

Characterization and Analysis Development for the ALPHA-g Barrel Scintillator

by

Gareth Smith

B.Sc., Carleton University, 2019

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF SCIENCE

in

The Faculty of Graduate and Postdoctoral Studies

(Physics)

THE UNIVERSITY OF BRITISH COLUMBIA

(Vancouver)

September 2021

© Gareth Smith 2021

The following individuals certify that they have read, and recommend to the Faculty of Graduate and Postdoctoral Studies for acceptance, the thesis entitled:

Characterization and Analysis Development for the ALPHA-g Barrel Scintillator

submitted by **Gareth Smith** in partial fulfillment of the requirements for the degree of Master of Science in Physics.

Examining Committee:

Makoto Fujiwara, Research Scientist, TRIUMF
Supervisor

Stan Yen, Adjunct Professor, Department of Physics and Astronomy, UBC
Examining Committee Member

Abstract

The ALPHA-g experiment is the current focus of the ALPHA collaboration; it aims to measure the trajectory of antihydrogen atoms in a gravitational field in order to verify the weak equivalence principle for antimatter. It includes a scintillator detector called the Barrel Veto which will measure particle time of flight in an effort to distinguish antihydrogen annihilation products from cosmic ray background. A new analysis framework for BV data is presented. A small-scale replica of the BV was constructed at TRIUMF and used to develop procedures for calibrating corrections for differences in path length and pulse amplitude. These corrections improve the time of flight resolution to approximately 170 ps, which should be sufficient for cosmic ray background rejection.

Lay Summary

Every normal type of particle has an opposite, called an antimatter particle, which are very rare in nature. The ALPHA-g experiment will test the effect of gravity on antimatter particles for the first time. To this end, it must detect the fragments created when an antimatter particle collides with normal matter. These fragments are difficult to distinguish from background particles called cosmic rays which are constantly descending from space. A special detector called the Barrel Veto was built for this purpose. Particles passing through this detector cause flashes of light, and by looking at the patterns and timing of these flashes, it is possible to tell the difference between antimatter fragments and cosmic rays. In my work, I prepared the detectors and analysis for the Barrel Veto, and demonstrate that it should be able to remove the majority of cosmic ray background from our data.

Preface

This thesis is composed of original, unpublished work by the author, based on the ALPHA-g experimental apparatus and the work of the ALPHA collaboration. The Barrel Veto apparatus described in chapter 5 was constructed at TRIUMF primarily by Nicolas Massacret, Robert Henderson, and Philip Lu. The data acquisition system was created by Konstantin Olchanski and Pierre Amaudruz at TRIUMF, and the analysis framework was implemented by Andrea Capra of TRIUMF and Joseph McKenna of Aarhus University. The work described in chapters 6-9 was contributed entirely by the author.

Table of Contents

Abstract	iii
Lay Summary	iv
Preface	v
Table of Contents	vi
List of Figures	ix
Acknowledgements	xii
1 Introduction	1
2 On antimatter	3
2.1 History	3
2.2 Motivation for the study of antimatter	5
2.2.1 CPT symmetry	5
2.2.2 Baryon asymmetry	6
2.2.3 Antimatter and gravity	6
2.3 Annihilation process	8
2.4 The advantages of antihydrogen	10
2.5 Antihydrogen in experiment	10
3 The ALPHA-g experiment	13
3.1 History of ALPHA	13
3.2 Physics goals	14
3.3 Antihydrogen production	15

Table of Contents

3.4	Antihydrogen release and detection	17
4	Antimatter detection	19
4.1	Ionization in a time projection chamber	20
4.2	Light production in plastic scintillators	23
4.3	Light collection in silicon photomultipliers	24
5	The ALPHA-g Barrel Scintillator	27
5.1	Motivation for the BV: Time-of-flight background discrimination	28
5.2	TOF time resolution requirement	32
5.3	Barrel veto design	33
5.4	Readout electronics design considerations: Time walk correction	36
5.5	Readout electronics design	38
6	Barrel Scintillator Analysis Algorithm Design	41
6.1	ADC waveforms	42
6.2	Calculation of TDC time	48
6.3	Event construction	49
6.4	Calculation of Z position	50
6.5	TPC track matching	53
6.6	Calculation of time-of-flight	54
7	The Vertical Slice	56
7.1	Motivation: Mysteries in the CERN TOF data	56
7.2	Implementation	58
7.3	Initial results	62
8	Time corrections and calibrations	67
8.1	Optimal TDC threshold	67
8.2	Channel-by-channel offset calibration	72
8.3	Time-walk correction	76
8.3.1	Form of the time walk correction	77

Table of Contents

8.3.2	Calibration schemes for time walk correction	78
8.3.3	Demonstration of time walk using oscilloscope	80
8.3.4	Calibration of time walk correction using data	84
8.3.5	Alternative time walk correction schemes	86
9	Final results and outlook	87
9.1	Final time resolution	87
9.2	Time of flight demonstration	89
9.3	Final position resolution	93
9.4	Final preparations and next steps	94
10	Conclusions	98
	Bibliography	100

List of Figures

2.1	First observation of a positron by Carl Anderson	4
2.2	Proton-antiproton annihilation interaction	9
3.1	Antihydrogen production process	17
4.1	Positron Emission Tomography operating principle	20
4.2	The TRIUMF TPC	22
4.3	ALPHA-g TPC track reconstruction	22
4.4	Molecular structure of polyvinyltoluene	24
4.5	Semiconductor band gap structure	25
5.1	Arrangement of ALPHA-g detector systems	28
5.2	Illustration of TOF-based background rejection principle . . .	31
5.3	Simulated annihilation signal acceptance	33
5.4	Reflections within the scintillator bars	34
5.5	Barrel veto SiPM arrangement	35
5.6	“Time walk”: The effect of pulse amplitude on threshold crossing time	37
5.7	Readout electronics for the BV	38
6.1	ADC waveform from cosmic ray event	43
6.2	Fitting of unsaturated ADC waveform	45
6.3	Fitting of saturated ADC waveform	45
6.4	Measured ADC waveform amplitude	46
6.5	Fitted ADC waveform amplitude	47
6.6	ADC waveform amplitude – measured vs. fitted	48
6.7	Z position of hits within bar	52

List of Figures

6.8	Time of flight distribution for cosmic rays at CERN	55
7.1	TOF: measured vs. predicted from geometry	57
7.2	Scintillator bar layout for vertical slice	58
7.3	Low voltage distribution and control board	59
7.4	Analogue sum discriminator (ASD) card front side	59
7.5	Analogue sum discriminator (ASD) card reverse side	60
7.6	Rear transition module (RTM)	61
7.7	Alpha-16 analogue to digital converter (ADC)	61
7.8	TRB3 time to digital converter (TDC)	62
7.9	Cosmic ray TOF distribution for vertical slice	63
7.10	End-to-end time differences for vertical slice	64
7.11	Crossed bars layout	65
7.12	End-to-end time difference using crossed bars	66
8.1	Multiplicity of TDC time values	68
8.2	Extraneous TDC time values	69
8.3	Cosmic ray count rates for varied threshold voltage	71
8.4	TDC time vs. approximate ADC waveform start time	72
8.5	Pulser-measured TDC channel offsets	73
8.6	Pulser-measured TDC channel offsets after channel-by-channel correction	75
8.7	Pulser-measured TDC channel offsets after swapping TDC inputs	76
8.8	Pulse shape used for time walk correction	78
8.9	Quantities needed to calibrate the time walk correction	79
8.10	ASD card test points for measuring time walk	81
8.11	Waveform measured at the first test point	83
8.12	Waveform measured at the second test point	83
8.13	Time correction needed for cosmic ray pulses of varying am- plitude	84
8.14	Scan of time walk correction parameter W	85
9.1	TOF distribution without corrections	88

List of Figures

9.2	TOF distribution with corrections	88
9.3	TOF distribution with corrections, with bars separated	90
9.4	TOF distribution with corrections, with bars swapped	92
9.5	TOF distribution with corrections, with bars swapped and separated	92
9.6	End-to-end time differences for vertical slice with corrections	93

Acknowledgements

Firstly, I would like to thank my supervisor Makoto Fujiwara for his enthusiasm and belief in my abilities; for providing an excellent environment for me to develop as a researcher by allowing me to find my own solutions; and for consistently providing guidance, encouragement, insight, and support.

Thanks as well to all the members of the ALPHA team at TRIUMF for their help in getting to grips with the many facets of the ALPHA-g experiment, as well as for creating such a vibrant work environment. In particular, I thank Andrea Capra for being such an excellent mentor, Lars Martin for his equal parts humour and wisdom, and Konstantin Olchanski for never missing a teaching opportunity.

More personally, I would like to thank my parents for their continual mental support, and for instilling in me the belief that I can achieve anything I set my mind to. And finally, I want to acknowledge all the friends I have made in Vancouver (and retained in Ontario), who have made the day-to-day of this degree both bearable and memorable.

Chapter 1

Introduction

Antimatter lies at the heart of one of the most prominent phenomena which remains unexplained by the Standard Model: the *baryon asymmetry problem*, the difference in the quantity of matter and antimatter in the observable universe. This makes it an attractive subject for particle physicists hoping to discover signs of new physics; however, the study of atomic antimatter has only recently been made possible at the Antiproton Decelerator (AD) facility of the European Organization for Nuclear Research (CERN). ALPHA-g is the newest venture of the ALPHA collaboration, one of the leaders in antihydrogen research at the AD. Building upon the successes of previous experiments, ALPHA-g aims to measure the trajectory of antihydrogen atoms in a magnetic field. This will be the first attempt to directly verify the weak equivalence principle for antimatter. The first operation of ALPHA-g is planned for the autumn of 2021, and commissioning of the various components of the experiment is fully underway at the time of writing.

The present thesis describes the efforts to commission one of these components, the barrel veto (BV), and to demonstrate its capabilities. The BV is composed of a series of bars of plastic scintillator which surround the experiment like a barrel. It is designed to precisely record the time of particles passing through it. Not only does this provide extra information about the antihydrogen atoms being studied, it is also crucial in distinguishing the true signal from the background of cosmic rays which permeate the experiment. The goal of this work is threefold: primarily, it is to demonstrate that the BV is capable of performing this function, and that no further hardware or electronics upgrades are required. In addition, this work will develop the analysis algorithms used to calculate the time of flight and other event information using the raw data coming from the BV. Finally, any calibrations

and corrections required for the proper operation of the BV will be identified and carried out.

The structure of this report is as follows. Chapter 2 provides a brief history of the study of antimatter, including an overview of the present experiments at the AD. Chapter 3 goes into more detail on the ALPHA-g experiment. Chapter 4 describes the particle detection techniques used by ALPHA-g, and chapter 5 gives an overview of the BV system and electronics. Chapter 6 explains the work done in developing an analysis algorithm for the BV, with reference to a cosmic ray data set taken previously. Chapter 7 shows the development of a vertical slice of the system created locally at TRIUMF, and the results it provided. Chapter 8 describes how this was used to calibrate a series of time corrections, which ultimately led to the final results presented in chapter 9.

Chapter 2

On antimatter

The fundamental goal of a physicist is to develop our understanding the universe around us and its composition. Since the days of the ancient Greeks [1], we have theorized that the universe is made of tiny indivisible particles which interact with each other to create the world we experience. Almost all of the particles which we encounter on a daily basis fall into the category of *matter*; however, there is a second much rarer class of particles called *antimatter* about which we know comparatively little.

Every matter particle¹ has a corresponding antimatter particle, which has identical mass and lifetime and interacts via the same forces, but has opposite charge and spin. When an antiparticle encounters its regular matter counterpart, they undergo an interaction called annihilation, whereby both particles are destroyed and their rest energy is released. Broadly, the goal of the ALPHA collaboration is to study these antimatter particles to gain a more complete understanding of our universe.

2.1 History

Antimatter has the distinction of being one of the first types of particles to be hypothesized before its discovery. The anti-electron or positron was first postulated by Dirac in 1927 [2]. While developing his theory describing the relativistic motion of an electron, now known as the Dirac equation, he noticed a mathematical oddity. As well as two positive-energy solutions corresponding to the two spin states of the electron, his equation also had two negative energy solutions. After some postulation, he eventually interpreted

¹The only known exceptions to this rule are particles which are their own antiparticle, such as the photon.

2.1. History

these solutions as a particle with equal mass but opposite charge to the electron – the positron.

The existence of this antiparticle was soon confirmed by Carl Anderson in 1932 [3], who used a bubble chamber to observe cosmic rays. Initially, it was impossible to differentiate between a positron curving in a magnetic field and an electron travelling in the opposite direction. By placing a lead sheet in the bubble chamber to absorb some of the energy of the particles, he was able to determine the direction of motion of these positrons and hence prove the existence of antimatter.




Figure 2.1: First observation of a positron by Carl Anderson in 1932. Taken from [3]. A lead sheet (placed horizontally, in the centre of the image) slows down a positron passing from the bottom of the image to the top. This is observed as a decrease in the radius of curvature due to the external magnetic field. The direction of curvature indicates the charge of the positron.

Following this discovery, antiparticles of many other types have been measured. Heavier antimatter particles are studied primarily in accelerator experiments, as only the lightest antimatter particles are readily available in nature.

2.2 Motivation for the study of antimatter

The symmetries between matter and antimatter are fundamental to all physical theories describing the interactions of particles. These symmetries are very well motivated, and have obtained a near-axiomatic status. Should they be broken, a reworking of our accepted laws of physics would be required. Exactly what this would pertain is still the subject of some debate, and is outside the scope of this thesis; however, this section will nonetheless outline the symmetries in question and how the work of the ALPHA collaboration is testing these symmetries.

2.2.1 CPT symmetry

The main symmetry which motivates the study of antimatter is called *CPT symmetry*. This refers to three fundamental transformations. Charge conjugation (C) refers to changing the sign of all quantum numbers of a particle; simplistically, replacing a particle with its antiparticle. Parity inversion (P) refers to changing the sign of all spatial coordinates, reflecting the particle in a three dimensional mirror. Time reversal (T) refers to letting the time coordinate of the system run backwards.

Initially, each of these transformations were thought to be fundamental symmetries which preserve the laws of physics. This is true for classical mechanics and electromagnetism. However, the weak interaction was found to break both C and P symmetries [4]. For a time it was thought that combining charge conjugation and parity inversion lead to an unbroken CP symmetry; however, CP symmetry violation has been observed in both kaons [5] and B mesons [6].

Combining charge, parity, and time reversal yields the CPT transformation. Thus far, there have been no observations of a process violating CPT symmetry; indeed, there are strong theoretical reasons for CPT to be conserved. It has been shown that any Lorentz invariant quantum field theory must also be invariant under CPT [7]. Lorentz invariance is generally taken axiomatically, and so should CPT violation be observed, the axioms of our theoretical framework would have to be re-examined.

2.2.2 Baryon asymmetry

Most models of the early universe result in matter and antimatter being created in equal amounts. This creates tension with our observations of the universe today, which indicate that it is entirely dominated by matter. A number of mechanisms have been suggested which could result in the observed imbalance between matter and antimatter; however, they require CP violating processes on a magnitude far greater than those which have been observed [8].

The theoretical search for such a mechanism leads naturally to an experimental study of the properties of antimatter. An observed discrepancy between any property of matter and antimatter would be highly useful in elucidating a mechanism for baryon asymmetry. Such a deviation from prediction has never been seen, but our tests on antimatter have thus far been limited, and the divergent property could take any number of forms. Further study of antimatter, then, could illuminate a mechanism for the overwhelming prevalence of matter and provide insight into the history of our universe.

2.2.3 Antimatter and gravity

Our understanding of gravity is at the centre of a second unsolved problem in modern physics. The electromagnetic, strong, and weak forces are all well described by quantum mechanics and quantum field theories. In this framework, they result from the exchange of certain force-mediating particles. However, this framework is not compatible with general relativity, our current theory of gravity. A quantum theory of gravity which is self-consistent has not been constructed; this is an active area of research.

The postulated particle which would carry the gravitational force is called the graviton; it is necessarily a spin-2 particle [9]. Furthermore, it must be its own antiparticle, similar to the photon. Thus CPT symmetry dictates that the gravitational interaction between two bodies composed of antimatter, i.e. an anti-apple falling on an anti-Earth, must be identical to the interaction between the corresponding matter bodies.

2.2. Motivation for the study of antimatter

For obvious reasons, this measurement is currently beyond human capabilities. However, we can test the *weak equivalence principle*, the statement that all bodies accelerate identically in a gravitational field, for antimatter. This principle can be verified by observing an antimatter particle in the Earth's gravitation field, i.e. an anti-apple falling on normal Earth. An observation of a difference in behaviour between a particle and antiparticle would have a huge effect on the efforts to write a theory of quantum gravity; it may require, for example, the introduction of a second spin-1 gravitational force carrying particle. Regardless, comparing the gravitational acceleration of antimatter particles to their matter counterpart is an obvious and necessary test to perform as we attempt to verify the fundamentals upon which a quantum theory of gravity can be built.

There are a number of arguments [10–12] which predict that the gravitational effect of antimatter particles is similar to that of matter. As an example, the reasoning given by Caldwell and Dvali [11] is as follows.

The mass of the proton is much larger than the mass of its constituent quarks; the extra mass comes from the binding energy of the quarks, which is stored in the gluon field. In order to calculate the inertial mass of the proton, the contributions from the binding energy must be summed up. In the language of Feynman diagrams, this can be thought of as a summation over an infinite number of QCD diagrams involving gluons and virtual quarks. In reality, the calculation is considerably more technical, and is typically carried out using lattice QCD, although this Feynman diagram description is theoretically sound in principle. From this point, the force of gravity acting on the proton can then be found by invoking the equivalence principle to obtain the gravitational mass, and replacing the proton with a point-like source of equal mass.

A second way to find the gravitational mass is to apply gravity to the summation of QCD diagrams directly, diagram by diagram. Caldwell and Dvali describe attaching a graviton leg directly at every point of each diagram, where the other end of the graviton line is thought to be attached to the Earth. The coupling of the graviton field to each diagram would determine the gravitational mass of the proton directly, with no need to call

2.3. Annihilation process

upon the equivalence principle. Note that this calculation is obviously not feasible to perform, especially with no current standard formulation of the graviton — this exercise is meant as a thought experiment. The equivalence of these two methods is a manifestation of the equivalence principle.

Now consider if the coupling of the graviton was different for matter particles and antimatter particles. This would also be true for off-shell particles, such as the virtual quark pairs present in many of the QCD diagrams of the proton. Calculating the gravitational mass using the second method, then, would necessarily give a different result than using the first method, since many of the gravitons legs would be attached to virtual antiparticles. Thus a gravitational force which acts differently on matter and antimatter would lead to a detectable difference between the gravitational and inertial mass of the proton. The fact that this has not been observed indicates that it is highly unlikely to find an anomalous gravitational force on antimatter, at least at the scale which is presently accessible by direct measurement.

However, the validity of these arguments is not certain, and it is important and exciting to test them experimentally. In a certain fashion, the presence of arguments predicting a null result from antimatter gravity experiments makes the possibility of an unexpected result all the more exciting. In the words of the late Swedish DJ Avicii, “there’s always haters”.

2.3 Annihilation process

The annihilation process is both a blessing and a curse when it comes to making measurements on antimatter. Any detector will necessarily be constructed of matter particles, and so a positron or antiproton will almost certainly interact with a detector via annihilation once it reaches a sufficiently low energy. Thus it becomes impossible to detect low energy antimatter particles directly.

