows: DIO units (red, EXT 0-1), Fastino (green, EXT 3-4), amplifier boards (blue, EXT 5-7, 10).

```

{
  "target": "kasli",
  "variant": "kasli-aegis-5tc1",
  "hw_rev": "v2.0",
  "base": "master",
  "core_addr": "192.168.1.70",
  "rtio_frequency": 125e6,
  "ext_ref_frequency": 10e6,
  "peripherals": [
    {
      "type": "dio",
      "ports": [0],
      "edge_counter": true,
      "bank_direction_low": "input",
      "bank_direction_high": "input"
    },
    {
      "type": "dio",
      "ports": [1],
      "edge_counter": true,
      "bank_direction_low": "input",
      "bank_direction_high": "input"
    },
    {
      "type": "fastino",
      "ports": [3]
    },
    {
      "type": "hvamp",
      "ports": [5]
    },
    {
      "type": "hvamp",
      "ports": [6]
    },
    {
      "type": "hvamp",
      "ports": [7]
    },
    {
      "type": "hvamp",
      "ports": [10]
    }
  ]
}

```

Figure F.2: JSON file used for gateway image compilation, corresponding to the setup of a typical Kasli used in AEGIS. The Kasli is used as master device with an external 10 MHz clock and controls several extension peripherals of different types, as given in Fig. F.1.

# Appendix G

## Issues of the custom amplifier boards

A schematic of the electronics circuit of one channel of the Sinara amplifier boards, custom-designed for AEGIS with a 20-fold amplification, is shown in Fig. G.1. Eight such channels are assembled on one board.

When first examining the output produced by the amplifier boards on an oscilloscope, large oscillations of the generated signal were observed, as shown in the photograph in Fig. G.2. The oscillations had a frequency on the order of MHz, depending on the cable length, and an amplitude of approximately 5 V, independent of the set voltage, which only changed the offset. This issue has been solved by adding a second compensation capacitor to the one already used in the circuit (4.7 pF, C45A in Fig. G.1). The value of the added capacitor is 100 pF.

A further issue that is removed by the successful implementation of the amplifier boards: When using the Fastino output directly without connecting an amplifier channel, the signal produces an overshoot during a potential change before settling on the desired voltage. An example of such an overshoot in a voltage ramp is shown in Fig. G.4. The magnitude of the overshoot depends on the voltage difference to be achieved (up to 0.5 V) and can lead to instabilities and potential losses of the plasmas contained in the traps, especially if they are amplified when using the amplifier channels. However, when leading the voltage through one of the amplifier channels (with the capacitor fix implemented as described above), no more significant overshoot is observed, as shown in an example voltage ramp in Fig. G.5, and the voltage settles smoothly.

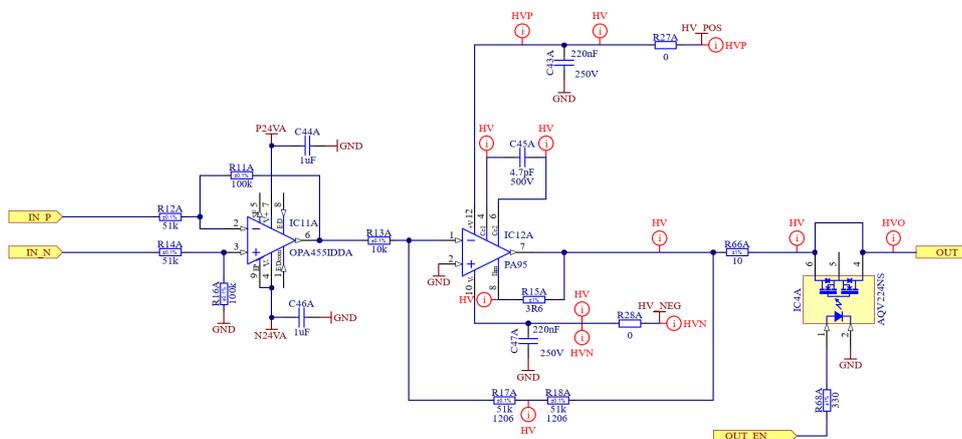


Figure G.1: Schematic of the electronic circuit of one channel of the custom-designed amplifier board.

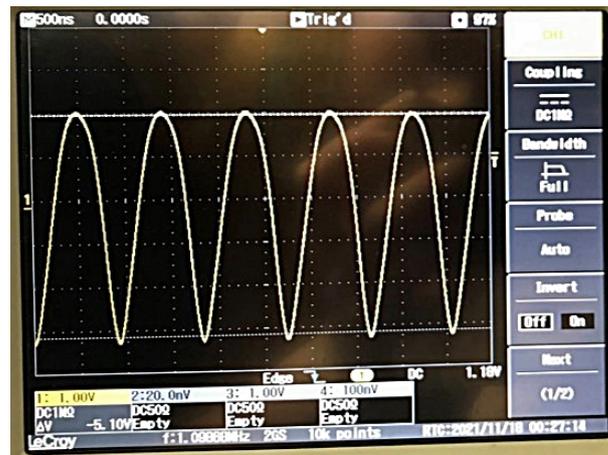


Figure G.2: Photograph of an oscilloscope showing the oscillations of the amplifier signal.

