

Sub-nanosecond Cherenkov photon detection for LHCb particle identification in high-occupancy conditions and semiconductor tracking for muon scattering tomography



Floris Keizer
Churchill College
University of Cambridge

This dissertation is submitted for the degree of

Doctor of Philosophy

June 2019

Abstract

The increase in luminosity during the LHC upgrade programme causes a challenging rise in track multiplicity and hit occupancy in the LHCb detector. In order to mitigate this effect, the use of photon detector hit time information is presented in the context of the Ring-Imaging Cherenkov (RICH) detectors. The application of a time gate in the FPGA of the digital readout board for the Upgrade Ia photon detector, which is being installed for LHC Run 3, is described. Data recorded during SPS charged particle beam tests using a 6.25 ns time gate show a reduction of up to a factor of four in asynchronous detector noise compared to the original 25 ns readout. A time-walk correction based on the time-over-threshold is proposed. Using the LHCb simulation framework, the intrinsic time resolution of the RICH detectors is demonstrated to be less than 10 ps. This is particularly relevant for the LHCb Upgrade II, which is scheduled for the year 2030 in preparation for a further order-of-magnitude rise in luminosity. Methods of time gating and scaling of the signal amplitude in the RICH reconstruction likelihood maximisation algorithm are presented. The results show that, considering improvements in the time-resolution only, a photon detector with an approximately 50 ps resolution can achieve today's particle ID performance in the high-luminosity LHC environment.

In the second part of this thesis, the first published semiconductor tracker for cosmic-ray muon scattering tomography is presented. The tracker uses silicon strip sensors from the ATLAS Semiconductor Tracker (SCT) with an $80\,\mu\text{m}$ pitch. A novel electronic readout system for the sensors is designed, based on a scalable, inexpensive, flexible, FPGA-based solution. A high-precision mechanical structure with integrated cooling is built to align the SCT modules. This alignment is fine-tuned in software, and the tracker performance is compared with a Geant4 simulation. A scattering angle resolution compatible with 1.5 mrad at the 4 GeV average cosmic-ray muon energy is obtained. Data are recorded for plastic, iron and lead samples using 45000 muons. Images are reconstructed using the Angle Statistics Reconstruction algorithm, and demonstrate good contrast between low and high atomic number materials.

Declaration

This dissertation is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the Preface and specified in the text. It is not substantially the same as any that I have submitted, or, is being concurrently submitted for a degree or diploma or other qualification at the University of Cambridge or any other University or similar institution. I further state that no substantial part of my dissertation has already been submitted, or, is being concurrently submitted for any such degree, diploma or other qualification at the University of Cambridge or any other University or similar institution. It does not exceed the prescribed word limit for the Degree Committee for the Faculty of Physics.

Floris Keizer



‘Natural Sciences’

Peter Keizer

Oil on canvas, 82 × 74 cm, June 2019

Preface

I joined the University of Cambridge in 2011 for my undergraduate degree in Natural Sciences at Churchill College, and started my PhD at the Cambridge High Energy Physics group in 2015. I am grateful to the teaching staff members at the Cavendish Laboratory and at Churchill College, especially Lisa Jardine-Wright, John Richer and Dave Green, who taught me amongst many things that ‘one only needs to comb half his beard, as the rest follows by symmetry’.

I would like to thank the members of the Cambridge HEP group, in particular Val Gibson, Susan Haines, Bart Hommels, John Hill and Maurice Goodrick. I am very grateful to Chris Jones, for many interesting conversations, for his kind advice and tremendous help with the LHCb simulation studies.

I gratefully acknowledge Andy Parker, for leading the muon scattering tomography project in collaboration with Chris Steer and AWE. The semiconductor tracker greatly benefitted from the technical expertise of Saevar Sigurdsson, who I thank for all his kind help, creativity and Icelandic humour. It was a pleasure to supervise undergraduate student Andrei Gorbach, without whom the image reconstruction software would not be the same.

Throughout my time in Cambridge there is one person I would like to specially thank: my supervisor Stephen Wotton. Steve and I met ten years ago, when I visited the Cavendish Laboratory as a high-school student from the Netherlands for a summer project on HiSPARC. Since then, Steve has taught me much of what I know during my undergraduate and master projects and, of course, this PhD research. Thank you so much, Steve, for giving me this opportunity, for your enthusiastic guidance and our friendship.

The test beam periods at CERN have been the most rewarding experience during my PhD. I sincerely thank the members of the LHCb RICH group for creating this friendly and highly motivating environment, especially Roberta Cardinale, Giovanni Cavallero, Christoph Frei, Silvia Gambetta, Philip Garsed, Thierry Gys, Mike McCann, Antonis Papanestis, Alessandro Petrolini, Didier Piedigrossi and Antonino Sergi. My gratitude extends to Sajan Easo for his advice on the simulation studies, and to Carmelo d’Ambrosio for his future perspective. I would also like to thank Michele Blago for our joint effort on the RICH photon detector performance at high rate. I will never forget our test beam night shift when the fire brigade suddenly entered the control room in what closely resembled space suits. From the FASER group, I thank Yosuke Takubo, Hidetoshi Otono and Jamie Boyd for their friendly collaboration.

Finally, a heartfelt thank you to my family and friends, especially to Nick Hills. *Astrid, Peter en Daan, dank jullie wel voor alle zorg en aanmoediging. Zonder jullie was deze droom nooit uitgekomen.*

Contents

I Sub-nanosecond Cherenkov photon detection for LHCb particle identification in high-occupancy conditions	5
Introduction	6
1 The upgrades of the LHCb experiment at the LHC	7
1.1 CERN and the Large Hadron Collider	7
1.1.1 The LHC and its main experiments	9
1.1.2 Future programme of the LHC and the LHCb experiment	10
1.2 The LHCb experiment	11
1.2.1 Flavour physics in a nutshell	11
1.2.2 Physics motivation for the LHCb detector Upgrade II	16
1.2.3 Overview of the LHCb detector	16
1.3 Tracking system upgrades	17
1.3.1 Vertex Locator	17
1.3.2 The Scintillating Fibre tracker and Mighty Tracker	20
1.3.3 Magnet side stations	22
1.4 Particle identification system upgrades	23
1.4.1 RICH detectors	23
1.4.2 TORCH detector	26
1.4.3 ECAL	28
1.5 Data processing upgrades	29
1.5.1 Upgrade Ia readout architecture	30
1.5.2 Time at the detector front-end electronics	31
1.5.3 Time at the event builder	32
1.5.4 Time at the High-Level Trigger	32
1.6 Chapter conclusion	34

2	Time information in the RICH Upgrade Ia photon detector	35
2.1	The RICH photon detector system	36
2.1.1	Multi-Anode Photomultiplier Tube (MAPMT)	36
2.1.2	The CLARO front-end ASIC	38
2.1.3	Photon Detector Module Digital Board (PDMDB)	39
2.2	Performance at high rate and high photon occupancy	41
2.2.1	Experimental setup	41
2.2.2	Readout board	43
2.2.3	Pulse-height spectrum analysis	44
2.2.4	Results at high rate and high photon occupancy	46
2.2.5	Dynode biasing scheme	48
2.2.6	Signal-Induced Noise (SIN)	51
2.3	SPS beam tests with a time gate in the FPGA	54
2.3.1	Test setup at the SPS charged particle beam facility	55
2.3.2	Implementation of a time gate in the PDMDB	57
2.3.3	Analysis and results	59
2.3.4	Time-over-threshold correction for the Upgrade Ia detector	63
2.4	Chapter conclusion	65
3	Sub-nanosecond time for particle identification in the HL-LHC	66
3.1	Time in the LHCb simulation framework	67
3.1.1	A 4-dimensional primary vertex generator	68
3.1.2	Particle propagation and photon dispersion	70
3.1.3	The RICH detector SmartID	73
3.1.4	Spillover	74
3.1.5	Pile-up	75
3.2	RICH Reconstruction in space and time	76
3.2.1	Reconstruction algorithm in space	76
3.2.2	Photodetector hit time distribution	78
3.2.3	Primary vertex time	80
3.2.4	Prediction of the Cherenkov photon time of arrival	81
3.2.5	Time in the RICH reconstruction software	86
3.2.6	The particle identification curve	88
3.3	Particle identification performance using time	90
3.3.1	Particle identification at the LHC Run 3 luminosity	90
3.3.2	Particle identification at the HL-LHC Run 5 luminosity	92
3.3.3	Time gating and time in the likelihood calculation	92
3.3.4	Photon objects and CPU resources	94
3.4	Implications for detector technologies	95
3.4.1	The Micro-Channel Plate (MCP) detector family	97
3.4.2	Silicon PhotoMultiplier (SiPM)	98
3.5	Chapter conclusion	99

II	A semiconductor tracker for muon scattering tomography	100
4	Apparatus for muon scattering tomography	101
4.1	Muon scattering tomography	102
4.1.1	Cosmic-ray muons	102
4.1.2	Multiple scattering in materials	103
4.1.3	Muon imaging detectors and applications	104
4.2	ATLAS Semiconductor Tracker module	105
4.3	Electronic readout design	107
4.3.1	Module adapter card	107
4.3.2	Trigger and readout	112
4.4	Software	114
4.4.1	Detector configuration interface	114
4.4.2	Data decoding	116
4.5	Characterisation of the sensor and front-end electronics	117
4.5.1	Noise scan	117
4.5.2	Three-point gain	118
4.5.3	Noise characteristic and the semiconductor tracker	121
4.6	Mechanics	122
4.6.1	Mechanical alignment of the SCT modules	122
4.6.2	Cooling of the tracker	124
4.6.3	Light-tight environment	124
4.7	Chapter conclusion	125
5	Performance of the MST tracker	126
5.1	Track identification	127
5.2	Software alignment of modules	128
5.3	Detector simulation and angular resolution	129
5.3.1	Geant4 detector description	129
5.3.2	Scattering angle resolution results	130
5.4	Image reconstruction	133
5.4.1	Angle Statistics Reconstruction	133
5.4.2	Cosmic-ray imaging results	134
5.4.3	Image recording time and tracker improvements	136
5.5	Chapter conclusion	137
A	The semiconductor tracker as a beam telescope	138
B	4-dimensional primary vertex distribution	143
C	Mechanics of the semiconductor tracker	147
	Glossary	151
	Bibliography	151

Part I

Sub-nanosecond Cherenkov photon detection for LHCb particle identification in high-occupancy conditions

Introduction

In this thesis, I present the potential benefits and technological challenges of the use of time information in the LHCb Ring-Imaging Cherenkov (RICH) detectors, in order to mitigate the high detector occupancy in the LHC upgrade programme.

It is important to understand how the hit time information could be used to improve the RICH detector performance, initially assuming that there are no technological restrictions. The standardised parameter to quantify this improvement is the particle identification efficiency, which will be studied in the LHCb simulation framework presented in chapter 3. The particle ID efficiency for a range of detector time resolutions will subsequently be explored in simulation, which can help identifying suitable detector technologies for the future upgrades.

Other LHCb sub-detectors also intend to use hit time information in the future. This can help to reduce the resources required for the RICH detector readout and event reconstruction. Furthermore, the technologies considered by other sub-detectors are often also of interest to the RICH detectors. The upgrade programme of the LHCb detector is therefore described in detail, with an emphasis on the time coordinate, in chapter 1.

The use of time information sets a challenge for the time resolution of the sensor and front-end readout electronics. The first tests of a time gate implemented in the readout electronics of the currently ongoing detector upgrade for LHC Run 3 are reported in chapter 2. This time gate is primarily intended to reduce asynchronous noise from the photon sensor. Additionally, the simulation studies in chapter 3 can help to determine the minimum width of the front-end time gate required to capture the signal, and to quantify the expected reduction in background hits produced in the particle beam collisions using the time gate.

Chapter 1

The upgrades of the LHCb experiment at the LHC


This chapter introduces the Large Hadron Collider beauty (LHCb) experiment at the Large Hadron Collider (LHC), with a focus on the upgrade programme of the experiment. Section 1.1 introduces the European Organisation for Nuclear Research (CERN) and its premier facility, the LHC. Section 1.2 highlights some of the particle physics concepts at the LHCb experiment along with an overview of the detector. The upgrades of the LHCb sub-detectors are outlined in sections 1.3 and 1.4, where the possible use of a precision hit-time coordinate and single-photon detection technologies are of particular interest to the Ring-Imaging Cherenkov (RICH) detector upgrades. The use of time in the readout and trigger decisions is discussed in section 1.5.

1.1 CERN and the Large Hadron Collider

Since its formation in 1954, CERN has become the host of the world's largest particle accelerator, the LHC. The LHC is housed in a 3.8 m-wide tunnel up to 175 m underground with a circumference of 27 km straddling the Franco-Swiss border near Geneva. In 1988, the ends of the former LEP tunnel were joined together to complete Europe's largest civil-engineering project at the time.

An impressive number of discoveries and measurements has been made at CERN. A small selection includes the first observation of anti-nuclei in 1965 [1], the discovery of neutral currents in 1974 [2], the Nobel-prize-winning discovery of W and Z bosons in 1983 [3, 4] and its precision measurements at the Large Electron-Positron collider [5], a measurement of the number of light neutrino families in 1989 [6], the first creation of anti-hydrogen in 1995 [7] and the discovery of pentaquarks in 2015 [8]. The discovery of the Higgs boson in 2012 led to the Nobel prize for Peter Higgs and Francois Englert for the development of

the Higgs mechanism [9, 10]. CERN has also driven technological developments in particle accelerators and detectors. The multi-wire proportional chamber, discovered at CERN in 1968 [11], marked the transition from manual particle detection using bubble or spark chamber photographs to the fast electronic era. Georges Charpak received the Nobel prize for this invention. Numerous other highlights range from the world's first website and web server, `info.cern.ch`, to the record-breaking ATLAS Barrel Toroid, a 100-tonnes superconducting magnet aligned to the millimetre [12]. To keep CERN at the forefront of particle physics, an ambitious upgrade programme has been developed for the LHC and its experiments. This programme is presented in the context of LHCb in section 1.1.2.



LHC - Large Hadron Collider // SPS - Super Proton Synchrotron // PS - Proton Synchrotron // AD - Antiproton Decelerator // CLEAR - CERN Linear Electron Accelerator for Research // AWAKE - Advanced WAKEfield Experiment // ISOLDE - Isotope Separator OnLine // REX/HIE - Radioactive Experiment/High Intensity and Energy ISOLDE // LEIR - Low Energy Ion Ring // LINAC - LINear ACcelerator // n-ToF - Neutrons Time Of Flight // HiRadMat - High-Radiation to Materials // CHARM - Cern High energy AcceleraTOR Mixed field facility // IRRAD - proton IRRADIation facility // GIF++ - Gamma Irradiation Facility // CENF - CErN Neutrino platForm

Figure 1.1: Overview of the CERN accelerator complex, showing the individual stages of acceleration. The proton beams in the Large Hadron Collider are brought together with a centre-of-mass energy of up to 14 TeV at the ATLAS, CMS, ALICE and LHCb experiments. Source: [13].

1.1.1 The LHC and its main experiments

The CERN accelerator complex is shown in figure 1.1. Several accelerator stages successively increase the beam energy. The protons reach 50 MeV in the linear accelerator (LINAC 2), followed by 1.4 GeV in the Proton Synchrotron Booster (PSB), 25 GeV in the Proton Synchrotron (PS), 450 GeV in the Super Proton Synchrotron (SPS) and 7 to 13 TeV in the LHC. The structure of the accelerator complex illustrates the progression of CERN, where each accelerator was first built for a programme of measurements, and later used as a stepping-stone for the next-scale accelerator. In January 2019, the conceptual design report for the Future Circular Collider (FCC), a 100 km accelerator reaching a 100 TeV collision energy, was submitted for the next update of the European Strategy for Particle Physics [14]. The FCC, which would be injected by the LHC after 2035, provides an exciting next phase in the tradition of increasing acceleration at CERN.

The LHC ring consists of two beam pipes with counter-rotating beams of protons that are accelerated using radio-frequency cavities. Dipole magnets bend the beams onto circular paths while quadrupole magnets focus the beams. The beams are brought into collision, creating a large number of proton-proton interactions, at the ATLAS, CMS, ALICE and LHCb experiments. In addition to these main experiments, the LHC facilities are shared by a variety of smaller experiments, such as the Forward Search Experiment (FASER) [15], which is referred to in chapter 4 of this thesis.

ATLAS [16] and CMS [17] are general-purpose detectors investigating a wide range of physics, from measurements on the Higgs boson to searches for new physics. The detectors were independently designed, with different technologies and magnet systems, in order to allow the cross-confirmation of new discoveries. ATLAS is the largest volume particle detector ever built, weighing 7000 tonnes and occupying $46 \times 25 \times 25 \text{ m}^3$. On the opposite side of the LHC ring, the CMS (Compact Muon Solenoid) detector is built around a solenoid magnet with a coil of superconducting cable that generates a 4 T field. The steel yoke that confines the field makes CMS the heaviest particle detector, weighing 14000 tonnes and occupying $21 \times 15 \times 15 \text{ m}^3$.

ALICE (A Large Ion Collider Experiment) [18] studies the collisions of heavy ions. For a period of time each year, heavy ions are accelerated by the LINAC 3 and the Low-Energy Ion Ring, before being injected into the PS. The ALICE detector, occupying $26 \times 16 \times 16 \text{ m}^3$, is designed to study strongly interacting matter at extreme energy densities. The ion-ion collisions create a quark-gluon plasma, which is observed by the ALICE detector as it expands and cools, forming the particles that constitute the matter of our universe today.

The LHCb experiment [19] specialises in studying the differences between matter and antimatter in the decay of beauty quarks. The LHCb detector, which is the main topic of this thesis, is the smallest of the four main detectors, weighing 5600 tonnes and occupying $21 \times 10 \times 13 \text{ m}^3$.

1.1.2 Future programme of the LHC and the LHCb experiment

The future programmes of the LHC and the LHCb experiment are documented in [20] and [21] respectively. This section provides a summary of these programmes using the timeline in figure 1.2. The LHC was successfully commissioned in 2010 with a 7 TeV proton-proton collision energy, rising to 8 TeV in April 2012. During the first long shutdown (LS1) in 2013 and 2014, maintenance was performed on the high-current splices between the superconducting magnets across the LHC ring. During LHC Run 2, which started in 2015, the collision energy was increased to 13 TeV and a total integrated luminosity of 9 fb^{-1} was recorded at the LHCb experiment, exceeding the design integrated luminosity of 7 fb^{-1} . The second long shutdown (LS2) started at the end of 2018 and is currently ongoing.

During LS2, short-term machine improvements and preparations towards the High-Luminosity LHC (HL-LHC) are being made. For example, the LINAC 2 is being replaced by the new LINAC 4. Instead of protons, the LINAC 4 will inject hydrogen ions (H^-) into the PSB. The electrons are stripped off in the PSB, leaving protons. The H^- ions coming from LINAC 4 are attracted to these protons, resulting in a more intense, concentrated beam. LS2 also provides an opportunity for detector upgrades in the experiments at the LHC. A complete revision, called Upgrade Ia, of the LHCb experiment is ongoing, which introduces an LHCb-wide, triggerless, data-driven readout architecture with all event selections made in software. In this way, the Upgrade Ia detector makes full use of the available luminosity. A total integrated luminosity of around 23 fb^{-1} is expected by the end of LHC Run 3.

The installation of the HL-LHC will take place during the third long shutdown (LS3) between 2024 and 2026. The HL-LHC facilitates a significant increase in statistical precision of measurements compared to the LHC and aims to reach a peak luminosity of $5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ with luminosity levelling, accumulating to 3000 fb^{-1} by around the year 2038.




Figure 1.2: Timeline of the LHC and LHCb experiment. The detector Upgrade Ia is currently ongoing, and will be consolidated during Upgrade Ib. Although the HL-LHC will start operation in 2027, the luminosity in the LHCb experiment will not be increased until LHC Run 5.

During LS3, the ATLAS and CMS experiments will undergo a major upgrade to accommodate the higher event multiplicities and detector occupancies in the HL-LHC. In order to spread the available resources for experiment upgrades, the LHCb and ALICE experiments will not be upgraded until LS4 in 2030. Nevertheless, the three-year shutdown provides an opportunity for the LHCb detector Upgrade Ib, which will be used to consolidate the Upgrade Ia detector and to install low-cost prototypes of future detectors.

Throughout LHC Run 4, the luminosity at the LHCb experiment is maintained at the Run 3 level using luminosity levelling, which limits the (peak) pile-up of events in the detector and the radiation damage near the interaction region. This can be achieved at the interaction point by dynamically varying the beam separation, the crossing angle, the focal length or the crab-cavity voltage, which affects the tilt of the proton bunch with respect to the crossing angle [20]. A total integrated luminosity of around 50 fb^{-1} is expected by the end of LHC Run 4.

During the fourth long shutdown (LS4) scheduled in 2030, the LHCb detector will be completely revised (LHCb Upgrade II) in anticipation of the luminosity increase of a factor of 7.5 from LHC Run 5 onwards. This step in luminosity provides a challenging increase in primary vertex and track multiplicity, detector occupancy and fluence. As described in this chapter, the majority of LHCb sub-detectors will implement detector hit time information in Upgrade II in order to cope with the high detector occupancies. The LHCb upgrade II is a major task, and the installation and commissioning will most likely require more time than the one year currently scheduled for LS4.

1.2 The LHCb experiment

The particle physics studies and future prospects of the LHCb experiment are briefly outlined in sections 1.2.1 and 1.2.2. A general overview of the LHCb detector and its structural changes in the upgrades is given in section 1.2.3.

1.2.1 Flavour physics in a nutshell

In 1928, Paul Dirac wrote at St. John's College in Cambridge: ‘The wave equation thus refers equally well to an electron with charge e as to one with charge $-e$ ’ [22]. This electron with positive charge, the positron, was discovered four years later by Carl D. Anderson [23]. Antimatter has taken a profound role in our understanding of particle physics ever since. However, the observed predominance of matter over antimatter has also led to a fundamental question: *why are we here?* The LHCb experiment is looking for the answer.

The predominance of matter over antimatter is believed to have arisen in the early evolution of the universe. The Sakharov conditions [24] state that in the early universe there must have been (i) baryon number violation ($n_B - \bar{n}_B \neq \text{constant}$), (ii) C and CP violation and (iii) departure from thermal equilibrium.

The Standard Model of particle physics can accommodate CP violating effects in the weak interactions of quarks. The unitary Cabibbo-Kobayashi-Maskawa (CKM) matrix [25] connects the weak eigenstates (primed) and the corresponding mass eigenstates:

$$\begin{pmatrix} d' \\ s' \\ b' \end{pmatrix} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} \begin{pmatrix} d \\ s \\ b \end{pmatrix} \quad (1.1)$$

The absolute values of the matrix elements show a hierarchical pattern with the diagonal elements close to unity, $|V_{us}|$ and $|V_{cd}|$ around 0.2, $|V_{cb}|$ and $|V_{ts}|$ of order 4×10^{-2} and $|V_{ub}|$ and $|V_{td}|$ of order 5×10^{-3} . The Wolfenstein parameterisation [26] expands each element as a power series in the small parameter $\lambda = |V_{us}| \approx 0.22$, thereby reflecting the CKM matrix hierarchy:

$$V = \begin{pmatrix} 1 - \lambda^2/2 & \lambda & A\lambda^3(\rho - i\eta) \\ -\lambda & 1 - \lambda^2/2 & A\lambda^2 \\ A\lambda^3(1 - \rho - i\eta) & -A\lambda^2 & 1 \end{pmatrix} + O(\lambda^4) \quad (1.2)$$

The four parameters A , ρ , λ and η govern all flavour-changing transitions in the quark sector, where the complex phase η is non-zero for CP violation. The unitarity condition $\sum_i V_{ij}V_{ik}^* = 0$ for $j \neq k$ provides six constraints. Each of these is the sum of three complex numbers closing a triangle in the complex plane. The $j = d$ and $k = b$ triangle applies directly to beauty decays and is shown in figure 1.3. The sides of the triangle are given by:

$$\left| \frac{V_{ud}V_{ub}}{V_{cd}V_{cb}} \right| = \sqrt{\rho^2 + \eta^2} \quad \text{and} \quad \left| \frac{V_{td}V_{tb}}{V_{cd}V_{cb}} \right| = \sqrt{(1 - \rho)^2 + \eta^2} \quad (1.3)$$

Additionally, the angles of the unitarity triangle are defined as:

$$\alpha \equiv \arg \left[-\frac{V_{td}V_{tb}^*}{V_{ud}V_{ub}^*} \right], \quad \beta \equiv \arg \left[-\frac{V_{cd}V_{cb}^*}{V_{td}V_{tb}^*} \right] \quad \text{and} \quad \gamma \equiv \arg \left[-\frac{V_{ud}V_{ub}^*}{V_{cd}V_{cb}^*} \right] \quad (1.4)$$

Measurements on the unitarity triangle are therefore directly related to the CKM matrix and the CP violation encoded in the Standard Model, and any discrepancy between the observables would be a strong indication of new physics.