Instead, an antimatter detector must allow for annihilation, and then detect the products of this interaction. To work backwards and understand the properties of an antimatter particle from its annihilation products, we must have a good understanding of the annihilation process itself.

2.3. Annihilation process

The simplest annihilation interaction is between a low-energy electron and a positron. Each of these has a rest energy of 511 keV, and so the annihilation releases around 1 MeV of energy. Momentum and angular momentum conservation require the creation of two or more particles, and the only particles which can be produced with this energy budget are photons and neutrinos. The most common electron-positron annihilation interaction produces two or more photons.

The annihilation interaction between a proton and an antiproton is far more complicated. The energy budget is now close to 2 GeV, which opens the door for a wider range of particles to be produced. Furthermore, the proton is not a point-like particle like the electron, but is composed of quarks. At high energies, one or more quarks from the proton and anti-proton will interact, while at low energies the protons will behave more as single entities. The details of these interactions are largely beyond the scope of this thesis.

The simplest interaction between a proton and antiproton results in the creation of three pions: a π^+ , a π^0 , and a π^- . This is essentially a rearrangement of the quarks of the proton and antiproton, as shown in figure 2.2.




Figure 2.2: The simplest diagram for a proton-anti-proton annihilation interaction at low energy. Note that this interaction is essentially a rearrangement of the constituent quarks into three mesons.

A typical proton-antiproton annihilation will create a number of mesons; the heavier varieties for the most part decay quickly to charged pions. There are many different sets of final states, with well studied branching ratios [13]. Excess energy from these interactions is carried away by the pions as momentum.

Annihilation of atomic antimatter follows this general prescription: the antiprotons and antineutrons of an antimatter atom annihilate with protons and neutrons, while the positrons annihilate with electrons. The final state includes both pions and photons.

2.4 The advantages of antihydrogen

With the advent of improved technologies for producing and controlling antiprotons, the study of the properties of antiprotons and atomic antimatter has begun to accelerate greatly in the last twenty years. While studying the properties of elementary antiparticles is important, there are a number of advantages to studying atomic antimatter. The hydrogen atom is one of the most well-understood systems in physics, with spectroscopic measurements reaching a truly monumental precision of one part in 10^{15} [14]. The CPT theorem predicts that this frequency is the same for antihydrogen, and so performing a similar measurement for atomic antihydrogen would be an unparalleled test of CPT symmetry.

Another advantage of studying the antihydrogen atom is that it is the simplest system of antimatter which is electrically neutral. While this makes it more difficult to contain, it removes residual electromagnetic fields as a background for an experiment studying the interaction of antimatter with other forces. In particular, the study of the effect of gravity on antimatter becomes easier to perform on an electrically neutral atom.

2.5 Antihydrogen in experiment

Antihydrogen was first created in the LEAR facility at CERN, which was operational from 1982 to 1996 [15]. LEAR used CERN's proton synchrotron

2.5. Antihydrogen in experiment

to produce a high intensity beam of antiprotons which was then accumulated and stored. In 1995 [16], this beam was passed through a Xe target, where antiprotons would occasionally lose energy due to electron-positron pair production. In some very rare cases, the positron would be created with a velocity matching that of the antiproton, and the two would become bound as an antihydrogen atom in the beam. These atoms made a unique signal as they annihilated in a detector.

These antihydrogen atoms were far too fast to perform measurements on. At the same time, progress was being made in slowing the antiprotons and confining them to Penning traps.

Antiproton and antihydrogen research really began to take off with the decommissioning of LEAR and the construction of the Antiproton Decelerator (AD) at CERN in 1999. Since then, a number of experiments have successfully contained and measured antiprotons and antihydrogen. In 2002, the ATHENA collaboration [17] and subsequently the ATRAP collaboration [18] were able to produce larger quantities of antihydrogen using trapped antiprotons and positrons. However, these atoms were still too short-lived to be able to perform spectroscopic measurements, either annihilating or becoming ionized in the electromagnetic fields.

The ALPHA collaboration in 2010 [19] and the ATRAP collaboration in 2011 [20] were able to confine antihydrogen produced in this way using magnetic minimum traps. This has allowed for a number of the transition frequencies of antihydrogen to be measured [21–24]. More recently, the ALPHA collaboration has been able to cool these antihydrogen atoms even further using laser cooling techniques [25]. Furthermore, the completion of the ELENA decelerator ring promises a supply of even lower energy antiprotons [26].

The list of properties of antiprotons and antihydrogen currently being studied at the AD is expansive [15]. The BASE experiment measures the magnetic moment of antiprotons trapped in a Penning trap [27]. ASACUSA studies antihydrogen and antiprotonic helium in flight, mainly performing spectroscopy [28]. GBAR creates and traps antihydrogen ions and releases them to study their gravitation acceleration [29]. The AEgIS group also

2.5. *Antihydrogen in experiment*

studies the effects of gravity on antihydrogen, by using interferometry on a beam of antihydrogen [30]. And finally the ALPHA collaboration continues to study the transition frequencies of antihydrogen, and is beginning to measure gravitational acceleration in ALPHA-g [31].

Chapter 3

The ALPHA-g experiment

3.1 History of ALPHA

The ALPHA collaboration was formed in 2005 by many members of the former ATHENA collaboration [17], in an effort to build from the successes of its predecessor. Whereas ATHENA was unable to capture the antihydrogen atoms it created, the original ALPHA experiment used a magnetic minimum trap to hold the antihydrogen atoms. In 2011, ALPHA succeeded in trapping cold antihydrogen atoms for up to 16 minutes [32].

With the ability to contain antihydrogen long enough for the majority of atoms to fall into the ground state, ALPHA began to use spectroscopy to measure transition frequencies. The ground state of antihydrogen in a magnetic field splits into four hyperfine states. Two of these states are trapable in a magnetic minimum trap, while the other two move towards higher magnetic fields and thus escape. In 2012, by injecting resonant microwave radiation, ALPHA was able to induce a transition between these states and observe the subsequently ejected antihydrogen [33].

At the same time, ALPHA used a silicon vertex detector to reconstruct the annihilation positions of antihydrogen atoms after the magnetic minimum trap is turned off. By examining the vertical distribution of these annihilations, a proof of concept measurement of the gravitational mass of antihydrogen was performed; this was an early test of the weak equivalence principle in antihydrogen, published in 2013 [31]. This laid the framework for the present ALPHA-g experiment.

For its second run, the ALPHA apparatus was upgraded to ALPHA-2. One primary upgrade was the introduction of a separate *catching trap*, responsible for slowing and cooling antiprotons into a plasma before they

are moved into the main apparatus. This allowed far greater access to the antihydrogen atoms once formed. In 2016, by stepping the incident radiation frequency across the hyperfine transition frequencies, ALPHA-2 was able to measure the hyperfine splitting to a precision of 4×10^{-4} [21]. Using a similar method, the Lyman-alpha 1s-2s [22] and 1s-2p [23] transition frequencies were also measured. More recently, the fine structure of antihydrogen has been observed [24], and laser cooling of antihydrogen atoms has been demonstrated for the first time [25], which opens the door for a new set of precision measurements.

The current focus of the ALPHA collaboration is split between an upgrade of the ALPHA-2 system called ALPHA-3, and a new experiment called ALPHA-g. The upgrades for ALPHA-3 include photon detection for fluorescent light, as well as plans to observe further transitions such as the 2s-4s transition. Meanwhile, the new ALPHA-g will build upon the previous proof-of-concept gravitational mass measurements with an entirely new trap and set of detectors [34]. Crucially, the new trap is oriented vertically to provide a much larger statistical sensitivity. In addition, the magnetic fields of the new apparatus are more homogeneous, and their behaviour is known much more precisely as they are turned off. This reduces the systematic error introduced by the magnetic forces which would otherwise dwarf the gravitational interaction.

3.2 Physics goals

The goal of the ALPHA-g experiment is to measure the gravitational acceleration g of antihydrogen in the Earth's gravitational field. The weak equivalence principle states that the trajectory of any object in a gravitational field is independent of the composition or internal structure of the object; on the Earth's surface, all objects will accelerate downwards at $g_0 = 9.806\,65\,\text{m s}^{-2}$.² ALPHA-g aims to verify that this is also the case for

²This is the value for *standard gravity*, the nominal acceleration due to gravity near the Earth's surface which was fixed in 1901 [35]. In reality, it varies significantly by altitude and latitude.

antihydrogen, and in doing so, provide evidence that the weak equivalence principle holds for antimatter. ALPHA-g aims to be the first precision test of this principle on antihydrogen, and in doing so provide a more solid experimental basis for the theoretical efforts surrounding gravity discussed in section 2.2.3.

3.3 Antihydrogen production

The ALPHA-g apparatus is built upon the success of multiple generations of antihydrogen experiment. The antihydrogen production process begins with the production of a cold plasma of antiprotons. The first steps are performed by CERN’s Antiproton Decelerator (AD) group. A proton beam from the proton synchrotron is collided with a target to produce a smorgasbord of particles, some of which are antiprotons. A fraction of these antiprotons are collected by electromagnetic fields and focused into a beam, which is sent to the AD. Here they are contained in a cooling ring by bending magnets, while electric fields are used to slow the antiprotons to 5.3 MeV. From there, they are transported to the new ELENA (Extra Low ENergy Antiproton) ring [26], which uses similar techniques to further slow the antiprotons to an energy of 100 keV.

From this point, the antiprotons are sent to the ALPHA apparatus. They are captured by a Penning-Malmberg trap called the *catching trap*. Here, a solenoid is used to create a uniform axial magnetic field. This confines the radial motion of the antiprotons, and causes them to orbit around the trap axis. A static electric potential applied at both ends of the trap by a set of cylindrical electrodes creates a potential barrier which confines the antiprotons in the axial direction. Once captured in the catching trap, the antiprotons are sympathetically cooled to further reduce their energy to below 100 eV. This is done using a cold plasma of electrons. Once this cooling is complete, the trap potentials are lowered to eject higher energy antiprotons. Finally, the plasma is compressed radially using the rotating wall technique. Electric fields are pulsed sinusoidally to cause the plasma to rotate, and thus contract radially due to conservation of angular momentum.

3.3. Antihydrogen production

Following this, the antiprotons are moved into the ALPHA-g apparatus by lowering the potential wall on one end of the trap. Further axial magnetic fields ensure radial confinement of the antiprotons.

Meanwhile, a cold plasma of positrons is prepared by an apparatus called the *positron accumulator*. A sodium-22 radioactive source emits positrons with a range of energies up to 545 keV. A solid neon moderator film is used to slow these positrons. Only a fraction of positrons emerge from the moderator without annihilating, but these have a low energy and can be focused into a beam. This beam is then passed through a number of stages of nitrogen gas, where positrons lose energy through collisions. The plasma is compressed using the rotating wall technique. They are further cooled by evaporative cooling, whereby the potential wall at one end of the trap is lowered, so that particles with higher energy are able to escape while lower energy particles remain contained. At this point they are sent to the ALPHA-g apparatus, using an analogous method to the antiproton technique.

Once at the ALPHA-g apparatus, both plasmas (antiprotons and positrons) are held by Penning-Malmberg traps. They are further cooled using the methods previously described, and compressed to be of similar sizes. With the plasmas held in two adjacent and opposite potential wells, the potential barriers are then lowered so that both plasmas combine in mixing region, as shown in figure 3.1. Some number of positrons will escape; however, this is thought to be an advantage, as it leads to further cooling which offsets the energy increase caused by the mixing. With the plasmas held in an overlapping position, antihydrogen atoms are allowed to form.

A fraction of the atoms produced have magnetic dipole moments properly aligned for trapping; these are held in a magnetic minimum trap constituting of two coils and an octopole magnet. The magnetic minimum trap remains energized over a number of cycles of the antihydrogen formation process, so that a large number of antihydrogen atoms can be collected. For further information on the antihydrogen production process, see e.g. [36].

Once sufficient antihydrogen has been collected, the atoms are further cooled by a laser cooling technique known as Doppler cooling. The 1S-2P transition of antihydrogen is excited using a pulsed Lyman-Alpha laser,

3.4. Antihydrogen release and detection

which is slightly below the transition frequency. Because of this detuning, light will be preferentially absorbed by antihydrogen atoms travelling towards the laser. Since this photon absorption process also transfers momentum, and photon reemission is an essentially isotropic process, the momentum of the antihydrogen atom will decrease on average. The implementation of laser cooling in ALPHA [25] was the first demonstration of the technique for antihydrogen atoms, and will be instrumental in lowering the antihydrogen kinetic energy sufficiently for the gravity measurement.

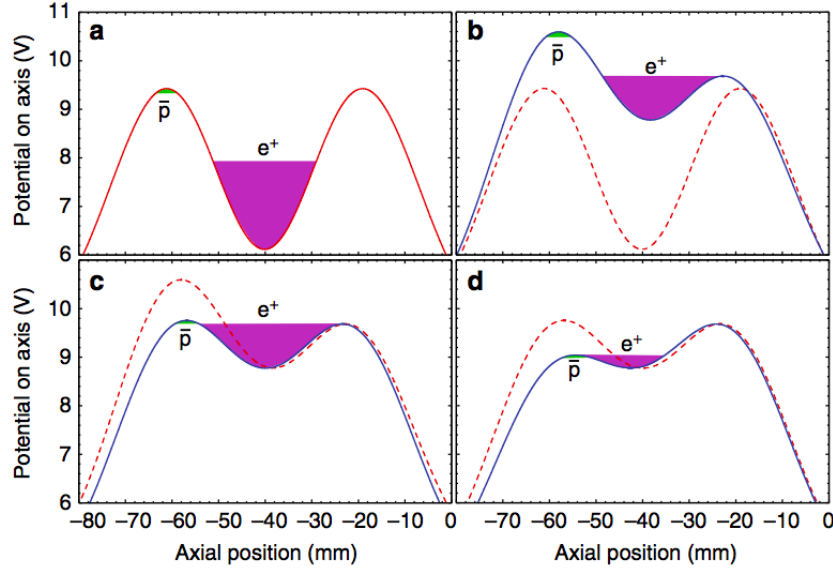


Figure 3.1: The on-axis potential during the antihydrogen production process. Note how the two plasmas are brought into contact by careful manipulation of the fields. The red dashed line shows the potential in the previous step. Taken from [36].

3.4 Antihydrogen release and detection

Once a sufficient number of cold antihydrogen atoms have been collected, they are released from the trap. This is done by a controlled power down of the magnetic minimum trap. The magnetic fields of the trap during this

3.4. Antihydrogen release and detection

process have been studied extensively, in order to minimize any “kick” which might be given to the antihydrogen atoms.

The atoms are then allowed to free fall, with their natural trajectories due to thermal kinetic energy being modified slightly by the presence of the Earth’s gravitational field. After travelling a short distance, they will annihilate on the walls of the trap. This process is described in section 2.3. The products of this annihilation will then pass through the ALPHA-g detectors described in the chapters 4 and 5.

The technicalities of measuring the strength of the gravitational interaction from the annihilation positions is considerably more nuanced than the naive “drop the anti-atoms and see where they land”. Firstly, the thermal energy of the anti-atoms is still much larger than the gravitational potential, and so under normal downwards gravity, many atoms would annihilate above the trap. Furthermore, the speed at which the magnetic field can be quenched is far from instantaneous; instead, one should consider a slow lowering of the gravitational potential barriers above and below the trap.

The simplest measurement will be performed with the magnetic potential barriers below and above the trap at the same energy. Since the barrier above the trap is (under normal gravity) at a higher gravitational potential, this upper wall will have a higher total height than the lower one when considering both the magnetic and gravitational potentials. Thus as the walls are lowered, anti-atoms will preferentially escape the trap in the downwards direction. The rough rule-of-thumb based on simulations within the collaboration is that four out of every five anti-atoms should escape downwards.

The height of the upper and lower potential barriers can also be varied individually. Another interesting measurement is when the lower potential barrier is augmented so that the total height of the two barriers is equal, including the gravitational potential. In this case, anti-atoms should escape the trap upwards and downwards in the same ratio. By measuring the distribution of annihilation positions with different configurations of the magnetic potential barriers, and comparing these measurements to simulations with different values for the gravitational interaction, the value of g for antihydrogen can be determined.

Chapter 4

Antimatter detection

Detecting antimatter particles is a surprisingly well studied problem in physics, for an important reason – it forms the basis of positron emission tomography (PET), a prolific medical imaging technique. The procedure for a typical PET study is to inject a biological compound into the patient which has been tagged with a radionuclide. This particle then undergoes a decay and releases a positron, which travels a distance of around 1 mm before annihilating with an electron in the body. This annihilation can produce two photons with opposite momentum. By detecting the two photons, the location of each annihilation can be reconstructed³, and active areas of the body can be identified [37].

The ALPHA-g detectors work by a similar principle: they allow the antimatter particles to annihilate, then track the annihilation products in order to reconstruct the annihilation vertex. The key difference, of course, is that the antimatter particles in question are antihydrogen atoms instead of positrons, and so the annihilation products are pions instead of photons. With more than two annihilation products, a more sophisticated tracking algorithm is required for the reconstruction of the annihilation sites.

This chapter will discuss the particle detection technology used in ALPHA-g to track the pions resulting from antihydrogen annihilations. For completeness, the principles of the time projection chamber will be outlined, although this detector is not the focus of the present work. More detail will be given on the plastic scintillators of the BV, and the silicon photomultipliers used to capture the resulting photons.

³More precisely, each annihilation can only be constrained to the line between the two photons. Identifying the active areas requires a large volume of data and sophisticated analysis. This is not the case for ALPHA-g; the vertex of each annihilation can be reconstructed exactly even in the two pion case, since the two pions are rarely collinear.




Figure 4.1: The principle of operation for Positron Emission Tomography. Positrons produced in active areas of the body annihilate and produce two gammas with opposite momentum, which are detected to reconstruct the annihilation site. Taken from [38].

4.1 Ionization in a time projection chamber

The time projection chamber (TPC) is a tried and tested detector which is used to record the path of a charged particle. It is typically filled with a gas mixture; a wide variety of mixtures can be used, but the ALPHA-g TPC is filled with a mixture of argon and CO_2 . When a charged particle traverses this gas, it primarily loses energy through Coulomb interactions with the large number of electrons of the argon atoms. This causes ionization of the argon atoms, creating electron-ion pairs. The gas density is low enough that typically only a fraction of the charged particle energy is deposited as the particle traverses the length of the TPC.

There are also a number of excitation interactions in the argon gas as well as ionizations, where an electron is raised to an excited state instead of

being completely freed from the argon atom. As these excited atoms decay, photons are produced. In a pure argon gas, these photons might be able to trigger additional ionizations via the photoelectric effect. For this reason, CO₂ is added as a *quenching gas*, in order to absorb these excess photons and avoid unwanted behaviour.

An electric field is applied to the gas using an array of anode and cathode wires. This prevents the electron-ion pairs from recombining; instead the electrons drift toward the anode wires and the ions toward the cathode wires. The electric field of an individual wire is proportional to $1/r$, as frequently verified by first year physics students. Consequently, as the electrons reach the anode wire, they experience a much larger electric field. The drift velocity increases greatly, and a threshold is crossed where each electron has enough energy to cause a subsequent ionization and produce a new electron ion pair. The number of free charge carriers increases exponentially in what is known as an *avalanche*. The electrons are collected on the anode wire, while the new ions are drifted back to the cathode, inducing an even greater charge on the anode wires. This charge is read out by a data acquisition system.