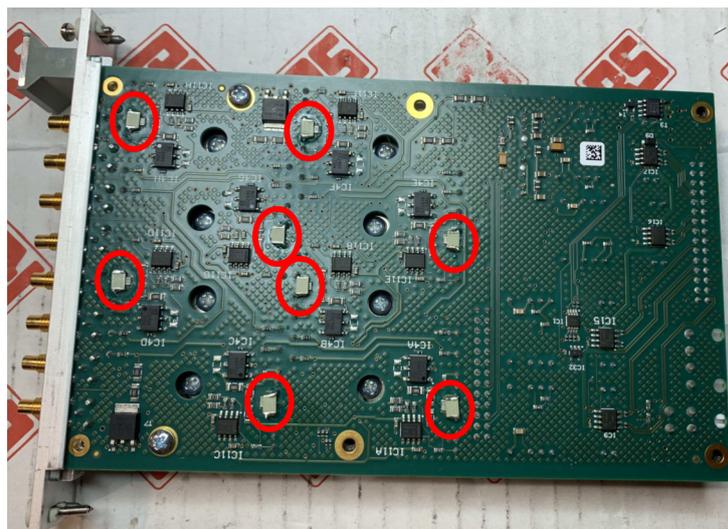


Figure G.3: Photograph of one amplifier board with the added capacitors. The locations of the corresponding capacitors (one per amplifier channel, i.e. eight per board) are indicated by the red circles.

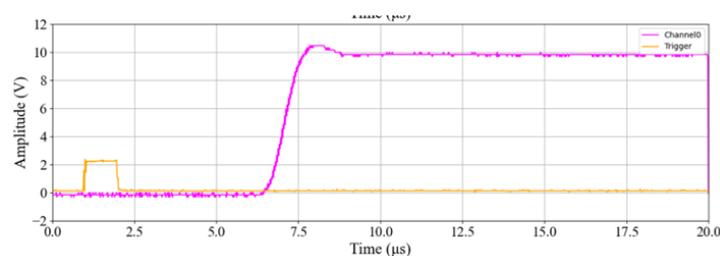


Figure G.4: Output voltage ramp from 0 to 10 V on an example Fastino channel, including a significant overshoot, as recorded by an oscilloscope.



Figure G.5: Output voltage ramp from 0 to 50 V on three example channels of Fastino and custom amplifier with no observed overshoot, as recorded by an oscilloscope (oscilloscope channels 2-4). The ramp happens subsequently to an external trigger pulse (oscilloscope channel 1).

# Appendix H

## Calibration results of the amplifier channels

Fig. H.1 shows the calibration results, i.e. the difference between expected and measured produced voltage, for each Sinara amplifier channel used in AEGIS, 88 channels in total. This result is plotted for a range of voltages spanning from minimum ( $-200\text{ V}$ ) to maximum ( $200\text{ V}$ ) of the possible values. To observe the relative precision for all voltage ranges, the same number of data points are obtained for every step in order of magnitude (i.e. equally many points between  $200$  and  $20\text{ V}$  as between  $20$  and  $2\text{ V}$ , etc., down to the range between  $0.2$  and  $0.02\text{ V}$ ). The identifier names of each channel are obtained as follows: The first four characters denote the name of the Sinara crate the amplifier board belongs to; then comes the number of the board followed by the number of the channel itself on the given board. As desired, the voltages can be tuned with an uncertainty below  $20\text{ mV}$  on all channels in all voltage ranges, reaching the level of a few millivolt on most channels.