The LHCb experiment strongly contributes to the precision bands in figure 1.3. The experiment extends the beauty physics results from B factories such as Belle and BaBar [27], by studying decays of a wide range of (heavy) beauty and charm hadrons produced at the high collision energy at the LHC. A theoretically clean example of a measurement at the LHCb experiment is the CP violation in the interference between the common final states with and without mixing before the decay [28]:

$$\begin{aligned} B^0 &\rightarrow J/\psi K_s^0 \\ B^0 &\rightarrow \bar{B}^0 \rightarrow J/\psi K_s^0 \end{aligned} \quad (1.5)$$




Figure 1.3: Unitarity triangle of the CKM matrix. The LHCb experiment contributes to the precision bands constraining the triangle. Any discrepancies would be a clear sign of new physics. The current status after LHC Run 1 and Run 2 (top) is followed by projections to 23fb^{-1} at the end of Run 4 (middle) and 300fb^{-1} at the end of Run 6 in the High-Luminosity LHC (bottom). Source: [21].

The J/ψ rapidly decays to a $\mu^+\mu^-$ pair and the K_s^0 decays to a $\pi^+\pi^-$ pair. A schematic of the LHCb event is shown in figure 1.4, following the approach in [29]. The LHC bunch-crossing collisions produce $b\bar{b}$ pairs mainly through the process of gluon fusion, where the gluons typically carry a different fraction of the beam momentum. This results in a boost of the centre-of-mass energy of the produced $b\bar{b}$ pair along the beam direction. The boost spreads the primary vertex (PV) and the decay vertices, allowing these to be resolved by the high-resolution LHCb tracking detectors.

Figure 1.4 shows the $\bar{B}^0 \rightarrow D^+\mu^-\bar{\nu}_\mu$ (or $B^0 \rightarrow D^-\mu^+\nu_\mu$) decay vertex, which marks the collapse of the wave function, meaning that the other B meson produced at the PV must be in the B^0 (or $\bar{B}^0$) state. This is called the flavour tag. The tagged particle subsequently propagates as the mass eigenstates and decays. The time-dependent decay rate asymmetry can be written as [28]:

$$\begin{aligned} \mathcal{A}_{J/\psi K_s^0} &\equiv \frac{\Gamma(\bar{B}_{t=0}^0 \rightarrow J/\psi K_s^0) - \Gamma(B_{t=0}^0 \rightarrow J/\psi K_s^0)}{\Gamma(\bar{B}_{t=0}^0 \rightarrow J/\psi K_s^0) + \Gamma(B_{t=0}^0 \rightarrow J/\psi K_s^0)} \\ &\approx S \sin \Delta m t - C \cos \Delta m t \end{aligned} \quad (1.6)$$

Here, Δm is the mass difference between the two mass eigenstates, and their decay-width difference has been neglected. The sine-term arises from interference between the decays in equation 1.5 with a common final state, and (in the limit of negligible higher order terms) the unitarity triangle observable $S \approx \sin 2\beta$ can be measured. The cosine-term arises either from direct CP violation in the decay or from CP violation in the B^0 - $\bar{B}^0$ mixing, which are both negligible. Figure 1.5 shows the measured asymmetry as a function of the decay time. The amplitude of the fit contributes to the $\sin 2\beta$ band in figure 1.3, while the period contributes to the Δm band.

Direct searches for new physics occur at the TeV mass range at the LHC. In order to increase this mass range, the *relativistic pathway* would be to build a larger collider with




Figure 1.4: Schematic of a $B^0 \rightarrow J/\psi K_s^0$ decay, where the B^0 meson is flavour-tagged at $t = 0$.




Figure 1.5: The time-dependent asymmetry provides unitarity triangle constraints on $\sin 2\beta$ and Δm . Source: [28].

higher centre-of-mass energy, such as the FCC. However, flavour physics also provides a sensitive probe via the *quantum pathway*, as the weak interactions in the mixing and decay of beauty and charm hadrons are mediated by gauge bosons that can be much heavier than the initial and final state particles. The LHCb precision measurements of rare and CP-violating decays are particularly sensitive to new physics through these loop contributions in flavour-changing neutral current transitions.

The photon and W and Z bosons have identical couplings to the three lepton generations in the Standard Model. This is known as lepton universality [30] and provides a strong test of the Standard Model. The LHCb experiment determines the ratio of purely leptonic or semi-leptonic processes that involve the same quark transition but leptons of different generations. This way, the quark CKM contributions cancel out in the ratio, along with most systematic experimental errors. Any deviation from unity in the ratio is a clear sign of new physics. The production of all beauty hadron species at the LHC facilitates a large variety of lepton universality tests.

Another part of the LHCb programme involves observations of exotic hadrons, which have internal structures more complex than the meson or baryon systems. For example, an unexpected narrow resonance in the mass spectrum of $\Lambda_b \rightarrow (J/\psi p) K^-$ was observed in 2015, which was consistent with the pentaquarks states [31].

Despite the wide search for CP violation, the obtained results are consistent with a sole source of CP violation encoded in the CKM mechanism. Studies of the dynamic generation of the baryon asymmetry in the universe suggest that this Standard Model prediction is too small to account for the observed matter-antimatter asymmetry [32]. As long as this question remains open, the LHCb experiment will keep testing the Standard Model with increasing precision.

The programme of measurements at the LHCb experiment relies upon a high detector performance. Events take place in a high (combinatorial) background environment, which requires highly efficient particle identification for flavour tagging as well as the determination of the final state. An impact parameter resolution on the micrometre scale is required to distinguish vertices and to measure the decay time. Physicists are therefore continuously finding ways to improve the detector performance, which is the topic of this thesis.

1.2.2 Physics motivation for the LHCb detector Upgrade II

As outlined in the physics case for the LHCb detector Upgrade II [21], the LHCb flavour physics programme would greatly benefit from the statistics that can be obtained during LHC Run 5 and Run 6. An example of the expected gain in the unitarity triangle measurements is given in figure 1.3. The precision bands at an integrated luminosity of 300 fb^{-1} put tight constraints on the Standard Model. Similar high gains can be expected in lepton universality tests and searches for rare decays and exotic states.

The unique opportunities to probe the Standard Model at high energy scales through the quantum pathway should be fully exploited. Although the complete revision of the LHCb experiment costs significant resources, the alternative, to push the particle physics frontier through the relativistic pathway by increasing the centre-of-mass energy, would be much more challenging.

The luminosity projections in the Upgrade II physics case contain the assumption that the excellent performance does not degrade under HL-LHC conditions. The increase in luminosity at the HL-LHC results in a higher track multiplicity, detector occupancy and fluence, which poses a challenge to maintain the high performance of the LHCb detector. Nevertheless, new contributions are foreseen in the tracking, such as the magnet side stations, and in the particle ID, such as the TORCH detector, which could improve the overall performance of the experiment.

1.2.3 Overview of the LHCb detector

The LHCb experiment is a single-arm forward spectrometer with an angular coverage of approximately 10 to 300 mrad in the horizontal plane and 250 mrad in the vertical plane, which is designed to detect the boosted beauty hadrons in the forward direction. The detector uses a right-handed coordinate system with the origin at the nominal interaction point, the x -axis pointing at the centre of the LHC ring and the z -axis along the forward beam direction.

During Upgrade Ia, the bottleneck of the level 0 (hardware) trigger is removed and the detector is read out in a data-driven fashion with event selection performed entirely in software. During Upgrade II, time information from the detector can be added and used to save CPU resources at the event reconstruction. The readout architecture is described in section 1.5.

The Upgrade Ia detector

The detector consists of planar sub-detectors, as shown in figure 1.6 for the Upgrade Ia design. The silicon Vertex Locator (VELO) is positioned closest to the interaction point and provides the high impact parameter resolution required for the decay time resolution and flavour tagging. The VELO tracks are combined with the segments from the Upstream Tracker (UT), which replaces the current Tracker Turicensis (TT), and the Scintillating Fibre (SciFi) tracker, which replaces the current three downstream tracking stations. Hence the charged particle track vector is reconstructed, along with a momentum measurement from the LHCb dipole magnet. The charged hadron ID in the momentum range of 10 to 100 GeV/c is performed by the Ring-Imaging Cherenkov (RICH) detectors. The RICH 1 detector is positioned close to the VELO and measures a higher track density than the RICH 2 detector located downstream of the LHCb magnet. The current first muon station is replaced by a neutron shield, which protects the SciFi photon detectors from the backscplash from the Electromagnetic Calorimeter (ECAL). The Hadronic Calorimeter (HCAL) and Muon System (M2-M5) form the remainder of the detector.

The Upgrade II detector

The Upgrade II detector is shown in figure 1.7. The sides of the LHCb dipole magnet are equipped with tracking stations, in order to track low-momentum charged particles exiting the detector. The inner part of the SciFi tracker, where the track multiplicity and fluence is highest, is replaced by a silicon tracker, and the combination of the SciFi and silicon tracker is renamed as the Mighty Tracker (MT). The Time of Internally Reflected Cherenkov Light (TORCH) detector, for charged hadron identification at 1 to 10 GeV/c, may be installed during detector Upgrade Ib. The TORCH detector would be positioned between the SciFi tracker (MT in Upgrade II) and the RICH 2 detector. The HCAL is replaced by up to 1.7 m of iron shielding in front of the muon system.

1.3 Tracking system upgrades

The future LHCb tracking system consists of the VELO, SciFi tracker, UT and magnet side stations, which are described in sections 1.3.1 and 1.3.3 and together provide tracking with a momentum resolution of $\delta p/p < 1.0\%$ [33]. The developments towards 4-dimensional tracking are discussed in this section, as well as the photon detector of the SciFi tracker, which is a candidate technology for the future RICH detector.

1.3.1 Vertex Locator

The VELO identifies the primary vertices and the secondary decay vertices. The tracker consists of two retractable halves that can be positioned closely to the beam interaction region, as shown in figure 1.8, and is exposed to a non-uniform fluence accumulating up to approximately $8 \times 10^{15} \text{ neq cm}^{-2}$ during LHC Run 3 and Run 4 [36]. The current silicon strip sensors are replaced by a hybrid pixel technology with a $55 \times 55 \mu\text{m}^2$ pitch in




Figure 1.6: The UpgradeIa LHCb detector contains the Vertex Locator (VELO), Ring-Imaging Cherenkov (RICH1) detector, Upstream Tracker (UT), dipole magnet, Scintillating Fibre (SciFi) tracker, RICH2 detector, electromagnetic calorimeter (ECAL), hadron calorimeter (HCAL) and Muon System (M2-M5). Source: [33].




Figure 1.7: The UpgradeII LHCb detector contains additional magnet side stations and silicon inner modules for the SciFi tracker, which is renamed as the Mighty Tracker (MT). The HCAL is replaced by iron shielding for the Muon System. The Time of Internally Reflected Cherenkov Light (TORCH) detector will be installed in Upgrade Ib. Source: [21].

Upgrade Ia [33]. The silicon pixels are read out by the VeloPix ASIC, which is based on the Timepix3 ASIC [37], at a 40 MHz rate without pixel hit time information. The hit occupancy in the inner regions will reach 900 Mhits/s yielding a total data rate of more than 3 Tbit/s.

The luminosity increase in LHC Run 5 poses a further challenge, with an order of magnitude increase in occupancy and data rate, and a radiation level of around 5×10^{16} neq cm⁻². Preliminary results from VELO simulation studies are shown in figure 1.9 and demonstrate that the PV mis-association level typically reaches 15 to 20% in this environment. Therefore, a new detector, with time information and smaller pixel sizes, needs to be installed during Upgrade II.

While a single technology is preferable from the R&D and commissioning point of view, a hybrid solution has also been proposed for the Upgrade II VELO. This hybrid consists of small, radiation hard pixels without timing or with low time resolution in the inner region and larger, less radiation hard pixels with a high time resolution in the outer region. As shown in figure 1.9, the PV mis-association is reduced to 6 to 8% using 50 to 100 ps time resolution in the outer region and 4 to 6% with an additional time resolution of 200 ps in the inner region of the hybrid [21]. The gains from time resolution are much greater than from the spatial resolution. For example, a reduction of the pixel size from 200 to 25 μ m in the outer region, with a 25 μ m pixel size in the inner region, improves the PV mis-association by 10%. However, a 70% improvement is expected when the time resolution in the outer region improves from 200 to 40 ps [35].




Figure 1.8: The Upgrade Ia VELO design. The sensors of the opposing detector halves form a diamond shape with the beam passing through the centre. Source: [34].




Figure 1.9: Preliminary estimates for the PV mis-association fraction after Upgrade II as a function of the outer pixel time resolution in the VELO hybrid design, where inner pixels have no time information (red) or 200 ps resolution (blue). Source: [35].

The NA62 Gigatracker provides an example of a current detector with high radiation hardness, high granularity and a time resolution better than 200 ps [38]. Its hybrid pixel matrix is made by bump-bonding two rows of five 100 μm -thick TDCpix ASICs to a 200 μm -thick planar silicon sensor with an area of $60.8 \times 27 \text{ mm}^2$ and a pixel size of $300 \times 300 \mu\text{m}^2$. Both the VeloPix and the TDCpix are based on 130 nm CMOS technology. In the new generation of TimePix ASICs, the TimePix4, 65 nm CMOS technology will be explored along with a time resolution of approximately 200 ps.

The TimeSpot project is a major R&D project for pixels with a time resolution below 100 ps, high-speed front-end readout electronics and real-time track reconstruction algorithms [39]. 3D silicon or diamond sensors are a promising development, where the bias and collection electrodes are fabricated in the vertical direction with respect to the silicon surface. These sensors are naturally fast and radiation hard, owing to the short paths taken by the charge carriers.

VELO time information would significantly reduce CPU resources by rejecting the combinatorial background at an early stage of the track reconstruction. This can be used to speed up the reconstruction in the online selection, to reduce ambiguities in track matching with the UT and MT in a high track multiplicity environment and to provide a PV timestamp to other sub-detectors that use time information.

1.3.2 The Scintillating Fibre tracker and Mighty Tracker

The SciFi tracker is being installed during detector Upgrade Ia, and consists of 6 planes, each covering a 30 m^2 area with 2.5 m-long scintillating fibres of 250 μm diameter [33]. The replacement of the inner part of the SciFi is anticipated for Upgrade Ib due to radiation damage [21]. The MT, shown in figure 1.10, is proposed for Upgrade II and consists of SciFi modules in the outer region and an area of approximately 17 m^2 of radiation-hard silicon modules in the inner and middle regions.

Silicon Photomultiplier (SiPM)

The SiPM used for light collection in the SciFi tracker is shown in figure 1.11. This sensor was developed by Hamamatsu¹ and consists of an array of Single-Photon Avalanche Diode (SPAD) pixels of 60 μm width. The pixels covering an area of $0.25 \times 1.62 \text{ mm}^2$ are combined into a single readout channel. This surplus of pixels per channel gives the SiPM photon-counting capabilities. An additional advantage of the SiPM is its peak quantum efficiency of about 45% at 450 to 500 nm [40].

The radiation environment at the LHCb experiment poses a challenge for the SiPMs. The neutron fluence strongly increases the dark-count rate (DCR), which is produced by random avalanches in the amplifying silicon region and forms the dominant source of noise in the

¹Hamamatsu Photonics: <http://www.hamamatsu.com>

SiPMs. Other contributions are from pixel-to-pixel crosstalk, with typically a 10% probability, and from after-pulsing. After receiving a neutron fluence of $6 \times 10^{11} \text{ neq cm}^{-2}$, which is about half of the expected fluence during LHC Run 3, the SiPMs produce a DCR of 14 MHz per channel at -40°C temperature. This DCR is acceptable for the SciFi tracker, which will be read out at 40 MHz and only retains signal clusters of more than 4 or 5 photoelectrons.

In the Upgrade II detector, the DCR can cause a high number of ghost tracks. Irradiation tests have shown that the DCR increases linearly with the neutron fluence [33], while the relation with temperature has an exponential dependence. Further cooling would therefore be a solution, and preliminary results show that the DCR in Upgrade II could be controlled using a cryogenic bath at 133 K [41]. Annealing of the SiPMs reduces the radiation damage and it has been considered to heat the detector to 40°C for one week during the annual shutdown period [33]. Alternatively, microlenses could be installed to focus the light onto a smaller SiPM area, thereby reducing the overall DCR. The bias voltage of the SiPMs can also be optimised and the time constant of the shaper in the front-end electronics can be minimised. Additional polyethylene-based neutron shielding could be installed upstream of the calorimeter system, in order to block high-energy neutrons from the calorimeter, which form the main source of SiPM radiation damage [42]. Despite all these possibilities, the lack of radiation hardness remains a strong disadvantage of SiPMs.

No photon hit time information below 25 ns is currently foreseen in the SciFi modules, due to the photon path length distribution and the decay time of excited states in the scintillating fibres [43, 41]. The discussion of the SiPM time resolution is therefore postponed until




Figure 1.10: The Upgrade II Mighty Tracker (MT) consists of SciFi modules in the outer region and silicon modules in the inner and middle regions. Source: [41].




Figure 1.11: The custom Silicon Photomultiplier (SiPM) of the Upgrade Ia SciFi tracker. Source: [40].

chapter 3. However, the SiPM signal shape is still important to improve the pixel recovery time and photon detection efficiency in the SciFi tracker. The recovery time constant is 20 to 100 ns, which is estimated to lead to an approximately 1% inefficiency in the highest-occupancy region during LHC Run 3 [33].

Silicon tracker modules

The addition of hit time information is foreseen in the silicon central region of the MT, in order to reduce the combinatorial background and to minimise the rate of incorrect matching of upstream and downstream track segments. A suitable technology may be the Low-Gain Avalanche Detector (LGAD). This n-in-p silicon detector contains a heavily doped p-layer just below the p-n junction, which results in a high internal electric field and an avalanche effect. The device therefore has internal gain and a high timing resolution. LGADs are used in the ATLAS High-Granularity Timing Detector, where a time resolution of 30 ps at the beginning of the sensor lifetime is expected [44]. The sensors have limited radiation hardness and a substantial zero-gain region. While this is acceptable for typically $1 \times 1 \text{ mm}^2$ pixels, the technology is less applicable to the VELO for example.

(HV-)CMOS (also called DMAPS) detectors are thin silicon devices with in-pixel analog electronics and on-chip digital readout electronics, which avoid the need for complex assembly processes. Current examples are the ATLASPix and MuPix chips, which have a hit time resolution of 13 ns and 7 ns respectively [45]. This time resolution is dominated by the rise time of the amplifier, as the collection from the depleted region typically occurs in less than 1 ns [46]. Although a factor of 2 to 3 is gained in the track time resolution using multiple detector planes, this is relatively slow compared to the PV time spread of 1 to 2 ns, as discussed in chapter 3. The radiation hardness up to a fluence of approximately $10^{15} \text{ neq cm}^{-2}$ is also lower than for the typical hybrid silicon technologies.

1.3.3 Magnet side stations

As shown by the green curve in figure 1.12, the magnet side stations would improve the low-momentum tracking capabilities of the LHCb detector, which enhances the determination of the multi-body final state and the flavour tagging [21]. A typical example of a decay involving a slow pion track is $D^{*+} \rightarrow D^0 \pi^+$. A proposal has been made to install a prototype in Upgrade Ia followed by full population during Upgrade Ib [47].

The stations consist of triangular extruded scintillator bars with embedded wavelength-shifting fibres read out by SiPMs. The technology considerations are therefore similar to the SciFi detector, and the same readout electronics can be used. As shown in figure 1.13, the SiPMs are located outside the magnet and the LHCb acceptance, in a lower radiation environment with less constraints for cooling and shielding. A spatial resolution of about 1 mm in the (horizontal) bending plane of the dipole magnet would provide a momentum resolution of $\delta p/p < 1\%$ [48].




Figure 1.12: The low-momentum pion tracking capabilities improve using the magnet side stations between the UT and SciFi tracker. Source: [48].




Figure 1.13: Magnet side stations viewed along the z -direction. The wavelength-shifting fibres are read out by SiPMs in a low-radiation environment outside the LHCb acceptance. Source: [47].

1.4 Particle identification system upgrades

Particle ID is of fundamental importance for the reconstruction of the invariant mass of beauty and charm hadrons with multi-body final states in LHCb. It plays a major role in flavour tagging and in the LHCb trigger decisions. Section 1.4.1 introduces the RICH detector. The TORCH detector concepts and technologies are described in section 1.4.2. The ECAL upgrades, including time information, are presented in section 1.4.3.

1.4.1 RICH detectors

The Ring-Imaging Cherenkov (RICH) detectors perform charged hadron identification at LHCb. An example of the critical role of the RICH detectors in the two-body hadronic decay $B \rightarrow h^+ h^-$, where h indicates a charged hadron, is given in [49].

Cherenkov radiation is the constructive interference of electromagnetic waves from atoms that are polarised by a charged particle travelling through the medium at a speed above the speed of light in the medium. The radiation is emitted at the characteristic Cherenkov angle θ_c with respect to the direction of motion of the charged particle:

$$\cos \theta_c = \frac{1}{n(\epsilon) \beta} \quad (1.7)$$

Here, $n(\epsilon)$ is the refractive index of the medium for a photon of energy ϵ and $\beta = v/c$ is the ratio of the charged particle speed to the speed of light in vacuum. The particle speed threshold for Cherenkov radiation is dictated by $|\cos \theta_c| \leq 1$ in equation 1.7. Different particle types cross the Cherenkov threshold at different momenta as shown in figure 1.14. At high momentum, $\beta \rightarrow 1$ and the Cherenkov angle reaches a limiting value. The principle of the RICH detector is to measure the Cherenkov angle, and therefore the particle speed,

and to determine the particle mass by combining this speed with a momentum measurement from the LHCb tracking system. Details on the RICH reconstruction algorithm are provided in chapter 3.

The RICH system consists of two detectors. The RICH 1 detector is located upstream of the LHCb magnet and contains a C_4F_{10} gas radiator for a particle momentum range of 10 to 65 GeV/c. The downstream RICH 2 detector contains CF_4 gas for a momentum range of 15 to 100 GeV/c. Cherenkov photons generated in the gas radiator are focused onto an array of single-photon-sensitive detectors by the RICH mirror system, which consists of a tilted spherical mirror followed by a plane mirror. While the design of the RICH 1 and RICH 2 detectors is very similar, the RICH 1 detector is more compact, with a higher track density and photon occupancy.

With the increased photon occupancy predicted for the LHCb upgrades, the reconstruction of events becomes less efficient as the overlap between Cherenkov rings from different particle tracks increases. In order to reduce this effect in the Upgrade Ia detector, the focal length of the spherical mirror is doubled and the mirrors are rearranged, as shown in figure 1.15. This results in an overall magnification at the detector area, which reduces the photon occupancy per pixel. During Upgrade II, lightweight composite precision mirrors will be placed inside the acceptance, which reduces the overall optical aberrations of the system. Optical aberrations give rise to the emission point error, as photons emitted along the particle trajectory take different optical paths to the photon detector plane. As shown in table 1.1, this error is significantly reduced in the upgrade detector.