The key feature of a TPC is that given a constant electric field, the electron drift velocity is constant. Thus the time taken for charge to reach an anode wire is proportional to the distance of the ionizing particle from the wire. The location along the wire can also be determined by segmenting the wire length using conducting pads, such that electrons arriving at the anode wire will also induce a charge on these pads which can be read out. In this way, a three dimensional reconstruction of the path taken by the ionizing particle can be achieved.

The design of a typical TPC is shown in figure 4.2. ALPHA-g uses an unconventional TPC design, where the time projection axis is the radial one, and the anode wires run parallel to the cylinder axis. This is illustrated in figure 4.3. This design naturally brings its own advantages and challenges, which lie outside the scope of this discussion.

Much of this brief overview of the the TPC was based on a lecture series given by Doug Bryman at TRIUMF [39], although there are many resources

4.1. Ionization in a time projection chamber

available on the topic. For a comprehensive treatment, see [40].




Figure 4.2: An illustration of the TRIUMF TPC. (1) Magnet iron; (2) coil; (3) outer trigger scintillators; (4) outer trigger proportional counters; (5) end-cap support frame; (6, 8) field cage wires; (7) central high-voltage plane; (9) inner trigger scintillators; (10) inner trigger proportional wire chamber; (11) TPC end-cap proportional wire modules. Taken from [40].




Figure 4.3: Electron track simulation in the ALPHA-g TPC. Note how the anode wires are aligned with the central axis and ring the circumference of the TPC. Taken from [41].

4.2 Light production in plastic scintillators

Another commonly used technique for detecting charged particles is to use scintillators to produce photons. A scintillator is a material which takes advantage of the excitation interactions caused by an incident charged particle, where electrons in the scintillator atoms are raised to an excited state. When the atoms fall back down to the ground state, photons are produced with energy equal to the difference. These photons can then be detected and the location of the excitation determined.

There are a number of materials which can be used as a scintillator, each with their own advantages and disadvantages [42]. For example, liquid argon is often used for its scintillating properties in dark matter search experiments, due in part to its large cross section for nuclear recoils. Inorganic crystal scintillators are frequently used in collider experiments due to their large light yield and radiation hardness. For ALPHA-g, the main consideration is a short scintillation time; the lifetime of the excited state needed to be short to reduce uncertainty on the interaction time and find an accurate time of flight. To this end, plastic scintillator is the material of choice.

The particular plastic [43] used by ALPHA-g is called polyvinyltoluene, although a variety of long chain organic molecules can be used as plastic scintillator. The key structure is the benzene ring, a hexagonal arrangement of six carbon atoms, which has a unique set of molecular bonds. Three of each carbon atom's four valence electrons form sigma bonds with their neighbours. The final valence electron of each atom becomes involved in a pi bond, and due to the nature of the ring structure, these orbitals become delocalized over the entirety of the benzene ring. This delocalized orbital has a set of excited states which are especially suited for producing scintillation light.




Figure 4.4: The molecular structure of polyvinyltoluene, the plastic used as a scintillator. The benzene ring, the structure responsible for the scintillation properties, is denoted by the hexagonal ring.

The main issue with this and many other scintillators is that the wavelength of light produced by the decay of the excited pi bonds will be absorbed by these same pi bonds. As a result, the attenuation length of light within the scintillator is naturally short. To circumvent this, the plastic scintillator is doped with a fluorescent emitter called a “fluor”. These particles will absorb the light emitted from the scintillator, and fluoresce light at a different wavelength which will be able to travel more easily through the plastic.

The development of scintillators for physics experiments is an area of study with far more nuance than can be presented here. The interested reader should refer to e.g. [44].

4.3 Light collection in silicon photomultipliers

There are a number of methods which can be used to detect light coming from a scintillator. Historically, the most common method is to use a photomultiplier tube (PMT). A photon will strike a thin photocathode and will produce a single photoelectron depending on the quantum efficiency, typically around 25%. The photoelectron will then be accelerated by a series of dynodes, each held at a progressively higher voltage. Upon hitting a dynode, each electron will free additional electrons, which will be accelerated themselves. This leads to an electron cascade which eventually creates a measurable voltage. However, PMTs have some disadvantages compared to modern detectors, such as the requirement for high voltage and their

large physical footprint. One especially relevant shortcoming of PMTs is that they do not perform well in magnetic fields, as the trajectory of the electrons between the dynodes will be impacted.

A newer technology for photon collection is the silicon photomultiplier (SiPM). These rely on the semiconducting nature of the silicon lattice. Energy levels of atoms within the lattice form continuous bands rather than discrete levels. Between the highest band which is fully occupied (the valence band) and the lowest unoccupied band (the conduction band) is a small energy gap. Due to thermal energy, a small number of electrons will be excited into the conduction band, producing electron-hole pairs. A charged particle passing through the silicon can further excite electrons and create more electron hole pairs; this is the mechanism by which a silicon detector works.




Figure 4.5: Band gap structure for insulators and semiconductors. Electrons can be excited into the conduction band of a semiconductor thermally or by an incident particle. Taken from [44].

Silicon can be doped with ions or other impurities, which adds extra energy levels either directly above the valence band or directly below the conduction band. This leads to either n-type silicon, which has extra weakly-bound electrons and a proclivity towards producing electrons, or p-type silicon, which has extra ions which accept electrons, and a proclivity towards

producing holes. The junction between a region of p-type silicon and a region of n-type silicon is of special interest. Electrons move from the n-type silicon to fill the holes in the p-type silicon, until the redistribution of charge creates a static potential which opposes this movement. As a result, a depletion region forms around the junction which is mostly free of thermal electron-hole pairs.

In a typical silicon detector, a reverse bias voltage is applied to counteract the natural potential built up, and thus increase the size of the depletion region. Above a certain voltage known as the breakdown voltage, the drift velocity of electrons is high enough that they have sufficient energy to cause additional excitations and create further electron hole-pairs. This causes an avalanche of electrons which are then detected at an anode, similar to in a gaseous detector. SiPM are typically held above the breakdown voltage, so that a single photon creating a single electron-hole pair will produce a measurable signal. However, this does lead to a rise in thermal noise, as a single thermal excitation will also trigger a breakdown. As a result, most SiPM consist of a number of individual cells which are isolated from each other, so that the majority of the detector remains active at all times despite thermal noise.

Further development of silicon technologies is an area of active research, both in physics and for commercial applications.

Chapter 5

The ALPHA-g Barrel Scintillator

The ALPHA-g experiment relies on two main detector systems. The first is the radial time projection chamber (TPC) which is responsible for the tracking of annihilation products. The second is called the barrel scintillator or barrel veto (BV), which provides precise timing information of the detected particles. It also provides a fast and precise trigger for the TPC. The BV is composed of 64 trapezoidal bars of plastic scintillator which are read out at both ends by arrays of SiPM. These detector systems are concentrically arranged around ALPHA-g trap; this layout is shown in figure 5.1. The barrel veto will be the focus of this chapter, and indeed of the rest of this report.




Figure 5.1: Arrangement of ALPHA-g detector systems. The TPC lies immediately outside the inner trap, and the BV surrounds the TPC. Surrounding the BV is the solenoid magnet (not pictured).

5.1 Motivation for the BV: Time-of-flight background discrimination

The ALPHA-g experiment is located at surface level at CERN and has minimal shielding from any external particle flux. As a result, there is a large background from cosmic rays incident on the detectors. Protons and alpha particles interact in the upper atmosphere, creating pions, which quickly decay to muons travelling near the speed of light. These muons have no hadronic interaction and so are able to penetrate the atmosphere and reach the surface, their lifetimes in the laboratory frame increased by relativistic time dilation. The cosmic ray flux at the surface is well understood: the angle of incident cosmic ray muons approximates a $\cos^2 \theta$ distribution, and the rate of incidence on a horizontal area is approximately $1 \text{ cm}^{-2} \text{ min}^{-1}$ [45]. This rate is sufficient to produce a significant background over the timescale

5.1. Motivation for the BV: Time-of-flight background discrimination

of a typical ALPHA-g data run. In particular, a cosmic ray passing through the TPC will leave a track by ionizing the gas, which can be indistinguishable from a pion resulting from an antihydrogen annihilation. Should a cosmic ray pass close to (or interact with) the wall of the trap, then the TPC alone does not have the capability to recognize this as an externally incident particle instead of an annihilation on the trap wall.

The main purpose of the BV is to measure the time of flight (TOF) of particles crossing the TPC, which is one of the main methods of identifying events caused by cosmic rays. The BV is not a true veto in the traditional sense – it is not used to limit the data rate by excluding background events in real time. Instead, it is used to distinguish between cosmic rays and true annihilation products on an event-by-event basis during offline data analysis. This background rejection algorithm remains to be written, but it will be based in machine learning and will rely on a number of input parameters. The TOF will be one (very important) such parameter, but the algorithm will also use the number of tracks, angles between tracks, and proximity to the trap wall, to name a few others.

To understand why TOF is a powerful tool for background rejection, consider how the TPC signatures of an antihydrogen annihilation can be mimicked by cosmic rays. Figure 5.2 illustrates the four most common cases:

- (a) A cosmic ray passing through the detector.
- (b) A cosmic ray interacting within the detector and producing one or more secondary particles.
- (c) An antihydrogen annihilation producing two pions on the trap wall.
- (d) An antihydrogen annihilation producing three or more pions on the trap wall.

A cosmic ray muon will enter the detector on one side, ionize gas in the TPC which will be recognized as a track, produce a second track on the opposite side of the TPC (or more in the case of showering), and then exit.

5.1. Motivation for the BV: Time-of-flight background discrimination

The time delay between the entry and exit is equal to the path length of the muon divided by its speed (for our purposes, the speed of light). The time difference between the entry and exit is called the time of flight (TOF).

On the other hand, consider the products of an antihydrogen annihilation. In the case of two pions, two tracks will be produced originating on the trap wall with opposite momenta. This could be indistinguishable from the tracks of a cosmic ray passing through the trap wall. Similarly, the multiple pion case could resemble a cosmic ray shower on the trap wall.

Where these two classes of events differ is in their TOF. In an annihilation event, the pions are produced at the annihilation vertex and travel through the TPC at roughly the same speed. As a result, they will exit the TPC at roughly the same time. The *effective TOF*, the largest time difference between two exit times, will be small. Meanwhile for a cosmic ray event, the exit hit(s) will occur a measurable time delay after the entry hit, proportional to the path length taken. By precisely measuring the entry and exit times of particles in the detector, the BV will be able to distinguish annihilation products from the cosmic ray muon background based on their effective TOF.

5.1. Motivation for the BV: Time-of-flight background discrimination

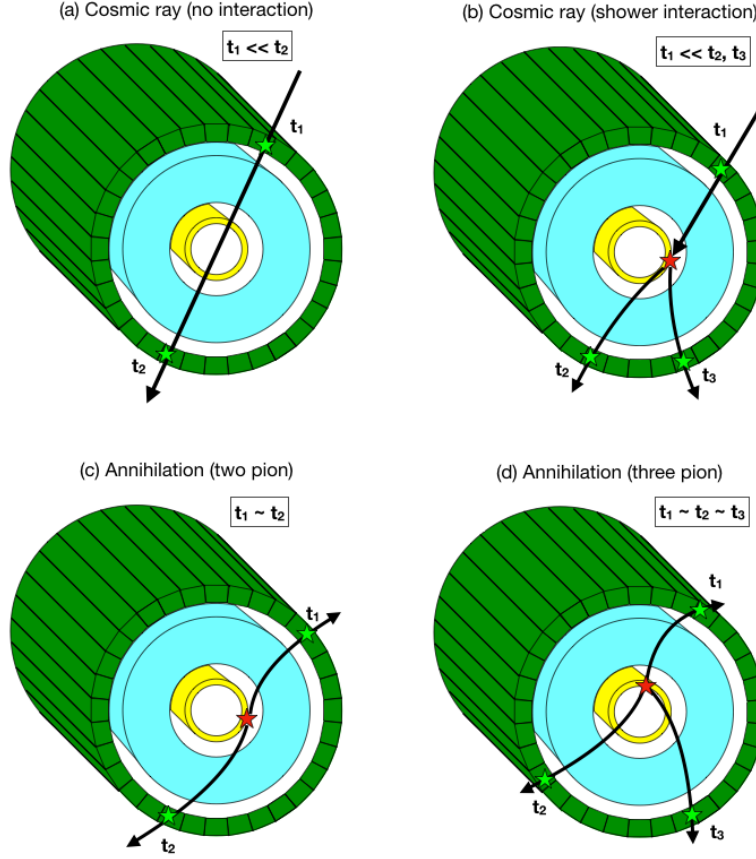


Figure 5.2: An illustration of the TOF-based cosmic ray background rejection principle. Cosmic rays hit the BV with a measurable time delay between hits, while antihydrogen annihilation produce multiple BV hits at the same time.

In addition to this important use as a background rejection tool, the BV will also perform the role of providing a fast trigger for ALPHA-g. Relying on the TPC to self-trigger is not ideal; there is significant variation in trigger time caused by variations in drift time within the gas, depending on the geometry of each event. The response of the BV to a particle is much faster than that of the TPC, since the propagation of light down the bars is much faster than electron and ion drift within the gas.

5.2 TOF time resolution requirement

Clearly, the BV must have some minimum time resolution to be able to effectively reject cosmic rays. An order of magnitude of the required time resolution is easy to calculate as follows. The ALPHA-g barrel veto has a diameter of approximately 45 cm. Consequently, a perpendicular cosmic ray muon travelling at the speed of light has a TOF of $45 \text{ cm}/c = 1.5 \text{ ns}$. Thus to distinguish external cosmic rays from annihilation events with an effective TOF close to 0 ns, the detector must be sensitive to a level of a few hundred picoseconds.

To get a more accurate account of the required time resolution, a Monte Carlo simulation was previously performed [46]. A sample of cosmic rays and a sample of antihydrogen annihilation products were simulated using Geant4 [47] in an approximation of the detector geometry. The effective TOF for each event in this simulation was convoluted with a Gaussian of predefined width to mimic the finite detector response time and electronics time resolution.

Partial results of this study are shown below in figure 5.3, which shows the annihilation event acceptance rate for a cut on the effective TOF which rejects 99% of the cosmic ray background. The study included the opening angle between tracks, i.e. the angle from one track to the other with respect to the centre of the detector. Large opening angles correspond to tracks close to the centre of the detector and the trap wall; these are the cosmic rays which more closely resemble annihilation events. We see that for most opening angles, an acceptance rate of around 97% is attainable. This is achieved with a TOF time resolution of 200 picoseconds in the BV detector. Any further improvement on this time resolution has marginal impact on the discriminating power of the the TOF.

Informed by this study, the BV must be able to provide the effective TOF of events within 200 ps. This report will cover the steps taken to attain this time resolution.

5.3. Barrel veto design

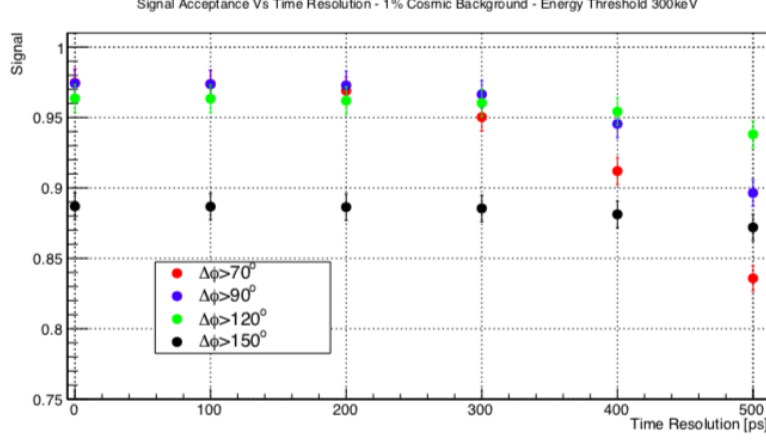


Figure 5.3: Simulated annihilation signal acceptance for 99% cosmic background rejection. Simulated TOF values were convoluted with a Gaussian corresponding to the specified time resolution. Cuts were made on total energy deposited in the TPC and the opening angle between tracks. A cut was then made on the TOF which rejected 99% of the simulated cosmic ray background, and the fraction of simulated annihilation events which pass this cut is plotted. Taken from [46].

5.3 Barrel veto design

The barrel veto was designed to measure the TOF of incident particles with this specified precision of 200 ps. It finished construction and was first tested in conjunction with the TPC at CERN in 2018.

The barrel veto is composed of a series of 64 bars of plastic scintillator. These are arranged in a cylinder, similarly to the slats of a barrel. Each bar has a slightly trapezoidal cross sectional shape so that they fit seamlessly together. This is to aid structural integrity and stability of the barrel and to ensure uniformity of the active volume. The length of each bar is 260.4 cm. When assembled, the inner diameter of the barrel is 44.6 cm and the outer diameter is 48.6 cm.

The bars are made of EJ-200, a polyvinyltoluene scintillator produced by Eljen Technologies [43]. This scintillator was designed with two properties

5.3. Barrel veto design

in mind: a long attenuation length, which makes it suitable for the 2.6 m bars, and fast timing. The rise time of the light pulse produced by incident ionizing radiation is quoted at 0.9 ns, and the decay time is 2.1 ns. The scintillation light produced has a wavelength which peaks at 425 nm, and the refractive index at this wavelength is 1.58.

Any ionizing radiation crossing a scintillator bar will produce scintillation light isotropically. This light will travel down the length of the bars, reflecting off the surface of the scintillator which has been treated to maximize reflectivity and facilitate transmission. Naturally, a photon reflecting multiple times before reaching the end of the scintillator bar will take a longer path. Thus although all the scintillation light is produced in a few ns, photons will reflect around inside the bar for far longer. As a result, the actual length of the light pulse at the end of the bars was measured to be around 40 ns. A simple diagram demonstrating this effect is provided in figure 5.4. The full effect of photon reflections will be investigated with future Monte Carlo simulations.




Figure 5.4: Reflections within the scintillator bars. Photons are produced isotropically and reflect off the bar walls. Photons taking longer path lengths are responsible for the long 40 ns pulse observed at the bar ends.

At both ends of the scintillator bar, the light signal is detected by MicroFJ-60035 silicon photomultipliers (SiPMs), which are produced by SensL [48]. These sensors are 6 mm by 6 mm in size. Each bar end is instrumented with 6 SiPMs assembled into an SiPM module, which provides approximately 50% coverage of the surface area at the end of the bar. This is a compromise between light collection efficiency and space conservation –

5.3. Barrel veto design

enough space must be left for structural support of the bars.

The chosen SiPMs have a peak light collection efficiency at a wavelength of 420 nm, which sufficiently matches the wavelength of the scintillation light. To increase light collection efficiency, 1 mm to 2 mm of RTV 615 silicone rubber was used to couple the scintillator bar to the SiPM module. This transparent rubber is used to prevent the presence of an air gap, which would otherwise cause some reflection at the interface. The inclusion of an RTV layer was shown in previous tests to increase the light transmission by a factor of 30%.

A diagram showing the arrangement of SiPM and their coupling to the bars is given in figure 5.5.