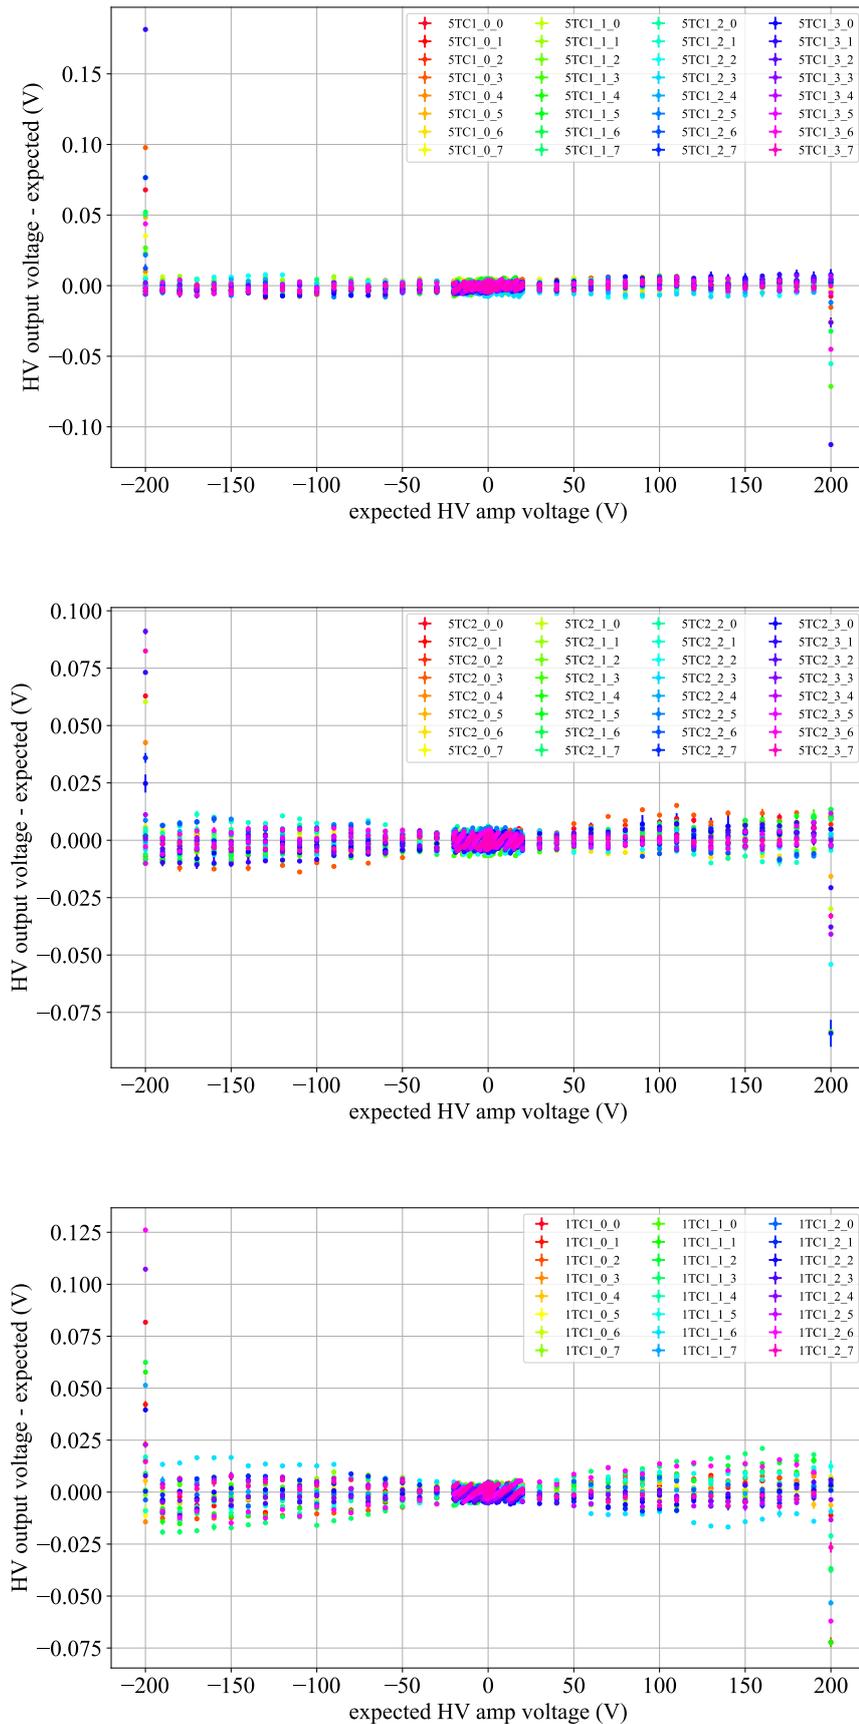


Figure H.1: Difference between the desired voltage on the amplifier channels and the measured output voltage versus the expected voltage after the amplifier calibration for all amplifier channels of all boards involved in the AEGIS apparatus. The legend identifies the channel numbers of the given boards.

# Appendix I

## CMOS images of electron and antiproton plasmas

Fig. I.1 shows images of the combined electron-antiproton plasma confined in the P Trap, taken with the downstream MCP/CMOS camera setup. The images are obtained for different amounts of space charge transferred from the C Trap to the P Trap. As explained in the caption, these images are used to determine the radius of the plasma. Details on the procedure are given in the main text. The determined radii are compared to that of the antiproton bunch arriving in the experiment from ELENA. An example image taken of such a bunch with the same camera is shown in Fig. I.2. It is analyzed in the same way as the mixed plasma.