The unknown energy of the emitted Cherenkov photons gives rise to a chromatic error in the dispersive gas radiator and gas enclosure quartz window. The improvements in the optical system reduce this spatial error. Additionally, a wide search for alternative radiator materials is ongoing, including gases, aerogel, meta-materials and photonic crystals. The




Figure 1.14: Reconstructed Cherenkov angles in the RICH 1 C_4F_{10} gas radiator. The different particle types separate into distinct bands. Source: [49].

	RICH 1			RICH 2		
	Current	Upgrade I	Upgrade II	Current	Upgrade I	Upgrade II
N_{pe}	32	42	60-30	24	22	30
Chromatic	0.84	0.58	0.24-0.18	0.48	0.31	0.10
Pixel	1.04	0.44	0.15	0.35	0.19	0.07
Emission	0.76	0.37	0.10	0.27	0.27	0.05
Total [mrad]	1.60	0.78	0.30-0.24	0.65	0.45	0.13

Table 1.1: The mean number of photoelectrons N_{pe} expected for saturated rings, and the Cherenkov angle resolution for the LHCb RICH detector and its upgrades. The N_{pe} and chromatic error ranges for Upgrade II come from the choice of optical filter. Source: [21].

chromatic error can also be reduced by choosing a photon detector with higher sensitivity in the green region of the optical spectrum or by reducing the quartz gas enclosure window thickness. The photon dispersion is discussed in more detail for the timing studies in chapter 3.

In order to achieve the 40 MHz readout rate of the UpgradeIa detector, the hybrid photon detectors, which contain encapsulated 1 MHz readout electronics in the vacuum tubes, are replaced by Multi-Anode Photomultiplier Tubes (MAPMTs) with modular readout electronics, which will be described in section 2.1. For the UpgradeII detector, a photon detector with increased granularity and photon hit time information is foreseen. A small prototype of this detector is being proposed for installation during Upgrade Ib [50]. A commonly used requirement for good pattern recognition is that the Cherenkov angle resolution σ_θ does not smear the ring over more than one pixel: $\sigma_\theta \cdot f \leq \sqrt{A_p}$. Here, f is the focal length of the mirror and A_p is the pixel area. The pixel size should therefore be reduced from 2.6 mm for MAPMTs to about 1 mm for the future upgrade detector, which improves the pixel error in table 1.1. The simulation studies presented in chapter 3 demonstrate that the addition of a photon hit time measurement is critical to reduce the high combinatorial background in Upgrade II.




Figure 1.15: The RICH1 optical system in the current (left), UpgradeIa (middle) and UpgradeII (right) configuration. The focal length of the spherical mirror is doubled during UpgradeIa, and the flat mirror is brought into the detector acceptance to reduce optical aberrations during Upgrade II. Source: [51].

1.4.2 TORCH detector

The TORCH detector, which has been proposed for installation during Upgrade Ib, would be the first large-scale system using sub-nanosecond timing in the LHCb experiment. The concepts and technology R&D of the TORCH detector are also relevant for time considerations in the RICH detector.

The TORCH detector enhances the charged hadron ID in the momentum range of 2 to 10 GeV/c [52]. The Cherenkov threshold for kaons is 9.3 GeV/c in the RICH1 radiator (C_4F_{10} with a refractive index of $n = 1.0014$) and 15.6 GeV/c in the RICH2 radiator (CF_4 with $n = 1.0005$). The TORCH detector combines time-of-flight and detection of internally reflected Cherenkov light techniques. The detector would be located upstream of RICH2 at about 9.5 m from the collision point, resulting in a time-of-flight difference of approximately 35 ps for 10 GeV/c kaons and pions. A particle track typically generates 30 Cherenkov photons in the TORCH radiator. A single-photon time resolution of approximately 70 ps therefore yields a time-of-flight resolution of $70/\sqrt{30} \sim 13$ ps, which is sufficient to discriminate kaons and pions based on the Rayleigh criterion [53].

The Cherenkov photons are trapped in the quartz plate by total internal reflection and collected at the periphery as shown in figure 1.16. A quartz block with a plated cylindrical mirror is used to focus the photons onto the detector plane. The Cherenkov angle θ_c and the photon path length d_γ through the quartz are reconstructed using the photon detector hit position and the track vector from LHCb tracking. The measured photon detector hit time of arrival t_{hit} is the sum of the primary vertex time t_{pv} , the time-of-flight and the




Figure 1.16: Schematic of the TORCH optics. Cherenkov photons are internally reflected and travel to the edge of the 10 mm-thick quartz plate, where they are focused onto the photon detector plane. Source: [52].

time-of-propagation:

$$t_{hit} = t_{pv} + \frac{d_{track}}{c} \sqrt{1 + \left(\frac{mc}{p}\right)^2} + \frac{d_\gamma n_g}{c} \quad (1.8)$$

Here, d_{track} is the reconstructed distance from the PV to the TORCH radiator interception point and p is the reconstructed track momentum. The time-of-propagation, which contains the group index n_g , has a contribution from the chromatic dispersion of quartz, because the Cherenkov photons are emitted over a range of energies. The Cherenkov angle measurement is used to correct for this chromatic error. Equation 1.7 can be rewritten as:

$$n(\lambda) = \frac{\sqrt{(mc)^2 + p^2}}{p \cos \theta_c} \quad (1.9)$$

For a *given* track mass hypothesis m , the refractive index can be calculated, which also provides a measurement of the photon wavelength. This information can then be used to calculate the group index ($n_g = n - \lambda \frac{dn}{d\lambda}$). Since the particle time-of-flight and photon time-of-propagation both increase with the particle mass, the particle ID performance is enhanced by measuring their sum. Similar to the RICH reconstruction algorithm described in chapter 3, the TORCH reconstruction will step through different track mass hypotheses and compare the measured photon hit time distribution to the predicted hit times in equation 1.8, in order to find the set of mass hypotheses that maximises the overall likelihood of the event.

The start time of the particles, which can be approximated to the PV time, is required for the time-of-flight measurement. In the first instance, the TORCH reconstruction assumes that all tracks are pions, and uses equation 1.8 to obtain the PV time. Any non-pion tracks will give outliers in this distribution, which can be rejected. Owing to the high pion multiplicity, the PV time can be determined with typically a few picoseconds resolution. This is the only possible way for TORCH to obtain the PV time during LHC Run 4.




Figure 1.17: Quantum efficiency of the Photek Phase-1 prototype MCP-PMT as a function of the collected integrated charge for different wavelengths. Source: [53].




Figure 1.18: Photek Phase-3 prototype MCP-PMT, with an active area of $53 \times 53 \text{ mm}^2$ in a $60 \times 60 \text{ mm}^2$ device. Source: [54]

After Upgrade II, however, a likely scenario is that the LHCb tracking system will adopt 4-dimensional tracking as discussed in section 1.3. In this case, the TORCH detector could save CPU resources by using the PV timestamp from the LHCb tracking system.

Micro-Channel Plate Photomultiplier Tube (MCP-PMT)

The MCP-PMT combines an excellent time resolution with a low dark-count rate, which suits the requirements of the TORCH detector. The operational principles of MCP detectors are explained in section 3.4.1. A peak quantum efficiency of the photocathode of 20 to 25% was measured for the TORCH MCP-PMT prototype [55], which is manufactured by Photek² and shown in figure 1.18. The device has a transit-time spread of 23 ps [56, 57]. A small-scale prototype of the full TORCH system, including the optics and electronics readout, achieved a $(83\text{--}115)\pm 6$ ps single-photon time resolution during 2018 beam tests [55], which is approaching the target time resolution of the system.

The main drawback of the MCP-PMT is its restricted lifetime, despite significant improvements using an atomic-layer deposition coating [58]. The measurements in figure 1.17, conducted over a period of 2.5 years using single-photon illumination from a blue LED, show a loss of quantum efficiency above approximately 3 C cm^{-2} [53]. An average integrated anode charge of at least 5 C cm^{-2} is expected in the LHCb experiment, with non-uniform occupancies across the TORCH photon detector plane [55].

1.4.3 ECAL

The electromagnetic calorimeter (ECAL) performs photon, electron and π^0 identification. The current Shashlik modules remain in the ECAL during LHC Run 3. The modules consist of alternating layers of lead and scintillator with integrated wavelength-shifting fibres read out by PMTs. The radiation-damaged inner part of the detector will be replaced during Upgrade Ib.

It is planned to use time information in order to mitigate the high combinatorial background from the pile-up of events in LHC Run 5. This can be achieved using dedicated silicon timing layers, following the approach of the CMS High-Granularity Endcap Calorimeter for example, which uses $320\text{ }\mu\text{m}$ -thick silicon sensors with a pixel area of 1.1 cm^2 and a time resolution of 10 to 50 ps [59]. The silicon planes could be embedded between absorber layers or behind a single converter plane at the front of the module. Simulation studies have shown a strong reduction in the number of overlapping primary vertices using three timing layers, as shown in figure 1.19.

Several technologies can provide a high intrinsic time resolution in the ECAL module, without the use of timing planes. A homogeneous crystal calorimeter with longitudinal segmentation offers a fast response time and a high radiation hardness and energy resolution, but is also the most expensive solution. Sampling calorimeters of the SpaCal (‘spaghetti’ of

²Photek Ltd.: <https://www.photek.com>




Figure 1.19: The mean number of incorrect primary vertices included in a one-sigma time window around the primary vertex of interest, using three timing planes with a resolution of 20 and 50 ps per plane. Source: [48].

fibres for scintillation and transportation of light) or Shashlik types, with a tungsten-alloy converter, provide an alternative option. The absorber plates in the SpaCal detector can be machined with grooves to host GAGG (Gadolinium Aluminium Gallium Garnet) crystal fibres [60], which are read out by PMTs. The cerium and magnesium co-doping in these fibres leads to a rise time of several tens of picoseconds [61]. The radiation-hardness of the wavelength-shifting fibres is a challenge for the Shashlik detector. Test beam studies with prototypes have demonstrated a shower time resolution of less than 100 ps [60]. However, this resolution may degrade for small energy depositions.

1.5 Data processing upgrades

In the LHCb Upgrade Ia detector, the full inelastic collision rate of 30 MHz will be processed in software. This new triggerless, data-driven design is motivated by the large amount of signal in the detector. This signal data rate after applying moderate transverse momentum and vertex displacement requirements would sum up to 133 GB/s for beauty hadrons, charm hadrons and light, long-lived particles such as K_S^0 and Λ_0 [62]. Since the bandwidth to storage is 2 to 5 GB/s, the trigger has to decide which signal to store according to the physics requirements, which can only be done in a flexible software trigger. Furthermore, the event processing is moved from the radiation environment at the detector front-end to remote computing farms, which reduces costs and allows expansion in the future.

The LHCb readout design is described in section 1.5.1. Hit time information is an important tool to reduce the order-of-magnitude increase in data rate and event reconstruction time during Upgrade II. Starting from the Upgrade Ia design, the use of time information is considered at different stages of the data processing in sections 1.5.2 to 1.5.4.

1.5.1 Upgrade Ia readout architecture

The UpgradeIa readout architecture is shown in figure 1.20. The sub-detector front-end electronics transfer the data through a radiation-hard optical link (the Versatile Link), using the GBT protocol at a 100 GBit/s rate. The data arrive at the event builder PCs through PCIe40 boards [63], which are the generic readout boards for all LHCb sub-detectors. Since the PCIe40 is embedded in the event builder PC, the data readout and the event building can be performed using the large PC memory as a data storage buffer. The PCIe40 also provides the bi-directional fast- and slow-controls communication channel to the front-end electronics through the Versatile Link.

The high-performance FPGA on the PCIe40 decodes and aligns the data frames according to their bunch crossing ID and can perform sub-detector-specific data processing. The data are buffered in the PC memory and all fragments that make up a single event are transferred to a single CPU on one of the 500 target PC-servers connected to the local-area network. The fragments are combined in the event builder CPU, with the possibility of applying a first stage of event selection. The events are then despatched to the event-filter farm, where the reconstruction and event selection takes place. The timing and fast-control supervisor maintains the overall system synchronisation. The supervisor is running in the FPGA on the PCIe40 and controls the data flow, synchronous and asynchronous commands and clock, timing and trigger signals.




Figure 1.20: The Upgrade Ia readout architecture. Data are sent from the front-end electronics to the event builder PCs, which despatch events to the event-filter farm for processing and storage. Source: [64].




Figure 1.21: The high-level trigger reduces the data rate from the event builder to storage and performs run-by-run calibration. Source: [65].

The software high-level trigger (HLT) is implemented at the event-filter farm. The events in the HLT are reconstructed in real-time with a maximum processing time of 13 ms imposed by the inelastic collision rate and the 2 to 5 GB/s data rate to storage. An important feature of the HLT is the automated alignment and calibration procedure that can be run for the VELO, UT, Muon System and the optical system of the RICH detectors. This allows event filtering to be based on fully reconstructed and calibrated quantities. Details of the HLT sequence are shown in figure 1.22. The VELO tracks are reconstructed first. While in the offline procedure all VELO tracks are extrapolated to the SciFi, which comprises the third ‘forward tracking’ stage, the online procedure contains an intermediate stage where (low-momentum) tracks are removed if there is no match with the UT, which reduces the CPU time. The total time taken in the tracking sequence is estimated to be 5.6 ms. The particle identification in the RICH and TORCH (Upgrade Ib) detectors relies upon a momentum measurement from the tracking system, which is obtained from the Kalman fit procedure at the seventh stage.

1.5.2 Time at the detector front-end electronics

A time gate at the detector front-end electronics can be used to reject background hits. As a practical implementation, the digital readout electronics would only register hits that arrive within a fixed-width time gate at a constant offset with respect to the LHC clock. This offset is sub-detector specific, and depends on the average particle time-of-flight from the nominal collision point to the sub-detector as well as the detector response time. The front-end gate needs to be sufficiently wide to accept all PVs, which are created with a time spread of $\sigma_t = 214$ ps as discussed in chapter 3, and to allow for the detector time resolution including sensor jitter, time walk or other sources. The minimum time gate is therefore 1 to 2 ns. The use of a front-end time gate in the RICH Upgrade Ia detector is discussed in chapter 2.

The front-end time gate has a direct impact on the data rate. For example, a decrease of a factor of eight in asynchronous detector noise can be expected by rejecting hits outside a 3 ns time window around the signal, provided the rise time and recovery time of the sensor are negligible. A smaller data reduction is expected for (background) hits that are caused by the bunch crossing, for example due to low-momentum curling tracks, signal spillover from previous bunch-crossings or scintillation in the detector, because these hits would need to arrive sufficiently out-of-time to be rejected by the front-end time gate. Figure 3.7 in chapter 3 will show that more than 90% of the hits, both signal and background, are correlated with the bunch-crossing and would arrive within the front-end time gate in the RICH detector.

The use of detector correlations in space or time, for example between photons in a Cherenkov ring or between hits in different tracking layers, is challenging at the front-end electronics, due to the modular nature of the electronics and the difficulty of implementing additional data paths between modules. Currently it would seem that the event builder would be a more practical alternative to exploit this possibility, as discussed in the next section.

1.5.3 Time at the event builder

The event builder is in a low-radiation location which allows the use of standard commercial technology. Since full information of each sub-detector event is available, the data can be pre-processed at the event builder farm, which reduces the data rate to the event-filter farm.

An example of track pre-processing is the 4-dimensional artificial retina algorithm, which is based on a neural network scheme allowing massively parallel and fast processing [66]. The algorithm, which can be executed on commercial FPGAs, first identifies stubs, which are coincidences between pairs of hits in adjacent planes compatible in space and time with tracks from the bunch-crossing area. A high time resolution at the pixel level is important for an efficient and selective stub definition. Stubs with similar time and space coordinates are subsequently combined into fully reconstructed tracks. The propagation of tracks to the event-filter farm saves bandwidth, because each track can be uniquely described by four space parameters and one time parameter rather than as a list of hits.

The 4-dimensional tracking algorithm would allow the event builder to take on a new, additional role to process detector hits to reduce the data rate to the event-filter farm. This can be done using a narrow time gate for the RICH1 detector hits, provided the extra step is made to reconstruct the primary vertex time from the tracks. As will be shown in chapter 3, the Cherenkov photons from a single PV arrive with a time spread of less than 30 ps, allowing a gate of approximately this width to be applied around the PV timestamp. During the LHC Run 5, approximately 50 PVs are generated in the bunch crossing. The width of the time gate would therefore vary between 30 ps (in the unlikely event that all PVs are created at the same time) and 1.5 ns (when the PVs are evenly spread in time with a 30 ps separation), which is similar to the front-end gate. Since the time spread from the sensor and electronics has not yet been taken into account and is likely to increase the required width of the gate, these cuts are not considered very feasible.

1.5.4 Time at the High-Level Trigger

Most of the detector hits are correlated with the bunch crossing and need to be transmitted to and processed by the event-filter farm. Here, the main challenge in the reconstruction and particle ID is to eliminate the combinatorial background caused by the high track multiplicity in the high-luminosity environment.

A scheme for the use of time information in the HLT is proposed in figure 1.22. The VELO track time would be obtained from 4-dimensional tracking in the event builder or reconstructed at the first HLT stage. Additional hit time resolution in the UT or the inner region of the MT would allow these trackers to be included in the 4-dimensional tracking algorithms, which would improve the track matching speed and efficiency. At the fourth stage, the time resolution of the PVs improves by a factor of $\sqrt{N_{tracks}}$ due to the large number of tracks emerging from each PV. Although the TORCH and RICH detectors could find the PV time as a first stage of their reconstruction algorithms, as

described for the TORCH detector in section 1.4.2, a PV timestamp from tracking would avoid this and improve the reconstruction speed. As discussed in detail in chapter 3, the RICH reconstruction algorithm greatly benefits from time information for the reduction of combinatorial background. As a consequence, the particle ID stage becomes significantly faster and more efficient.




Figure 1.22: The upgradeIa High-Level Trigger. Stages 2, 5 and 6 are not applied offline and intended to speed up the HLT. The improvements with time information after Upgrade II are illustrated in green.

1.6 Chapter conclusion

The increasing luminosity of the LHC upgrade programme provides a challenging, order-of-magnitude increase in the track multiplicity, hit occupancy and fluence at the LHCb detector. A common trend for the LHCb sub-detectors is the use of hit time information to mitigate this effect. Some of the emerging technologies for the detector upgrades are summarised in table 1.2. In order to retain the precise timing coming from the sensors, significant improvements in the readout electronics are also required.

Towards the future upgrades, detector hit time information can be used at various stages of the readout and trigger design. For example, a 4-dimensional FPGA-based tracking algorithm can be implemented in the event builder. For the RICH detector, a two-stage approach is proposed:

- A time gate of 1 to 2 ns width implemented in the digital logic at the detector front-end electronics, which is discussed in chapter 2.
- Sub-nanosecond time information implemented in the RICH reconstruction algorithm at the high-level trigger, which is discussed in chapter 3.

The front-end time gate is most effective at cutting asynchronous sensor noise, which lowers the overall bandwidth requirements. The sub-nanosecond time gate in the RICH reconstruction algorithm would reject combinatorial background, which improves the reconstruction speed and particle ID efficiency.

Technology	Time resolution [ps]	Sub-detector
SiPM	~ 200	SciFi, RICH
MCP-PMT	~ 30	TORCH, RICH
MAPMT	~ 300	RICH
TimePix ASIC	< 200	VELO
3D silicon sensor	< 100	VELO
LGAD	~ 30	ECAL, MT, UT
HV-CMOS	$\sim 10^3$	MT
SpaCal	< 100 (shower)	ECAL

Table 1.2: Overview of the time resolution of some of the technologies that are considered for the LHCb detector upgrades.

Chapter 2

Time information in the RICH Upgrade Ia photon detector

The Multi-Anode Photon Multiplier Tubes (MAPMTs) and the electronic readout system of the Upgrade Ia detector are described in section 2.1. The hit occupancy per MAPMT pixel reaches 25% at 40 MHz in the central region of the RICH 1 detector, which demands a high count-rate capability at the front-end [67]. It is essential for the RICH detector to maintain a high photon detection efficiency in these conditions. Furthermore, a low noise level is important for the particle identification performance of the RICH pattern recognition algorithms. The studies in this chapter address these parameters.

The photon detection efficiency of the MAPMTs is expected to degrade at sufficiently high exposure due to saturation effects in the dynode chain and instabilities in the biasing network. The photon detector was therefore tested at high photon rates and occupancies using a pulsed laser, as presented in section 2.2.

Signal-Induced Noise (SIN), where a noise hit arrives up to about $1\,\mu\text{s}$ after a photon is detected, was subsequently found by the collaboration for a subset of MAPMTs and pixels. These hits can be effectively reduced in the data by a time gate in the FPGA-based digital front-end electronics, which was implemented and tested for the first time during the October 2018 tests in the charged particle beam at the SPS facility in the CERN North Area. These measurements and the application of time information in the RICH Upgrade Ia detector are described in section 2.3.

2.1 The RICH photon detector system

An elementary cell (EC) of the photon detector is shown in figure 2.1. The MAPMTs, described in section 2.1.1, are directly plugged into a passive baseboard, which contains the resistor network for biasing the dynodes and routes the anode signals to the CLARO ASICs on the Front-End Boards (FEBs) for pulse shaping and discrimination, as described in section 2.1.2. The digital signals are then routed through the backboard to the Photon Detector Module Digital Board (PDMDB). This FPGA-based board, described in section 2.1.3, provides the configuration interface for the CLAROs, collects the hits from all channels and prepares the data for transmission across the Versatile Link. An aluminium case provides mechanical support, heat dissipation and an electrical ground connection.




Figure 2.1: Exploded view of an elementary cell of the RICH photon detector. The backboard connects to the digital readout electronics. Source: [68].

2.1.1 Multi-Anode Photomultiplier Tube (MAPMT)

The MAPMTs, developed by Hamamatsu and shown in figure 2.2, perform single-photon detection in the RICH detector. Two versions of different dimensions are used: the R13742 ‘1-inch’ and the R13743 ‘2-inch’ MAPMT, which are custom modifications of the R11265 and R12699 MAPMT respectively. Both types contain 8×8 pixels, with a pixel area of $2.88 \times 2.88 \text{ mm}^2$ and $6.0 \times 6.0 \text{ mm}^2$ for the 1-inch and 2-inch MAPMTs respectively [70]. The 2-inch MAPMTs are used in the low-occupancy regions of the RICH2 detector, as shown in figure 2.3.

The MAPMT has a UV-glass window, which enhances its radiation hardness. A super-bialkali photocathode is deposited onto this window. The quantum efficiency (QE) curve for the R11265 MAPMT is shown in figure 2.4. A peak QE of at least 30 to 35 % around a 350 nm wavelength is guaranteed [70]. Photoelectrons liberated at the photocathode are focused onto one of the 64 dynode chains using a wire mesh. This mesh generates a stronger




Figure 2.2: Photograph of the 1-inch MAPMT, looking through the UV-glass front window, showing the focusing wire mesh and the first dynode stage.




Figure 2.3: Layout of the RICH2 photon detector. 1-inch MAPMTs are used in the central, high-occupancy region, while the outer region is equipped with 2-inch MAPMTs. Source: [69].




Figure 2.4: QE of the R11265 MAPMT. The super-bialkali (SBA) photocathode and UV-glass window are used for the MAPMTs for the RICH detector. Source: [68].

focusing field at the top and bottom rows. A 12-stage (1-inch MAPMT) or 10-stage (2-inch MAPMT) dynode chain is used for electron amplification, each with a nominal gain of 10^6 at 1.0 kV bias voltage. The potential difference between the dynodes accelerates the electrons, which bombard the dynode and liberate more electrons. This repeated stop-and-start leads to a relatively long transit time of 5.1 to 5.3 ns, and a nominal transit-time spread of 350 ps and 280 ps for the 1-inch and 2-inch MAPMTs respectively [70]. The photocurrent is collected at the anode pins protruding from the back of the MAPMT. The MAPMTs are plugged into a bespoke baseboard, with a different design for the 1-inch and 2-inch MAPMTs.