Figure 5.5: Layout of SiPM and electronics at the end of the scintillator bars. Note the arrangement of six SiPM per bar end. Not shown is the RTV layer or the signal cables.

The breakdown voltage of these silicon photomultipliers is approximately 24 V. They are operated at a reverse bias voltage of 30 V, corresponding to an overvoltage of 6 V. They provide a fast output and a standard output; the ALPHA-g BV uses the fast output. In response to a single photon, the fast output produces a pulse with a full width half maximum (FWHM) of 3.0 ns.

5.4 Readout electronics design considerations: Time walk correction

The combination of scintillator bars and SiPMs does an excellent job of producing an electronic signal when an ionizing particle crosses the detector, but designing readout electronics to record these signals has its own challenges. Clearly the readout electronics must be capable of recording signal time at the 100 ps level to provide an accurate time of flight. However, a common factor in these types of measurements which must be accounted for is called the *time walk* or *time skew* correction, and the readout must provide enough information to perform this correction.

Consider that the length of the light pulse observed by an SiPM after an incident particle is around 40 ns due to reflections within the bar. As discussed previously, the time of flight needs to be known to within around 200 ps. Thus it becomes necessary to construct a very precise and well-defined start time for the relatively long SiPM pulses.

The simplest method is to define a fixed threshold voltage, and record the time that the electronic pulse coming from the SiPM crosses this threshold. However, this method has a fundamental problem: not all pulses are of equal magnitude. A particle which loses more energy in the scintillator will produce more scintillation light, and a particle crossing closer to the SiPM will have less light attenuated in the scintillator. Given two pulses of different magnitude, even if they have the same shape and begin at the same time, the larger pulse will reach the threshold first. This effect is illustrated in figure 5.6.




Figure 5.6: “Time walk”: The effect of pulse amplitude on threshold crossing time

There are two common methods to work around this issue.⁴ The first is to simply set the threshold voltage low enough such that it will be crossed by the first photon reaching the sensor. This will effectively time the start of the signals. However, SiPMs are known to have a *dark* background rate, where an electron-hole pair is created by thermal excitations in the silicon. Setting the threshold voltage too low will lead to the thermal noise being recorded. As a result, the data acquisition system must be capable of handling high rates from the SiPMs, and another system must be implemented to distinguish the true signal events from this thermal background.

The second method, and the one implemented for the ALPHA-g BV, is to fix the threshold higher, and then subtract a correction time value in analysis in order to take the time walk effect into account. This correction value is amplitude dependent, and so it is necessary to measure the amplitude of each pulse.⁵

⁴A third method to minimize the time walk effect is the use of a technique called *constant fraction discrimination* (CFD). Here, the pulse from the SiPM is split into two. One pulse is inverted, delayed, and then added again to the first pulse. The new pulse will cross 0 V at a time which does not depend on the pulse amplitude. This method was considered for ALPHA-g, but ultimately not used. Compared to digitizing the full waveforms with an ADC, using CFD adds no redundancy while adding an extra element to the system.

⁵There are other quantities which are roughly proportional to pulse amplitude and can

Thus the requirements for the readout electronics are twofold. Firstly, it must be able to precisely record the time that the pulses from the SiPMs crosses a set threshold, such that the TOF can be reconstructed with a resolution of 200 ps. Secondly, it must independently record the amplitude of each pulse, which will be used to correct for the time walk effect.

5.5 Readout electronics design

The readout electronics of the BV contain a number of custom-built printed circuit boards and other components. This section will describe the general function and implementation of these components. For reference, a diagram detailing the readout electronics system is included in figure 5.7.

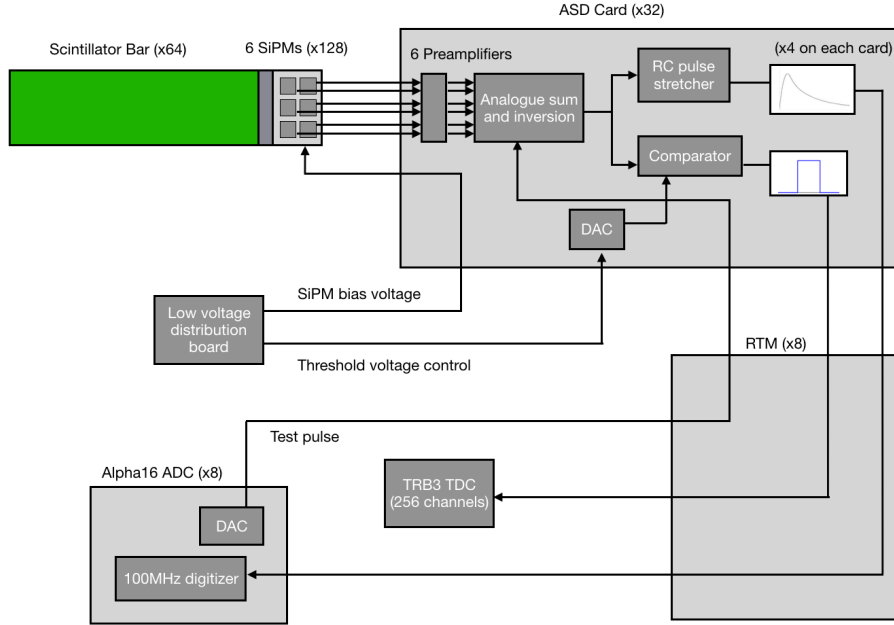


Figure 5.7: A schematic overview of the readout electronics for the BV. The ASD cards are located on the BV itself, while all other components are in a designated electronics rack.

be used instead. For example, the ALICE experiment uses *time over threshold* [49], and the integrated charge of the pulse is also commonly used.

The first stage of the readout electronics is a custom *SiPM board*, designed to hold the 6 SiPMs in place on the end of each bar. These boards collect the fast output of each SiPM and carry it away from the bar via a short (a few cm) ribbon cable.

The ribbon cable leads to a set of purpose built printed circuit boards called *analogue sum discriminator cards*, or ASD cards. These are mounted directly next to the ends of the SiPMs. Each ASD card reads out the signal from four bar ends, for a total of 32 ASD cards, 16 at each end of the BV. They contain much of the readout electronics for the BV.

After a set of preamplifiers, the first stage of electronics on each board is an analogue sum. This sums together the signals produced by the 6 SiPM at the end of a bar. Following this is an amplification stage, where this analogue signal is increased in amplitude to around 1 V to 2 V but keeps the same timing. The pulse is then inverted to give a positive waveform.

At this point, the signal is split into two pathways. The first pathway crosses a comparator, which compares the signal to a threshold voltage. The comparator produces a LVPECL digital signal which is on only when the SiPM signal exceeds the threshold voltage. This voltage is programmable and is generated by a digital-to-analogue converter (DAC) on the ASD card. It is controlled by a signal from an external low voltage distribution board, which is also responsible for providing the voltage needed for pulse amplification, as well as the bias voltage of the SiPM.

The second pathway leads to an RC filter. This acts as a low-pass filter, which removes the high frequency components of the signal pulse. This results in a stretching of the pulse. An analysis was performed to determine the transfer function of this pulse stretcher in the frequency domain. The results confirmed that the pulse is stretched by a factor of approximately 5, which was the design intention.

Both the stretched analogue pulse and the digital signal from the comparator are carried out of the ASD card via a flex cable. The signals are taken to a third custom board called the *rear transition module* (RTM). One RTM board is sufficient for two pairs of ASD cards, and so a total of 8 RTM boards are required. They are located in a VME crate in the electronics

rack. The purpose of the RTM is to further distribute the signals.

The digital signal is converted from LVPECL to LVDS in the RTM. From there, it is carried to a time-to-digital converter (TDC). ALPHA-g uses a TRB3 TDC produced by GSI [50]; a single TDC with 256 channels is used. The TDC is responsible for recording the precise time of the signal. It uses four FPGAs to create a delay line, where the signal is passed in sequence through a number of delay elements, each of which takes a certain amount of time. After the signal arrives, it propagates through the delay chain until the rising edge of the next clock cycle. By counting the number of delay elements the signal has propagated through, a precise time measurement can be made. The RMS time resolution of the TDC is quoted at 14 ps.

The stretched analogue signal is passed through the back of the VME crate to an analogue-to-digital converter (ADC). There is one ADC directly behind each RTM, for a total of 8. The ADCs used are ALPHA-16 boards, a repurpose of the GRIF-16 ADCs previously used by the GRIFFIN experiment at TRIUMF [51]. They run at 100 MHz, taking a sample of the signal every 10 ns. Sampling the full waveform in this way allows for the measurement of the pulse amplitude required for the time walk correction. Since the original pulses have a length of around 40 ns, and the sample length of the ADC is 10 ns, the need to stretch the pulses by a factor of 5 is apparent.

One final system implemented in the readout electronics is a calibration pulser. An external pulse can be generated by a DAC in each ADC. This is passed to the RTM by a lemo cable, and then distributed to each ASD card through the flex cables which also carry the signals. The pulse is injected into the signal pathway on the ASD cards directly following the first preamplification stage. The calibration pulser allows for testing and calibration of the full system.

The electronics for the BV were fully implemented by 2018, when testing and calibration runs were taken. However, it was not clear that the system was sufficient to distinguish cosmic background events using TOF. This report will show that the BV should be capable of performing the precision TOF measurements needed for cosmic background rejection, and the methods used to attain this precision in analysis will be discussed.

Chapter 6

Barrel Scintillator Analysis Algorithm Design

This chapter will discuss the design and implementation of the analysis algorithm for the data produced by the barrel veto. The goal of this analysis is to use the raw data provided by the ADC and TDC, directly following the event building by the ALPHA-g data acquisition system, to determine an effective TOF for each event which will be used to discard cosmic ray background events. This code was written and tested using results from the first barrel scintillator commissioning data run, taken at CERN during the summer of 2018. The plots in this chapter pertain to data from this run. This run used the full ALPHA-g detector system, with the barrel scintillator placed around the TPC and inside the solenoid magnets. In this section, results for a run with the magnetic field disabled are shown, although the algorithm for when the magnetic field is enabled is also described.

The general steps taken by the analysis algorithm are as follows:

1. The approximate start time and the pulse amplitude are extracted from the waveforms saved by the ADC.
2. A precise time value is determined using the information saved by the TDC.
3. ADC waveforms and corresponding TDC time values are associated together as *end hits*.
4. End hits on opposite ends of the same bar are associated together as *bar hits*.

5. The axial position of the bar hit is calculated using the TDC time values.
6. Bar hits are associated with corresponding tracks in the TPC.
7. The time of flight between each pair of bar hits is computed.

These steps will be further explained in the following sections.

6.1 ADC waveforms

Upon receiving a trigger signal, the ADC saves a waveform for all channels. A time range of 700 bins is sampled, corresponding to 7 μ s. The data acquisition system includes a buffer, so that some of these samples are from before the arrival of the trigger signal, in order to record the full waveform. There are numerous possible trigger conditions which are used by ALPHA-g for different runs. For the cosmic ray tests used to develop this analysis, two conditions were used for different runs: either a coincidence of multiple TPC channels, or a coincidence of two or more bars in the BV.

The dynamic range of the ADC is ± 1 V. It records values as 16 bit integers, and so the values range between $\pm 2^{15}$, or ± 32768 . The first step in analysis was to convert these raw *ADC values* to a voltage, by multiplying by 1 V/32768.

An example of a waveform recorded by the ADC after a cosmic ray event is shown below in figure 6.1. The detectors and electronics which lead to this distinct shape are discussed in the previous chapter. Although each SiPM and electronics channel will behave slightly differently, the differences between channels are not sufficient to create noticeable differences in the pulse shape.

6.1. ADC waveforms

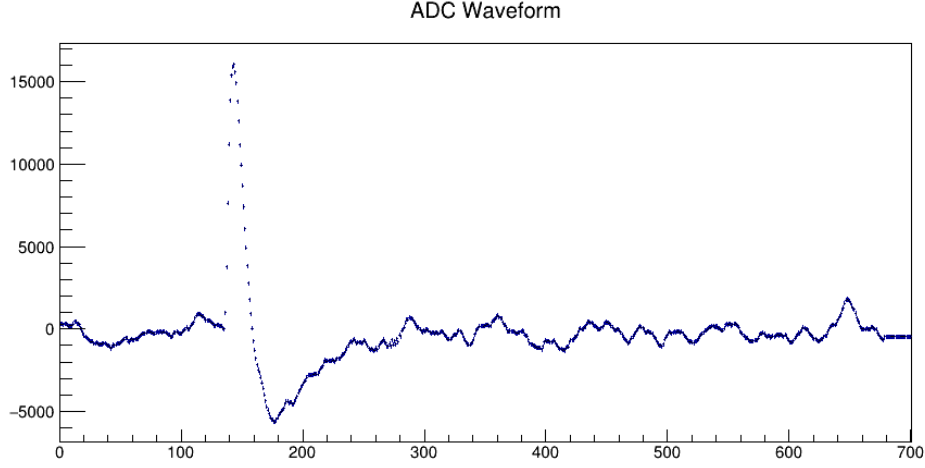


Figure 6.1: A sample ADC waveform from a cosmic ray event.

There are two quantities which the analysis is required to extract from each waveform: the approximate start time of the pulse, and a pulse amplitude to be used to correct for time walk. The pulse analysis algorithm implemented follows these steps:

1. Calculate the baseline voltage. This is done by averaging the value of the first 100 samples. The data collection window was configured to start sufficiently early before the trigger signal, such that there is guaranteed to be no pulse within the first 100 samples.
2. Determine the window where the waveform is above the baseline by a predefined threshold, currently set to 1400, or 40 mV. This was chosen to be just above the noise level for all channels.
3. Within this window, fit the waveform with a skewed gaussian function of the following form:

$$f(x) = \begin{cases} Ae^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2} & x > \mu \\ Ae^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma-k(\mu-x)}\right)^2} & x < \mu \end{cases} \quad (6.1)$$

This function was determined empirically to be an excellent fit for the

waveform produced by an incident cosmic ray. It has only four free parameters: A, μ, σ, k .

4. Subtract the baseline from the maximum of the fit function (the parameter A) to find the amplitude of the fitted pulse.
5. Interpolate the time where the fit function crosses the predefined threshold to get approximate start and end times for the pulse.

One issue which was faced while developing this algorithm is the saturation of the ADC, where pulses were larger than the ± 1 V dynamic range. When this occurred, the ADC records its maximum value of 32768. This is a fundamental issue which cannot be avoided by increasing the dynamic range or modifying the electronics. Particles hitting the bar closer to the SiPM module will necessarily lead to more photons being collected in a shorter period of time, producing a larger amplitude pulse. Increasing the dynamic range of the ADC would allow a larger fraction of the pulses to be properly recorded, in exchange for a loss of sensitivity, or an increased data rate if a more powerful ADC was used. It is not feasible to design readout electronics which can record the largest waveforms without saturating. Instead, a different solution was implemented.

This problem was circumnavigated by making a change to the waveform fitting procedure: any bins with a value of 32768 were ignored by the fitter. Thus only the unsaturated part of the waveform was fit, and the fit was able to estimate the true height of the waveform as if there was no ADC saturation. Figure 6.2 shows an example of an unsaturated fitted waveform, and figure 6.3 shows an example of a saturated fitted waveform.

6.1. ADC waveforms

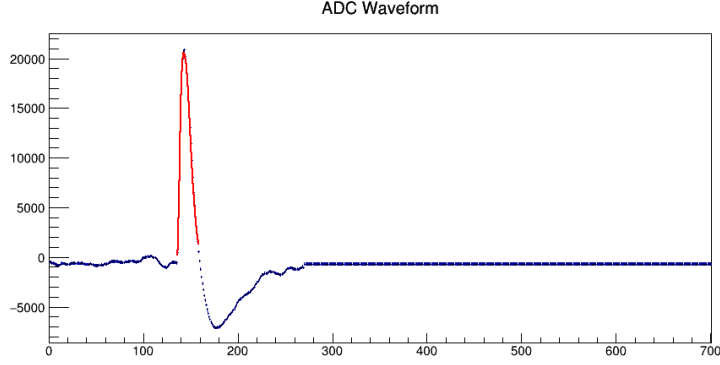


Figure 6.2: An unsaturated ADC waveform from a cosmic ray event, fit with equation 6.1.




Figure 6.3: A saturated ADC waveform from a cosmic ray event, fit with equation 6.1. The fit was modified to exclude points with the maximum ADC value, in order to ignore the flat top and circumvent the limited range of the electronics.

To demonstrate the validity of the waveform fitting procedure as a method for addressing ADC saturation, figure 6.4 shows a histogram of the measured amplitude for a sample of the cosmic ray dataset, and figure 6.5 shows the fitted amplitude for this sample. In the first figure, the peak around 32768 corresponding to ADC saturation is clearly visible.⁶ After the fitting proce-

⁶The amplitude is calculated as the difference between the maximum and the baseline, and so variation in the baseline is responsible for the width of the peak.

6.1. ADC waveforms

ture, the distribution closely matches an exponential decay, with no visible artefacts at the transition between the saturated and unsaturated regions.

Theoretically, light is attenuated⁷ in the bar according to some attenuation length λ , so that the amplitude of a pulse originating a distance x from a photosensor is given by

$$A = A_0 e^{-x/\lambda} \quad (6.2)$$

Assuming a uniform flux of cosmic rays along the length of the bars thus leads to an exponential amplitude distribution. An exponential was fit to figure 6.5, and matches the data well.




Figure 6.4: The measured amplitude of ADC waveforms for a sample of cosmic rays. The amplitude is measured from the baseline to the maximum. Note the large peak around 32768 due to ADC saturation.

⁷The value of λ is cited by [43] as 380 cm, although the effective value is expected to be lower due to photons reflecting inside the bar and not taking the shortest path.

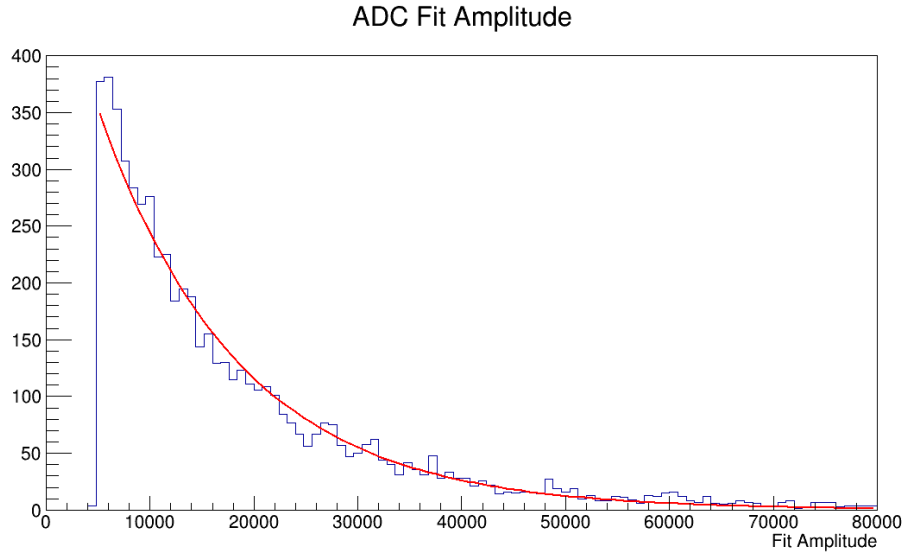


Figure 6.5: The fitted amplitude of ADC waveforms for a sample of cosmic rays. Note how the distribution extends far further on the X axis than figure 6.4, with no artefact around 32768. An exponential fit is shown.