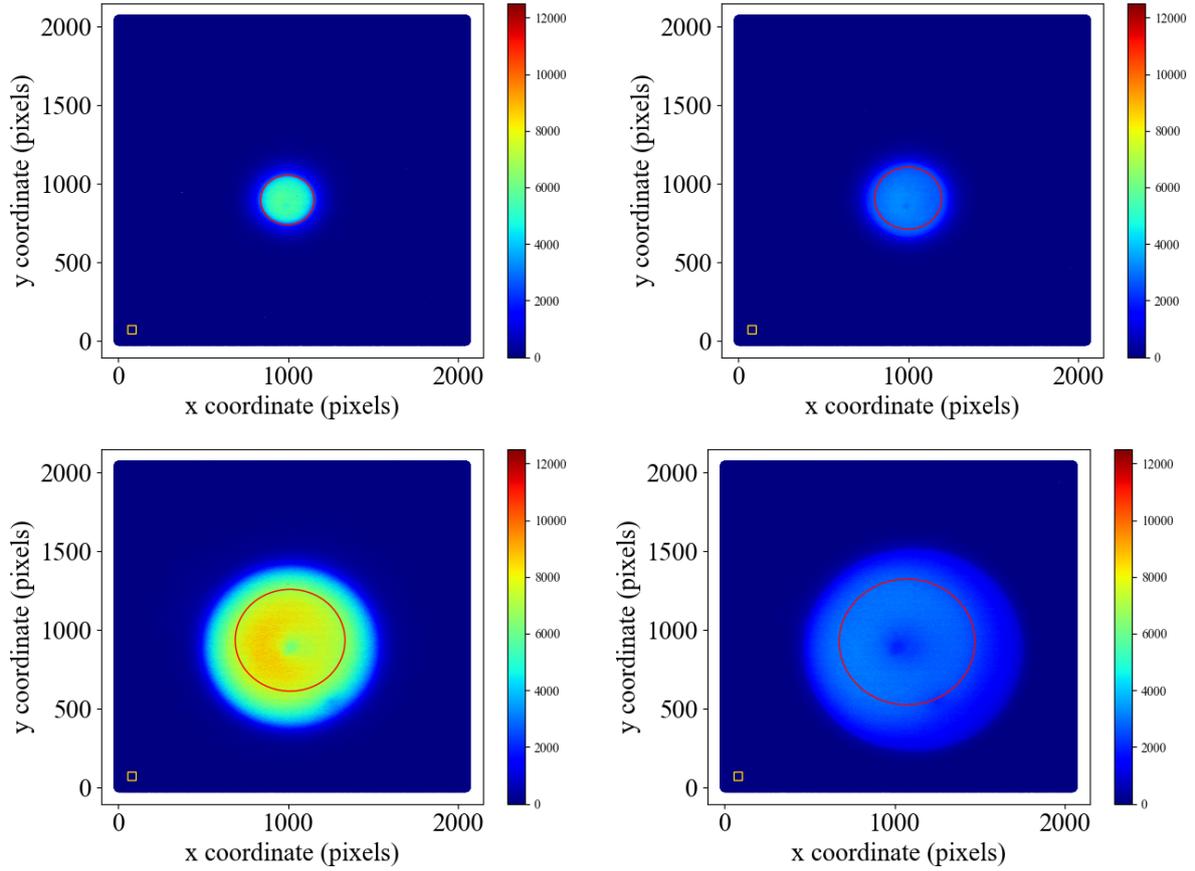


Figure I.1: CMOS images taken with the downstream MCP of the mixed electron-antiproton plasma for different amounts of electron space charge transferred from the C Trap to the P Trap. The pixel intensity is given by the color bar. The red circle indicates the pixel coordinates included by the average plasma radius, as determined from the FWHM method described in the main text. The yellow square borders the pixels used to obtain the average background intensity. Top left: 25 V of transferred space charge. Top right: 50 V of transferred space charge. Bottom left: 95 V of transferred space charge. Bottom right: 145 V of transferred space charge.

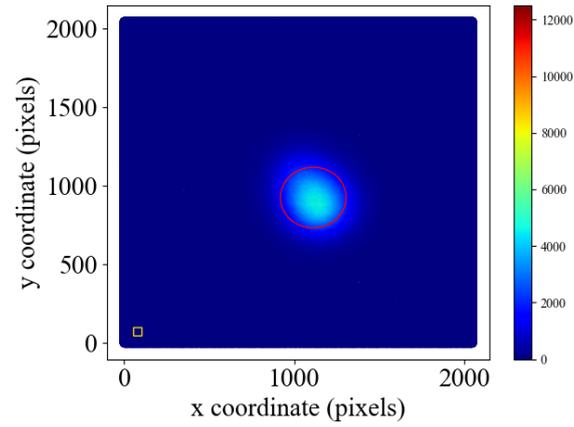


Figure I.2: Example CMOS image taken with the downstream MCP of the incoming ELENA antiproton bunch. The pixel intensity is given by the color bar. The red circle indicates the pixel coordinates included by the average radius, as determined from the FWHM method described in the main text. The yellow square borders the pixels used to obtain the average background intensity.

# Appendix J

## Potential reshaping between two $\bar{\text{H}}$ production cycles

Given that both electrons and antiprotons are recycled in the antihydrogen production procedures as described in Sect. 7.1, they end up in the recycling trap potential upstream from the launch location. In order to use them again for the next production cycle, they need to be transported from there to the dump trap, which is then in turn reshaped to form the launch trap for the subsequent transfer towards the Ps target. Optionally, new electrons and/or antiprotons can be added to the recycled plasma if the capture and cooling trap is shaped in between, as shown in Fig. 6.2. Fig. J.1 visualizes the reshaping operations performed on the low-voltage potential to return from the voltage configuration used for recycling and launch to the dump trap potential, inside which the antiprotons cool on the electrons before being transported to the launch trap as depicted in Sect. 7.1.