2.1.2 The CLARO front-end ASIC

The CLARO is an 8-channel ASIC designed in $0.35\ \mu\text{m}$ CMOS technology. Eight CLAROs, four on either side, are mounted on a FEB. Four FEBs are connected to the 1-inch MAPMT baseboard and read out a total of 256 channels. In the case of the 2-inch MAPMT, two FEBs are used in order to reduce the signal path lengths between the anodes and the CLARO inputs, which lowers the input capacitance and noise.

The CLARO contains a charge-sensitive amplifier with an attenuation network at its input. This 2-bit configurable attenuation can reduce the input signal by up to a factor of eight, in order to compensate for the spread in gain of the MAPMT pixels. The amplifier output is DC-coupled to a discriminator with a 6-bit configurable threshold ranging up to an input-equivalent charge of $2\ \text{Me}^-$ in $30\ \text{ke}^-$ steps at zero attenuation or $16\ \text{Me}^-$ in $240\ \text{ke}^-$ steps at maximum attenuation. An additional offset bit can be used to lower the entire range of the discriminator thresholds by approximately $1\ \text{Me}^-$, which facilitates studies of the noise pedestal. An equivalent noise charge of $4\ \text{ke}^-$ was measured for the CLARO when operated inside the EC [71], which is three orders of magnitude smaller than the nominal signal amplitude. The average crosstalk in the system was 1% with a maximum value of 2.5%.

The CLARO is optimised for a low power consumption such that, despite the closely packed system, the combined heat output of the MAPMT biasing network and CLAROs is conducted through the EC to the actively cooled column mechanics. The power consumption depends on the photon occupancy, with 0.7 mW at a low occupancy and 2.5 mW at a 25% occupancy at 40 MHz rate. The power consumption of a disabled channel is negligible. Triple modular redundancy is implemented in the CLARO configuration registers and the CLARO radiation hardness has been tested at radiation levels that exceed the expected level for the RICH detector [72].

The CLARO operates asynchronously. The amplifier has a 6 ns peaking time with a 25 ns recovery time and a recovery time from saturation of less than 50 ns [73]. The RMS jitter introduced by the CLARO was measured to be 113 ps for signals around $300\ \text{ke}^-$ and 9 ps for large signals of $4.5\ \text{Me}^-$ [74], which is negligible compared to the MAPMT transit-time spread and the CLARO time walk.




Figure 2.5: Time walk arises from the variation of (MAPMT) signal amplitudes at a fixed (CLARO) threshold. A measurement of the time-over-threshold allows a time-walk correction.




Figure 2.6: Two PDMDBs are used to configure the CLARO ASICs and to collect and transmit the data from four ECs with 1-inch MAPMTs. Source: [75]

Figure 2.5 illustrates the principle of time walk. The ‘delay’ is the time at which the threshold is crossed, and the time walk is defined as the variation in this delay with the signal amplitude. The CLARO input bias current can be increased using an external resistor on the FEB, which reduces the time walk by a factor of 2 in a compromise with an increased power consumption of 1.5 mW per channel at low occupancy [74]. Two common techniques to correct for the time walk are to measure the time-over-threshold or to apply a constant-fraction discriminator, where the threshold is set to a constant fraction of the peak amplitude of the signal. While the constant-fraction discriminator is not implemented in the CLARO, a time-over-threshold correction is possible in the digital readout electronics, as discussed in section 2.3.4.

2.1.3 Photon Detector Module Digital Board (PDMDB)

The PDMDB provides the interface between the CLARO ASICs and the Versatile Links to the PCIe40 readout board. The PDMDB for the 1-inch MAPMT readout is shown in figure 2.6. The FPGAs capture the digital signals from the ECs, format the data and pass it on to the Data Transmission Module (DTM). This plug-in module uses the GBTX transceiver ASIC [76] in wide-frame mode to transmit the data to the PCIe40 card via the Versatile Link.

Two different PDMDB versions are used. For the 1-inch MAPMT readout, the board is equipped with three FPGAs and three DTMs, and two boards are used to read out four ECs, as seen in figure 2.3. Therefore, each board processes $4 \text{ ECs} \times 4 \text{ MAPMTs} \times 64 \text{ pixels} / 2 \text{ PDMDBs} = 512 \text{ channels}$ [75]. The PDMDB for the 2-inch MAPMT readout contains

two FPGAs and two DTMs. Here, one board is used to read out the four ECs with a total of $4 \text{ ECs} \times 1 \text{ MAPMT} \times 64 \text{ pixels} = 256 \text{ channels}$.

The XC7K70TFBG676-1C FPGA of the Xilinx¹ Kintex-7 series is the smallest, most cost-effective FPGA that meets the IO bank requirements for the connectivity to the CLAROs and the GBTX. The FPGA was chosen for its high flexibility, which allows the firmware to be updated to benefit from future developments and potential bug-fixes. An example is the recent use of time in the readout electronics, which is discussed in section 2.3. However, the radiation environment is a disadvantage for the use of FPGAs, due to the risk of single-event upsets (SEUs). The FPGA therefore contains triple modular redundancy and its logic footprint is minimised by design, with much of the architectural and synchronisation logic moved to the PCIe40 board. The Kintex-7 was selected to take advantage of its built-in configuration memory correction logic which reduces the need for scrubbing since most SEU-induced bit errors are automatically recovered. The use of partial reconfiguration further reduces the scrubbing time to a few seconds by selectively reconfiguring only the affected region of the FPGA. Its robustness against radiation has been tested up to a total ionising dose of 1 Mrad [77], which is five times higher than the expected maximum dose in the RICH 1 detector.

The Trigger and Control Module (TCM) contains the bi-directional master link for fast commands to control and synchronise the data acquisition and for the exchange of configuration and monitoring data. The master GBTX ASIC controls the Versatile Link and distributes an LHC-synchronous clock to the front-end electronics. The Slow-Control Adapter (SCA) ASIC acts as a bridge between the master link and the configuration protocols for the FPGAs, CLAROs and DTM GBTXs. Additionally, it provides programmable digital IOs, ADCs for temperature and voltage monitoring and the DACs for the CLARO test pulse.

The CLARO outputs are asserted asynchronously to the LHC clock and are captured and synchronised in the PDMDB FPGA. Three 160 MHz clocks, one for each DTM, are distributed across the PDMDB and used to sample the CLARO signals at the FPGA IOs. As described in more detail in section 2.3.2, this sampling occurs at a 320 MHz bit rate using both clock edges. The clock phase can be adjusted per PDMDB or per FPGA to compensate for variations in signal latency across the photon detector plane. This latency is the sum of the time-of-flight of particles and photons through the detector, the MAPMT response time and the CLARO response time and time walk. The grouping of MAPMTs with similar gain to the same PDMDB, in order to be able to apply the appropriate bias voltage for example, will also be beneficial for the time resolution of the system.

¹Xilinx inc.: <https://www.xilinx.com>

2.2 Performance at high rate and high photon occupancy

In this section, I present measurements of the behaviour of the photon detector system at high photon occupancies of up to 45% and high rate up to 20 MHz, which are similar conditions to the central region of the RICH1 detector in the LHCb experiment. The results, which are published in [78], motivated the individual biasing of the eleventh and twelfth dynodes of the 1-inch MAPMT.

A pulsed laser light source illuminated the EC, as described in section 2.2.1. Rather than the PDMDB, an earlier type of readout board was used for these studies, which is described in section 2.2.2 and will also be used for the semiconductor tracker in part 2 of this thesis. The pulse-height spectrum of an MAPMT pixel is explained in section 2.2.3 and the results are presented in section 2.2.4. Additional tests, subsequently performed by the RICH group, revealed a systematic source of signal-induced noise, which is discussed in section 2.2.6.

2.2.1 Experimental setup

The EC for these studies was mounted on a cold bar as shown in figure 2.7. Unlike in the figure, only one EC was mounted at a time, and this EC was exchanged for the comparison between the 1-inch and 2-inch MAPMTs. A single power supply channel was connected across the biasing network on the baseboard and the supply current was recorded as an indicator of the photoelectron current in the MAPMTs.




Figure 2.7: Two elementary cells, one with four 1-inch MAPMTs and one with one 2-inch MAPMT, are mounted on a cold bar.




Figure 2.8: Overview of the test setup. A veto is used to downscale the data rate with respect to the laser pulse rate.

The EC was illuminated by a pulsed laser diode with a wavelength of 405 nm and a pulse width of 60 ps FWHM [79]. The light intensity and pulse rate were set by the external driver. The laser was coupled to an optical fibre, and the other end of the fibre was aligned with the centre of the EC inside a light-tight box. Figures 2.9 and 2.10 show a map of the average number of photoelectrons per pixel per laser pulse, which is referred to as the hit occupancy, for the ECs with 1-inch and 2-inch MAPMTs respectively. This hit occupancy was estimated from the pulse-height spectrum as explained in section 2.2.3. The non-uniform pattern of illumination allowed a high occupancy to be achieved at the centre of the EC while staying in the linear regime of the laser driver. While no data are available on the linearity of the laser, the intensity is expected to vary with the laser rate at high output power, and a careful calibration would be required in order to be able to compare data recorded at different rates. The linearity of the laser intensity is further discussed in section 2.2.5. The non-uniform illumination allows a direct comparison of a range of occupancies in the same MAPMT.

An FPGA-based digital readout board with Ethernet connectivity to the DAQ PC was used in the test setup. Two boards read out one EC with 1-inch MAPMTs, and a third board, with different FPGA logic, was used to provide the readout triggers and is called the trigger controller. A schematic of the trigger flow is shown in figure 2.8. The synchronous output signal from the external laser driver was downscaled to a 20 kHz rate by a NIM veto module and connected to the trigger controller. The controller distributed the trigger




Figure 2.9: Occupancy (%) map for an elementary cell containing four 1-inch MAPMTs at the ‘low intensity’ setting.



Figure 2.10: Occupancy (%) map for the 2-inch MAPMT. The EC was illuminated non-uniformly.

simultaneously to both readout boards, and discarded any new readout triggers that arrived while the readout boards were processing the event.

For any given laser rate, the readout rate was downscaled to 20 kHz in order to keep the data rate within the limit of the 100 MBit/s Ethernet connection between the digital boards and the DAQ PC. With the designed FPGA input logic, two photon signals in consecutive clock cycles cannot be distinguished from a single photon signal extending into the second clock cycle. This places an upper limit on the laser pulse rate of approximately 20 MHz since most photon signals return to zero within 25 ns.

2.2.2 Readout board

The readout board shown in figure 2.12 contains a Xilinx Spartan-6 XC6SLX25 FPGA. Owing to its versatility, relatively low cost and external USB and Ethernet connectivity, these boards were used for many tests of the Upgrade Ia elementary cell and its prototypes [80]. Additionally, the wide range of configurable input-output (IO) logic standards facilitates the interfacing to different types of front-end electronics, such as the ASICs of the semiconductor tracker modules in part 2 of this thesis.

Two readout boards, without the IO expander, are directly coupled to the EC backboard as shown in figure 2.7. The FPGA firmware is designed for the readout to work autonomously once configured at the start of a run. Configuration data are sent as a single Ethernet packet with custom protocol and containing a unique target identifier in the packet payload. The signals from the CLAROs are captured and buffered in FPGA RAM. The events are formatted into multi-event data packets and the appropriate Ethernet headers are added. The packets are then forwarded, employing fragmentation where needed, to the external transceiver. For these studies, the Ethernet interface was operated in 100-baseTX mode, but the hardware supports 1000-baseT mode with minor firmware changes.

A readout board with a different version of FPGA firmware was used to distribute the readout trigger. In this configuration, which is referred to as the ‘trigger controller’, the readout board is controlled through the USB interface. An additional IO expander and hardware




Figure 2.11: Plug-in for the trigger controller.




Figure 2.12: The FPGA-based readout board with Ethernet and USB connectivity (right) and the IO expander (left).

plug-in, shown in figure 2.11, can be connected to the trigger controller to discriminate LEMO signals from scintillator trigger detectors. Although not used for the high-rate tests in this section, this plug-in is important for the trigger controller in the beam tests in section 2.3 and the semiconductor tracker in part 2.

2.2.3 Pulse-height spectrum analysis

The pulse-height spectrum characterises the gain, photon detection efficiency and noise behaviour. The procedure to record the pulse-height spectrum of an MAPMT pixel is outlined in this section.

Since the MAPMT readout is binary, the pulse-height spectra are derived from the integral distribution obtained in a threshold scan. This scan records 2.5×10^5 readout triggers at each step of the 6-bit CLARO discriminator threshold. The CLARO offset bit is set in order to study the noise pedestal and first photoelectron peak of the spectrum. The fraction of events containing a hit during the threshold scan is plotted in figure 2.13 for three MAPMT pixels with different nominal gains and hit occupancies. Pixel 1 was closest to the centre of the laser spot and pixel 64 furthest. At low threshold, the CLARO outputs are always high, followed by a drop in the fraction of events above the narrow noise pedestal. At high threshold, the hits are dominated by the photoelectron signal. The bins immediately after the pedestal drop are a measure of the total number of events with at least one photon in the pixel. Neglecting the events with more than one photon, it gives an estimate of the pixel occupancy. A sufficiently high gain is required to separate the signal from noise and




Figure 2.13: Threshold scans of three MAPMT pixels with different hit occupancy. A function based on three Fermi PDFs and an exponential decay is fitted. The occupancy is estimated from the plateau after the noise pedestal drop at low thresholds, as indicated in the plot.

to create a plateau in the region used to estimate the occupancy, as shown in figure 2.14. This is guaranteed by an amplitude ratio of at least 1.3 between the first photoelectron peak and the valley between the pedestal and the first photoelectron peak [70].

A probability density function f was fitted to each threshold curve and the fitted function differentiated to obtain the pulse-height spectrum [80]. The function with seven free parameters is based on three Fermi PDFs, representing the pedestal and first and second photoelectron peaks, as well as an exponential to describe the right-hand side tail of the pedestal:

$$f(\mu; \sigma_{0,1}; x_{0,1}; A; k) = \frac{\mathcal{P}(0, \mu)}{1 + e^{(x-x_0)/\sigma_0}} + \frac{\mathcal{P}(1, \mu)}{1 + e^{(x-x_0-x_1)/\sigma_1}} + \frac{\mathcal{P}(2, \mu)}{1 + e^{(x-x_0-2x_1)/(\sqrt{2}\sigma_1)}} + \frac{A}{k} e^{-k(x-x_0)} \quad \text{for } x > x_0 \quad (2.1)$$

Here, $\mathcal{P}(N, \mu)$ is the Poisson probability of N photoelectrons for an average number of μ photoelectrons. The mean x_i and width σ_i are for the noise pedestal ($i = 0$) and first photoelectron peak ($i = 1$), from which the second photoelectron peak properties are derived. The exponential tail of the pedestal is caused by dark counts and photoelectrons that skip the first dynode stage of the MAPMT, leading to a reduced signal amplification. The exponential has normalisation A and coefficient k .

The pulse-height spectra of three pixels are shown in figure 2.14. The noise pedestal has a smaller width than the threshold step. The height of the first photoelectron peak is used as a measure of the occupancy and the photon detection efficiency, while the position of the peak




Figure 2.14: The pulse-height spectra obtained by differentiating the fitted function in figure 2.13. The positions of the first photoelectron peak reflect the nominal gain of the pixels, which was obtained from Hamamatsu specifications.

characterises the gain of the pixel. The second photoelectron peak would only be within the threshold scan range for the lowest-gain pixel shown in red. However, the occupancy in this pixel was approximately 20%, resulting in a probability of two photoelectrons of only $\mathcal{P}(2, 0.20) \approx 2\%$.

2.2.4 Results at high rate and high photon occupancy

The pulse-height spectra were recorded and compared for (i) a range of laser pulse rates between 100 kHz and 20 MHz, (ii) two laser intensity settings, referred to as ‘low’ and ‘high’ intensity, (iii) the 1-inch and 2-inch MAPMTs and (iv) 0.9 kV and 1.0 kV bias voltages.

During LHC Run 3 and 4, the peak occupancy can reach 25% for the 1-inch MAPMTs in the central region of the RICH 1 detector and less than 10% for the 2-inch MAPMTs in the RICH 2 detector [67]. The occupancy map for the 1-inch MAPMTs in the test setup was shown for the low intensity setting in figure 2.9. An occupancy of nearly 30% was reached in the centre of the EC, which increased to 45% for the high intensity setting. The 2-inch MAPMTs were tested at a laser intensity setting selected to give an occupancy of up to 20% as shown in figure 2.10. Although these occupancies are higher than in the RICH detector, the laser rate was 20 MHz compared to the 40 MHz LHC collision rate.

Figure 2.15 compares the pulse-height spectra of a typical 1-inch MAPMT pixel at laser pulse rates from 100 kHz to 20 MHz at the two intensity settings at 1.0 kV bias voltage. The first photoelectron peak exhibits no changes in gain or photon detection efficiency with the laser rate. As expected, the hit probability is higher at the high laser intensity setting. Close to the noise pedestal, at threshold DAC values between 30 and 40, the hit probability




Figure 2.15: Pulse-height spectra for a typical 1-inch MAPMT pixel at 1.0 kV bias voltage and a CLARO input attenuation of a factor of four, showing a range of laser rates from 100 kHz to 20 MHz at two light intensity settings.

is observed to increase at the 20 MHz laser rate for both laser intensities and at 1.0 kV, as well as 0.9 kV in figure 2.16. This effect could be caused by a small number of photon signals that return to zero after 50 ns, thereby causing a hit during the next laser pulse. Since this tail of the signal is only observed at low thresholds, it can be effectively eliminated by an appropriate choice of threshold during operation.

The pulse-height spectra at 0.9 kV bias voltage are shown in figure 2.16. Here, the laser rate affects the first photoelectron peak, reducing the bin content by approximately 20%, which indicates a loss of photon detection efficiency. This is caused by instabilities in the dynode voltages, induced by the photocurrent inside the MAPMT, as discussed in the next section. The observed effect is representative for all pixels and is more pronounced at increased photon occupancy.

The 2-inch MAPMT contains 10 dynodes with a total resistance of the biasing network of $2.5\text{ M}\Omega$, compared to $3.0\text{ M}\Omega$ for the 12 dynodes in the 1-inch MAPMT. The photocurrent therefore makes up a smaller fraction of the total current through the biasing network for the 2-inch MAPMT, which results in a greater stability of the dynode voltages. The typical pulse-height spectra in figure 2.17 were recorded at 0.9 kV bias voltage. A reduction of approximately 10% in bin content at the first photoelectron peak was observed at high laser rates.




Figure 2.16: Pulse-height spectra for a typical 1-inch MAPMT pixel at 0.9 kV bias voltage and CLARO input attenuation of a factor of two. At this bias voltage, the amplitude of the first photoelectron peak drops at high rate.

2.2.5 Dynode biasing scheme

This section discusses the dynamic behaviour of the biasing network that caused the observed reduction in photon detection efficiency at high rate. A schematic of the MAPMT biasing network is shown in figure 2.18. An increase in photoelectron current through the dynode chain, due to a higher laser rate or light intensity, can change the potential differences between the individual dynodes. This effect was studied using a SPICE [81] simulation in [67]. Here, the authors study the stability of the eleventh and twelfth dynode as a function of the photocathode current as shown in figures 2.19 and 2.20. The reduction in V_{11} and V_{12} at high photocathode current degrades the collection efficiency at the anodes and accounts for the observed reduction in photon detection efficiency at 0.9 kV bias voltage in figure 2.16. Although this reduction in V_{11} and V_{12} is more pronounced at increased bias voltage, it is compensated by an increase in the photoelectron multiplication through the dynode chain. As a result, no loss in photon detection efficiency was observed at 1.0 kV bias voltage.

The authors in [67] proposed a modified powering scheme where V_{11} and V_{12} are biased with an independent voltage source. This powering scheme, which is shown in figure 2.18, provides a constant voltage at 16% and 10% of the overall bias voltage for dynode 11 and 12 respectively, which recovered the high-rate performance as shown in figure 2.23. This solution was only applied to the 1-inch MAPMTs, as these showed a stronger rate-dependence than the 2-inch MAPMTs and are operated in the regions of the detector with the highest occupancy.

The power supply current was monitored as a function of the laser pulse rate at the high laser intensity setting, and converted into a photocathode current using the 1.0 kV curve from the SPICE simulation results in figure 2.21. This photocathode current is plotted as a function of the laser pulse rate in figure 2.22. The photocathode current increases linearly up to approximately 100 pA at 40 MHz rate. The linearity of the photocathode current with the laser rate indicates no variation of the laser light intensity with the pulse rate, which was an important assumption discussed in section 2.2.3.

The maximum rating for the average anode output current of the 1-inch MAPMT is 100 μ A [70]. For the nominal gain of 10^6 at 1.0 kV, the photocathode current in figure 2.22 shows that this maximum rating is reached at 40 MHz, suggesting a very high anode current at the pixels with highest occupancy. As this is likely to accelerate the ageing of the MAPMTs, it is beneficial to start the operation during the experiment at low gain, which is achieved using a bias voltage of 0.8 to 0.9 kV.




Figure 2.17: Pulse-height spectra for a typical pixel of the 2-inch MAPMT at 0.9 kV bias voltage and CLARO input attenuation of a factor of two.




Figure 2.18: 1-inch MAPMT dynode biasing scheme, with the possibility to individually bias dynodes 11 and 12. Although the photocathode is negatively biased, the charge polarity has been inverted everywhere for the purposes of the simulation.




Figure 2.19: SPICE simulation results for the voltage across dynode 11 as a function of the photocathode current at total bias voltages between 0.8 kV and 1.1 kV.




Figure 2.20: Identical plot to figure 2.19, for dynode 12. The dynode numbering is defined in figure 2.18. Source: [68].




Figure 2.21: SPICE simulation results for the power supply bias current as a function of the photocathode current at total bias voltages between 0.8 kV and 1.1 kV. Up to 1.0 kV, the relationship is approximately linear. Source: [68].




Figure 2.22: Photocathode current as a function of the laser rate during the tests presented in section 2.2.4. The photocathode current was obtained from the measured bias current at 1.0 kV using the SPICE simulation results in figure 2.21.




Figure 2.23: Pulse-height spectrum for a typical 1-inch MAPMT pixel at 0.9 kV bias voltage without (left) and with (right) individual biasing of dynodes 11 and 12 as illustrated in figure 2.18. Source: [82].

2.2.6 Signal-Induced Noise (SIN)

In February 2018, the RICH group installed a detector module containing four ECs behind the detector plane inside the RICH 2 detector, in order to test the UpgradeIa detector performance in the LHCb environment using Cherenkov light passing through the gaps between the hybrid photon detectors. During these tests, Signal-Induced Noise (SIN) in the MAPMTs was discovered, which manifests itself as a noise hit arriving typically between a few hundred nanoseconds and a few microseconds after a signal hit in the same pixel. Due to this relatively long time scale, there is a risk of pile-up of SIN hits at high photon occupancy and high rate in the experiment. In this section, the characteristics of SIN are described and the presented high-rate tests, which were recorded around September 2016, are revisited in the light of SIN.

The SIN behaviour is shown in figure 2.24, where three overlapping datasets of $0.6 \mu\text{s}$ were recorded, each with a different offset with respect to the photon signal. The number of hits in the MAPMT in each 25 ns bin is plotted over the total range of $1.4 \mu\text{s}$. The signal peak at approximately 100 ns is followed by a minimum at around 150 ns and an increase in the number of hits due to SIN. The results for four different MAPMTs are shown, ranging from low SIN for the red curve to very high SIN with a long time constant for the blue curve.

As shown in figure 2.25, the SIN hits are concentrated around the first row, pixels 1 to 8, and especially the last row, pixels 57 to 64. These rows are distinguished by their high focusing electric field in the wire mesh behind the photocathode, which was seen in figure 2.2, and




Figure 2.24: The number of hits per 25 ns over a time period of $1.4 \mu\text{s}$, for four different MAPMTs. The signal at 100 ns is followed by a tail of SIN hits with different time constants. Source: [83].