Figure 6.6 plots the measured and fitted amplitude against each other. The two values are perfectly correlated in the unsaturated region, as expected. In the saturated region, the ADC maximum remains fixed while the fitted amplitude continues to increase. These three plots give a degree of confidence that the fitting procedure is working as expected, and that this is an adequate solution to the waveform saturation problem.

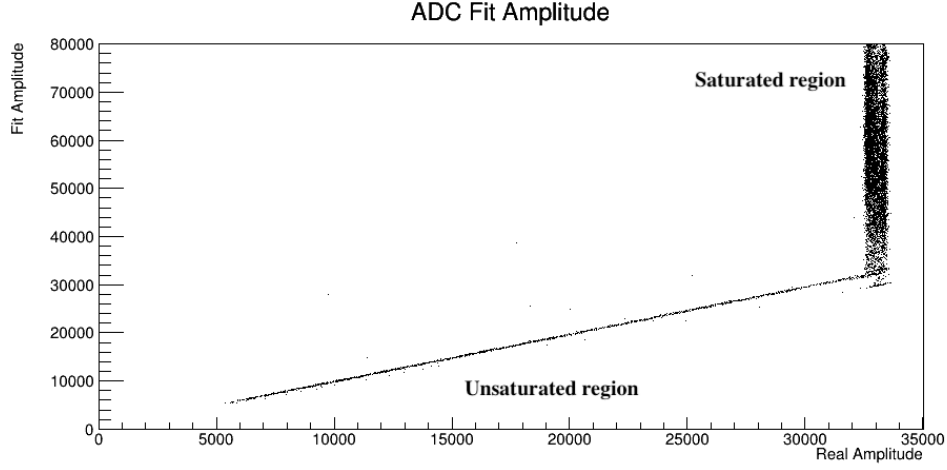


Figure 6.6: The real/measured ADC amplitude vs. the amplitude of the fit, for cosmic ray events. Notice the exact correspondence in the unsaturated region. In the saturated region, the measured amplitude stays constant due to saturation, while the fit amplitude continues to increase. The width of the measured amplitude in the saturated region is due to different baselines.

6.2 Calculation of TDC time

The TRB3 TDC records its time values as a group of three numbers, which are described in the TRB3 manual provided by GSI [52, 53]. The *coarse time counter* corresponds to cycles of a 200 MHz clock, and so an increase of 1 in the coarse time counter is equal to a delay of 5 ns. The main 200 MHz clock is built into the TDC, although the on-board clock is synced to an external clock input to ensure stability. The coarse time counter is stored in 11 bits, and so has a range of $2^{11} \times 5 \text{ ns} = 10.24 \mu\text{s}$.

To further increase this range to around 45 minutes, an *epoch counter* is included, which increments each time the coarse time counter rolls over. The epoch time counter is the most coarse timing element, and is used to generate an absolute time value. Almost all events generated by an incident particle have the same epoch time, and so this value is largely irrelevant for the purpose of finding the time difference between any two hits.

The *fine time counter* is the most precise timing element. It is generated

6.3. Event construction

by a delay line, which is a series of electronic components which each take a fixed amount of time to propagate a signal. Determining precisely how far along the delay chain the signal has propagated at the start of the next clock cycle allows the TDC to perform very precise timing measurements on the order of a few tens of picoseconds. The fine time counter records the number of elements the signal has propagated through by the next clock cycle. The simplest way to interpret this is to assume that each element in the delay line takes the same amount of time (the linear approximation). A more precise time value can be obtained by measuring the time taken for a signal to propagate through each delay element⁸, although the linear approximation works well enough for our purposes.

Using the linear approximation, the final time value recorded by the TDC for each event is calculated as:

$$t_{\text{final}} = \frac{c_{\text{epoch}}}{f_{\text{epoch}}} + \frac{c_{\text{coarse}}}{f_{\text{coarse}}} + \frac{\frac{c_{\text{fine}} - c_{\text{low}}}{c_{\text{high}} - c_{\text{low}}}}{f_{\text{coarse}}} \quad (6.3)$$

Here, the coarse frequency is 200 MHz, the epoch frequency is $1/10.24 \mu\text{s} = 97\,656.25 \text{ Hz}$, and the fine time was determined to range between $c_{\text{low}} = 17$ and $c_{\text{high}} = 450$. This was done simply by plotting a histogram of the fine time counter, and observing the edges of the distribution.⁹

6.3 Event construction

After extracting the necessary information from the ADC and TDC data flow, the analysis algorithm groups together the information into suitable data structures.

An *end hit* holds five quantities which are measured when light creates a signal in one SiPM module at the end of a bar: the time measured by

⁸The full calibration is performed by collecting a large sample of random data, and creating a histogram of the fine time counter. Some fine time values will have more counts since the corresponding delay elements take more time; the time for each delay element is inversely proportional to the number of counts.

⁹The limits of the fine time counter were observed to vary by a few units between channels. Currently the same linear approximation is used for all channels, but using a channel-by-channel approximation would yield slight improvements.

the TDC, the pulse amplitude and the approximate time measured by the ADC, and the bar number and end. One end hit is constructed for each ADC waveform. An established channel mapping is used to find the TDC channel corresponding to channel of the ADC waveform. In the case that there are multiple TDC hits in the corresponding channel within the event window, the first TDC time value is used as the time for the end hit. Care was taken to ensure that the thresholds were set high enough to minimize extra hits caused by background noise – this will be discussed later.

A *bar hit* is a grouping of two end hits, one for each end of a given bar. When a particle crosses the BV and produces a scintillation signal, all the information of this interaction coming from both ends of the bar will be recorded as a bar hit. Any end hits without a corresponding hit on the opposite bar end are discarded.

6.4 Calculation of Z position

Using the time values recorded by the TDC, it is possible to calculate the Z position (i.e. position axially along the bar) at which the interaction occurred for each bar hit. It must be assumed that light travels down the bar at a fixed effective speed, such that the time taken for light to reach the SiPMs increases linearly with the distance of the interaction from the sensor. This is certainly true to the first order, although the degree to which reflections and attenuation within the bar cause deviations from a fixed effective speed is not yet known. This will be studied in the future using ray tracing simulations.

The time measured by the TDC can be split into the sum of three major components: the time t_{int} of the interaction, the time t_{prop} taken by the photons to propagate down the bar, and the time t_{elec} taken for the signal to propagate through the electronics to the TDC. We can write the photon propagation time in terms of the distance d from the photosensor using the fixed effective speed approximation $t_{\text{prop}} = d/c_{\text{eff}}$, so that:

6.4. Calculation of Z position

$$t_{\text{TDC}} = t_{\text{int}} + \frac{d}{c_{\text{eff}}} + t_{\text{elec}} \quad (6.4)$$

If we define the Z axis to be centred around the centre of a bar of length L , then the distance d_{top} that light must travel to reach the top photosensor from an interaction at z_{int} is given by:

$$d_{\text{top}} = \frac{L}{2} - z_{\text{int}} \quad (6.5)$$

The distance travelled by light to reach the bottom photosensor is:

$$d_{\text{bot}} = \frac{L}{2} + z_{\text{int}} \quad (6.6)$$

Assuming that the electronics delay is the same between channels, we can subtract the TDC times for the bottom and top photosensors to obtain:

$$t_{\text{bot}} - t_{\text{top}} = \left(t_{\text{int}} + \frac{d_{\text{bot}}}{c_{\text{eff}}} + t_{\text{elec}} \right) - \left(t_{\text{int}} + \frac{d_{\text{top}}}{c_{\text{eff}}} + t_{\text{elec}} \right) \quad (6.7)$$

$$= \frac{d_{\text{bot}} - d_{\text{top}}}{c_{\text{eff}}} \quad (6.8)$$

$$= \frac{\left(\frac{L}{2} + z_{\text{int}} \right) - \left(\frac{L}{2} - z_{\text{int}} \right)}{c_{\text{eff}}} \quad (6.9)$$

$$= 2 \cdot \frac{z_{\text{int}}}{c_{\text{eff}}} \quad (6.10)$$

Thus we can calculate the Z position of an interaction with the bar using the two TDC time values in the following way:

$$z_{\text{int}} = \frac{c_{\text{eff}}}{2} (t_{\text{bot}} - t_{\text{top}}) \quad (6.11)$$

This calculation is carried out for each bar hit. Initially, the effective speed of light was assumed to be given by the refractive index; the scintillator bars have a refractive index of 1.58 [43], which would imply a speed of light of $0.63c$. The Z position distribution was found using this refractive index and is shown in figure 6.7.

6.4. Calculation of Z position

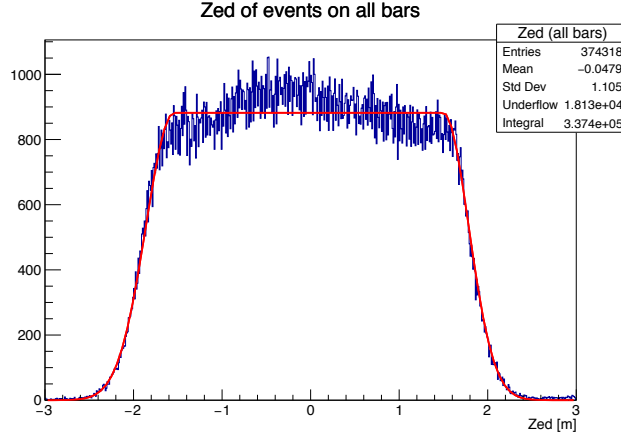


Figure 6.7: The distribution of the Z position of hits within all bars. The width of this distribution corresponds to the length of the bars. This study used an effective speed of light of $v_{\text{eff}} = 0.63c$, which leads to a slightly inflated apparent bar length.

This distribution was fit with a flat top function with Gaussian edges.¹⁰ The width of the Gaussians at the edges is 300 mm, which is consistent with a time resolution of 1.6 ns at the ends of the bars. The width of the flat top function is 3014 mm, which is considerably longer than the true length of the 2604 mm bars. The conclusion drawn from this is that the effective speed of light is significantly smaller in the bars than that naively expected given the refractive index, due to the effect of reflections with the bar discussed in section 5.3. The effective speed of light which is consistent with the observed width of the $t_{\text{bot}} - t_{\text{top}}$ distribution is $0.55c$. This effective speed of light was used in all subsequent analysis. In the future, it will be verified and further understood using ray tracing simulations.

¹⁰There is a clear non-uniformity of the flat top distribution, with more events occurring near the centre of the bar. The cause of this is not understood. The non-uniformity was not reproducible in the tests carried out at TRIUMF described in the following chapters. It will be interesting to see if this pattern is seen in cosmic ray data taken with the BV in the fall of 2021.

6.5 TPC track matching

In most cases, a charged particle with enough energy to produce scintillation light in the barrel will also produce a track in the TPC. Matching each bar hit with the corresponding track in the TPC is essential for understanding and properly treating more complicated events with multiple tracks and bar hits. This is accomplished using the following procedure:

1. Generate a 3-vector describing the position of each bar hit.
2. For each TPC track, extrapolate it until it intersects with the BV. Create a 3-vector describing the intersection position.
3. Match each bar hit 3-vector with the nearest track intersection 3-vector, if applicable.

The first stage is to find a 3-vector for each bar hit. The procedure for finding the Z coordinate is described in section 6.4, where the centre of the TPC is the origin of the coordinate system. The phi coordinate is determined using the bar number: each of the 64 bars covers an angle of 5.625° , and the hit position is set to the centre of the bar. The radial coordinate is set to the inner edge of the bar, at 223 mm.

As part of the TPC analysis, tracks through the gas chamber are reconstructed and saved for each event. They are either straight lines or helical, depending on the state of the magnetic field. In this analysis, each track is extrapolated out of the gas volume until it reaches the radius of the barrel veto. The point where it intersects this radius is saved as a 3-vector.

Finally, the two sets of 3-vectors are matched together. For each 3-vector corresponding to a bar hit, the geometric distance is calculated to each TPC track intersection point, and the closest track is selected. The same TPC track can be matched to multiple bar hits; indeed, this is expected for high angle particles which pass through the edges of multiple bars.

If the difference in phi or difference in Z between the two 3-vectors exceeds a fixed maximum value, then the barrel scintillator hit is said to have

no matching TPC track and is discarded. Otherwise, the TPC track intersection point is saved into the bar hit data structure. For each hit, the difference in Z and phi between the two 3-vectors was plotted (not shown). The maximum differences were chosen in order to fully encompass the main peak in each of these distributions. Currently, the maximum permitted phi difference is 17° (three bars) and the maximum Z distance is 300 mm.

6.6 Calculation of time-of-flight

The final step in the analysis algorithm is to calculate the time of flight between any two bar hits. The calculation is similar to that of the Z position. For a given bar hit, the mean of the TDC time measured at the two ends is:

$$\frac{t_{\text{bot}}^A + t_{\text{top}}^A}{2} = \frac{\left(t_{\text{int}}^A + \frac{d_{\text{bot}}^A}{c_{\text{eff}}} + t_{\text{elec}}\right) + \left(t_{\text{int}}^A + \frac{d_{\text{top}}^A}{c_{\text{eff}}} + t_{\text{elec}}\right)}{2} \quad (6.12)$$

$$= t_{\text{int}}^A + t_{\text{elec}} + \frac{L}{2c_{\text{eff}}} \quad (6.13)$$

where L is the total length of the bar. Now suppose that an event has two bar hits, on different bars. Subtracting the average time of the two bar hits yields the time of flight:

$$\frac{t_{\text{bot}}^A + t_{\text{top}}^A}{2} - \frac{t_{\text{bot}}^B + t_{\text{top}}^B}{2} = \left(t_{\text{int}}^A + t_{\text{elec}} + \frac{L}{2c_{\text{eff}}}\right) - \left(t_{\text{int}}^B + t_{\text{elec}} + \frac{L}{2c_{\text{eff}}}\right) \quad (6.14)$$

$$= t_{\text{int}}^A - t_{\text{int}}^B \quad (6.15)$$

This calculation makes the assumption that the effective speed of light is constant within the bar. It makes the additional assumption that this speed is the same for all bars, and that all bars are the same length and have the same electronics delay.

The time of flight is calculated in this way between every pair of bar hits.

6.6. Calculation of time-of-flight

This will be used in the future as a main input into a cosmic ray background discriminator.

The time of flight distribution at this stage is shown below in figure 6.8. This plot includes all events with two or more bar hits matched to separate TPC tracks, and includes the TOF between all combinations of hits within these events.




Figure 6.8: Time of flight distribution for cosmic rays at CERN

At first glance, this seems reasonable; the largest possible TOF for a particle travelling diagonally across the length of the BV is around 10 ns. However, it is devoid of any complex features which might be expected due to the geometry of the BV. Understanding and verifying this observed TOF distribution is the main motivation for setting up a test system at TRIUMF, which will be discussed in the following section.

Chapter 7

The Vertical Slice

A *vertical slice* of a system is a small scale replica often used for testing or benchmarking. It has one copy of each component of the system so that functionality may be tested without requiring the entire system, which might have hundreds of copies of each component. This section describes the implementation of a vertical slice of the BV at TRIUMF.

7.1 Motivation: Mysteries in the CERN TOF data

The necessity of further testing and troubleshooting of the BV system became apparent after further analysis of the commissioning data taken at CERN – there were a number of indicators that the response to cosmic rays was not fully understood.

One particular issue arose when comparing the geometry of each cosmic ray event to its TOF. Only events with the simplest topology were considered: those with two hits on the BV and two tracks in the TPC, with the hits and tracks successfully matched together. These correspond to a single cosmic ray crossing the BV, the TPC, the central trap volume, and then exiting through the TPC and BV. With this cut in place, it is simple to calculate the geometric distance between the two hits on the barrel scintillator, using the positions determined either through extrapolating the TPC tracks or using the barrel scintillator itself. Using this geometric distance and the assumption that the cosmic ray is travelling at the speed of light, an expected time of flight was calculated. However, when comparing this expected time of flight to the time of flight measured by the barrel scintillator,

7.1. Motivation: Mysteries in the CERN TOF data

the two did not correlate at all, as shown in figure 7.1.

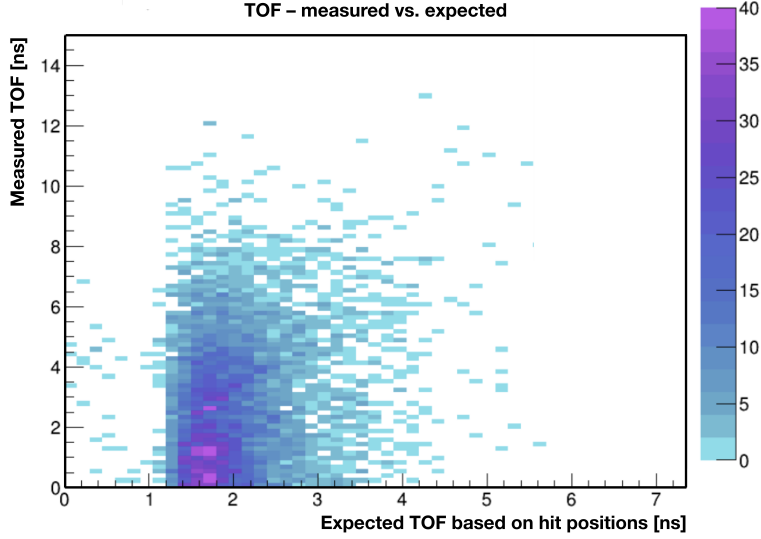


Figure 7.1: Measured TOF plotted against TOF expected from the geometry of hits. Data was taken at CERN using cosmic rays. Expected TOF was calculated as $(\text{distance between hits})/(\text{speed of light})$. Notice how there is no apparent correlation.

The most obvious feature of this plot is that the two variables are not correlated at all, as mentioned above, where we would expect a full correlation. Another feature of note is that the expected TOF has a minimum value below which there are no events. This is due to the requirement that there are two distinct tracks in the TPC, which necessarily have some distance between them (otherwise they would be merged into one track), and that the BV hits are matched to these two separate tracks. This cut-off should be present in the measured TOF as well; however, the distribution extends all the way down to zero on the Y axis. These major discrepancies made it clear that the time of flight values being measured had some issues. In order to investigate, a vertical slice of the system was set up at TRIUMF to perform further tests.

7.2 Implementation

The TOF system at TRIUMF is a replica of the full BV and its associated electronics, but with only the minimum number of each component necessary to function. Two shorter 70 cm scintillator bars were used in place of the full 64 2.6 m bars. The electronics associated with the bars are an exact copy of the electronics for the full BV, as summarized in figure 5.7. Each of the four bar ends was instrumented with an ASD card, which was powered by a single low voltage distribution board. Only one RTM board was needed for the two pairs of ASD cards, and subsequently only one ADC. Finally, an identical TRB3 TDC was used, and the test pulser line was set up to provide calibration pulses in the same fashion. For the majority of the following tests, the two bars were placed parallel to each other and stacked directly on top of each other.

Included in this section are some annotated photographs of the vertical slice.




Figure 7.2: Layout of the scintillator bars in the vertical slice.