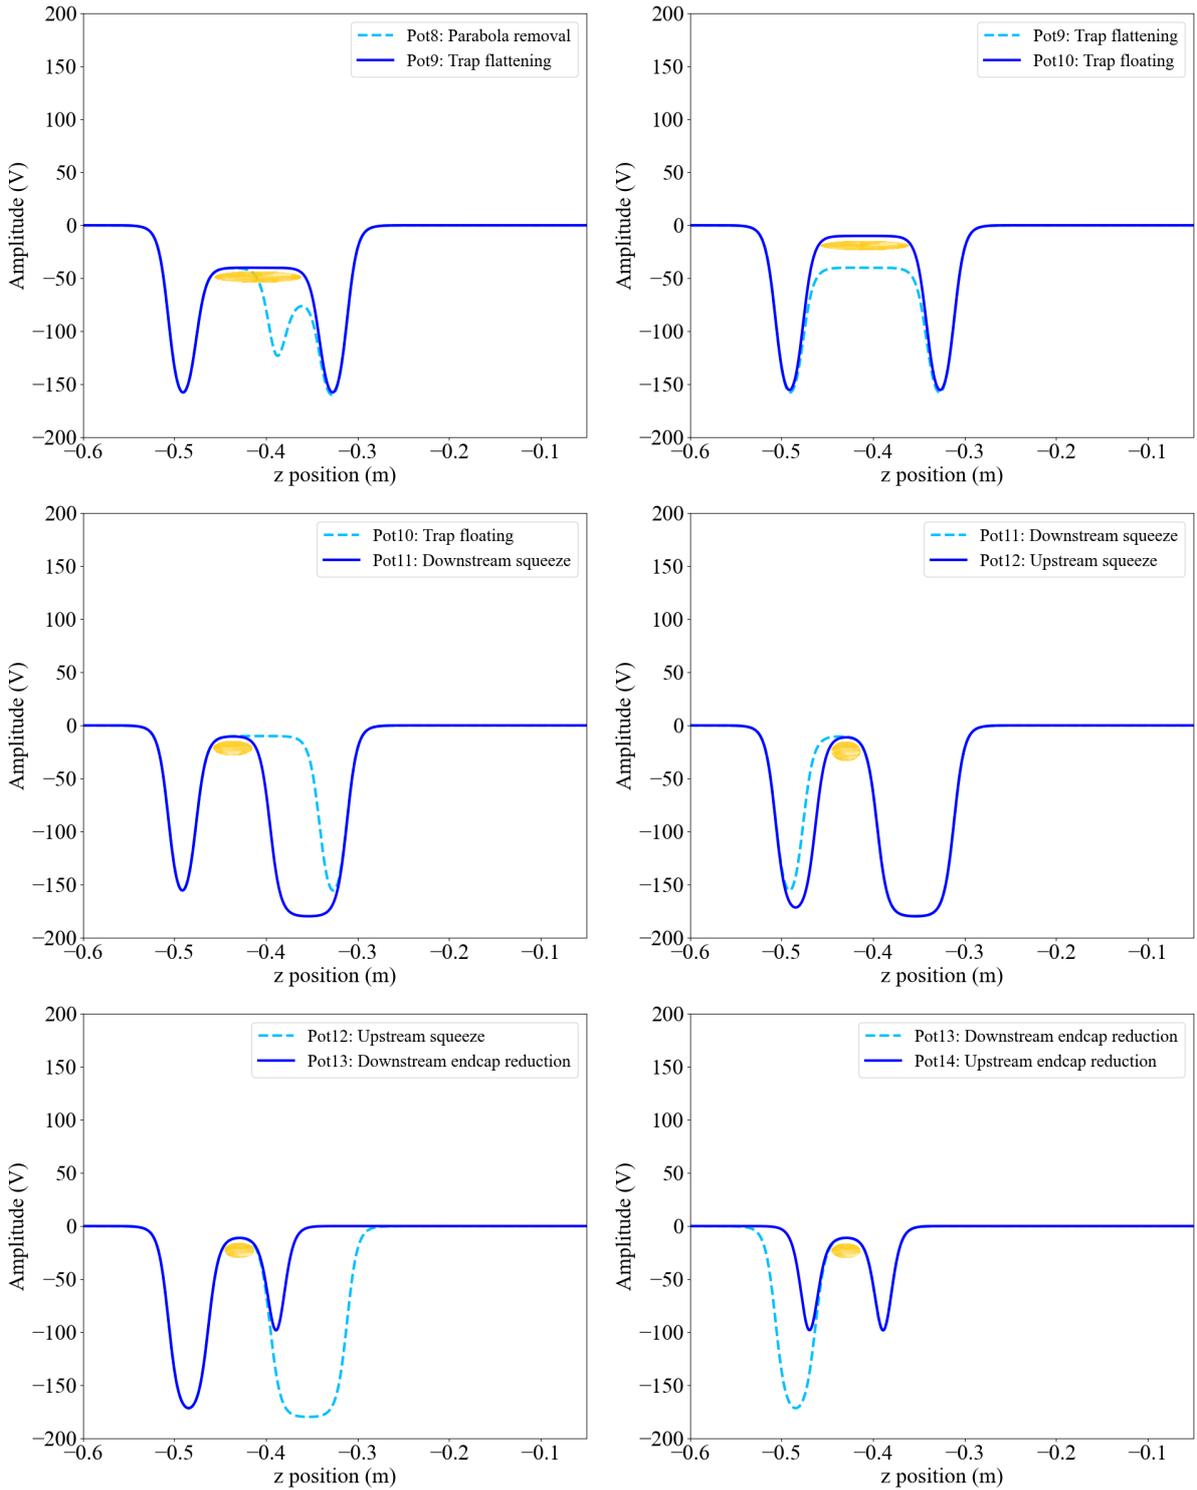


Figure J.1: Low-voltage potential reshaping operations to return from the recycling and launch trap to the dump trap configuration. The initial potentials are plotted by the cyan dashed line; the final potentials by the blue solid line. Assuming that both electrons and antiprotons are previously recycled to be included in the next  $\bar{H}$  production round, the location of the mixed plasma in the final potential is shown by the yellow ellipse (not to scale). Top left: Following the removal of the parabolic transfer potential, the launch and recycling traps are flattened. Top right: The trap well is deepened by floating the floor to an only slightly negative voltage. Middle left: Inward extension of the downstream endcap. Middle right: Inward extension of the upstream endcap. Bottom left: Downstream endcap reduction. Bottom right: Upstream endcap reduction.

# Appendix K

## Scan over the target electrode potential

Fig. K.1 shows a spectrum measured by PMT 24, as an example, for a scan over the voltage applied on the A0 electrode,  $V_{A0}$ , from  $-60$  to  $-130$  V in a set of runs with antiprotons only, i.e. without positrons and thus without antihydrogen formation but with the same procedure as for  $\bar{H}$  formation otherwise employed. Plotted is the range between  $60$  and  $150$   $\mu\text{s}$  from the start of the acquisition. The signals obtained in the runs with different  $V_{A0}$  are individually plotted. Fig. K.2 shows the same data sets, with the plotted time range zoomed to the time around the first distinctive structures, between  $62$  and  $80$   $\mu\text{s}$ . The different signal components and their evolution is discussed in the main text in Sect. 7.2.2.