Figure 2.25: The distribution of SIN hits in a 1-inch MAPMT. No SIN is observed in the central region, while the rows adjacent to the (labelled) bias pins contain highest SIN.

their proximity to the dynode bias voltage pins of the MAPMT, which are labelled in figure 2.25. A small number of MAPMTs display very high SIN, where hits in neighbouring pixels can also induce noise. SIN has a low dependence on temperature and resistor values in the biasing network [84]. However, as will be demonstrated in section 2.3, the number of SIN hits reduces with the bias voltage and is negligible for the 2-inch MAPMTs.

The presented high-rate tests were revisited after SIN was discovered. Figure 2.26 shows the ratio of the hit occupancy at 20 MHz to 100 kHz laser rate for the four 1-inch MAPMTs in the EC at 1.0 kV bias voltage. The average ratio is approximately 1.0 for each MAPMT, which is consistent with the observations made for the pulse-height spectra. The pixels 57 to 64 show a 15% increase in the occupancy at high rates for MAPMT 2 and 3. Due to the non-uniform illumination pattern during the tests, these peripheral pixels had a low occupancy of typically 5% as shown in figure 2.9. The *absolute* increase in hit occupancy was therefore less than 1%. As shown in figure 2.2, the dynode structure for the edge rows of pixels 1 to 8 and 57 to 64 are similar, and different to the other (central) rows. However, pixels 1 to 8 did not show any rate-dependence, despite the estimated photon occupancy of up to 45% at pixel 1 for the results in figure 2.26.

MAPMT 3 showed the largest SIN effect in the results in figure 2.26 but is amongst those with the smallest effect in the randomly selected sample used in the subsequent studies. This can be seen in figure 2.24, where the red curve (labelled ‘FA0044’) shows the measurement for MAPMT 3, and the other curves show MAPMTs that were not used at the time. An example of an MAPMT with very high SIN, subsequently recorded using a uniform low-intensity illumination in a similar setup to the presented high-rate tests, is shown in figure 2.28. This MAPMT reaches a ratio of approximately 10 for the hit occupancy at 20 MHz to 100 kHz laser rate at the pixels 1 to 8 and 57 to 64. The SIN effect was therefore initially not observed in the high-rate tests due to the random choice of MAPMTs.

Although studies to establish the fraction of MAPMTs with high SIN are still ongoing, first indications are that the fraction is small enough that these MAPMTs can be substituted or placed in the relatively low-occupancy regions of the RICH detectors. Further efforts in the collaboration are being made to simulate the effect of different levels of SIN on the particle ID performance. Additionally, a time gate of less than 25 ns in the FPGA of the PDMDB can significantly reduce the number of SIN hits in the data, as discussed in the next section.




Figure 2.26: Ratio of the hit occupancy at 20 MHz to 100 kHz laser rate for the four 1-inch MAPMTs at 1.0 kV bias voltage. The illumination was non-uniform as described in figure 2.9.




Figure 2.27: Laser light spot during the high-rate tests. For each 1-inch MAPMT, the pixels 57 to 64 are located in the low-occupancy region on the outside of the EC.




Figure 2.28: Example of the ratio of the hit occupancy at 20 MHz to 100 kHz laser rate for an MAPMT with very high SIN. This MAPMT has not been used for the high-rate tests. The pile-up of SIN hits increases the occupancy by approximately a factor of ten for the outer rows. Source: [83].

2.3 SPS beam tests with a time gate in the FPGA

The emergence of SIN has prompted intensive activity within the RICH group in order to understand and to mitigate its effect. One possible approach could be to adopt time gating to suppress the out-of-time SIN. In this section, the feasibility of applying a time gate in the electronic readout that has been designed for the RICH Upgrade Ia detector is demonstrated. These measurements, performed during the October 2018 beam test period, show a reduction of SIN hits in the data whilst retaining nearly all signal from the SPS charged particle beam. The test setup at the CERN North Area in Preveessin is described in section 2.3.1, the implementation of the time gate in section 2.3.2 and the analysis and results in section 2.3.3. Section 2.3.4 discusses a time-over-threshold correction that is beneficial for the implementation of a time gate in the RICH detector Upgrade Ia.

Throughout the course of my PhD, I participated in six beam test periods of the Upgrade Ia system. A paper of the 2015 beam tests has been published in [80]. For the two test periods in 2017, I installed the semiconductor tracker, which is described in part 2 of this thesis, in order to provide tracking information for photon detection efficiency measurements of the MAPMTs. The use of the tracker during the beam tests is described in appendix A.




Figure 2.29: Schematic of the lens used to generate Cherenkov radiation and to focus it onto the MAPMTs. The beam passes through the centre of the lens. Source [80].




Figure 2.30: Image of the aligned test system including the crossed scintillator trigger detectors, the Cherenkov lens and a column of 1-inch MAPMTs.

2.3.1 Test setup at the SPS charged particle beam facility

The charged particle beam from the SPS facility is mainly composed of protons and pions with a momentum of $180 \text{ GeV}/c$, and has an approximately Gaussian transverse intensity profile with a standard deviation of 18 mm . The beam arrives in spills of approximately 4.8 s at 50 s intervals, with a particle rate of typically around 500 kHz during a spill.

The beam passes through the borosilicate glass planoconvex lens of the RICH detector setup shown in figures 2.29 and 2.30 [80]. The lens, with a diameter of 15 cm and a thickness at the centre of 27 mm , acts as both the Cherenkov radiator and the focusing device. Beam particles enter the lens at the centre of the spherical side. The generated Cherenkov photons undergo total internal reflection at the flat surface, followed by a reflection from the reflective coating deposited on the spherical surface of the lens, which focuses the photons onto the MAPMTs at the detector plane. As shown in figure 2.29, a black tape mask was used to select photons produced in the first 13 mm of the lens, which reduces the error due to the unknown emission point of the Cherenkov photons.

During the tests in 2018, two columns of four ECs detected arcs of the Cherenkov ring, as shown in figures 2.36 and 2.37. These two columns, one containing 1-inch MAPMTs and the other 2-inch MAPMTs, were placed outside the beam acceptance with the beam passing between the columns. A crossed scintillator detector pair with a $5 \times 5 \text{ mm}^2$ area was used, in coincidence with another crossed pair behind the lens and MAPMT columns, to trigger




Figure 2.31: The readout scheme for the RICH test setup in the particle beam. The trigger controller performs the discrimination and coincidence of the signals from the scintillator detectors. On receipt of the trigger, the MiniDAQ stores the data frame from the PDMDB.

the system. The alignment of the beam with the centre of the lens was achieved using a remotely controllable translation table as described in more detail in appendix A.

The trigger and readout scheme of the system is illustrated in figure 2.31. The PDMDBs are connected through data-driven Versatile Links to the MiniDAQ, which provides the 40 MHz system clock. A trigger controller, as described in section 2.2.2, is used to discriminate the signals from the scintillator detectors. The FPGA on the trigger controller performs the coincidence of the signals and distributes the trigger to the tracking system and the MiniDAQ. On receipt of the trigger, the MiniDAQ stores the corresponding data frame arriving from the PDMDB via the Versatile Link.

Unlike in the LHCb experiment, the time of arrival of the beam particles at the test setup was asynchronous. Therefore, the subset of the beam particles that arrived approximately synchronously with the system clock needed to be selected in order to be able to test the implementation of a time gate in the PDMDB. The strategy adopted for the trigger controller was to create a synchronous time gate of approximately 6.25 ns and only trigger the system if a charged particle arrived within this gate. As shown by the dashed line in figure 2.31, the clock signal from the MiniDAQ was connected to an input of the trigger controller FPGA in order to achieve the synchronisation of the gate.

The FPGA logic for the trigger controller gate is shown in figure 2.32. The first register is clocked by the 40 MHz system clock (SCLK) and creates a toggled clock (TCLK), which toggles on each rising edge of the system clock owing to the inverter at the output of the register. The TCLK signal is input to the second register, which is clocked by a 160 MHz clock (FCLK) derived from the system clock using a phase-locked loop. This register switches on the rising edge of FCLK and therefore has the effect of delaying TCLK by one clock period of FCLK, resulting in DTCLK. TCLK and DTCLK are input to an XOR gate, which asserts a high GATE signal when its two inputs are not equal. The coincidence of the trigger scintillators is only made if the fourth scintillator detector signal arrives while the GATE signal is high.

The signal timings presented in figure 2.32 are in the absence of any signal delays introduced by the signal paths inside the FPGA. In practice, depending on the synthesis of the VHDL firmware, delays can change the relative phase of TCLK and DTCLK and therefore change the width of the GATE signal by typically up to a few nanoseconds. Since the width of the gate is fixed for a given synthesis, the test strategy is valid, although the gate is not reproducible once the firmware is re-synthesised. In order to estimate the width including signal delays, the number of scintillator detector counts was recorded at the trigger controller. While the three un-gated scintillator detectors recorded $(3.23 \pm 0.28) \times 10^5$ counts, the gated detector recorded 4.7×10^4 counts, which suggests a width of around 3.7 ns. Although the width of the gate introduces some uncertainty, this scheme was relatively quick to implement on the short timescale during the test period.

2.3.2 Implementation of a time gate in the PDMDB

This section describes the implementation of a hit time gate for the CLARO signals arriving at the PDMDB. The logic exploits the deserialiser, which is optimised to be able to operate at gigabit rates and embedded in every input-output (IO) logic block of the Kintex FPGA [85].

The operation of the deserialiser is shown in figure 2.33. The FPGA deserialiser samples the CLARO signals using both edges of the 160 MHz clock and shifts the sampled data at 320 MBit/s into an 8-bit shift register. This byte is used to address the 256-bit memory of a lookup table, which is a readily available resource with a small logic footprint in the general purpose logic of the FPGA. The value of the bit at this memory location is presented at the output of the lookup table on each 40 MHz system clock edge. The lookup table can therefore be programmed to detect specific signal patterns arriving from the CLARO channel.

Two patterns were used in the lookup table during the October 2018 beam test period. These patterns, referred to as the ‘3.125 ns level’ logic and the ‘6.25 edge’ logic, are shown in table 2.1. The edge (low to high transition) in the ‘6.25 edge’ logic was required in the third or fourth 3.125 ns bin of the 25 ns system clock period, and during the beam test the length of the cable transmitting the clock signal from the MiniDAQ to the trigger controller was adjusted in order to ensure that the CLARO signal transition occurred in this region. An alternative, programmatic method for the RICH Upgrade Ia detector to align the CLARO transition with the edge pattern in the lookup table would be to use the bit-slip feature of the IO deserialiser, which offsets the 8-bit window in steps of one bit (3.125 ns at 320 MBit/s).

The ultimate limit of this method of time gating in the FPGA of the PDMDB is imposed by the 1 GHz sampling rate at the Kintex 7 IOs. Although this would allow the width of the gate to be adjusted to 1 ns, it would also mean that each byte from the deserialiser corresponds to 8 ns rather than 25 ns in the current implementation. Additional general-purpose logic resources would therefore be required in order to maintain the time-over-threshold correction discussed in section 2.3.4. The use of FPGAs will allow these different implementations to be explored. The front-end time gate in the LHCb experiment should accept all primary vertices, which have a time spread of about 1 ns as discussed in chapter 3. This front-end gate is therefore within the capabilities of the existing Upgrade Ia PDMDB.




Figure 2.32: Schematic of the FPGA logic in the trigger controller that creates the 6.25 ns GATE signal synchronous to the 40 MHz system clock (SCLK). The ideal timings in the absence of signal delays are shown.




Figure 2.33: Schematic of the implementation of a time gate in the FPGA. The CLARO signals at each IO are sampled using a 320 MHz clock and deserialised. The lookup table allows flexible hit pattern recognition and time gating.

Logic	Lookup pattern
3.125 ns level	1XXX XXXX : Eight versions with the required bit in different locations
6.25 ns edge	000 1 XXXX 00 1X XXXX
Time-walk correction	0 1 1 1 1 1 X X 0 1 X XXX 0

Table 2.1: Patterns in the FPGA lookup table for different types of logic. X can be either 0 or 1. The ‘3.125 ns level’ and ‘6.25 ns edge’ logic were implemented during the October 2018 beam test. The ‘time-walk correction’ logic for the UpgradeIa detector is proposed in section 2.3.4.

2.3.3 Analysis and results

This section presents the beam test results for the signal detection efficiency and the reduction in the amount of recorded SIN hits using the 3.125 ns level logic and 6.25 ns edge logic in the FPGA.

The CLARO signal detection efficiency averaged over all channels is shown as a function of time in figure 2.34. Each point in the distribution was recorded by programming the 3.125 ns level logic with a different position of the required bit in the pattern of the lookup table. The CLARO thresholds were set to the minimum between the noise pedestal and the first photoelectron peak in the pulse-height spectrum. The trigger was gated using the scheme presented in section 2.3.1, and the uncertainty in the time of the trigger smears the leading and trailing edges of the signal by an estimated amount of 3.7 ns. As will be shown in figure 2.38, time walk introduces a further spread of up to 6 ns for the leading edge and up to 15 to 20 ns for the trailing edge. These spreads and the channel-to-channel variations in the FPGA result in a peaked structure, rather than a plateau, of the distribution. The distribution falls rapidly to zero after 25 ns due to additional logic intended to avoid artificial consecutive hits in the experiment. Without this cut-off, the distribution would show a longer tail determined by the characteristics of the CLARO amplifier.

The recorded number of SIN hits with the 3.125 ns level logic was found to be similar to the ungated readout. This was due to the width of the CLARO signals, which is typically between 15 and 30 ns as shown in figure 2.38. Hence, a hit could be registered by the level logic even if the leading edge of the signal arrived before the 3.125 ns bin. The 6.25 ns edge




Figure 2.34: The CLARO signal time distribution using the gated trigger in the beam test setup. Each data point was recorded using a different required bit in the 3.125 ns level logic in the FPGA.




Figure 2.35: Hit time distribution for pixel 62 of a 1-inch MAPMT on the Cherenkov ring. The baseline (B), signal (S) and signal-induced noise (N) regions are used to quantify the benefit of time gating.




Figure 2.36: Hit time distribution for three selected pixels on the Cherenkov ring during the particle beam tests. The 6.25 ns edge logic significantly reduces the noise baseline.




Figure 2.37: Hit time distribution for two selected pixels on the Cherenkov ring, with highest SIN observed in pixel 62.

logic was therefore introduced, with the low to high transition in the third or fourth 3.125 ns bin as shown in figure 2.34. Since the hit time distribution peaks in the sixth 3.125 ns bin, a small number of signal transitions occurs after the fourth bin and is therefore not registered by the 6.25 ns edge logic.

The hit time distributions of six selected pixels on the Cherenkov ring are shown in figures 2.36 and 2.37 for the 2-inch and 1-inch MAPMT column respectively. These plots show the number of hits detected by the logic as a function of the 25 ns readout slot. Readout limitations necessitated that, in order to cover the desired time window of $2.2\ \mu\text{s}$, the data were recorded as four different datasets, each with a different time offset with respect to the trigger. The signal arrives at the start of the second data set. Comparing the red and the black curves, the 6.25 ns edge logic significantly reduces the background while maintaining a high signal efficiency. A figure of merit for this improvement in signal-to-noise will be introduced in the following discussion. The SIN is highest for pixel 62 in figure 2.37, where it peaks between 400 and 500 ns after the signal, and lowest for the 2-inch MAPMT.

The hit time distribution of the pixel in figure 2.35 is divided into three regions: the baseline (B), the signal (S) and the signal-induced noise (N) regions. The baseline region is populated by particles that are outside the trigger acceptance, SIN hits from previous beam particles and dark counts for example. The integral over this 550 ns-wide baseline region was scaled to account for the different range of the signal (25 ns) and the noise (1625 ns) regions and subtracted from the signal and noise integrals.

The signal and the signal-to-noise ratio are shown with and without a time gate and at 0.9 kV and 1.0 kV bias voltage in table 2.2. The SIN generally reduces at lower bias voltages, which is reflected in the higher S/N ratio at 0.9 kV. At 1.0 kV bias voltage, the S/N ratio improves by a factor of 3.13 to 3.96 with the introduction of the 6.25 ns edge logic. This is slightly lower than the expected factor of four due to a 10.5% loss in signal caused by the trigger time spread and time walk. The signal loss increases to 33.6% at 0.9 kV due to additional time walk. The MAPMT pixel gain is typically a factor of 2.5 lower at 0.9 kV than at 1.0 kV [86], which reduces the signal amplitudes and increases time walk as shown in figure 2.38. This strongly motivates a time-over-threshold correction, which is discussed in the next section.

The results in table 2.2 can be used to estimate the rate of SIN in the LHCb environment. Let q be the number of 25 ns bins for the SIN effect to diminish after the signal arrives. The observed SIN in a given LHCb bunch crossing b is then the sum of the SIN hits induced by the signal in bunch crossings $b - 1$ up to $b - q$. Under the reasonable assumption of a constant signal level at each LHCb bunch crossing, this sum is the same as the integral over the noise region in figure 2.35, because here the noise approximately returns to the baseline near the end of the range. Therefore, the rate of SIN can be estimated for the pessimistic

Bias [kV]	Readout logic	Signal [$\times 10^3$]	Signal/Noise	Improvement factor
1.0	Original (25 ns)	44.67 ± 0.21	$1.07 - 1.36$ ¹	$3.13 - 3.96$ ¹
	6.25 ns edge	40.01 ± 0.20	4.25 ± 0.05	
0.9	Original (25 ns)	41.96 ± 0.20	4.53 ± 0.05	2.56 ± 0.06
	6.25 ns edge	27.90 ± 0.17	11.60 ± 0.25	

Table 2.2: Comparison between the traditional 25 ns readout and the 6.25 ns edge logic at 0.9 kV and 1.0 kV bias voltage for 1-inch MAPMT pixel 62. The improvement factor is approaching the expected factor of four at 1.0 kV. At 0.9 kV, the S/N ratio is generally higher, but the improvement factor is lower due to the signal loss from time walk.

case of 1.0 kV bias voltage, the traditional 25 ns readout, the highest SIN pixel 62 and the highest occupancy of 30% in a small region of the RICH 1 detector:

$$SIN \text{ rate} \approx occupancy \cdot \frac{N}{S} = 0.25 \pm 0.03 \text{ hits/pixel/bunch} \quad (2.2)$$

The S/N improvements presented in table 2.2 suggest that this 25% maximum SIN occupancy can be reduced to 3% using 6.25 ns edge logic at 0.9 kV bias voltage, provided a suitable time-walk correction can be applied. These numbers should be considered as estimates, and for example do not apply to pixels with a significantly longer SIN time scale, such as the blue curve in figure 2.24.

2.3.4 Time-over-threshold correction for the Upgrade Ia detector

The results in the previous section have shown that the SIN effect can be reduced by lowering the bias voltage, and that the SIN data rate can be reduced by applying a time gate in the FPGA. The time-walk correction proposed in this section would allow both techniques to be combined. The FPGA is the only possible place to implement the time-walk correction in the Upgrade Ia detector, owing to the flexibility of its programmable logic. Closer to the sensors, the implementation of, for example, a constant-fraction discriminator in the CLAROs would require a redesign of the ASIC, and closer to the PCIe40 cards, the pulse-width information required for a time-over-threshold (ToT) correction cannot be available due to the limitation in the available bandwidth.

¹These are limits due to an incomplete data set from the beam tests. At 0.9 kV, data were recorded with the 25 ns, 6.25 ns edge logic and the 3.125 ns level logic. The 3.125 ns level logic gives a S/N ratio of 5.73, which is a factor of $s = 1.26$ higher than the ratio for the 25 ns readout. At 1.0 kV, the 25 ns readout data set was incomplete. The improvement factor of 3.13 therefore arises from the S/N ratio of the 3.125 ns level logic compared the 6.25 ns edge logic. From the observations at 0.9 kV bias voltage, this is known to be an underestimate. The improvement factor of 3.96 was obtained using the scale factor $s = 1.26$ at 0.9 kV bias voltage, which is an overestimate because the CLARO signals are expected to be longer at 1.0 kV than at 0.9 kV bias voltage, increasing the probability of a hit in the 3.125 ns level logic and thereby reducing the scale factor s .

The lookup table in the FPGA logic presented in section 2.3.2 could be enhanced with a time-walk correction without increasing the logic footprint. The 6.25 ns edge logic during the beam tests required the CLARO transition to occur in the third or fourth 3.125 ns bin to span the trigger time spread. In table 2.1, the envisaged lookup patterns for a time-walk correction in the LHCb experiment are shown. These patterns look for the CLARO transition in a single 3.125 ns bin, which is achievable given the signal time spread of approximately 1 ns in the LHCb experiment, as discussed in chapter 3, rather than the estimated 3.7 ns uncertainty on the trigger gate during the beam tests. Depending on the ToT, the lookup table checks for the transition in the second or third 3.125 ns bin. In the proposed pattern, the ToT criterion is 18.75 ns, corresponding to six bins of 3.125 ns. At a lower ToT, the signal delay is expected to delay the CLARO transition from the second to the third 3.125 ns bin. This initial choice for the ToT criterion is based on figure 2.38. Choosing the blue and red curves for CLARO thresholds between 210 ke⁻ and 450 ke⁻, small input signals arrive after around 13 ns and large input signals arrive after around 7 ns, resulting in a time walk of approximately 6 ns. A signal delay of approximately 10 ns corresponds to a ToT of 15 to 20 ns. Therefore, provided the phase of the FPGA sampling clock can be fine-tuned to align the boundary of the IO sampling bins with this 10 ns signal delay, signals with a ToT of more than 18.25 ns are expected to arrive in the second 3.125 ns bin, while the signals with a smaller ToT are expected in the third 3.125 ns bin.




Figure 2.38: The time-over-threshold (‘width’) and the time taken to reach the threshold (‘delay’) as a function of the input signal amplitude at different threshold settings in the CLARO low-power mode. Source: [73].

The proposed scheme is a first step towards the implementation of a time-walk correction in the Upgrade Ia detector, and requires further study. This can be performed independently of the commissioning of the upgrade, owing to the programmable nature of the FPGA. Optimisations of the pattern in the lookup table are foreseen, for example in the ToT criterion. Furthermore, a study of the signal detection efficiency could lead to a more relaxed condition, such as 0X1XXXX0, where the transition is allowed in the second *or* third 3.125 ns bin, which would be similar to the beam test implementation. In hardware, the possibility of operating the CLAROs in ‘timing mode’ could be explored [74], although this would require the modification of a resistor on the FEB for each CLARO and would alter the power consumption and thermal behaviour at the front-end. Ultimately, the particle identification performance is the most important parameter, and simulation studies can help to find the optimal compromise between the signal efficiency and noise reduction in the detector.

2.4 Chapter conclusion

In this chapter, a time gate in the FPGA of the digital readout board was used to reduce noise hits from the Upgrade Ia MAPMT. The MAPMT performance at high rate and high photon occupancy, as well as its noise characteristic, were described.

The LHCb RICH Upgrade Ia MAPMTs are operated at a hit occupancy per pixel of up to 25% at 40 MHz in the central region of the RICH 1 detector. The high count-rate capability at the front-end was studied up to approximately 45% hit occupancy at 20 MHz rate using a pulsed laser light source. The high photon exposure led to dynamic instabilities in the biasing network, which could be mitigated by individually biasing the last two dynodes in the MAPMT dynode chain.

In February 2018, the RICH group discovered a source of signal-induced noise in the MAPMTs, where a noise hit can typically arrive between a few hundred nanoseconds and a few microseconds after a signal hit in the same pixel. The emergence of this noise has prompted intensive activity within the group in order to understand and to mitigate its effect. One promising approach is the use of a time gate to suppress out-of-time noise hits, which was implemented in the FPGA on the PDMDB and tested at the SPS beam facility. The results demonstrated that the 25% maximum noise occupancy at 1.0 kV MAPMT bias voltage can be reduced to 3% using a 6.25 ns time gate at 0.9 kV, provided the proposed time-walk correction is applied.