7.2. Implementation




Figure 7.3: Low voltage distribution and control board




Figure 7.4: Analogue sum discriminator (ASD) card front side




Figure 7.5: Analogue sum discriminator (ASD) card reverse side

7.2. Implementation




Figure 7.6: Rear transition module (RTM)




Figure 7.7: Alpha-16 analogue to digital converter (ADC)




Figure 7.8: TRB3 time to digital converter (TDC)

7.3 Initial results

Following the implementation of the vertical slice at TRIUMF, a similar analysis procedure was developed to the one described in the previous chapter. There are two key differences between the analysis for the full BV and the analysis for the vertical slice. Firstly, there was no matching to TPC tracks implemented, since the bars were set up in absence of a TPC. Furthermore, with only two bars, there was no need to cut on event topology to focus only on the “simple” events. Any event with coincidence between the four bar ends was included in the analysis.

It is also worth noting that the two physical quantities we can measure with the BV, the Z position of cosmic rays and their TOF, become useful diagnostic tools with only two bars. With the two bars directly stacked, the Z position measured for vertical cosmic rays should be the same in both bars. In practise, only a small difference between the two Z measurements which is expected from cosmic rays travelling with a diagonal trajectory.

7.3. Initial results

Consequently, the difference between the two Z measurements can serve as a test of the Z position resolution of the bars.

In addition, with the two bars directly adjacent to each other, the time of flight between the two hits should be very small. A cosmic ray traveling the distance of around 2 cm at the speed of light takes 66 ps. Again, only very small variations are expected due to diagonal trajectories. As a result, the TOF measurement with the two bars together provides a means of determining the TOF resolution of the BV.

The initial time of flight distribution taken using cosmic rays is shown in figure 7.9.

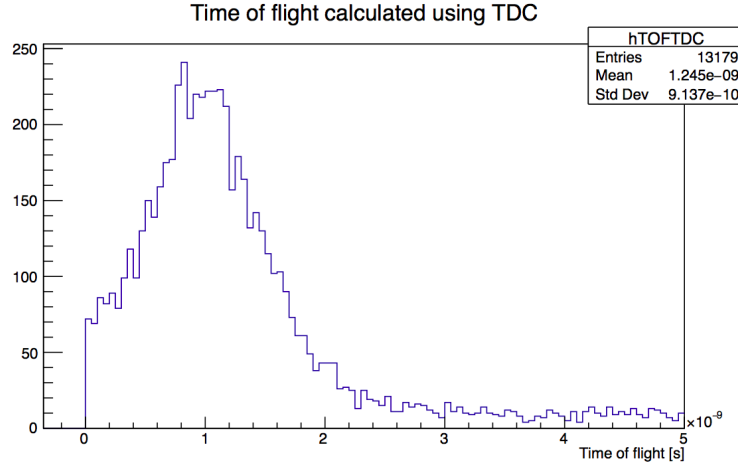


Figure 7.9: Unsigned TOF distribution for cosmic rays in the vertical slice. Note the 1 ns offset of the peak from zero, the large width of the distribution, and the number of *grass* events outside the main peak.

There are three features of note in this plot. Firstly, the mean TOF is around 1 ns, which is far larger than expected – this indicates an offset between channels which will be addressed in the following section. Secondly, the standard deviation of the distribution is also around 1 ns, which is far larger than the required 200 ps resolution. Steps taken to improve this resolution will also be discussed in the following section. Finally, there are many *grass* events which have TOF extending far beyond the main distribution.

7.3. Initial results

As described in section 6.4, the time difference between the two ends of a bar is proportional to the Z position of the event along the bar. For coincidence events between the two bars, the time difference between the bar ends for each bar are plotted against each other in figure 7.10.




Figure 7.10: End-to-end time difference comparison between bar A and bar B of vertical slice. This time difference is proportional to the Z position of the hit, and so we expect the two quantities to be very well correlated.

The two time differences (and therefore the two Z positions) are very well correlated, as expected. The width of this distribution along both axes is around 10 ns. For a bar length of 70 cm, this corresponds to an effective speed of light of $v_{\text{eff}} = 0.47 \times c$. This conflicts slightly with the speed of $v_{\text{eff}} = 0.55 \times c$ measured in the full BV, although both are smaller than the speed of light given by the refractive index of the scintillator. It is not surprising that the effective speed of light is different for the two detector systems, given the significant length difference between the two sets of scintillator bars.

It is also worth noting that there is an offset from zero of about 1 ns for both bars. This is further evidence for a time offset between channels.

Another useful test which can be performed is to cross the scintillator

bars, as shown in figure 7.11.



Figure 7.11: Crossed bar layout for the vertical slice. In this layout, coincidence events are localized to the overlap of the two bars.

Setting up the scintillators in this way and using a coincidence trigger between the two bars effectively limits the active area to the intersection of the two bars, an area roughly 2 cm by 2 cm in size. With this layout, the Z position of all events is known to be approximately the same, and so the width of the Z position distribution gives an approximation of the Z position resolution of the bars. The time difference distribution (which is proportional to the Z position) is shown in figure 7.12.

The standard deviation is 1.3 ns, which is similar to the time resolution found using the TOF. Using $v_{\text{eff}} = 0.47 \times c$, this corresponds to a distance of 18 cm. The other features of the TOF plot also reappear here: namely, the offset from zero and the *grass* events.

7.3. Initial results




Figure 7.12: End-to-end time difference using crossed bars. Shown is bar B, although bar A is similar. The width of this distribution is due to the time resolution, as well as the width of the overlap area between the bars.

Since the vertical slice is a replica of the true BV, all the features discussed in this section should also affect the cosmic ray data taken with the full BV. However, the full data were further obfuscated by the large number of channels in the BV. Constructing the vertical slice was very useful in identifying these issues in a smaller-scale environment. For example, there is no guarantee that the offset observed in figures 7.9 and 7.10 is the same for all pairs of bars in the BV. This would result in a smearing out of the TOF depending on which of the bars were involved in an event. This would explain the effects discussed in section 7.1. The next section will focus on how the vertical slice was used to address these problems and fully prepare the BV.

Chapter 8

Time corrections and calibrations

8.1 Optimal TDC threshold

After the initial setup, the first step in calibrating the vertical slice was to set the threshold voltage described in section 5.5. With the threshold voltage set to a reasonable value, each ADC waveform should correspond to exactly one time value produced by the TDC. However, with the default threshold voltage set too low, electronic noise caused the threshold to be crossed multiple times in the event window, leading to a number of TDC hits per ADC waveform. This was observed in the CERN data; difficulty in selecting the correct time value is suspected to be a cause of the strange behaviour discussed in section 7.1.

This behaviour was quickly reproduced with cosmic rays in the vertical slice. The default threshold voltage set by the run control software was 0 V, which is obviously low enough to trigger on the slightest electronic noise. There is no record of the threshold voltage used in the CERN run, so it seems likely that it was also set to 0 V.

The trigger used was a logical OR across all ADC channels, such that the DAQ would create an event whenever the waveform recorded by any ADC channel reached a certain amplitude, recording all ADC and TDC data within a set window. Figure 8.1 shows the number of TDC time values recorded per ADC waveform in the vertical slice with the default threshold voltage.




Figure 8.1: The number of TDC time values corresponding to each ADC waveform. Note how for one ADC waveform, the corresponding TDC channel can record up to 25 time values. This is due to the default threshold being below the electronic noise level.

To further investigate these extraneous time values, the time of each TDC hit was plotted for each ADC waveform. The approximate ADC time as determined in section 6.1 was used as the X axis. The results of this are shown in figure 8.2 and have some interesting ramifications.

The first feature to note is that there is a good correlation between the two times, at least within the main group of data. This indicates that the TDC time values and the approximate times generated by fitting the ADC waveforms are in agreement. The width of the line is due to the low time resolution of the ADC.

Another obvious feature is that there are two clear lines of data. This indicates that for each ADC waveform, two time values are being recorded by the TDC: the true time value, and a second time value approximately




Figure 8.2: Plot of TDC time value vs. approximate start time of ADC waveform. Each ADC waveform corresponds to multiple TDC time values, i.e. there are multiple values of Y for each value of X. Data is from cosmic rays in the vertical slice.

200 ns later. The cause of this is not fully understood, although it is very likely to be due to an electronic reflection within the signal cables. This behaviour was also seen in the data from the full BV. A sufficient solution is to simply use the first of these two time values in analysis and discard the second, although this is obviously not ideal and this behaviour should be investigated in the future.

The third interesting feature is that there is a group of events which happen around 200 ns earlier in both the ADC waveform and the TDC time value (these are the events on the far left of figure 8.2). This is related to the time chosen by the DAQ to be the start time of the event; for these events, the event trigger and TDC start happen 200 ns later in relation to the time of the SiPM pulse. Since we are interested only in the difference between times, the DAQ event start time is not relevant to this work, and so this 200 ns shift in start time was deemed interesting but unimportant.

8.1. Optimal TDC threshold

The final and most obvious feature of this plot is the wash of background points away from the two main lines. These are the extraneous TDC time values caused by the threshold voltage being set below the electronics noise level.

Removing these extra data points was done by reconfiguring the threshold voltage to be just higher than the noise coming from the SiPMs. This voltage was determined via a scan: the threshold voltage was increased from 0 to 30 mV in steps of 1 mV. At each step, five minutes of cosmic ray data was taken. The number of counts on each channel was recorded.

Each ASD card is designed to take data from four bar ends, but only one bar was connected to each ASD card. Consequently, during this scan data was recorded from four channels connected to the SiPMs and scintillator bars, as well as from twelve channels which were not connected to SiPMs. The TDC still recorded time values in these channels, due to electronic noise or interference surpassing the threshold voltage.

The results of the scan for the four SiPM-connected channels and the twelve disconnected channels are shown in figure 8.3.

For the connected channels, as the threshold is increased, the number of TDC events decreases until all the channels converge to the same value. This value is presumably the true rate of cosmic rays passing through the 70 cm bars multiplied by the detection efficiency: 28000 events per 5 minutes, or 93 events per second.¹¹

In the disconnected channels, there is an initial increase in the count rate as the threshold is increased. This is possibly due to the higher threshold voltage creating more electronic interference, although this effect is not well understood and not hugely relevant.

In both cases, all background events stopped across all channels at a threshold voltage of 25 mV. The threshold voltage was set to this value for all channels in the vertical slice. It is also worth noting that as the threshold is increased further past this value, the number of TDC events

¹¹Actually, this number is twice the cosmic rate due to the double counting seen in figure 8.2. This is still much higher than the true cosmic rate cited in section 5.1, and it is unclear why the flux measured was so high.

8.1. Optimal TDC threshold

remains constant and does not decrease, implying that no cosmic ray events are being excluded by the threshold.




Figure 8.3: Cosmic ray count rates during scan of threshold voltage. Threshold voltage is in Volts. Four channels were connected to SiPMs (left), while twelve were left unconnected (right). Counts are due to electronic noise exceeding the threshold voltage. Note how all count rates converge to the same value after 25 mV, implying that the threshold has passed the noise level but does not exclude any cosmic ray signal.

After setting the threshold voltage, another set of cosmic ray data was taken and figure 8.2 was reproduced, shown in figure 8.4. The random wash of TDC time values was successfully removed. Now, each ADC waveform is accompanied by exactly two TDC time values: the true one and an additional later time value due to reflections. This second time value is ignored in analysis.




Figure 8.4: Time of TDC time values vs approximate ADC waveform start time, after setting TDC thresholds. Note that the background wash of events has disappeared, and now exactly two time values are recorded for each event. The second is due to reflections in the electronics and is ignored.

8.2 Channel-by-channel offset calibration

Another problem which needed to be addressed is the offset in the TOF plot (figure 7.9). As described in section 7.3, the time of flight between the two bars should be very close to zero when the bars are placed directly next to each other. Instead, the mean TOF is close to 1 ns.

As shown in this section, the reason for this offset was discovered to be a systematic delay which is different for each channel. This delay could in theory be introduced at any point along the electronics path between the SiPM and the TDC.

In order to determine the characteristic delay for each channel, data was taken using the calibration pulser described in section 5.5. A square pulse was generated by a DAC on the ADC card and distributed via the RTM to each channel simultaneously. The pulse is injected on the ASD card,

directly after the preamplification stage following the SiPM. Although only four SiPMs were connected, the pulse was injected into all four channels on each of the four ASD cards. The injected pulse then followed the same electronics pathway as a real signal from an SiPM would.

The TDC time value for all hits was computed. For each hit on channel N , the time difference between the time value on channel N and the time value on channel 8 was found. This should be zero if all channels are properly calibrated. The reason for using a time difference here is that there is no stable reference or trigger time which is as precise as the TDC. All the analysis is carried out using time differences within the TDC as described in chapter 6, and so using an absolute event time would introduce an external factor into this calibration.

These time differences are shown in figure 8.5.




Figure 8.5: Time offsets between TDC channels in the vertical slice, using the calibration pulser to send a pulse to all channels simultaneously. All time differences are relative to channel 8, hence why the values for channel 8 are zero. The two peaks in channel 6 are not understood; this channel was never connected to an SiPM.

8.2. Channel-by-channel offset calibration

The width of the peaks within each channel is around 80 ps. However, the time value recorded for simultaneous pulses differs by up to 2 ns between channels.

The median time difference¹² for each channel was recorded to a file and used for calibration of the vertical slice. For any given TDC time value in all following analysis, a channel-by-channel time correction is applied. The new time value is calculated as

$$t' = t_{\text{TDC}} - \Delta t_{\text{chan}} \quad (8.1)$$

where Δt_{chan} is the mean value in the corresponding bin of figure 8.5.

With this correction in place, the time difference between channels using simultaneous pulses is shown in figure 8.6.

At least for simultaneous test pulses, the correction has removed any time delays between channels. It is worth pointing out that channel 6 exhibits a behaviour not present in any other channel, where some time values are recorded 3 ns to 4 ns before the main peak. This channel also showed a large amount of noise, for example firing on occasion even when the pulser was disconnected. The second peak in channel 6 is not fully understood, but since it was never connected to an SiPM it never interfered with these results.

Considerable effort was put into understanding the nature of the channel-specific delays. Numerous pulser runs were taken, replacing or swapping in turn each of the electronics components – the ASD cards, the signal cables, the RTM, and the TDC channels used. The results were conclusive. The ASD cards, RTM, and signal cables all had no perceptible effect on the channel-specific delays, and so it can be concluded that path length differences in any of these components are negligible.

Recall from section 5.5 that the TRB3 TDC has a total of 256 channels. In the vertical slice with only 4 ASD cards, only 16 of these channels were in use. This is seen in figure 7.8, where only one signal cable is connected

¹²The median was used because it produced more acceptable behaviour for the defective channel 6. The mean could be used instead.

8.2. Channel-by-channel offset calibration

to one of 12 possible This caused the offsets to change entirely. To see this, compare figure 8.5 to figure 8.7, which was taken using a different set of inputs on the TDC.




Figure 8.6: Time offsets between TDC channels in the vertical slice, using the calibration pulser to send a pulse to all channels simultaneously, after the channel-by-channel offset correction was applied. Note how the channels are brought into line with each other compared to figure 8.5. The variation across all channels is now around 80 ps.




Figure 8.7: Time offsets between TDC channels in the vertical slice, using the calibration pulser to send a pulse to all channels simultaneously, using a different set of TDC inputs to figure 8.5. Note how the delay for all channels take on entirely new values.

The conclusion reached is that these offsets are due entirely to internal path length differences in the TDC. It is possible that there is a more elegant way to calibrate these, but none is suggested in the manual provided by GSI [52, 53]. For the purposes of the BV, it is sufficient to correct for this effect in analysis, using the method described above.

8.3 Time-walk correction

With the systematic delays between TDC channels dealt with, the most important remaining correction necessary for improving the TOF sensitivity is the time walk correction described in section 5.4. Since the pulses have a long rise time and the readout uses a fixed threshold voltage, pulses with smaller amplitude will cross the threshold later. Thus an amplitude-dependent correction term must be added to disentangle the amplitude dependance of the hit time. The size of this correction for any given pulse amplitude must be determined using data; this section describes the process used to calibrate this correction.

8.3.1 Form of the time walk correction

Applying an amplitude-dependent correction term is a common practise in analysis for fast TOF detectors. Examples of this correction being applied are given in e.g. [54–60].

The vast majority of sources use a correction term given by

$$\delta t = \frac{W}{\sqrt{A}} \quad (8.2)$$

where A is the pulse amplitude, and W is a parameter which must be fit using data. This form arises by assuming that the SiPM signal shape is given by

$$V = A \left(\frac{t}{t_R} \right)^2 e^{-2\left(\frac{t}{t_R}-1\right)} \quad (8.3)$$

i.e., a quadratic rise and an exponential decay, both determined by a rise time t_R . Assuming the behaviour at the start of the pulse is quadratic leads to

$$V = A \left(\frac{t}{t_R} \right)^2 \quad (8.4)$$

$$\Rightarrow \sqrt{\frac{V}{A}} \cdot t_R = t \quad (8.5)$$

Thus for a time measurement taking place at $t = t_{thr}$ when the voltage reaches the threshold $V = V_{thr}$, the true start time $t = 0$ can be recovered by subtracting

$$t_{thr} = \sqrt{\frac{V_{thr}}{A}} \cdot t_R \quad (8.6)$$

from the measured time. Here, $\sqrt{V_{thr}}$ and t_R combine into the parameter W in equation 8.2.

This shape of equation 8.3 was fit to a cosmic ray waveform in figure 8.8. It is a good match at least for the rising edge of the pulse; the trailing

8.3. Time-walk correction

edge of the pulse deviated from the exponential decay due to the negative tail introduced by the amplification and stretching electronics. With this functional form being a reasonable match for the leading edge of the observed cosmic ray waveforms, the form for the time walk correction given by equation 8.2 was comfortably applied.¹³




Figure 8.8: The pulse shape given by equation 8.3, fit to a cosmic ray waveform similar to the one shown in figure 6.1. Observe that the match is a reasonable fit at least on the leading edge, and so the form of the time walk correction in equation 8.2 is justified.

8.3.2 Calibration schemes for time walk correction

Finding the value of W for the time walk correction requires precise knowledge of three quantities for a large number of events: a precisely defined start time t_0 which is amplitude independent, the exact time t when the waveform crosses the threshold voltage, and the amplitude A of the waveform. These quantities are illustrated in figure 8.9. Once they are known,

¹³As discussed in the following section, the waveform recorded by the ADC is not the waveform seen by the comparator. The pathway to the ADC includes an additional RC pulse stretcher, which removes the high frequency components of the pulse, stretching it temporally. However, this transformation preserves the functional form of the pulse, and so showing that the waveform recorded by the ADC has the proper functional form implies that the pulse seen by the comparator does as well.

8.3. Time-walk correction

$\delta t = t - t_0$ can be plotted against $1/\sqrt{A}$, and the slope will be W . In any further analysis, $\delta t = \frac{W}{\sqrt{A}}$ can be used to determine the time correction applied for each waveform based on its amplitude.




Figure 8.9: Quantities needed to calibrate the time walk correction. There are three quantities which must be known for a large number of events: the amplitude of the pulse A (here a), the true amplitude-independent start time t_0 , and the measured time where the pulse crosses the threshold t .