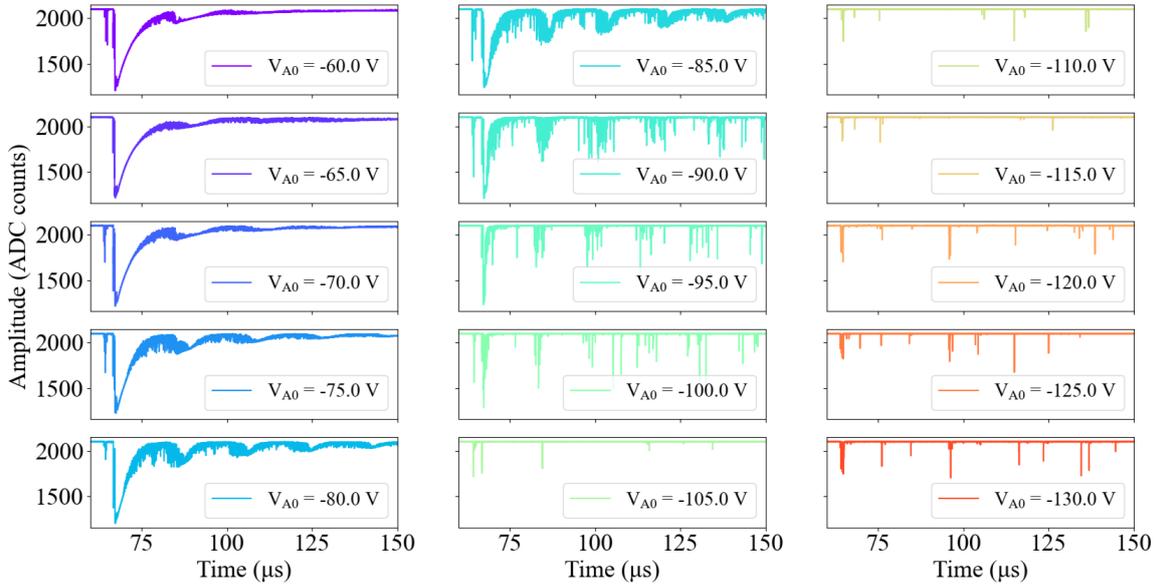


Figure K.1: Signal distributions observed on the digitized PMT 24 as a function of time for a scan over the target electrode voltage in  $\bar{H}$  production runs without positrons. The time is given as the time from the start of the acquisition. Plotted is the range between  $60$  and  $150$   $\mu\text{s}$ .

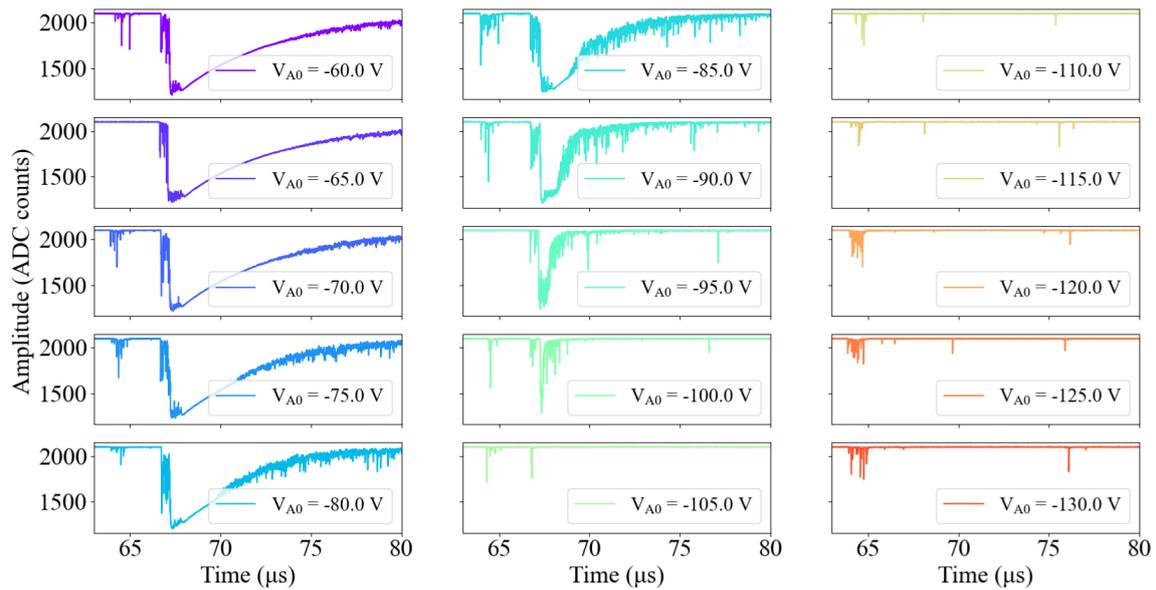


Figure K.2: Signal distributions observed on the digitized PMT 24 as a function of time for a scan over the target electrode voltage in  $\bar{H}$  production runs without positrons. The time is given as the time from the start of the acquisition. Plotted is the range between 62 and 80  $\mu$ s.

# Acknowledgements

Everyone who knows me is likely aware of the fact that I am not exactly doing physics for the love of the subject or a passion for fundamental research: Dedicating one's life to the deepening of human knowledge of the origin and flow of our world is an immense privilege and a marvellous thing; it just might not be mine. That, however, does not mean that I am not grateful for what working in the field has blessed me with. There have been many blessings. They certainly include a knowledge of my personal strengths, weaknesses and boundaries, and the confidence that most things can apparently be achieved with an open mind, team spirit and a mixture of discipline and an organizing ability. Working in physics allows one to be extremely creative and free in choosing one's way of working and the topics of interest. That's pretty amazing and surely not the case in other areas of work. Furthermore, despite the wonderful team spirit I have experienced here, autonomy, personal responsibility and one's own initiative are really fostered in the physics environment - skills that are very precious to possess. I have worked hard and challenged myself, and I have developed a "just try it" attitude towards most things, which I think will be very useful in the future and which I think I really owe to physics because that's the one field where I've been constantly confronted with new things I had no clue about, "being thrown into cold water" and having to deal with it - and it has usually worked out in the end somehow. Almost all things can be learned rather quickly, and it really is not such a big deal to start something new. I have also understood that perfection is impossible and certainly unnecessary. Instead, it is always worth looking beyond one's own nose, and I am so grateful that I have had the opportunities as well as the desire to do that. So I am very thankful for all the things I have learned along the way (some physics, too, which can definitely be fascinating, but mainly more important things). The one thing I am most grateful for, though, is having become part of the physics community. I have made wonderful friends along the way, and generally, the physics world really genuinely seems to be a place of open-minded acceptance: Everyone can be as they are and they will be appreciated for it. We're all weird in our own ways - certainly not only us physicists, but most of us here are aware, which helps. This spirit is an incredibly valuable feature in such a diverse community, which I am proud to be a part of. So, in that sense, first of all thank you to physics.



This work has been sponsored by the Wolfgang Gentner Programme of the German Federal Ministry of Education and Research (grant no. 13E18CHA), for which I am very grateful.

Generally, I truly appreciate CERN's work towards international exchange and collaboration as well as the member states' contributions, which allow for such instructive and adventurous opportunities for young people.