Chapter 3

Sub-nanosecond time for particle identification in the HL-LHC

In this chapter, the potential of time information to improve particle identification in the future LHCb RICH detector is demonstrated without constraints on the time resolution imposed by the photon sensor or readout electronics. This is particularly relevant for the LHCb Upgrade II. As discussed in chapter 1, the majority of LHCb sub-detectors expect to use time information to mitigate the increased hit occupancy during the HL-LHC Run 5 and Run 6.

The results in this chapter demonstrate the exceptionally high intrinsic time resolution of less than 10 ps for the RICH detectors, owing to the prompt Cherenkov radiation and the focusing mirror system. This provides a powerful tool for background suppression and improved particle ID performance using photon detectors and readout electronics with a high time resolution in the future.

The benefits of time information are explored using the LHCb simulation software, which is based on the flexible, configurable and object-oriented Gaudi framework [87]. The simulation contains a detailed description of the proton-proton interactions at the bunch crossing, the particle propagation through and interaction with the sub-detectors, the photon sensor, the digitisation processes at the front-end electronics and the event reconstruction. The Upgrade Ia detector description and the LHC Run 3 and Run 4 luminosity are simulated. Other potential detector improvements in Upgrade II, such as the mirror alignments or pixel granularity discussed in section 1.4.1, are not included.

The developments required to incorporate time into the simulation are described in section 3.1. The 3-dimensional RICH reconstruction algorithm is presented in section 3.2. This is followed by the development of an accurate prediction of the Cherenkov photon hit time of arrival on the RICH photon detector plane, which is exploited in the proposed

4-dimensional RICH reconstruction algorithm. The particle ID performance using the time coordinate and the implications for future detector technologies are discussed in sections 3.3 and 3.4 respectively.

3.1 Time in the LHCb simulation framework

The developments of the LHCb simulation framework in the time domain are discussed in this section. A novel 4-dimensional primary vertex generator is presented in section 3.1.1. The generated particle tracks and photons are then propagated through the LHCb detector and produce Cherenkov photons in the RICH detector. Although the chromatic error of the RICH detector has been studied in detail in the spatial domain, its evaluation in the time domain required a correction to the photon group velocity, which will be discussed in section 3.1.2. The storage of hit time information in the reconstruction is described in section 3.1.3. The simulation of spillover is discussed in section 3.1.4. An approximation of the conditions at the HL-LHC Run 5, with a high multiplicity of interactions within the same bunch crossing (pile-up), is introduced in section 3.1.5.




Figure 3.1: The LHC bunches are modelled as 3-dimensional Gaussian PDFs propagating at the bunch-crossing angles θ and ϕ towards the interaction point. The bunches are shown at three different times around the global $t = 0$.

3.1.1 A 4-dimensional primary vertex generator

A novel simulation tool has been developed to generate a 4-dimensional PV distribution function and to randomly sample PVs from this function. The tool, which replaces an existing tool where the PV time was simply fixed at $t = 0$, is the first step of the LHCb simulation and is invoked for each PV.

The LHC proton bunches are modelled as 3-dimensional Gaussian PDFs, as shown in figure 3.1. The PDFs are denoted f_+ and f_- for the bunches travelling forward and backward respectively along the z -direction and can be expressed as:

$$f_{\pm}(\mathbf{x}, t) = \frac{1}{\sqrt{(2\pi)^3 |\boldsymbol{\Sigma}_{\pm}|}} \exp \left\{ \left[-\frac{1}{2} (\mathbf{x} - \boldsymbol{\mu}_{\pm})^T \boldsymbol{\Sigma}_{\pm}^{-1} (\mathbf{x} - \boldsymbol{\mu}_{\pm}) \right] \right\} \quad (3.1)$$

Here, $\mathbf{x} = (x, y, z)$ is the position vector. Each bunch is rotated in the LHCb coordinate system by the half-crossing angle θ about the x -axis and ϕ about the y -axis. The time coordinate t enters the equation through the positions $\boldsymbol{\mu}_{\pm}$ of the centres of the bunches:

$$\boldsymbol{\mu}_{\pm} = \begin{bmatrix} -ct \cos \theta \sin \phi \\ -ct \sin \theta \cos \phi \\ \pm ct \cos \theta \cos \phi \end{bmatrix} \quad (3.2)$$

The covariance matrices $\boldsymbol{\Sigma}_{\pm}$ of the forward and backward travelling bunches are given by:

$$\boldsymbol{\Sigma}_{\pm} = \mathbf{R}_y(\mp\phi) \cdot \mathbf{R}_x(\pm\theta) \cdot \mathbf{C} \cdot \mathbf{R}_x^T(\pm\theta) \cdot \mathbf{R}_y^T(\mp\phi) \quad (3.3)$$

Here, $\mathbf{R}_x$ and $\mathbf{R}_y$ are the rotation matrices and $\mathbf{C}$ is defined as:

$$\mathbf{C} = \begin{bmatrix} \sigma_{\perp}^2 & 0 & 0 \\ 0 & \sigma_{\perp}^2 & 0 \\ 0 & 0 & \sigma_{\parallel}^2 \end{bmatrix} \quad (3.4)$$

The standard deviation $\sigma_{\parallel}$ is parallel to the beam. $\sigma_{\perp}$ is perpendicular to the beam and can be expressed in terms of the Twiss beta parameter β^* at the interaction point and the emittance ϵ [88]:

$$\sigma_{\perp} = \sqrt{\epsilon \beta^*} \quad (3.5)$$

The product of the bunch PDFs, $f_+ \cdot f_-$, describes the 4-dimensional PV distribution function, which is omitted here for clarity and presented in appendix B.

For the UpgradeIa conditions in this study, $\beta^* = 2.79$ m, $\epsilon = 0.509$ nm, $\sigma_{\parallel} = 90$ mm, $2\phi = 0.385$ mrad and $\theta = 0$. Using equation 3.5, the width of the Gaussian bunches is therefore $\sigma_{\perp} = 38 \mu\text{m}$. The PV distribution is truncated at 4σ in the space coordinates.




Figure 3.2: Markov chain decision tree to randomly sample 4-vector $\mathbf{x}$ from function $f(\mathbf{x})$.

A Markov chain sampler is used to randomly sample from the primary vertex distribution function. The implementation of the Markov chain is shown schematically in figure 3.2. The chain starts at the origin. A Gaussian perturbation is then applied with a standard deviation of typically 10% of the width of the function. The sample of the function is the value $\mathbf{x}$ obtained after n repetitions of the chain, where $n = 1000$ was chosen for this study. The Markov chain sampler was tested using various functions. The result of a triangle function test is shown in figure 3.3.




Figure 3.3: Triangle function test of the Markov chain sampler used for the primary vertex generator.




Figure 3.4: 4-Dimensional PV distribution obtained using a Markov chain to sample from the distribution function of two colliding Gaussian bunches.

Figure 3.4 shows the PV distribution for ten thousand bunch crossing events. The observed time distribution is symmetric around zero with a standard deviation of $\sigma_t = 214$ ps. The PVs are distributed along the z -direction with $\sigma_z = 57$ mm. Due to the small bunch crossing angle, the width of the PV distribution, $\sigma_{x,y} = 36 \mu\text{m}$, is similar to the width of the individual bunches.

The main parameters determining the PV distribution are the LHC bunch length and the beam crossing angles. The generator imports these quantities, as well as the beta parameter, emittance and bunch width, from the configuration files of the simulation to facilitate modifications for possible future running conditions. Luminosity levelling is taken into account through the crossing angle and beta parameter, but the function in appendix B would need to be updated if the levelling strategy changes to a crab cavity or beam offset for example. The presented PV generator has become the default tool for LHCb upgrade studies, and has also been used for the timing studies by the ECAL group [89] discussed in section 1.4.3.

3.1.2 Particle propagation and photon dispersion

This section describes the particle propagation through the LHCb detector, with the aim to quantify the spread in the photon time of arrival at the detector plane due to chromatic dispersion in the RICH detector. All particles and photons generated during a bunch crossing and subsequent interactions are propagated using the Geant4 simulation toolkit [90] interfaced to the GAUSS software framework [91]. GAUSS v51r0 was used, which contains a detailed model of all UpgradeIa sub-detector geometries and materials. The track speed, energy loss and interactions with the sub-detectors and magnetic field are simulated. In this study, the particle identification performance is explored for both RICH detectors using the ‘long tracks’ that are reconstructed in all LHCb tracking stations. In addition to the reconstructed track parameters, the GAUSS framework stores a full Monte-Carlo (MC) history of the event, which can be accessed from the RICH reconstruction algorithms.

When a charged particle passes through the RICH detector, Cherenkov photons are produced along the particle track segment in the gas radiator, provided the Cherenkov threshold is exceeded. The detected photons were emitted over an energy range of 1.75 to 7.0 eV, reflected by the mirror system and focused onto the photon detector plane. The photons therefore travel up to 5.4 m through the RICH radiator and 6 to 8 mm through the gas enclosure quartz window. The photon dispersion along these paths produces a spread in the photon time of arrival at the RICH detector plane, which is simulated in this study.

In order to verify the photon time-of-flight calculations performed by the Geant4 toolkit, the GAUSS ‘step-action’ procedure was enabled. This configurable step-action prints requested parameters at each particle propagation step in the Geant4 simulation. Here, it was found that the speed of the photon at a given step was used to compute the time-of-flight for the

next step. This method can result in significant errors when the photon travels from the gas enclosure quartz window into the nitrogen gas in front of the photon detector plane for example. The solution in GAUSS and Geant4 is under investigation within the RICH group. In this study, a calculation of the group velocity at each step in the C_4F_{10} and CF_4 radiators, the gas enclosure quartz window and the nitrogen gas was introduced. This calculation will be discussed in the remainder of this section.

The group velocity for the propagation of photons of energy ϵ in eV is given by:

$$v_g = \frac{c}{n_g} \quad (3.6)$$

Here, c is the speed of light in vacuum and n_g is the group index:

$$n_g = n(\epsilon) + \epsilon \cdot \frac{dn(\epsilon)}{d\epsilon} \quad (3.7)$$

The refractive index n of the gases is estimated using the Lorentz-Lorenz equation [92]:

$$n = \sqrt{\frac{1+2r}{1-r}} \quad (3.8)$$

Here, the quantity r is defined in terms of the Sellmeier fit parameters $F_{1,2}$ and $E_{1,2}$, the molar density ρ/M , the Bohr radius a_0 , and Avogadro's number N_A :

$$r = \left(\frac{F_1}{E_1^2 - \epsilon^2} + \frac{F_2}{E_2^2 - \epsilon^2} \right) \cdot \frac{4\pi a_0^3}{3} \cdot N_A \cdot \frac{\rho}{M} \quad (3.9)$$

The coefficients for the different gas types are tabulated in table 3.1. These values were originally obtained from a fit to data from the DELPHI experiment and were tuned to agree with data taken during the LHCb RICH detector beam tests [93]. The molar density of the gases is stated at 20°C temperature and atmospheric pressure.

The refractive index of the quartz is described by the Sellmeier approximation:

$$n = \sqrt{1 + \frac{B_1}{1 - D_1^2 \epsilon^2} + \frac{B_2^2}{1 - D_2^2 \epsilon^2}} \quad (3.10)$$

Here, the coefficients $B_1 = 0.6962$, $B_2 = 0.4079$, $D_1 = 0.0552 \text{ eV}^{-1}$ and $D_2 = 0.0938 \text{ eV}^{-1}$ are obtained from [94] and converted into units of energy.

Gas type	$F_1 [\text{eV}^2]$	$E_1 [\text{eV}]$	$F_2 [\text{eV}^2]$	$E_2 [\text{eV}]$	$\rho/M [\times 10^{-6} \text{ mole/cm}^3]$
C_4F_{10}	12305.3	18.938	~ 0	~ 0	69.86
CF_4	7626.16	20.003	~ 0	~ 0	41.59
N_2	921.28	13.414	3569.60	23.215	41.58

Table 3.1: The Sellmeier fit parameters and molar density for the RICH detector gases that contribute to photon dispersion.

Figure 3.5 shows the group indices of C_4F_{10} and quartz. The variation in this group index gives rise to the chromatic error, which is smaller for low-energy photons. This effect is quantified for the gas radiators and quartz window in table 3.2, which shows the chromatic error, in picoseconds, for typical photon path lengths in the RICH 1 and RICH 2 detectors. These numbers were calculated using the parameters in table 3.1. Photons emitted at the limits of the simulated Cherenkov radiation energy spectrum, at 1.75 and 7.0 eV, develop a total time-of-flight difference of 16.5 and 18.3 ps by the time they reach the RICH 1 and RICH 2 photon detector respectively. The QE of the Upgrade Ia MAPMTs, which was shown in figure 2.4, restricts the detected energy range to be between 2.1 and 6.2 eV, which reduces the maximum total time-of-flight difference to approximately 11 ps. In the future, this contribution could be reduced to only a few picoseconds by using a green-shifted photon detector with a high QE in the low photon-energy range. In addition to the quartz and gas radiators in table 3.2, the photons travel through 39 cm of N_2 gas between the RICH 1 gas window and the photon detector, which introduces a time-of-flight difference of less than 0.3 ps. The effect of the 5 % admixture of CO_2 in the RICH 2 gas volume is expected to be minor [95].

The smaller chromatic error for a green-shifted photon detector comes at the cost of reduced photon yield. The number of Cherenkov photons dN generated with energy $d\epsilon$ over a track path length dx in the radiator is given by the Frank-Tamm equation [96]:

$$\frac{d^2 N}{d\epsilon dx} \propto \left(1 - \frac{1}{n(\epsilon)^2 \beta^2}\right) \quad (3.11)$$

The small refractive index at low photon energies reduces this photon yield.




Figure 3.5: Group index of the C_4F_{10} gas radiator and the gas enclosure quartz window as a function of the photon energy.

Energy range [eV] (nm)	RICH 1		RICH 2	
	Quartz (6 mm)	C_4F_{10} (2.4 m)	Quartz (8 mm)	CF_4 (5.4 m)
1.75-7.0 (180-710 nm)	11.1	5.4	14.8	3.5
2.1-6.2 (200-600 nm)	6.9	3.8	9.1	2.5
1.55-3.1 (400-800 nm)	1.0	0.7	1.3	0.5

Table 3.2: Maximum photon time-of-flight difference in picoseconds due to chromatic dispersion, for typical path lengths in the RICH detectors and different detected photon energy (wavelength) ranges.

3.1.3 The RICH detector SmartID

The photon detector description is implemented in the RICH digitisation software package named BOOLE [97]. This is the final stage of the simulation before the RICH reconstruction algorithms are started, which are identical for the simulation and the experiment. The RICH SmartID provides the interface between the photon detector and the RICH reconstruction.

BOOLE provides the flexibility to implement different detector technologies. In this study, BOOLE v32r2 was used, which contains the UpgradeIa MAPMTs. Therefore, all detector properties, except the hit time resolution, are for the MAPMTs. By design, the simulation does not include electronic noise, cross talk or time spread introduced by the MAPMT or readout electronics. Although there are ongoing developments to include these effects in order to study the consequences of SIN and time gating on the particle ID performance during LHC Run 3, the absence of these detector-specific properties is preferred for the UpgradeII hit time studies in this chapter as it allows the conclusions to be independent of the final technology choice. A range of technologies can subsequently be explored, by smearing the hit times in the RICH reconstruction algorithms by the photon detector time resolution, which will be discussed in section 3.4.

The photon detector hits are stored in the RICH SmartID, as shown in figure 3.6. The UpgradeIa SmartID contains 32 bits to specify for example which RICH detector, photon detector panel, column, row and MAPMT pixel was hit in the event. In the experiment, any event information that is required in the RICH reconstruction algorithms needs to be stored in the SmartID.

In order to store the photon detector hit time information, the RICH SmartID was extended from 32 to 64 bits. Out of the 32 added bits, 20 bits are allocated to the hit time coordinate, with the remaining bits available for other studies. A time-to-digital converter (TDC) is emulated in software and the floating-point hit time generated in GAUSS is encoded as an




Figure 3.6: The RICH SmartID in the LHCb simulation framework. The SmartID and reconstruction algorithms (green) are identical for simulation (red) and experiment (yellow). The photon detector hit time is directly written into the SmartID from GAUSS, circumventing the detector-specific BOOLE digitisation stage.

integer and written into the 20-bit field. A range from -50 to 150 ns was chosen, which is sufficiently long to accommodate spillover as discussed in section 3.1.4. The resulting time resolution in the SmartID is $200 \text{ ns}/2^{20} = 0.2 \text{ ps}$, which is smaller than the time spread from photon dispersion.

As the number of bits has an important impact on the bandwidth requirements of the electronic readout, fewer bits are envisaged to store the hit time in the future upgrade detector. As seen in section 3.1.1, the standard deviation of the PV time distribution is 214 ps , meaning that a time gate at the front-end electronics of 1 to 2 ns is sufficient to collect all signal from the bunch crossing. For a TDC with 30 ps resolution, this would correspond to 5 or 6 bits in the readout electronics. The grouping of channels with the same timestamp may help to further reduce the associated increase in data rate.

3.1.4 Spillover

Spillover is defined as the activity from other bunch crossings that increases the overall background level in the bunch crossing of interest. The sources of spillover can be categorised as (a) originating from track or photon activity or (b) originating from photon detector activity. In category (a), a photoelectron is generated in the bunch crossing period of interest, but the origin of the photon lies outside this bunch crossing. Examples are out-of-time tracks generating photons, delayed scintillation and photons undergoing multiple reflections. In category (b), a photon detector pixel registers a signal above threshold due to a photoelectron generated outside the bunch crossing of interest. This process depends on the time response of the photon sensor and readout electronics. An example is the SIN




Figure 3.7: Photon detector hit time distribution at the BOOLE digitisation stage, showing five simulated bunch crossings (labelled 1-5) with signal (S), background (B) and reflected signal (R) peaks for the RICH 1 and RICH 2 detectors.

described in section 2.2.6. In general, the amount of spillover in category (b) is a compromise with the efficiency of capturing the photon sensor signal in the readout electronics. In order to keep this study as technology-independent as possible, only spillover due to (a) photon activity is simulated.

Spillover is not restricted to previous bunch crossings, but can also be caused by early beam activity from the next bunch crossings. Therefore, five bunch crossings are simulated, as shown in the BOOLE digitisation hit time distribution in figure 3.7. This distribution contains hits from the RICH 1 and the RICH 2 detectors. The signal (S) and background (B) peaks are labelled for bunch crossings 1 to 5. The signal typically arrives 50 ns later at the RICH 2 detector, which is further from the interaction point. The hit time distributions are discussed in more detail in section 3.2.2. The spillover hits in the region between 10^3 and 10^4 on the vertical scale in figure 3.7 gradually increase and peak at the third bunch crossing. This bunch crossing is the one of interest in this study, as it is most representative of a continuous stream of bunches. The contribution from spillover is around three orders of magnitude smaller than the signal amplitude. The spillover hits are also relatively spread out in time, without peaks, as seen in the tail after the S5 peak for RICH 2. This background can be effectively eliminated using a time gate.

3.1.5 Pile-up

The LHCb simulation generates conditions equivalent to a luminosity of $2 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$, corresponding to the LHC Run 3. Following the LHC programme outlined in section 1.1.2, this luminosity will increase by a factor of 7.5 in the high-luminosity LHC era. Although these conditions are beyond the capabilities of the current simulation, an approximate scheme was introduced to simulate the pile-up of events, as illustrated in figure 3.8. At the photon detector hit clustering stage of the RICH reconstruction, the hits from seven subsequent bunch crossings are overlaid onto one event. The tracking information is only stored for the event of interest, shown in the solid frame in figure 3.8. The potentially useful tracking information from the other six events, shown in the dashed frames, is lost




Figure 3.8: In order to estimate the effect of pile-up in the Upgrade Ia framework, six events (dashed frames) are superimposed as background (marked in red) onto the event of interest (solid frame).

and these photons are assigned as background since they cannot be associated with their tracks.

While the proposed pile-up scheme gives a realistic estimate of the hit pattern at the RICH photon detector plane, the simulated resolutions of each LHCb sub-detector still correspond to the Upgrade Ia detector. This, combined with the incomplete tracking information, means that the proposed simulation of pile-up conditions is a worst-case scenario. For a more accurate simulation of the Upgrade II conditions, the geometry descriptions and reconstruction software of the sub-detectors would need to be revised, which will take place in a process of gradual refinement over the coming years.

The background introduced by pile-up arrives around the same time as the signal, and is more challenging to eliminate than the spillover hits discussed in the previous section. In fact, a sub-nanosecond time gate is always necessary to mitigate pile-up, and this gate reduces the spillover hits to a negligible level. Therefore, spillover was switched off when simulating pile-up, in order to increase the simulation speed.

3.2 RICH Reconstruction in space and time

The RICH reconstruction algorithm uses the reconstructed charged particle trajectories to predict where photons are incident on the photon detector plane for a given choice of particle mass hypothesis. Using this prediction, the likelihood of the observed photon detector hits is calculated and the set of mass hypotheses that maximises the overall likelihood is searched for. The reconstruction is explained in section 3.2.1. The photon detector hit time distributions are discussed in section 3.2.2. In order to fully exploit the high intrinsic RICH detector time resolution, an accurate prediction of the Cherenkov photon hit times is computed using reconstructed parameters in section 3.2.4, which requires the PV time discussed in section 3.2.3. The extension of the RICH reconstruction algorithm into the time domain is presented in section 3.2.5. The particle ID curves, which are used to compare the simulation results, are introduced in section 3.2.6.

3.2.1 Reconstruction algorithm in space

The objective of the reconstruction algorithm is to find the most likely particle types for the set of tracks in the RICH detector at each bunch crossing, in regions of high track density with overlapping ring images. The reconstruction algorithm computes the Cherenkov emission angle for the photons associated with a track, and combines this angle with the knowledge of the track momentum to obtain the particle mass. This procedure can be broadly divided into three stages: pixel and track processing, photon reconstruction and particle identification.

In the first stage, the photon detector hits stored in the RICH SmartID are clustered. The coordinates of the clusters are calculated in the global and local RICH coordinate systems, assuming the position of the hit was at the centre of the photon detector pixel.

From the tracking information, the track segments within the RICH radiator volumes are reconstructed. In order to reduce the number of combinations considered in subsequent reconstruction stages, clusters are masked if these are outside the spatial acceptance associated with the segment and the Cherenkov emission angle for all particle types. This acceptance is computed by ray-tracing a fixed number of Cherenkov photons at the expected Cherenkov angle range to the detector plane [98].

In the second stage, the segments and detected hit clusters are combined into candidate Cherenkov photon objects. The photon objects contain the Cherenkov emission angles for each cluster, assuming that the cluster originated from each segment in the event in turn. This step is essentially the inverse of ray tracing, and is computationally intensive due to the large number of cluster and segment combinations. A full analytical solution of the RICH optics is used, which reconstructs the photon trajectory through the RICH optical system, taking into account the mirror and photon detector alignment [99, 100]. The mid-point of the segment is taken as the emission point, as the exact emission point is unknown and can be anywhere along the particle trajectory through the radiator.

In the third and final stage, the reconstruction algorithm maximises the overall likelihood associated with the assignment of particle types to the segments in the event. The algorithm identifies pion, kaon, electron, muon, proton and deuteron particle types. The reconstruction supposes that there are N sources of hits in the photon detector, and sets the event hypothesis $\mathbf{h} = (h_1, \dots, h_N)$, where h_j is either the mass-hypothesis for track j or a parameter associated with background j . As shown in [99], the event likelihood $\mathcal{L}(\mathbf{h})$ is the product, over all hit pixels, of the Poisson probabilities to obtain a hit in the pixel given the expected signal $\nu_i(\mathbf{h})$. This expected signal in pixel i is given by:

$$\nu_i(\mathbf{h}) = \sum_{j=1}^N a_{ij}(h_j) \quad (3.12)$$

The expected number of detected photoelectrons from track source j with hypothesis h_j in pixel i is:

$$a_{ij}(h_j) = \epsilon_i \cdot \lambda_j(h_j) \iint_{Pixel\ i} f_{h_j}(\theta, \phi) d\theta d\phi \quad (3.13)$$

Here, ϵ_i is the photon detector pixel efficiency and λ the expected total number of Cherenkov photons emitted by a track. The integral gives the fraction of these photons arriving at pixel i . The scaling of the background parameter, which is associated with the event hypothesis h_j where j is background, is discussed in section 3.2.6.