In fact, t_0 does not necessarily have to be the true start time of the pulse. It can be any stable reference time, as a constant added to t_0 would be cancelled out when taking the time difference in the TOF calculation. Many experiments use a third scintillator or other similar detector to provide a common start time for the TDC [55, 56, 59]. In the case of the BV, there is no easy method to devise a stable reference time or common start time. The time recorded by the TDC necessarily has an event-by-event jitter due to a trigger time which is not stable.¹⁴ This is removed when finding the difference between two times in the event, and so is not an issue for the

¹⁴The trigger in this case is based on the ADC. When the signal entering an ADC channel exceed a certain threshold, the DAQ will save all data from a set time window. The jitter is caused by the low time resolution of the ADC, among other DAQ-related reasons.

TOF measurement. However, it does make it impossible to devise a stable t_0 using the TDC in order to calibrate the time walk correction.

At this point, it was considered if the ADC can be used to calibrate the time walk correction, as the ADC does record the full waveform. However, there are two reasons why it is not suitable for finding the time correction. Firstly, the time resolution of the ADC is too imprecise. Each time bin has a width of 10 ns, and so even with interpolation, it cannot be used to find the time correction with any degree of precision. The second reason is that there is an RC pulse stretcher before the signal is sent to the ADC, as described in section 5.5. This is not present in the pathway leading to the threshold comparator and the TDC. This means that the time correction found using the ADC would not be for the same pulse as the time correction which needs to be applied to the TDC.

Some time was spent trying to precisely find the transfer function of the RC stretcher. The intention was to find the time correction term using the ADC, and then use the inverse transfer function to determine the true time correction which should be applied to the TDC time. However, this approach was met with a fundamental issue. The RC stretcher acts as a low pass filter, removing all high frequency components. This is not a reversible transformation – the high frequency components were set to zero, and so a multiplication in frequency space cannot effectively recover them. Consequently, it is difficult to take the full waveform recorded by the ADC and revert it to the unstretched pulse which passes through the comparator.

8.3.3 Demonstration of time walk using oscilloscope

Instead of using the ADC, an oscilloscope was used to measure the signal directly on the ASD cards in a cosmic ray run. This was not intended to be used for calibration – it is too imprecise, and is not replicable for the full BV. However, the oscilloscope was able to show that equation 8.2 is a suitable form for the time walk correction.

The signal was measured simultaneously at two test points. The first test point is directly before the comparator which compares the signal to

8.3. Time-walk correction

the generated threshold voltage and sends the resulting discriminated digital signal to the TDC. The second test point is after the RC pulse stretcher, before the signal is sent to the ADC. These test points are shown in figure 8.10. The reader should also refer to figure 5.7 to see where these points fit into the electronics pathways.




Figure 8.10: Test points on the ASD card used for measuring the time walk effect. These were measured using an oscilloscope during a cosmic ray run. Test point 1 was used to record the threshold crossing time and true start time for each pulse, directly before the comparator. Test point 2 was used to record the amplitude of the pulse after the RC stretcher.

An analysis was done using around 1000 waveforms measured in this

way. The waveforms at the two test points were digitized simultaneously using two channels of the oscilloscope. There were significant oscillations in both waveforms due to the coupling of the oscilloscope to the electronics. To smooth the pulses, a standard Savitzky-Golay filter was applied. For each data point, the filter fits a small subset of adjacent points with a low degree polynomial, and uses the value of this polynomial to correct the data point in question. It repeats this for each data point, resulting in a smoother waveform free from local oscillations.

The sample rate of the oscilloscope was 2 GHz, leading to time bins of length 0.5 ns. In order to obtain a finer time resolution, a curve was interpolated from the smoothed data points. The resulting curve is shown for a waveform from the first test point in figure 8.11 and for a waveform from the second test point in figure 8.12.

The time when the first signal crosses the threshold voltage of 25 mV was subtracted from the time when it crosses a value of 1% of the threshold voltage, to give an approximation of the time correction which should be applied. This is shown as the difference between the two purple dashed lines in figure 8.11. The pulse amplitude was found using the second, stretched signal. This is the dotted purple line in figure 8.12. Then these two quantities were plotted against each other. Figure 8.13 shows this plot for 29 V and 30 V SiPM bias voltages. A fit of the form given in equation 8.2 was applied to each of these plots. These fits were a reasonable approximation of the data, providing further evidence that this is a suitable form for the time walk correction.

8.3. Time-walk correction




Figure 8.11: Waveform measured using an oscilloscope at the first test point shown in figure 8.10. The yellow line is the measured waveform, while the red line is the curve resulting from the application of a Savitzky-Golay filter and interpolation. The intersection of this curve with the value of the threshold voltage (25 mV) and 1% of this value (0.25 mV) are marked with purple dashed lines.




Figure 8.12: Waveform measured using an oscilloscope at the second test point shown in figure 8.10. The yellow line is the measured waveform, while the blue line is the curve resulting from the application of a Savitzky-Golay filter and interpolation.

8.3. Time-walk correction




Figure 8.13: Time walk correction for cosmic rays at 29 V (left) and 30 V bias voltage, determined using an oscilloscope. The X axis is the amplitude of the pulse at test point 2, i.e. the dotted purple line in figure 8.12. The Y axis is the delay between when the pulse at test point crosses 1% of the threshold voltage and 100% of the threshold voltage, i.e. the time difference between the two purple dashed lines in figure 8.12. A fit of the form given in equation 8.2 was applied to each.

8.3.4 Calibration of time walk correction using data

The value of the parameter W in equation 8.2 was determined roughly using the fit to the oscilloscope data. However, given the imprecise nature of the oscilloscope analysis, this parameter value was not used in the final correction. Instead, the parameter was derived from data as follows.

A scan was performed where the analysis was repeated with a range of different values for the parameter W . The width of the TOF distribution

8.3. Time-walk correction

was used as a measure of the time resolution, as described in section 7.3. The value of W which minimized the width of the TOF distribution for cosmic rays was used in the final correction. The results of this scan are shown in figure 8.14. The final value of the parameter W which was used to correct the vertical slice data is 1.75×10^{-9} . This will possibly have to be recalculated for the full BV, which will be discussed in section 9.4.




Figure 8.14: Scan of time walk correction parameter W . The TOF distribution with the two bars parallel was found for each value of W . The value of W which minimized the width of this distribution was used.

Including the channel-by-channel correction and the time walk correction, the final time of an end hit is computed as

$$t' = t_{\text{TDC}} - \delta t_{\text{chan}} - \frac{W}{\sqrt{A}} \quad (8.7)$$

8.3.5 Alternative time walk correction schemes

In some experiments, the time walk correction is not applied to each end hit, but is instead applied to the average time of both ends of the bar [58, 59]. This is typically done using the geometric mean of the pulse amplitudes A_{top} and A_{bot} at both ends:

$$\delta t = \frac{W}{\sqrt[4]{A_{\text{top}} \cdot A_{\text{bot}}}} \quad (8.8)$$

This scheme was considered but ultimately discarded. While the TOF measurement uses the average time, the position measurement along the bar uses the difference between both end times, and so this scheme does not allow for an accurate position measurement.

Another method involves using the calibration pulser to calibrate the time walk correction. The advantage of this is clear – since the pulses are injected with known timing, the trigger time (TDC start time) t_0 will be stable. The pulse amplitude can be gradually increased and a simple plot of time vs. amplitude can be used to find W and calibrate the correction. However, this method requires that the waveform injected into the electronics exactly matches the shape of the pulse produced by the SiPM. This should be possible using the current test pulser setup. However, only square wave test pulses were implemented into the run control code at the time of this study, and so this method was not used. However, it remains an interesting option to explore in the future, especially with regards to calibrating the full BV.

Chapter 9

Final results and outlook

This chapter presents the performance improvements of the BV resulting from the calibrations described in chapter 8. It also discusses the final steps which need to be taken prior to the operation of the BV in the fall of 2021.

9.1 Final time resolution

Following the calibration of the threshold voltage, the implementation of TDC channel delays, and the time walk correction, the TOF resolution improved dramatically. Further cosmic ray data was taken with the two bars at TRIUMF positioned next to each other. With a distance of no more than 5 cm between the bars, the TOF of these events should be less than 150 ps, and the width of the distribution will give an indication of the TOF resolution of a pair of bars. The TOF of these events was calculated twice, once without the channel-by-channel and time walk corrections implemented, and once with these corrections. These distributions are shown in figure 9.1 and 9.2.

9.1. Final time resolution




Figure 9.1: The TOF distribution for the vertical slice with the bars parallel and touching, after the thresholds were set and before the channel-by-channel and time walk corrections were applied.




Figure 9.2: The TOF distribution for the vertical slice with the bars parallel and touching, after the thresholds were set and with the channel-by-channel and time walk corrections applied. Note that the channel-by-channel corrections have shifted the mean closer to zero, and the time walk correction has significantly lowered the width of the distribution.

These figures should also be compared to figure 7.9. The main difference between figures 7.9 and 9.1 is the calibration of the threshold voltage to be above the background noise level. The result of this is clear and conclusive: the removal of the “grass” events outside of the main TOF peak. These events were caused by noise in the SiPMs or electronics triggering the TDC at random intervals, leading to uniformly random TOF measurements. Their disappearance confirms the necessity and effectiveness of calibrating the threshold voltages.

The difference between figures 9.1 and 9.2 is the implementation of the channel-by-channel and time walk corrections. Since only two bars were used, and these same two bars are obviously present in all events, applying the channel-by-channel by channel correction applies the same offset to every event. Thus this correction serves only to shift the distribution, and does not change its shape or width. The time walk correction is responsible for the narrowing of the distribution.

Indeed, the standard deviation σ of the TOF distribution in figure 9.1 is 400 ps. Applying the time walk correction improves this time resolution to 169 ps in figure 9.2. This is within the 200 ps time resolution requirement quoted in section 5.2. To this end, the work performed in this thesis has demonstrated that the BV will most likely have the TOF resolution required to be an effective veto of cosmic ray background events.¹⁵

9.2 Time of flight demonstration

Aside from measuring the time resolution and comparing to the requirement set by simulation, this section presents another way to verify the efficacy of the BV to perform cosmic ray discrimination.

¹⁵The final caveat is that this demonstration was performed with shorter 70 cm bars. The longer bars may have an inferior TOF resolution due to increased variance in the photon propagation time. However, it is not clear whether this variance constitutes one of the main uncertainties in the TOF resolution; this is another reason to carry out simulations of the light propagation within the bars. To verify the bar length does not heavily impact the TOF resolution, these tests should be repeated with full-length bars in the vertical slice when possible.

9.2. Time of flight demonstration

The TOF spectrum presented in the previous section was recorded with the two bars touching each other; i.e., at an effective distance of around 2 cm between the centre points. By moving the bars further apart, the TOF for cosmic rays passing through the bars is seen to increase. This is shown in figure 9.3, where the distance between the two bars was increased to 50 cm, around the diameter of the BV.¹⁶




Figure 9.3: The TOF distribution for the vertical slice with the bars parallel and separated by approximately 50 cm, with corrections applied. Notice the increase in TOF corresponding to the separation between bars.

The main feature to note when comparing this to figure 9.2 is the increase in the mean TOF of 1.5 ns. Multiplying this by the speed of light accurately returns a distance of 50 cm as expected. Notably, the two TOF distributions have no overlap. This is very encouraging as it indicates that given an event from one of the two distributions, it is possible to accurately determine which distribution it was taken from.

The TOF distributions for cosmic ray background and antihydrogen annihilation are not expected to be quite so simple. For example, while the effective TOF for an annihilation event will be small, it will not be exactly

¹⁶The distance between bars was not measured to be exactly the BV diameter. Instead, the two bars were placed above and below the ALPHA-g prototype TPC which is kept at TRIUMF.

9.2. Time of flight demonstration

zero, and so some additional width will be added to this distribution due to the physics and the geometry of the events, in addition to the intrinsic time resolution seen here. Similarly, cosmic rays are not guaranteed to pass across the diameter of the BV; cosmic rays grazing the edge of the detector will naturally have a shorter TOF. Nonetheless, the independence of the two distributions seen here is a good indicator that the TOF measurements taken by the BV will be very useful in rejecting cosmic ray background events.

An interesting feature of the distribution in figure 9.3 is the long tail extending down towards lower TOF values. The slightly increased tail toward higher TOF is understood to be due to cosmic rays striking both bars at a diagonal, and thus having a longer path length than those striking perpendicularly. The tail towards smaller TOF values, however, is not currently understood.

As a further check, the order of the bars was swapped, taking the top bar and putting it on the bottom and vice versa. Since the cosmic ray flux is entirely coming from above, this changed the sign on the TOF distribution. This was repeated for both measurements at 2 cm distance and at 50 cm distance, and is shown in figures 9.4 and 9.5.

Comparing figures 9.2 and 9.4 we can see the mean TOF changes by about 0.15 ns. This is consistent with the distance between the two bars being 2.4 cm, which is accurate when considering the centre points. Comparing figures 9.3 and 9.5, we see a change in mean TOF of 3.1 ns. This is consistent with a distance between the bars of just less than 50 cm. To summarize, swapping the order of the bars changes the sign of the TOF distribution as expected.

The only holdover here is a common offset of around 550 ps which is visible on all four of these plots; that is, the zero of each plot is shifted to 550 ps. This will be addressed briefly in section 9.4, although the origin of this offset is still not understood and will be the subject of further investigation.

9.2. Time of flight demonstration




Figure 9.4: The TOF distribution for the vertical slice with the bars parallel and touching, corrections applied, with the top and bottom bars swapped from figure 9.2. Note the slight shift of the distribution, since the cosmic ray flux originates from above.




Figure 9.5: The TOF distribution for the vertical slice with the bars parallel and separated by approximately 50 cm, corrections applied, with the top and bottom bars swapped from figure 9.3. Note the sign change on the distribution, since the cosmic ray flux originates from above.

9.3 Final position resolution

Along with TOF measurements, recall that an additional functionality of the BV is to provide supplementary position data for hits. As described in section 6.4, the Z position of a hit along the bar can be found using the times recorded by the SiPM at both ends. The Z position is proportional to the time difference between both ends of the bar, and so the end-to-end time difference is presented in this section for clarity.

Using the setup of two parallel touching bars, the Z position recorded by a cosmic ray passing through both bars should be approximately the same. Plotting the end-to-end time difference of the two bars against each other should give a straight line with a slope of 1; this was shown initially in figure 7.10. After applying the channel-by-channel and time walk corrections, this plot is presented again in figure 9.6.




Figure 9.6: End-to-end time difference comparison between bar A and bar B of vertical slice, after applying corrections. The distribution is much more well-centred compared to figure 7.10.

The most notable change is that the time difference for both bars is now more accurately centred around zero, whereas before there was an offset due to the time delays between channels. This shift is entirely due to the

channel-by-channel correction. Once again, these corrections do not change the shape or size of the distribution but only add an offset, following the same logic as the previous section.

The diagonal width of this line distribution is around 360 ps, which corresponds (via equation 6.11) to a Z position resolution of around 2.5 cm. This is around double the time resolution for the TOF. This is expected: the diagonal width of the line distribution is the addition and subtraction of four TDC time values, while the TOF is calculated in equation 6.14 by adding and subtracting four TDC time values, and then dividing by two.

Interestingly, the time walk correction does not appear to decrease this width. A natural explanation for this is that for two parallel touching bars, the events on each of the two bars happen at the same Z position, and so at the same distance from the respective sets of SiPMs. As a result, the amplitudes of the pulses at each end are approximately the same for both bars. Thus for each event, the same corrections are being applied to both bars, and so the time walk correction serves to move points along the line, but does not stretch or narrow the width of the line distribution.

This manifests itself as a shortening of the line distribution due to the time walk correction. In figure 7.10, the end-to-end time difference of each bar spans a range of 11 ns, while in figure 9.6 this range is decreased to 9 ns. This is because the pulses with the smallest amplitude occur for interactions close to the end of the bar opposite the sensor. These have the highest time delay between the interaction and the light collection; applying the time walk correction decreases the recorded time delay compared to larger amplitude events. This causes the ends of the distribution to contract inwards. It also helpfully decreases the size of the tails of the distribution, making it easier to determine the endpoints.

9.4 Final preparations and next steps

The results obtained from the vertical slice at TRIUMF have indicated that the BV should have the capability to use TOF to identify cosmic ray background events. However, there remain a number of steps which must be

taken before this capability is realized.

Currently, the BV is fully functional at CERN. Cosmic ray waveforms and test pulse waveforms have been observed on all channels.¹⁷ However, the BV is not currently capable of performing TOF measurements since the issues outlined in section 7.1 have not been addressed. To prepare the BV for full operation, the calibration steps taken for the vertical slice must be replicated for the full BV.

The first step which must be taken is a calibration of the threshold voltages as described in section 8.1. The procedure for carrying this out will be mostly identical to the one used on the vertical slice. The threshold voltage for all channels will be simultaneously stepped across the interval used previously, and cosmic ray data will be taken with each value of threshold voltage. Whereas one value of threshold voltage was used for all channels in the vertical slice, it is expected that the threshold voltage should be set independently on a channel by channel basis for the BV. With so many more channels, it is unlikely that the one value will be suitable for all channels.

The second calibration which will be performed is the measurement of the time offset for each TDC channel. This procedure will be the same as the one used for the vertical slice. The test pulser will be used to send a pulse to all ASD cards simultaneously, and the delay between channels will be measured. These delays will be recorded, saved, and subtracted from the times reported by the TDC in future data runs.

Finally, the time walk correction must be calibrated. This is more difficult, as the same procedure cannot be used: with more than two bars, the TOF distribution will not be a gaussian centred around zero, and so minimizing the width of this distribution cannot be used to calibrate the correction. Initially, the parameter used for the vertical slice will be used for the BV.

Should a more accurate calibration of the time walk correction be required, one option is to use the test pulser. By varying the amplitude of the test pulse sent through the electronics, it is possible to observe the time delay for smaller and smaller pulses due to time walk. This can then be

¹⁷Almost all TDC channels are working

used to calibrate the correction curve. Ideally, the test pulse would also be sent directly to the TDC by a separate pathway to be used as a reference time, in order to keep the analysis timed entirely by the TDC. Care must be taken to use a pulse shape which is close to the one produced by the SiPM.

Aside from preparing the BV for operation at CERN, there is one final issue which must be resolved using the vertical slice. The astute reader may have noticed that the TOF distribution in figure 9.2 is offset from zero by about 550 ps. This offset should have been removed by the TDC channel-by-channel calibration using the test pulse, but clearly some offset remains. The cause of this is still unknown, and will be the subject of further investigation.

The importance of understanding this offset in the TOF distribution cannot be overstated. It is a small issue for the vertical slice, since there are only two bars. It becomes a much larger issue for the BV with the full 64 bars. If this offset is not understood, any pair of two bars could have a unique offset. The difference in offset between pairs of bars could be large enough to smear out any information that could be gained from the TOF measurement; it makes comparing TOF measurements between two different pairs of bars meaningless.

The fact that this offset does not appear in test pulse data limits its possible sources. It cannot be due to the data acquisition or system of electronics, since the test pulse is injected directly after the first amplification stage on the ASD cards. The only electronics possibility is if the test pulse is in fact not completely simultaneous between channels, leading to an incorrect calibration of the TDC channel delays. Another possibility is that it is caused by differences in the SiPMs themselves. The gain of the SiPMs differs substantially, with some SiPMs producing larger amplitude pulses than others. Unless properly taken into account when performing the time walk correction, this could lead to artificial delays between channels. Regardless, understanding this offset in the TOF distribution is one of the last remaining issues to be solved using the vertical slice.