Firstly, I'd like to thank Erika, who has become a sort of mentor for me over the years, for agreeing to supervise this thesis even though it's not your topic. I certainly would not be in this position now and I'm not even sure I would have continued these studies for as long if I hadn't entered your group. Thank you for asking me whether I was "insane" when I suggested I'd like to leave academia and might not pursue a PhD. I haven't regretted my decision, so I suppose it was a good one. I appreciate your honest feedback and the sharing of your personal experience. Furthermore, thank you for always showing me what I don't know yet and for teaching me to finally start asking questions. I have learned a lot from you and you have shown me how a group can be led in a really motivating, positive way. You and your group have played a major role in creating the physics environment I have enjoyed working in so much, and I've been very happy to kind of rejoin during the past months.

To the soul (and speaker) of modern AEGIS, Rug, I would like to say thank you for pouring so much of you (sometimes maybe too much) into the experiment and for inspiring everyone on the team to give it their best, leading by example. Thank you for having my back when I needed it, and thank you for teaching me so many things (even though I am able to understand a maximum of 10% when talking to you on-topic, that's a crucial boost every time). Thanks for your positive feedback and the constructive comments that have always meant good progress, and for your open advice on many different topics. Thank you also to Ale for your generosity when it comes to Rug's time and for the general kind support.

Michael, I would like to thank you for the opportunity to conduct my PhD work at CERN, and for giving such a brilliant guided tour during our seminar visit, prompting me to consider applying at CERN in the first place. Thanks for the trust you have placed in us newer team members to work on such integral parts of the experiment you have built up, and to freely transform them. I have very much enjoyed the open-minded spirit of AEGIS and the opportunities to learn and grow. Additionally, thank you for being so resourceful in every way, always finding a creative way to sustain us students and to accommodate our personal situations, and for offering possibilities to continue in the field wherever possible. In particular, I also appreciate the opportunities that are given in AEGIS, mainly through you, to young members to broaden one's horizon in the form of conference/school visits; it is not a given that these endeavours are always met with such support to this extent. I am immensely grateful to the wonderful "CR-now" team of AEGIS, whom it has been a true pleasure to work with and whom I am so glad to have met. The atmosphere created in the control room starts from the heart, from the best technical coordinator out there (who can apparently seem intimidating at first but is one of the kindest people): I cannot imagine AEGIS without you, Stef, and somehow I am sure you will always remain connected. Having wished for a "hands-on" thesis, it was a very welcome start to my time at CERN when I tried to set up some cryo pump or other, we couldn't find a matching connector and your immediate solution was "Well, then we obviously just make one ourselves". I have learned a great deal from you and I want to sincerely thank you for your patience, your interest in everyone's well-being (particularly maintained by your supply

of free coffee - not sure my version deserves the name "coffee" though...) and your drive to teach us parts of your vast knowledge. Thank you for creating such a warm working environment. I look forward to visiting the retired (as if...) version of you in Salzburg! If Stef is the heart of the AEGIS team (and we have Michael as the brain behind it and Rug as its soul), Benji would be the (very much alive!) spirit of AEGIS. Research and science certainly aren't only a job for you but you truly live them. It is amazing to me how you jump on any subject or detail and hammer it (rationally) until you get to the bottom. You are a true role model for a researcher in this regard but also in another one: Your kindness and willingness to help seem endless, and it's unbelievable how friendly, positive and supportive you remain even under the biggest stress, always appreciating any questions (even those that feel stupid to the person asking for the n-th time...) and being open-minded to all ideas, no matter how absurd. Your open and creative (and funny - I will really miss your... imaginative... jokes!) spirit is vital to the atmosphere in AEGIS and I am grateful to have learned from you in many ways; there could be no better new physics coordinator. A particular thanks also for your prompt and detailed feedback! Coming from the official coordinators to those they have tried to coordinate, I would like to thank Marco for sharing this crazy journey, for mutual support and for all the computer stuff I've learned from you (and for making the CIRCUS work in the way it does!). Also Lisa I want to thank for the mutual support and for sharing our worries and doubts so honestly, which is very valuable and grounding, but also for the fun times (and for the Stef Bingo and mugs, which are priceless!). A huge thank you to Gunn for your hands-on attitude, your resourceful help and for everything I have learned from you on electronics; obviously, nobody in AEGIS will ever forget your name (seeing as we read it about a thousand times per day...), and your team-oriented character has been super beneficial to the group (as have your selfie skills!). Furthermore, I want to thank Tasso for the breath of fresh air you bring, for your dedication and for your truly honest communication (and for ALPACA, which is saving all of our lives when it comes to analyzing data!), and Kuba for your rationality, your programming skills and your special humour (and for your beautiful CIRCUS art, which I love!). A particular thanks goes out to the fellow members of the cryo team, most importantly to our organizer Gosia (also for all the fun stories that brighten our days and for your lovely, refreshing bluntness) and to our champion Valts (not only for your immensely valuable love of craning or for never wasting our food but also for *your* kindness (!), your sense of team spirit and your eagerness to always jump up as the first person to help when anybody needs it). Thanks to everyone who has contributed to the positive atmosphere in AEGIS, also including Fredrik (for your ideas from thinking outside the box), Matthias (for being a constant of tranquility in the chaos that AEGIS can also be), and Tim (for your calm willingness to help with whatever comes up), and to those stopping by the control room occasionally, in particular Antoine, Franz & Bergi, Georgy, Gianni, Giovanni, Luca, Łukasz, the Prelz, Roberto, Theo and Tomasz, for offering valuable advice, manpower and new prospects. Thank you to Yasemin, Stephanie and Mikołaj for your friendly, efficient help whenever needed. Thank you to the CERN community in general and to all the kind and inspiring people I have met in my time in the Geneva area, both on and off site.