The probability density function of the emitted signal photons is taken to be a Gaussian in the angle θ , with expected Cherenkov angle $\theta_c(h_j)$ and Cherenkov angle resolution $\sigma(\theta)$, and uniform in the angle ϕ :

$$f_{h_j}(\theta, \phi) = \frac{1}{(2\pi)^{3/2}\sigma(\theta)} \exp \left[-\frac{1}{2} \left(\frac{\theta - \theta_c(h_j)}{\sigma(\theta)} \right)^2 \right] \quad (3.14)$$

The initial hypothesis assumes all tracks are pions, which are the most abundant particle type of tracks in the LHCb experiment. The change in the log-likelihood, $\Delta \ln \mathcal{L}$, is then computed by changing the particle type for one track at a time, leaving all other track hypotheses unchanged. The algorithm then identifies which change in mass hypothesis, amongst all tracks, gives the largest increase in the overall event likelihood. The mass hypothesis for that track is updated to the preferred type. This procedure is repeated until all tracks are set to the optimal mass hypotheses, and no further improvement in the likelihood is obtained, as defined by the $\Delta \ln \mathcal{L}$ threshold required to change a mass hypothesis. The extension of equation 3.14 into the time domain is discussed in section 3.2.5.

3.2.2 Photodetector hit time distribution

The RICH 1 hit time distribution is shown in figure 3.9. The particles from a bunch crossing around $t = 0$ take approximately 7 ns to reach the RICH 1 detector. Hits arriving at $t < 7$ ns are therefore exclusively from spillover. The background peak between 7 and 10 ns is due to tracks and photons that reach the photon detector plane directly, as shown by the blue dotted path in figure 3.11. The Cherenkov photon signal arrives at around 13 ns after the bunch crossing due to the additional path length through the RICH optical system. The RICH radiator gas is separated from the photon detectors by a 6 mm-thick quartz window positioned at 39 cm from the detector plane. Photons that are reflected between the MAPMT window and the gas enclosure quartz window therefore travel an additional 78 cm resulting in a 2.6 ns delay. This is the origin of the reflected signal peak at around 16 ns.

The RICH 2 hit time distribution is shown in figure 3.10. The Cherenkov signal peak arrives around 53 ns after the bunch crossing due to the large separation between the downstream RICH 2 detector and the beam interaction point. The typical Cherenkov photon path length in the mirror system is 5.4 m in the RICH 2 detector compared to 2.4 m in the RICH 1 detector. The signal and background peaks are therefore more spread out, and the background peak from the next bunch crossing (B2 in figure 3.10) arrives within the 25 ns window around the signal peak. The spacing between the gas enclosure quartz window and the RICH 2 photon detectors is only 40 mm. Additionally, the photon angle of incidence at the RICH 2 photon detector plane is typically closer to the normal, due to the longer focal length of the spherical mirror. No reflected photon signal peak was therefore observed.

A more detailed understanding of the sub-nanosecond structure of the hit time distributions is critical for the application of time in the future RICH detector, and will be the focus of sections 3.2.3 and 3.2.4.




Figure 3.9: RICH 1 photon detector hit time distribution showing background (B), signal (S) and reflected signal (R) peaks.




Figure 3.10: RICH 2 hit time distribution containing signal (S1) and background from two bunch crossings (B1 and B2).




Figure 3.11: Geometry of the RICH detector showing in blue the path for the tracks and photons resulting in background peak B in figure 3.9. The typical Cherenkov signal path is shown in red.




Figure 3.12: PV and track origin time from MC information.




Figure 3.13: The signal peak of the RICH 1 hit time distribution after subtraction of the PV time using equation 3.15.




Figure 3.14: The same plot as in figure 3.13, for the signal peak of the RICH 2 detector.

3.2.3 Primary vertex time

The width of the signal peaks in the hit time distributions in figures 3.9 and 3.10 are dominated by the PV spread. The time gate of a few nanoseconds at the detector front-end, as discussed in chapter 2, therefore cannot be reduced, as this would result in a loss of detected PVs. In the RICH reconstruction algorithm however, the PV time can be calculated and subtracted, opening up the possibility of resolving the pile-up time structure with sub-nanosecond time information.

The PV time in this study is obtained from MC information without error. It is envisaged that a high-resolution PV timestamp will be available as an input to the RICH reconstruction in the future LHCb experiment, from the Upgrade II VELO for example. Otherwise, analogous to the TORCH detector, the RICH reconstruction algorithm could generate the PV time from the data before the particle identification is performed. Further studies will be required to evaluate the best method to obtain the PV time in the RICH detector. For example, it may be possible to use the early arrival of background directly from the PV, through the blue dashed path in figure 3.11, as a reference time for the readout of the signal at the detector front-end.

The track origin time and the PV time are shown in figure 3.12 for tracks passing through the RICH detectors. In most events, the track origin corresponds to the PV. A small number of secondary tracks is generated in the first few nanoseconds after the bunch crossing, as the tracks propagate through the VELO. However, these secondary tracks ultimately also originate from the PV, and the time-of-flight difference associated with the different particle mass between the primary and secondary vertex is negligible. The PV time is therefore used in this study, as its resolution is expected to be $\sqrt{N_{Tracks}}$ higher than the track origin timestamp, where N_{Tracks} is the number of tracks emerging from the PV.

The PV time spread can be removed in order to demonstrate its contribution to the RICH 1 and RICH 2 signal peak widths. This is done by subtracting the time t_0 from the hit time distributions in figures 3.9 and 3.10, where t_0 is given by:

$$t_0 = t_{pv} - \frac{z_{pv}}{c} \quad (3.15)$$

Here, the z_{pv} coordinate is taken from MC without error. The time dependence on the x and y -coordinates of the PV is negligible, as $\sigma_{x,y} = 38 \mu\text{m}$ compared to $\sigma_z = 57 \text{ mm}$ and the tracks are boosted in the forward z -direction in the LHCb experiment. The difference between the detected hit times and the prediction t_0 is plotted in figures 3.13 and 3.14 for signal selected using MC information. A small background is present due to Cherenkov photons that are scattered or spilled over from previous bunch crossings. The width of the peak is approximately 30 ps for the RICH 1 detector and 300 ps for the RICH 2 detector, which is a significant reduction from the approximately 2 ns when the PV time spread was included.

3.2.4 Prediction of the Cherenkov photon time of arrival

In this section, an accurate prediction of the Cherenkov photon signal time of arrival is formulated, which allows the intrinsic time resolution of the RICH detector to be determined. Additionally, this predicted hit time forms the basis for the application of a sub-nanosecond time gate or the scaling of the signal amplitude in the $\Delta \ln \mathcal{L}$, as discussed in section 3.2.5.

The width of the hit time distributions in the absence of PV spread, which was shown in figures 3.13 and 3.14, is dominated by variations in the track and photon paths through the RICH detector. The following photon hit time prediction t_{pred} has therefore been implemented in the RICH reconstruction algorithms:

$$t_{pred} = t_{pv} + \frac{d_{pv,A}}{c} \sqrt{1 + \left(\frac{mc}{p}\right)^2} + \frac{d_{A,E}}{c} n \cos \theta_c + [d_{E,M1} + d_{M1,M2} + d_{M2,HIT}] \frac{n}{c} \quad (3.16)$$

Here, referring also to figure 3.15, $d_{pv,A}$ is the distance from the PV to the RICH entry point, $d_{A,E}$ from the RICH entry point to the photon emission point, $d_{E,M1}$ from the photon emission point to the spherical mirror reflection point, $d_{M1,M2}$ from the spherical to the flat mirror reflection point and $d_{M2,hit}$ from the flat mirror to the photon detector hit. These are calculated using the reconstructed points stored in the photon object, along with the reconstructed Cherenkov angle θ_c . The reconstructed track momentum p is provided by the LHCb tracking system and the track mass m is the mass hypothesis in the RICH likelihood maximisation algorithm. The refractive index n for the gas radiator is taken at the average photon energy of 4.25 eV and 4.40 eV for the RICH 1 and RICH 2 detector respectively. The photon emission point is taken at the mid-point of the track segment.




Figure 3.15: A track emerging from the primary vertex is deflected by the LHCb magnet and intercepts the RICH 2 detector, producing a Cherenkov photon shown in red. The track segment and points E, M1 and M2 are reconstructed parameters.

A track curvature correction for the RICH 2 detector

$d_{pv,A}$ in equation 3.16 is assumed to be a straight line. While this is a good approximation for RICH 1, the RICH 2 detector is positioned downstream of the LHCb magnet, which was shown in figure 1.6, and the tracks reaching the RICH 2 detector have followed a curved trajectory through the magnetic field. A correction to the RICH 2 track path length is therefore applied, based on a geometrical argument.

To first order, the path length of the curved trajectory is approximated as the path length associated with a ‘kick’ from the centre of the LHCb magnet, as shown in figure 3.16. In this figure, M is the interception point between the reconstructed track vector at the RICH 2 detector and the xy -plane at the coordinate z_M of the centre of the magnet. The original path length, $L = L_1 + L_2$, where L_1 and L_2 are before and after the LHCb magnet respectively, becomes L' :

$$L' = \frac{L_1}{\cos \alpha} + \frac{L_2}{\cos \gamma} \quad (3.17)$$

Here, the angle γ , which is in the (horizontal) bending plane of the dipole magnet, is computed in the reconstruction algorithm as the difference between the gradient of the reconstructed track vector and the ratio of the x and z -coordinates at the point of entry R of the track in the RICH 2 detector:

$$\gamma = \left. \frac{dx}{dz} \right|_R - \left. \frac{x}{z} \right|_R \quad (3.18)$$

As shown in figure 3.17, the track curvature typically causes a time-of-flight difference of up to 50 ps at low momentum, corresponding to a flight path difference of around 1.5 cm, which is small compared to the distance of about 15 m between the interaction point and the RICH 2 detector. The small angle approximations $(\cos \alpha)^{-1} \approx (1 + \alpha^2/2)$, $(\cos \gamma)^{-1} \approx (1 + \gamma^2/2)$ and $\alpha \approx \gamma \cdot L_2/L_1$ can therefore be used to simplify equation 3.17:

$$\frac{L'}{L} = 1 + \frac{\gamma^2}{2} \cdot \left(\frac{L}{L_1} - 1 \right) \quad (3.19)$$

Using similar triangles, equation 3.19 can be written in terms of the RICH entry window




Figure 3.16: Geometry for the first-order correction of the track path length for particles reaching the RICH 2 detector, which had a curved trajectory in the LHCb magnetic field. The small angles α and γ are drawn large for clarity.




Figure 3.17: Difference between the detected and predicted Cherenkov photon hit times as a function of track momentum for the RICH2 detector.




Figure 3.18: Identical plot to figure 3.17, after a correction is applied to the track path length for the curvature in the LHCb magnetic field.

coordinate z_R and the LHCb magnet centre coordinate z_M . Substituting equation 3.18, the path length correction becomes:

$$\therefore \frac{L'}{L} = 1 + \frac{s}{2} \left(\frac{z_R}{z_M} - 1 \right) \left(\frac{dx}{dz} \Big|_R - \frac{x}{z} \Big|_R \right)^2 \quad (3.20)$$

Here, a scale factor s has been introduced to optimise the correction, which is an over-estimate of the true curved track path length. The prediction of the photon hit time of arrival in figure 3.18 includes the path length correction with a scale factor of $s=0.575$, and is accurate across the full momentum range. This approach was inspired by a similar correction applied in the LHCb Outer Tracker [101].

The intrinsic resolution of the RICH detectors

The final result using the prediction in equation 3.16 is shown in figure 3.19. This is a key plot in this study, as it shows that the time of arrival of the RICH Cherenkov signal can be predicted to within 10 ps using the PV time and reconstructed parameters in the RICH detector. This high predictability of the hit time can be exploited to improve the particle ID in high-occupancy environments, as presented in section 3.3.

The positive tail of the distribution in figure 3.19 arises from the photon dispersion discussed in section 3.1.2. This is an irreducible error because the photon energy is not measured in the experiment. However, a photon sensor with enhanced QE in the green optical spectrum can reduce this dispersion effect. The RICH2 distribution shows a small offset from zero, suggesting a systematic underestimate of the signal time of arrival. This could for example have arisen from the scale factor in the curvature correction for the RICH2 detector, and in practice it can be corrected by adding a constant term to the signal hit time prediction.




Figure 3.19: The difference between the detected hits and the predicted photon time of arrival, which demonstrates that the intrinsic resolution of the RICH detectors is less than 10 ps.




Figure 3.20: The difference between the RICH2 detector hits and the predicted photon time of arrival, for all MC particle types, when the track mass in the prediction is the default assumed pion mass.




Figure 3.21: Same plot as figure 3.20, for the MC electron type. The track mass in the prediction is the (default assumed) pion mass for the red curve and the (correct) electron mass for the black curve.

The hit-time prediction in equation 3.16 is for a particular mass hypothesis, which is updated for each change of the particle type in the $\Delta \ln \mathcal{L}$ algorithm. This mass assignment results in a minor improvement in the predicted particle time-of-flight with respect to the default assumed pion mass. Figure 3.20 shows the difference between the hit time and the predicted time based on the pion mass, for the MC particle types in the RICH 2 detector. The incorrect mass assignment for the electrons causes a spread of approximately 10 ps, as shown in figure 3.21, while the effect is negligible for the other particles. As explained in section 3.2.6, any cut based on figures 3.20 or 3.21 has to be conservative in order to avoid removing all photons associated with the same track.

The emission point error in the time domain

For a straight particle track in the RICH detector, Cherenkov radiation is produced as a plane wave propagating at the Cherenkov angle away from the track. For the ideal optical system, free of aberrations, this wavefront would be focused onto the RICH detector plane. In practice however, the radiation emitted from different parts of the track takes different paths through the optical system and is affected differently by optical aberrations, giving rise to the spatial ‘emission point error’ discussed in section 1.4.1. In the time domain, the emission point error is negligible, as shown in figure 3.23, which plots the difference in the predicted signal hit times in equation 3.16 when the emission point is assumed to be at the extreme positions EP1 and EP2 on the RICH segment in figure 3.22. This emission point error strongly peaks around $\Delta t = 0$, with a range of approximately 2 ps for the RICH 1 detector and 5 ps for the physically larger RICH 2 detector.




Figure 3.22: Geometry of the RICH detector, showing the reconstructed Cherenkov photon paths from emission points EP1 and EP2 close to the entry and exit points of the gas radiator volume.




Figure 3.23: The Cherenkov photon emission point error in time, showing the difference in the predicted photon time of arrival for emission points 1 and 2 in figure 3.22.

3.2.5 Time in the RICH reconstruction software

Two implementations of the hit time information in the RICH reconstruction software are discussed in this section: a time gate around the predicted signal time of arrival and the extension of the likelihood maximisation algorithm into the time domain.

Time gating

The use of a 3.125 ns time gate in the Upgrade Ia PDMDB was discussed in chapter 2. The effect of this hardware gate will be simulated by defining a fixed time interval synchronous to the LHC bunch clock. All photon detector hits arriving outside this gate are masked in the reconstruction. Although this hardware gate at the front-end is an effective method to reduce the data rate and to save CPU resources in the reconstruction, the minimum width of the hardware gate is limited by the PV time distribution. An additional, more restrictive time gate can be applied in software, using reconstructed information on the PV time and the track and photon path lengths.

A software time gate is introduced in three stages of the reconstruction in order to reduce the CPU resources. While the hardware time gate is defined with respect to the bunch crossing, the software gates are centred on the hit time prediction in section 3.2.4. The difference between the three types of software gate is that the hit time prediction gets more accurate, allowing a smaller width of the gate.

The first time gate (‘type 1’) is applied prior to the construction of the photon object, which was discussed in section 3.2.1. At this stage, the association of photons to tracks has not yet taken place, and the reconstructed path lengths are therefore not available. Nevertheless, the PV time can be used as a first-order prediction of the signal hit time. The time distributions with the subtracted PV time were shown in figures 3.13 and 3.14. Based on these distributions, a conservative choice for the type 1 gate would be 70 ps for the RICH1 detector and 500 ps for the RICH2 detector. This background reduction is applied prior to the large combinatorics of the track-to-photon association for the photon objects, and therefore significantly saves CPU resources.

The second time gate (‘type 2’) is introduced at the photon object reconstruction stage, in the final stage before the object is stored. The detailed photon hit time prediction in equation 3.16 can therefore be used, albeit with the mass default assumed to the pion mass. As shown by the black curve for all particle types in figure 3.20, this allows a time gate of approximately 20 to 30 ps.

The third time gate (‘type 3’) is implemented in the $\Delta \ln \mathcal{L}$ algorithm. The track path length and the hit time prediction excluding the second term for the track time-of-flight in equation 3.16 are now stored separately in the photon object. The track time-of-flight is then calculated in the $\Delta \ln \mathcal{L}$ algorithm using the mass hypothesis. The time gate can therefore be reduced to the approximately 10 ps width shown in figure 3.19. Although this

time gate gives the best reconstruction performance, its implementation requires more CPU resources. The type 2 gate is therefore preferred unless a time resolution on the order of 10 ps can be achieved in the future detector.

The extension of the likelihood maximisation algorithm into the time domain

As an alternative implementation of time in the reconstruction, the likelihood maximisation algorithm, which was presented in section 3.2.1, can be extended into the time domain. The principle is to add the time coordinate to the probability density function of signal photons in equation 3.14, using the hit time prediction from equation 3.16. The association of detector hits to particle tracks can therefore be performed in space and time. Furthermore, the overall time resolution of the system can be taken into account, similar to the Cherenkov angle resolution in the spatial domain.

There are two subtle differences in the role of the space and time coordinates in the RICH reconstruction. Firstly, time can only indirectly contribute to the particle identification through background rejection and the association of photons to tracks. This is due to the negligible time-of-flight difference between different particle types for the relatively high-momentum tracks in the RICH detector, as was shown in figure 3.20. At lower momentum, the time-of-flight difference would no longer be negligible, which is the principle for the particle ID by the TORCH detector described in section 1.4.2. In the RICH detector, the spatial measurement of the Cherenkov angle will remain the only way to distinguish particle types. Secondly, the uncertainty in the time domain introduced by the sampling of the photon signals in the readout electronics cannot be neglected. In the spatial domain, the pixel size is small compared to the Cherenkov ring. The pixel error therefore contributes in the same way as the chromatic error and the emission point error to the Cherenkov angle resolution for the Gaussian PDF in equation 3.14. In the time domain, however, the width of the sampling bins in the electronic readout system is not necessarily smaller than the sensor time resolution, and the signal PDF in the reconstruction algorithm cannot be assumed to be Gaussian distributed.




Figure 3.24: Method of scaling the signal amplitude in the $\Delta \ln \mathcal{L}$ calculation using time. The scale factor is the red shaded area under the curve, and takes into account the predicted signal time of arrival (t_{pred}) and the time resolution of the photon sensor (σ_{sensor}) and readout electronics (Δt_{bin}).

The use of time in the Cherenkov photon PDF is illustrated in figure 3.24. The time spread from the photon sensor, shown in blue, is assumed to be Gaussian distributed around the predicted photon time of arrival t_{pred} with a width σ_{sensor} , which is assumed to be greater than the intrinsic time resolution of the RICH detector of less than 10 ps. This spread is introduced at the RICH reconstruction algorithms, and would need to be measured for the photon sensors of the future experiment. The spatial PDF $f_{h_j}(\theta, \phi)$ in equation 3.14 is then scaled by the integral over the width of the electronic readout bin Δt_{bin} around the detected hit time t_{hit} , as shown in red in figure 3.24. The signal PDF $g_{h_j}(\theta, \phi, t)$ in space and time is therefore given by:

$$g_{h_j}(\theta, \phi, t) = f_{h_j}(\theta, \phi) \cdot \left(\frac{1}{(2\pi)^{3/2} \sigma_{sensor}} \int_{t_{hit}-\Delta t_{bin}/2}^{t_{hit}+\Delta t_{bin}/2} \exp \left[-\frac{1}{2} \left(\frac{t - t_{pred}(h_j)}{\sigma_{sensor}} \right)^2 \right] dt \right) \quad (3.21)$$

If the future readout electronics would be much faster than the sensor, $\Delta t_{bin} \ll \sigma_{sensor}$, then $g_{h_j}(\theta, \phi, t)$ would become approximately Gaussian in space and time. Alternatively, if $\Delta t_{bin} \gg \sigma_{sensor}$, then all signal would be included in a single readout bin, and the time contribution would integrate to approximately 1, leaving only the spatial signal PDF. This would be equivalent to a time gate of the width of the electronic readout bin. The future implementation therefore depends on the time resolution of the available photon sensor technologies and readout electronics. The particle ID performance using this signal PDF will be evaluated in section 3.3.3.

3.2.6 The particle identification curve

The particle ID curves such as those in figure 3.25 provide a standardised and highly sensitive probe of the performance of the RICH detector, and will be used to compare the different methods of using the hit time information in section 3.3.

As was discussed in section 3.2.1, the reconstruction algorithm first assumes that all tracks are pions, and changes the mass hypothesis to a kaon if the associated $\Delta \ln \mathcal{L}(\pi - K)$ is above threshold. For each particle ID curve in figure 3.25, the kaon ID efficiency (kaon identified as a kaon) is plotted against the pion misID efficiency (pion misidentified as a kaon) for a range of $\Delta \ln \mathcal{L}(K - \pi)$ thresholds. The efficiencies are obtained by comparing the outcome of the reconstruction to MC information. At low $\Delta \ln \mathcal{L}(K - \pi)$ thresholds, more tracks will have their mass hypothesis changed from the pion initial condition to a kaon, resulting in a high kaon ID efficiency but also a high pion misID efficiency, and vice-versa at high $\Delta \ln \mathcal{L}(K - \pi)$ thresholds. Each curve shows the average over all track momenta for ten thousand simulated events.

Figure 3.25 shows the particle ID curves for five different background scaling parameters in the absence of a time gate. This parameter scales the expected contribution from background hits for each pixel, as defined by the event hypothesis h_j for background j in section 3.2.1. The background scaling parameter needs to be optimised each time the amount




Figure 3.25: Particle ID curves showing the pion misID efficiency against the kaon ID efficiency for a range of background scaling parameters in the RICH reconstruction. The black dot indicates the ideal reconstruction performance. Each point on the plot represents a different delta-log-likelihood threshold ($\Delta \ln \mathcal{L}(K - \pi)$ or DLL).




Figure 3.26: Particle ID curves for 1.0 and 2.0 ns time gates compared to no time gate. Since the PV time spread has a standard deviation of 214 ps, the 1.0 ns gate rejects a small number of tracks, resulting in a degradation of the particle ID performance.

of background changes in the RICH photon detector, for example due to a change in width of a time gate. In the example of figure 3.25, the value 0.20 yields the best performance. Comparing this particle ID curve to the curves for the adjacent steps of the scaling parameter, at 0.15 and 0.25, gives an indication of the error on the curve of typically 1% due to the background scaling parameter. This is sufficiently small to expose the clear trends in section 3.3. Each curve entering the comparisons in particle ID performance in section 3.3, first went through the stage to optimise its background scaling parameter.

The reconstruction performance is generally robust to small fluctuations in the photon yield for a given track. If for example the time spread introduced by the photon sensor causes some signal to be rejected by a time gate, the particle ID performance may still be high, as will be shown in section 3.3.3. If, however, the hardware gate or type 1 software gate removes signal, this results in a dramatic loss in the particle ID efficiency, as observed for the 1.0 ns time gate in figure 3.26. This is due to the high intrinsic resolution of the RICH detector, which means that if the gate rejects, say, 5% of the signal, it has removed *all* photons for 5% of the tracks, rather than 5% of the photons for all tracks. This is avoided for the results in section 3.3 by keeping a safe margin in the width of the gates.