In the longer term, further study should be performed to characterize the travel of photons down the bars, from the site of particle incidence to the photosensors at each end. In this thesis, it was demonstrated that the

9.4. *Final preparations and next steps*

effective speed of light in the bar is meaningfully slower than the speed of light given by the refractive index quoted by the manufacturer. The current hypothesis is that this is due to reflections with the bar. Ideally this effect should be modelled using Monte Carlo simulations, and confirmed using the actual bars.

Chapter 10

Conclusions

The three main goals of this work were to develop the analysis algorithms for the BV, to identify and implement any corrections and calibrations necessary, and to demonstrate that the BV is capable of generating accurate TOF information for use in cosmic ray background rejection. All three of these goals were met to some degree.

The analysis was implemented in chapter 6. It is able to find an effective TOF for potential cosmic ray events, and match these events to tracks in the TPC. It also calculates the Z position for each bar hit. A background rejection algorithm remains to be written based on machine learning using a number of inputs including the TOF recorded by the BV.

The steps taken to calibrate the BV and corrections were described in chapter 8. These include a calibration of the threshold voltage for the SiPMs, a calibration of the channel path length time delays within the TDC, and a calibration of the amplitude-dependent time walk correction. These steps were performed on a vertical slice of the system which was set up at TRIUMF, and will be replicated for the full system.

After the implementation of these calibrations and corrections, the final results shown in chapter 9 using the vertical slice demonstrate that the BV should be able to record particle TOF to a resolution of less than 200 ps. It was shown that this was sufficient to fully differentiate between an event with zero TOF and a particle crossing the width of the BV. Previous simulations had shown that this time resolution would be sufficient for cosmic ray background rejection; this study supports that claim.

The cosmic ray background is one of the main limiting factors on the detection of antihydrogen annihilations at ALPHA-g. The work done here will be instrumental in the first ALPHA-g run this year, where these methods

will be used to mitigate this background. In turn, this will improve the accuracy with which the measurement of gravity acting on antihydrogen can be performed, which will lead to an improvement of our fundamental understanding of antimatter.

Bibliography

- [1] Sharon Bertsch McGrayne, George F. Bertsch, and James Trefil. Atom, Jan 2021. URL <https://www.britannica.com/science/atom>.
- [2] 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*, 117(778):610–624, Feb 1928. doi: 10.1098/rspa.1928.0023.
- [3] C. D. Anderson. The apparent existence of easily deflectable positives. *Science*, 76(1967):238–239, 1932. doi: 10.1126/science.76.1967.238.
- [4] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson. Experimental test of parity conservation in beta decay. *Physical Review Letters*, 105:1413–1415, Feb 1957. doi: 10.1103/PhysRev.105.1413. URL <https://link.aps.org/doi/10.1103/PhysRev.105.1413>.
- [5] J. H. Christenson, J. W. Cronin, V. L. Fitch, and R. Turlay. Evidence for the 2π decay of the K^0 meson. *Physical Review Letters*, 13(4): 138–140, 1964. doi: 10.1103/physrevlett.13.138.
- [6] B. Aubert et al. Measurement of CP -violating asymmetries in b^0 decays to CP eigenstates. *Physical Review Letters*, 86:2515–2522, Mar 2001. doi: 10.1103/PhysRevLett.86.2515. URL <https://link.aps.org/doi/10.1103/PhysRevLett.86.2515>.
- [7] Gerhart Luders. Proof of the tcp theorem. *Annals of Physics*, 281(1):1004–1018, 2000. ISSN 0003-4916. doi: 10.1006/aphy.2000.6027. URL <https://www.sciencedirect.com/science/article/pii/S0003491600960275>.

- [8] Dietrich Bodeker and Wilfried Buchmuller. Baryogenesis from the weak scale to the grand unification scale, 2021.
- [9] Don Lincoln. Gravitons, Oct 2012. URL https://www.fnal.gov/pub/today/archive/archive_2012/today12-10-19_NutshellReadMore.html.
- [10] L. I. Schiff. Sign of the gravitational mass of a positron. *Physical Review Letters*, 1:254–255, Oct 1958. doi: 10.1103/PhysRevLett.1.254. URL <https://link.aps.org/doi/10.1103/PhysRevLett.1.254>.
- [11] Allen Caldwell and Gia Dvali. On the gravitational force on anti-matter. *Fortschritte der Physik*, 69(1):2000092, 2021. doi: 10.1002/prop.202000092. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/prop.202000092>.
- [12] Ulrich Jentschura. Antimatter gravity: Second quantization and lagrangian formalism. *Physics*, 2:397–411, 09 2020. doi: 10.3390/physics2030022.
- [13] Yang Lu and R.D. Amado. Branching ratios in proton antiproton annihilation from skyrmion dynamics. *Physics Letters B*, 357(3):446–450, 1995. ISSN 0370-2693. doi: 10.1016/0370-2693(95)00904-Y. URL <https://www.sciencedirect.com/science/article/pii/037026939500904Y>.
- [14] Christian G. Parthey, Arthur Matveev, Janis Alnis, Birgitta Bernhardt, Axel Beyer, Ronald Holzwarth, Aliaksei Maistrou, Randolph Pohl, Katharina Predehl, Thomas Udem, Tobias Wilken, Nikolai Kolachevsky, Michel Abgrall, Daniele Rovera, Christophe Salomon, Philippe Laurent, and Theodor W. Hänsch. Improved measurement of the hydrogen $1s - -2s$ transition frequency. *Physical Review Letters*, 107:203001, Nov 2011. doi: 10.1103/PhysRevLett.107.203001. URL <https://link.aps.org/doi/10.1103/PhysRevLett.107.203001>.
- [15] M. Hori and J. Walz. Physics at cern’s antiproton decelerator. *Progress in Particle and Nuclear Physics*, 72:206–253, 2013. ISSN 0146-6410.

- doi: 10.1016/j.ppnp.2013.02.004. URL <https://www.sciencedirect.com/science/article/pii/S0146641013000069>.
- [16] G. Baur, G. Boero, A. Brauksiepe, A. Buzzo, W. Eyrich, R. Geyer, D. Grzonka, J. Hauße, K. Kilian, M. LoVetere, M. Macri, M. Moosburger, R. Nellen, W. Oelert, S. Passaggio, A. Pozzo, K. Rhrich, K. Sachs, G. Schepers, T. Sefzick, R.S. Simon, R. Stratmann, F. Stinzling, and M. Wolke. Production of antihydrogen. *Physics Letters B*, 368(3):251–258, 1996. ISSN 0370-2693. doi: 10.1016/0370-2693(96)00005-6. URL <https://www.sciencedirect.com/science/article/pii/0370269396000056>.
 - [17] M. Amoretti, C. Amsler, G. Bonomi, A. Bouchta, P. Bowe, C. Carraro, C. L. Cesar, M. Charlton, M. J. Collier, M. Doser, and et al. Production and detection of cold antihydrogen atoms. *Nature*, 419(6906):456–459, 2002. doi: 10.1038/nature01096.
 - [18] G. Gabrielse, N. S. Bowden, P. Oxley, A. Speck, C. H. Storry, J. N. Tan, M. Wessels, D. Grzonka, W. Oelert, G. Schepers, T. Sefzick, J. Walz, H. Pittner, T. W. Hänsch, and E. A. Hessels. Driven production of cold antihydrogen and the first measured distribution of antihydrogen states. *Physical Review Letters*, 89:233401, Nov 2002. doi: 10.1103/PhysRevLett.89.233401. URL <https://link.aps.org/doi/10.1103/PhysRevLett.89.233401>.
 - [19] G. B. Andresen, M. D. Ashkezari, M. Baquero-Ruiz, W. Bertsche, P. D. Bowe, E. Butler, C. L. Cesar, S. Chapman, M. Charlton, A. Deller, and et al. Trapped antihydrogen. *Nature*, 468(7324):673–676, 2010. doi: 10.1038/nature09610.
 - [20] G. Gabrielse, R. Kalra, W. S. Kalthammer, R. McConnell, P. Richerme, D. Grzonka, W. Oelert, T. Sefzick, M. Zielinski, D. W. Fitzakerley, M. C. George, E. A. Hessels, C. H. Storry, M. Weel, A. Müllers, and J. Walz. Trapped antihydrogen in its ground state. *Physical Review Letters*, 108:113002, Mar 2012. doi: 10.1103/PhysRevLett.108.

113002. URL <https://link.aps.org/doi/10.1103/PhysRevLett.108.113002>.
- [21] M. Ahmadi et al. Observation of the hyperfine spectrum of antihydrogen. *Nature*, 548(7665):66–69, Aug 2017. ISSN 1476-4687. doi: 10.1038/nature23446. URL <https://doi.org/10.1038/nature23446>.
- [22] M. Ahmadi et al. Observation of the 1s–2s transition in trapped antihydrogen. *Nature*, 541(7638):506–510, Jan 2017. ISSN 1476-4687. doi: 10.1038/nature21040. URL <https://doi.org/10.1038/nature21040>.
- [23] M. Ahmadi et al. Observation of the 1s–2p lyman- α transition in antihydrogen. *Nature*, 561(7722):211–215, Sep 2018. ISSN 1476-4687. doi: 10.1038/s41586-018-0435-1. URL <https://doi.org/10.1038/s41586-018-0435-1>.
- [24] M. Ahmadi et al. Investigation of the fine structure of antihydrogen. *Nature*, 578(7795):375–380, Feb 2020. ISSN 1476-4687. doi: 10.1038/s41586-020-2006-5. URL <https://doi.org/10.1038/s41586-020-2006-5>.
- [25] C. J. Baker et al. Laser cooling of antihydrogen atoms. *Nature*, 592(7852):35–42, Apr 2021. ISSN 1476-4687. doi: 10.1038/s41586-021-03289-6. URL <https://doi.org/10.1038/s41586-021-03289-6>.
- [26] W Oelert, J Hangst, Michael Holzscheiter, P Lichard, G Gabrielse, and R Hayano. Elena: An upgrade to the antiproton decelerator, 09 2009.
- [27] C. Smorra et al. A parts-per-billion measurement of the antiproton magnetic moment. *Nature*, 550(7676):371–374, Oct 2017. ISSN 1476-4687. doi: 10.1038/nature24048. URL <https://doi.org/10.1038/nature24048>.
- [28] Masaki Hori, Hossein Aghai-Khozani, Anna Sótér, Andreas Dax, and Daniel Barna. Laser spectroscopy of pionic helium atoms. *Nature*, 581(7806):37–41, May 2020. ISSN 1476-4687.

- doi: 10.1038/s41586-020-2240-x. URL <https://doi.org/10.1038/s41586-020-2240-x>.
- [29] P. Pérez et al. The gbar antimatter gravity experiment. *Hypersfine Interactions*, 233(1):21–27, Aug 2015. ISSN 1572-9540. doi: 10.1007/s10751-015-1154-8. URL <https://doi.org/10.1007/s10751-015-1154-8>.
- [30] Claude Amsler et al. Pulsed production of antihydrogen. *Communications Physics*, 4(1):19, Feb 2021. ISSN 2399-3650. doi: 10.1038/s42005-020-00494-z. URL <https://doi.org/10.1038/s42005-020-00494-z>.
- [31] C. Amole et al. Description and first application of a new technique to measure the gravitational mass of antihydrogen. *Nature Communications*, 4(1):1785, Apr 2013. ISSN 2041-1723. doi: 10.1038/ncomms2787. URL <https://doi.org/10.1038/ncomms2787>.
- [32] G. B. Andresen et al. Confinement of antihydrogen for 1,000 seconds. *Nature Physics*, 7(7):558–564, Jul 2011. ISSN 1745-2481. doi: 10.1038/nphys2025. URL <https://doi.org/10.1038/nphys2025>.
- [33] C. Amole et al. Resonant quantum transitions in trapped antihydrogen atoms. *Nature*, 483(7390):439–443, Mar 2012. ISSN 1476-4687. doi: 10.1038/nature10942. URL <https://doi.org/10.1038/nature10942>.
- [34] W. A. Bertsche. Prospects for comparison of matter and antimatter gravitation with alpha-g. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 376(2116):20170265, 2018. doi: 10.1098/rsta.2017.0265. URL <https://royalsocietypublishing.org/doi/abs/10.1098/rsta.2017.0265>.
- [35] David B Newell and Eite Tiesinga. The international system of units (si):. *National Institute of Standards and Technology Special Publication*, 330, Aug 2019. doi: 10.6028/nist.sp.330-2019.

- [36] M. Ahmadi and others. Antihydrogen accumulation for fundamental symmetry tests. *Nature Communications*, 8(1):681, Sep 2017. ISSN 2041-1723. doi: 10.1038/s41467-017-00760-9. URL <https://doi.org/10.1038/s41467-017-00760-9>.
- [37] Ramsey Badawi. Introduction to pet physics: The physical principles of pet, Jan 1999. URL https://depts.washington.edu/imreslab/from%20old%20SITE/pet_intro/intro_src/section2.html.
- [38] Wei Jiang, Yamn Chalich, and M. Jamal Deen. Sensors for positron emission tomography applications. *Sensors*, 19:5019, 11 2019. doi: 10.3390/s19225019.
- [39] Doug Bryman. Physics 531 lecture 9: Geiger, mwpc, drift and gas microstrip chambers ii, 02 2021. University of British Columbia. Lecture notes.
- [40] Walter Blum, Werner Riegler, and Luigi Rolandi. *Particle detection with drift chambers*. Springer, 2010.
- [41] Andrea Capra. Alpha: Precision measurements of antihydrogen, 02 2019. Winter Nuclear Particle Physics Conference 2019, Banff.
- [42] Doug Bryman. Physics 531 lecture 12: Scintillators and photodetectors i, 03 2021. University of British Columbia. Lecture notes.
- [43] Eljen Technology. General purpose plastic scintillators: Ej-200, ej-204, ej-208, ej-212. <https://eljentechnology.com/products/plastic-scintillators/ej-200-ej-204-ej-208-ej-212>, 2016. [Online; accessed 4-May-2021].
- [44] Glenn F. Knoll. *Radiation Detection and Measurement*. John Wiley, 2000.
- [45] Particle Data Group. Review of Particle Physics. *Progress of Theoretical and Experimental Physics*, 2020(8), 08 2020. ISSN 2050-3911. doi: 10.1093/ptep/ptaa104. URL <https://doi.org/10.1093/ptep/ptaa104>. 083C01.

- [46] Andrea Capra. Time application of the barrel detector for alpha-g. TRIUMF, 2016.
- [47] S. Agostinelli et al. Geant4: a simulation toolkit. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 506(3):250–303, 2003. ISSN 0168-9002. doi: 10.1016/S0168-9002(03)01368-8. URL <https://www.sciencedirect.com/science/article/pii/S0168900203013688>.
- [48] SensL. J-series high pde and timing resolution, tsv package datasheet. <https://sensl.com/downloads/ds/DS-MicroJseries.pdf>, 12 2017. [Online; accessed 4-May-2021].
- [49] F. Carnesecchi. Performance of the alice time-of-flight detector at the lhc. *Journal of Instrumentation*, 14(06), 2019. doi: 10.1088/1748-0221/14/06/c06023.
- [50] A. Neiser et al. TRB3: a 264 channel high precision TDC platform and its applications. *JINST*, 8:C12043, 2013. doi: 10.1088/1748-0221/8/12/C12043.
- [51] A.B. Garnsworthy, C.J. Pearson, D. Bishop, B. Shaw, J.K. Smith, M. Bowry, V. Bildstein, G. Hackman, P.E. Garrett, Y. Linn, J.-P. Martin, W.J. Mills, and C.E. Svensson. The griffin data acquisition system. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 853:85–104, 2017. ISSN 0168-9002. doi: 10.1016/j.nima.2017.02.040. URL <https://www.sciencedirect.com/science/article/pii/S0168900217302243>.
- [52] Grzegorz Korcyl, Ludwig Maier, Jan Michel, Andreas Neiser, Marek Palka, Manuel Penschuck, Pawel Strzempek, Michael Traxler, and Cahit Ugur. A users guide to the trb3 and fpga-tdc based platforms. <http://jspc29.x-matter.uni-frankfurt.de/docu/trb3docu.pdf>, 02 2021. [Online; accessed 4-May-2021].

- [53] C Ugur, E Bayer, N Kurz, and M Traxler. A 16 channel high resolution (11 ps rms) time-to-digital converter in a field programmable gate array. *Journal of Instrumentation*, 7(02):C02004–C02004, feb 2012. doi: 10.1088/1748-0221/7/02/c02004. URL <https://doi.org/10.1088/1748-0221/7/02/c02004>.
- [54] W. Braunschweig, E. Konigs, W. Sturm, and W. Wallraff. A large area time-of-flight system with a resolution of 0.5 ns fwhm. *Nuclear Instruments and Methods*, 134(2):261–266, 1976. ISSN 0029-554X. doi: 10.1016/0029-554X(76)90280-9. URL <https://www.sciencedirect.com/science/article/pii/0029554X76902809>.
- [55] T. Tanimori, T. Kageyama, K. Nakamura, M. Yoshioka, J. Haba, H. Kawai, M. Kobayashi, T.S. Nakamura, and N. Sasao. A test of 150 cm \times 20 cm wide time-of-flight scintillation counters. *Nuclear Instruments and Methods in Physics Research*, 216(1):57–65, 1983. ISSN 0167-5087. doi: 10.1016/0167-5087(83)90330-7. URL <https://www.sciencedirect.com/science/article/pii/0167508783903307>.
- [56] Chong Wu, Yuekun Heng, Yuda Zhao, Xiaojian Zhao, Zhijia Sun, Jinjie Wu, Zhenghua An, Li Zhao, Linli Jiang, Fengmei Wang, Shengtian Xue, and Yifang Wang. The timing properties of a plastic time-of-flight scintillator from a beam test. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 555(1):142–147, 2005. ISSN 0168-9002. doi: 10.1016/j.nima.2005.09.029. URL <https://www.sciencedirect.com/science/article/pii/S0168900205018449>.
- [57] S. Nishimura, M. Kurata-Nishimura, K. Morimoto, Y. Nishi, A. Ozawa, T. Yamaguchi, T. Ohnishi, T. Zheng, M. Chiba, and I. Tanihata. Systematic studies of scintillation detector with timing resolution of 10 ps for heavy ion beam. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 510(3):377–388, 2003. ISSN 0168-9002. doi:

- 10.1016/S0168-9002(03)01824-2. URL <https://www.sciencedirect.com/science/article/pii/S0168900203018242>.
- [58] J.W. Zhao, B.H. Sun, I. Tanihata, S. Terashima, L.H. Zhu, A. Enomoto, D. Nagae, T. Nishimura, S. Omika, A. Ozawa, Y. Takeuchi, and T. Yamaguchi. Reaching time resolution of less than 10ps with plastic scintillation detectors. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 823:41–46, 2016. ISSN 0168-9002. doi: 10.1016/j.nima.2016.03.106. URL <https://www.sciencedirect.com/science/article/pii/S0168900216301516>.
- [59] Wen-Jian Lin et al. Plastic scintillation detectors for precision time-of-flight measurements of relativistic heavy ions. *Chinese Physics C*, 41(6):066001, Nov 2017. doi: 10.1088/1674-1137/41/6/066001.
- [60] A. Akindinov et al. Performance of the ALICE Time-Of-Flight detector at the LHC. *European Physics Journal Plus*, 128:44, 2013. doi: 10.1140/epjp/i2013-13044-x.