Other than Erika, I also want to thank the other members of the UHH-PD Detector team for warmly welcoming me back whenever I've been on Campus Bahrenfeld, in particular Christoph for always (literally!) keeping the door open for me, Bianca for the wonderful hospitality, for the mutual support and for making my office hours in Hamburg so enjoyable, and additionally Annika, Chin-Chia, Constanze, Georg, Jörn, Lea, Massi, and Thore for the extended lunch walks and/or for your advice and the fun times bouldering or at

the institute. Thank you to Mee Joo for the nice visit and for your administrative help. Thanks also to the extended UHH-PD group in building 68 for the shorter lunch walks and enjoyable breaks, for your hospitality and for the fun activities outside of physics.

There are many people who have directly or indirectly contributed to me succeeding in somehow pulling off a PhD in physics but there is a group of people without whom I would have never made it past the first semester of the Bachelor's programme (without exaggeration, as that was initially the plan...). So I want to deeply thank Küchensprache and all its important extensions for having become my physics family and better friends than I could have ever wished for here, and for making me fall in love with Hamburg. I started studying physics simply out of curiosity for something new and thanks to a lack of better ideas, but as stated above, I stayed in it because of the people I met along the way, and that all started with you guys. I am very grateful to have met you and to have carried each other through (parts of) this mad adventure, actually having fun while studying physics (or, you know, hanging out in Pflanzen un Blumen or at Bossli's/Tossli's instead... but somehow we still made it). Generally, I am very grateful to have become a part of the physics community of the UHH, which I think is just awesome.

(Even) Less physics-related, I want to say that I am super thankful to everyone who has made me feel so naturally welcome when coming back home to Germany, both during short trips while still at CERN and when coming back to stay. People had warned me that things would not be the same or the way I expected them to be when coming back, but either I had a realistic image from the start or my people are uncommonly wonderful (or both). Either way, both my friends and my (extended) family deserve a huge thank you for their patience and loyalty. In general, I truly appreciate all the kind people who have stayed in contact - digitally or through fantastic visits, proving that true friendship does not rely on being in the same place. Of course I already knew this, thanks to my wonderful American family, who have somewhat directly contributed to this work by ensuring that this writing part is much less of an obstacle than it could have been, but whom I am so grateful to in many ways beyond.

One person, who is at the same time family and a really good friend, has been a special constant and a rock for me in the past (uhm... however many...) years. Thank you, Tobi, for deciding to embark on this PhD adventure with me and for giving it a try; I know it hasn't always been easy, but I am so proud of us for (mostly) managing to make each other a priority despite the distance. Thank you for always being there for me, for listening, for your moral support, and for making everything immediately better with a (sometimes even remote) hug. Thank you for being the kind, funny, genuinely honest person that you are, for choosing to spend a lot of your time with me, and for balancing me out in so many ways - specifically in our studies, you are the one wanting to dig deeper and deeper and understand every detail (yes, you are the one that is a scientist at heart), while I tend to stay focused on the overall picture, a combination which I think can be very fruitful and which has apparently proved to work rather well. Directly on this work, I certainly owe you a thank you for your Python support, for your probing questions and for your profound proofreading - also here your eye for details comes in very handy. But I think the most important support I've received from you is the warm fluffy feeling I get from knowing that, no matter how badly everything is going with the thesis, I can always turn to you and feel loved and positive. I look forward to our next steps together (hopefully this time in the same place, although maybe not in a single small room the entire time anymore...), and I love you.

Above all, I cannot put into words the thankfulness I feel when I think of my family, most importantly my wonderful parents and grandparents. There is nothing more promoting

than the knowledge that there are loving people who will have your back independent of the outcome of any of your endeavours and who will support you in any of your choices. With this security in the back of my mind, I have always been free to explore, and I think that is an invaluable fortune. Being able to come back to the most loving, comfortable home whenever I've wanted to has brightened any frustrating times, even when I've been far away. I am endlessly grateful to you from the bottom of my heart and I love you so much!

♡

Am allerwichtigsten ist die Dankbarkeit, für die ich keine ausreichenden Worte finde und die ich empfinde, wenn ich an meine Familie denke, vor allem an meine wundervollen Eltern und Großeltern. Es gibt nichts Fördernderes als das Wissen, dass es Menschen gibt, die einen lieb haben, die immer hinter einem stehen, unabhängig vom Ausgang eines Unterfangens, und die einen immer in allen Entscheidungen unterstützen. Mit dieser Sicherheit im Hinterkopf war ich immer frei, Dinge auszuprobieren, und ich glaube, dass das ein unbezahlbares Glück ist. Dass ich jederzeit zurückkehren konnte zu dem liebevollsten, gemütlichsten Zuhause, das man sich wünschen kann, hat jede frustrierende Zeit aufgehellt, selbst wenn ich weit weg war. Ich bin euch von ganzem Herzen unendlich dankbar und ich habe euch sooo lieb!

# Eidesstattliche Versicherung / Declaration on oath

Hiermit versichere ich an Eides statt, die vorliegende Dissertationsschrift selbst verfasst und keine anderen als die angegebenen Hilfsmittel und Quellen benutzt zu haben.

Sofern im Zuge der Erstellung der vorliegenden Dissertationsschrift generative Künstliche Intelligenz (gKI) basierte elektronische Hilfsmittel verwendet wurden, versichere ich, dass meine eigene Leistung im Vordergrund stand und dass eine vollständige Dokumentation aller verwendeten Hilfsmittel gemäß der Guten wissenschaftlichen Praxis vorliegt. Ich trage die Verantwortung für eventuell durch die gKI generierte fehlerhafte oder verzerrte Inhalte, fehlerhafte Referenzen, Verstöße gegen das Datenschutz- und Urheberrecht oder Plagiate.

Hamburg, den 05.06.2025



---

Unterschrift der Doktorandin