3.3 Particle identification performance using time

The results for the particle ID including hit time information are presented for the LHC Run 3 and Run 4 conditions in section 3.3.1 and for the HL-LHC Run 5 and Run 6 in section 3.3.2. The performance with time gating is compared to the likelihood maximisation in space and time in section 3.3.3. The reduction in the number of reconstructed photon objects using time gating is presented in section 3.3.4, which is highly relevant for the CPU resources.

3.3.1 Particle identification at the LHC Run 3 luminosity

The particle ID performance using time gates of various widths is compared at the LHC Run 3 luminosity in figure 3.27. Using the time gate in the PDMDB described in section 2.3, a performance similar to the 2 ns curve can be expected. Although this hardware gate is primarily intended to reduce the MAPMT signal-induced noise contribution to the data bandwidth, the result shows that the additional rejection of out-of-time photons is advantageous for the particle ID efficiency.

The sub-nanosecond time gates are applied in software, as described in section 3.2.5, using a type 1 gate of 500 ps, type 2 gates of 50 to 200 ps and a type 3 gate of 10 ps. A clear trend of increasing particle ID performance with decreasing width of the time gates can be observed. The 10 ps time gate reaches a kaon ID efficiency of 99% at a pion misID efficiency of 2%. It can therefore be concluded that the particle ID performance without time is dominated by the error due to background photons and photon-to-track mis-association, rather than the Cherenkov angle resolution. This is an important consideration for choosing the strategy for




Figure 3.27: Particle ID curves at the LHC Run 3 and Run 4 luminosity, for different widths of the time gate. The 2 ns gate can be applied in hardware, while the other gates are implemented in the reconstruction software.




Figure 3.28: Particle ID curves at the HL-LHC Run 5 and Run 6 luminosity, using the approximation for pile-up presented in section 3.1.5. For comparison, the blue curve shows the Upgrade I conditions (low luminosity) without a time gate.

future detector upgrades. Although there may be improvements in the spatial resolution, these would require the pixel error, chromatic error and emission point error in table 1.1 to all improve. The use of time information directly targets the number of background photons, and is expected to play a major role in the future enhancement of the particle ID performance.

Although significant hardware upgrades are not currently planned, the trend towards improved particle ID performance in figure 3.27 motivates the use of time already in Upgrade Ib. A possibility would be to replace the readout electronics in part of the detector in order to implement a type 1 time gate of 500 ps using a PV timestamp from TORCH. This would improve the particle ID efficiency, as shown by the black curve, and remove additional sensor noise from the data.

3.3.2 Particle identification at the HL-LHC Run 5 luminosity

In order to simulate the particle ID performance at the Run 5 luminosity, pile-up is simulated as described in section 3.1.5. This pile-up increases the photon detector occupancy by a factor of seven, which is detrimental to the particle ID performance shown in figure 3.28, where a 200 ps time gate results in a kaon ID efficiency of 78 to 79% at a pion misID efficiency of 1%, compared to a 95 to 96% kaon ID efficiency without pile-up. The implementation of hit time information in Upgrade II is therefore critical to be able to operate the detector in the high-luminosity environment.

For comparison, the particle ID performance without time in the Run 3 conditions is shown in blue in figure 3.28. This performance would be met using a time gate of 50 to 100 ps in the high-luminosity environment, assuming that time would be the only improvement made during Upgrade II. As discussed for the Cherenkov angle resolution in section 1.4.1 and the photon detector in section 3.4.2, other improvements in the spatial domain are foreseen, which would lower the required time resolution. Following the same trend as in section 3.3.1, an excellent particle ID efficiency, better than achieved in the current detector, would also be possible in the high-luminosity environment using a time gate of 10 ps.

3.3.3 Time gating and time in the likelihood calculation

For the comparison in this section, a sensor time spread of $\sigma_{sensor} = 40$ ps is chosen, which approximately corresponds to the current state-of-the-art time resolution of an MCP detector. The extension of the likelihood maximisation algorithm into the time domain, as described by equation 3.21, is simulated with $\Delta t_{bin} = 40$ ps, 20 ps and 5 ps. The particle ID performance is then compared for time gates of 200 ps, 160 ps, 120 ps, 80 ps and 40 ps.

The use of time in the likelihood maximisation algorithm does not improve the particle ID efficiency. Compared to the particle ID curve for $\Delta t_{bin} = 40$ ps, which is shown as the red curve labelled ‘40 ps scale’ in figure 3.29, the curve for $\Delta t_{bin} = 20$ ps gives the same performance and for $\Delta t_{bin} = 5$ ps a few percent lower performance, which for clarity are

not shown in the figure. This is understood to be the result of the indirect contribution of time to the particle identification, through background rejection and photon-to-track association, as was mentioned in section 3.2.5. The extension of the likelihood algorithm is therefore a compromise. On the one hand, the time information improves the photon-to-track association by giving a lower weight to out-of-time detector hits. On the other hand, the hit-time spread is introduced by the photon detector, and is not related to the physical process associated with the particle ID. Instead, the particle ID is performed in the spatial domain using the Cherenkov angle information. It is therefore possible for a photon on the Cherenkov ring, with a ‘strong contribution’ to the spatial particle ID, to produce a hit in the tail of the photon sensor time distribution, resulting in a low weight in the 4-dimensional likelihood maximisation algorithm. The results suggest that this effect outweighs the advantages of a better photon-to-track association, even in the pile-up environment. A time gate is therefore the preferred use of time in the RICH reconstruction.

Figure 3.29 compares the particle ID performance for a range of time gates. The optimal width of the gate is found at 120 ps, which corresponds to $3\sigma_{sensor}$. The particle ID efficiency is reduced at narrower time gates because more signal is rejected, and at wider gates because less background is rejected. The time gate of $3\sigma_{sensor}$ gives approximately the same performance as the extension of the likelihood maximisation into the time domain with an electronic readout bin size of 40 ps. However, the 120 ps time gate is easier to implement in the reconstruction software and could provide a less stringent requirement on the readout bin size.




Figure 3.29: Particle ID curves for a sensor time resolution of $\sigma_{sensor} = 40$ ps and a range of time gates between 200 and 40 ps. Additionally, the red curve shows the performance of the likelihood maximisation in space and time for an electronic readout bin size $\Delta t = 40$ ps.

3.3.4 Photon objects and CPU resources

The number of reconstructed photon objects for different time gates is presented in figures 3.30 and 3.31. This is an important parameter for the CPU resources and reconstruction speed: each reduction in combinatorial background saves calculations and reduces costs.

The total number of photon objects per event, including both signal and combinatorial background, is plotted against the width of the applied time gate at the luminosity conditions of Run 3 in figure 3.30 and including pile-up for Run 5 in figure 3.31. The time gates are implemented as the type 2 gate described in section 3.2.5, and are applied after the spatial cuts in the reconstruction software. The strong reduction in the number of photon objects between 100 ps and 1 ns is due to the combinatorial background from tracks from one PV and photons from other PVs. The number of photon objects is approximately constant for time gates with a width above 1 ns, because the background outside the bunch crossing time interval of approximately 1 to 2 ns is small and expected to be more uniformly distributed in space, making the spatial cut more effective. The gating becomes less effective below 100 ps, because the majority of the photon objects are from tracks and photons originating from the same PV, which can be extremely difficult to resolve. However, the results in sections 3.3.1 and 3.3.2 have shown that a high particle ID efficiency can be achieved despite this irreducible combinatorial background.

The type 1 gate can be applied before the photon object gets reconstructed, as discussed in section 3.2.5. For a width of 100 ps in the RICH 1 and 500 ps in the RICH 2 detector, the number of photon objects is reduced by a factor of 6.73 ± 0.06 and 2.36 ± 0.03 respectively,




Figure 3.30: The number of photon objects per event as a function of the applied time gate in the RICH reconstruction, for the LHC Run 3 and Run 4 luminosity.




Figure 3.31: Identical plot to figure 3.30, but including pile-up to simulate the high-luminosity LHC environment.

where the statistical error will be dominated by systematic errors from the various future detector resolutions. Since the calculations for the photon objects are amongst the most computationally demanding in the reconstruction [98], the type 1 gate is expected to be the most effective at saving CPU resources.

MC information can be used to distinguish signal and background photon objects in figures 3.30 and 3.31. The 20 ps time gate reduces the number of background objects by 74% and 92% compared to no time information for the RICH 1 and RICH 2 detectors respectively during LHC Run 3. In the high-luminosity environment including pile-up, this background reduction increases to 94% and 97% respectively. The ratio of the number of signal to background photon objects using the 20 ps time gate decreases from 0.43 without pile-up to 0.27 with pile-up. This demonstrates that although the background increases by an order of magnitude in the high-luminosity environment, the cuts in space and time become more efficient, resulting in only a moderate increase in the signal-to-background ratio.

3.4 Implications for detector technologies

Figure 3.32 shows the particle ID curves for sensors with a standard deviation of the Gaussian hit time distribution ranging from 150 ps to zero. For each curve, the $3\sigma_{\text{sensor}}$ time gate described in section 3.3.3 is applied. The trend shows that a hit time resolution of the order of 20 ps (with a 60 ps time gate) is required to maintain the LHC Run 3 performance in the high-luminosity environment. Currently, the micro-channel plate (MCP) detector, discussed in section 3.4.1, is the only available single-photon detector that can match this time resolution requirement. A second candidate technology, the silicon photomultiplier (SiPM), is discussed in section 3.4.2.

Table 3.3 provides an overview of the MAPMT, SiPM and MCP detector characteristics. For the MAPMTs, which were described in section 2.1.1, the combination of the relatively large pixel size and the maximum anode current would lead to saturation and fast ageing in the high-luminosity environment. However, each technology has its drawbacks, and if future developments could reduce the MAPMT pixel size by a factor of 2 to 3, the MAPMT could still be a candidate for the Upgrade II detector.




Figure 3.32: Particle ID curves for different sensor time resolutions, each with a $3\sigma_{sensor}$ time gate and in the presence of pile-up. The curves demonstrate the improvement from time resolution only. A resolution of approximately 150 ps is expected for the MAPMT and SiPM, while 20 to 40 ps can be achieved by MCP detectors.

	MAPMT [68, 70, 86]	SiPM [40, 102]	MCP [55, 103]
Time res. (FWHM) [ps]	~ 300	~ 200	~ 30
Pixel size [mm]	2.88	~ 1	~ 1 (anode), 10^{-2} (channel)
Peak quantum efficiency	30 to 35% at 350 nm	45% at 400 nm	$\sim 30\%$ at 350 nm
Dark-count rate	< 5 Hz/channel	$\mathcal{O}(10^5 \text{ to } 10^7)$ Hz/mm ²	$\mathcal{O}(1)$ Hz/mm ²
Radiation hardness	UV glass window	Lattice defects	UV glass window
Gain ageing (50% loss)	$\sim 5 \times 10^3$ C	SPAD	$\mathcal{O}(10)$ C cm ⁻² (ALD)
Max. anode current	$\sim 100 \mu\text{A}$ (total)	Quenched	$\mathcal{O}(0.1) \mu\text{A}/\text{cm}^2$
Bias voltage [V]	$\mathcal{O}(10^3)$	10-100	$\mathcal{O}(10^3)$
Robust in B-field	RICH 1 shielding (< 5 mT)	Not affected	Micro-channel (< 2 T)

Table 3.3: Comparison of the MAPMT, SiPM and MCP single-photon detectors. The advantages of the technologies are shown in green, while the challenges are in orange. Details are provided in the main text.

3.4.1 The Micro-Channel Plate (MCP) detector family

The Photek MCP-PMT developed for the TORCH detector was discussed in detail in section 1.4.2. Many of the considerations made by the TORCH collaboration on for example the photon time resolution, QE, spatial resolution, radiation hardness, sensitivity to magnetic field, cooling requirements, gain and DCR are comparable to the requirements in the future upgrade of the RICH detector. The MCP-PMT is therefore a highly attractive candidate technology, primarily due to its exceptional time resolution of typically 30 ps.

The principle of an MCP detector is shown in figure 3.33. A photoelectron generated at the photocathode is accelerated in the electric field inside the device and reaches a micro-channel. As the electron travels through this channel, multiple collisions liberate additional electrons and provide signal amplification similar to the dynode chain of the MAPMT. In the example in figure 3.33, two MCPs are stacked in a ‘Chevron’ orientation, which produces a high gain and blocks the path for ion feedback from the second MCP to the first MCP. The resolution of the detector improves with a reduced separation between the plates and an increased inter-plate bias voltage [104]. The photocurrent from the MCP is collected at the anodes.

The flexibility in the choice of photocathode, MCP and anode has given rise to a family of MCP detectors. In addition to the TORCH MCP-PMTs, two particularly interesting varieties are the hybrid detector described in [105] and the Large-Area Picosecond Photon Detector (LAPPD) [106]. The charge collection for the hybrid detector is performed by the TimePix4 ASIC described in section 1.3.1, which provides a spatial resolution of better than $10\ \mu\text{m}$ and a TDC with a bin size of 195 ps in a radiation-hard design. The single-photon time resolution is further enhanced by the charge collection over multiple anodes, and may




Figure 3.33: Schematic of the cross section of an MCP detector, showing the photocathode, two micro-channel plates in chevron orientation and the anodes. The photoelectron current is illustrated in red. Similar to the MAPMT, a resistor network is used to (negatively) bias the photocathode and MCPs.

benefit from future improvements in the next generation of TimePix ASICs. The MCP for the LAPPD consists of a large-area glass capillary array with an emissive atomic layer deposition (ALD) coating. The detector in [106] has an active area of $195 \times 195 \text{ mm}^2$.

The main challenge for the application of MCPs in the RICH detector is their limited lifetime, as shown by the gain ageing and maximum anode current in table 3.3. Although the ALD coating improves the lifetime, this is still work in progress. For example, the results in [103] show that the ALD coating lowers the current saturation levels and increases the gain recovery time after reaching saturation.

3.4.2 Silicon PhotoMultiplier (SiPM)

The SiPM was discussed in the context of the SciFi detector in section 1.3.2 and is included in table 3.3. This device offers a number of attractive features: a high, green-shifted QE, small pixel size, photon-counting capability, low bias voltage and low sensitivity to a magnetic field. Furthermore, the silicon technology provides a cost-effective solution.

A single-photon time resolution between 175 and 330 ps FWHM was measured for a range of SiPMs from different manufacturers in [102]. Although the simulation studies in this thesis demonstrate that this time resolution alone is insufficient to maintain the LHC Run 3 particle ID performance in the high-luminosity environment, other benefits of the SiPM may relax the requirements for the time resolution. For example, the green-shifted, high photon sensitivity reduces the chromatic error, and the reduced pixel size combined with the Upgrade II optical system improves the Cherenkov angle resolution. These would allow a tighter spatial cut around the Cherenkov ring, providing additional background reduction. Simulation studies of the effect of the SiPM pixel size and enhanced QE on the particle ID performance are ongoing in the RICH group. These will ultimately be combined with the time resolution studies and the Upgrade II geometry to provide a detailed assessment of the SiPM performance in the future upgrade detector.

The main challenge for the SiPM is the radiation environment and the dark-count rate. For the SciFi tracker, a DCR of 14 MHz per channel after a fluence of $6 \times 10^{11} \text{ neq cm}^{-2}$ was observed at -40°C temperature. This background rate is at the same level as the peak photon occupancy in the RICH1 detector. A fluence of $3 \times 10^{12} \text{ neq cm}^{-2}$, further increasing by an order of magnitude in the high-luminosity environment, is expected in the RICH detector. Therefore, as was discussed in section 1.3.2, a mixture of cooling, annealing and bias voltage optimisation, as well as microlenses and neutron shielding would be required to operate SiPMs in the future RICH detector.

3.5 Chapter conclusion

Time is a powerful tool for background reduction in the LHCb RICH reconstruction, which will be critically important in the high-luminosity LHC environment. The high intrinsic time resolution of the RICH detector of less than 10 ps was demonstrated in this chapter and is limited by photon dispersion in the gas radiator and quartz window. The use of time information can strongly improve the particle identification performance. Using a hypothetical photon detector hit time resolution of 10 ps, an approximately 99 % kaon ID efficiency at a 2 % pion misID efficiency can be achieved for both the LHC Run 3 and the HL-LHC Run 5 conditions. Additionally, the order-of-magnitude reduction in combinatorial background using a time gate is expected to significantly reduce the CPU resources required for the RICH reconstruction, which is particularly important for the future real-time analysis.

The developments of the LHCb simulation framework for the timing studies in this chapter included a novel 4-dimensional PV generator, calculations for the photon dispersion effect in the RICH detector, the propagation of time information through the RICH reconstruction algorithm, the formulation of an accurate prediction of the Cherenkov photon hit time of arrival based on reconstructed parameters and the extension of the likelihood maximisation algorithm into the time domain. These improvements have laid the foundation for further timing studies. It was shown that the excellent particle identification of the RICH detector expected for the LHC Run 3 (without time information) can be maintained in the HL-LHC era by only using a hit time gate of 50 to 100 ps. The inclusion of other future detector improvements, such as a smaller pixel size, may relax this required time resolution. Some of the next steps for the simulation studies will therefore be the detailed description of a candidate photon sensor with readout electronics, such as the SiPM and MCP detectors, as well as the implementation of the Upgrade II optical system.

The advantages of a 3 ns time gate in the currently ongoing RICH detector Upgrade Ia were presented in chapter 2. The results in this chapter strongly motivate the use of sub-nanosecond time information in the future upgrades of the RICH detectors. A small prototype with high time resolution is being proposed for installation during Upgrade Ib, followed by a full-scale detector with time information during Upgrade II in the year 2030.

Part II

A semiconductor tracker for muon scattering tomography

Chapter 4

Apparatus for muon scattering tomography

Cosmic-ray muon-scattering tomography (MST) exploits highly penetrating muons as a probe to passively image objects with a high density and a high atomic number. In this chapter, a semiconductor tracker for high-resolution MST is presented, which employs silicon strip detector modules from the ATLAS Semiconductor Tracker (SCT) [107]. This is the first published use of a semiconductor tracker in the field of MST [108]. The $80\text{ }\mu\text{m}$ silicon strip pitch allows much smaller separation between the inner and outer detector planes while still achieving a similar angular resolution to previously reported MST detectors. The tracker therefore opens up the field of compact and portable MST systems capable of imaging objects that are difficult to reach or dangerous to move. MST and its applications are introduced in section 4.1 and the SCT modules are described in section 4.2.

The more open geometry of the MST system and the negligible radiation environment compared to the ATLAS experiment allow simplification of the electronic readout system. A novel, low-power, compact, portable and scalable electronic readout system was therefore designed using the FPGA-based readout board described in section 2.2.2. The system, which is presented in sections 4.3 and 4.4, was also subsequently adopted by the Forward Search Experiment (FASER) at the LHC. The analysis of data obtained during the FASER quality assurance program is used to characterise the excellent performance of the system, as described in section 4.5. The mechanical design, which is critical to the performance of the system, integrates active cooling and precision alignment of the SCT modules, as presented in section 4.6.

4.1 Muon scattering tomography

This section provides a brief introduction to MST. The properties of cosmic-ray muons, which provide a naturally occurring source of radiation are discussed in section 4.1.1. The multiple Coulomb scattering in materials is described in section 4.1.2, followed by an overview of cosmic-ray muon imaging detectors and applications in section 4.1.3.

4.1.1 Cosmic-ray muons

Cosmic rays are high-energy particles, predominantly protons, that reach the upper atmosphere from outer space. These primary particles interact with the atmosphere, producing a cascade of secondary particles. The short-lived, charged mesons in this cascade decay into more stable muons with a lifetime of $2.2 \mu\text{s}$, which is sufficient to reach sea level. Besides photons and neutrinos, muons are the most abundant secondary cosmic radiation at sea level. Owing to their charge, muons can be detected through ionisation in particle detectors. This source of highly penetrating radiation allows muon scattering tomography to be performed anywhere without the need for on-site particle accelerators.

The angular distribution and energy spectrum of muons at sea level is determined by the production spectrum, the energy loss in the atmosphere and the particle decays. Measurements of the energy spectrum at different angles from the zenith are shown in figure 4.1. The average muon momentum is 3 to 4 GeV/c, with a tail of the distribution up to approximately 1 TeV/c. The average momentum of the muons increases at large angles, due to a longer path length through the atmosphere, which filters out low-momentum muons and allows high-momentum pions to decay to high-momentum muons before the pions interact. This is often referred to as hardening of the muon spectrum. Additionally, the spectrum varies with altitude and geophysical location.




Figure 4.1: Measured muon momentum spectrum at sea level for a range of angles from the zenith. Source: [109].




Figure 4.2: The MST system contains two tracking stations positioned around a sample.

Muons reach the Earth's surface at a rate of approximately $1 \text{ min}^{-1} \text{ cm}^{-2}$ through a horizontal detector [110, 111]. The flux is greatest at the zenith and falls approximately as $\cos^2 \theta$, where θ is the angle from the zenith. The low muon flux compared to a particle beam or X-ray source means that the MST image recording time is relatively long. A decrease in this recording time generally causes a degradation of image resolution and contrast.

4.1.2 Multiple scattering in materials

Muons travelling through matter undergo multiple Coulomb scattering, where the overall deflection depends on the density, atomic number, path length and muon momentum. The scattering angle distribution is described by Molière's multiple scattering theory [112], where the central 98 % is approximately Gaussian [113] with width:

$$\theta_0 = \frac{z \cdot 13.6 \text{ MeV}}{p \beta c} \sqrt{\frac{x}{X_0}} \left[1 + 0.038 \cdot \ln \left(\frac{x}{X_0} \right) \right] \quad (4.1)$$

Here, p and βc are the momentum and velocity of the muon, $z = 1$ is the charge number of the muon and x and X_0 are the thickness and radiation length of the object respectively. The radiation length is the mean distance over which an electron loses all but $1/e$ of its energy and can be written in approximate form using Dahl's expression [114]:

$$X_0 = \left[\frac{A \cdot 716.4 \text{ g cm}^{-2}}{Z(Z+1) \ln(287/\sqrt{Z})} \right] \left[\frac{1}{\rho} \right] \quad (4.2)$$

Here, ρ , A and Z are the density, atomic weight and atomic number of the material. The approximate $\rho \cdot Z^2$ dependence for large Z characterises the sensitivity to dense objects with a high atomic number, which is a unique feature of MST. The scattering angle in equation 4.1 is inversely proportional to the muon momentum. The cosmic-ray momentum spectrum, which was shown in figure 4.1, therefore adds a spread to the measurement.

	Material properties			θ_0 [mrad] at three muon momenta		
	Z	A	ρ [g/cm ³]	0.3 GeV/c	3 GeV/c	30 GeV/c
Acetal	7	14	1.4	19.3	1.8	0.2
Aluminium	13	27	2.7	35.0	3.3	0.3
Silicon	14	28	2.3	34.4	3.2	0.3
Iron	26	56	7.9	83.4	7.9	0.8
Lead	82	207	11.3	156.2	14.7	1.5
Uranium	92	238	19.1	214.5	20.3	2.0

Table 4.1: The scattering angle θ_0 at three different cosmic-ray muon momenta, for a 5 cm-thick sample of a range of materials. The angle is calculated using the stated values of atomic weight A , atomic number Z and density ρ . In these studies, acetal, iron and lead samples will be imaged.

The scattering angle θ_0 for a 5 cm-thick object of different compositions is calculated using three muon momenta in table 4.1. The measured scattering angle has an additional contribution from the tracker resolution. The high-resolution semiconductor tracker is therefore expected to have a high object material and feature discrimination, which can lead to a decrease in image recording time.

4.1.3 Muon imaging detectors and applications

Cosmic-ray muon imaging can be broadly divided into two categories: muography based on attenuation and muon scattering tomography based on multiple Coulomb scattering. This section gives an overview of the applications and technologies for these techniques.

Muography

Although cosmic-ray muons do not interact strongly with matter, and can penetrate more than 100 m deep in rock [115], the wide muon momentum spectrum results in significant attenuation in large objects. This effect is exploited in muography, where the muon flux is measured and compared with the expected flux in order to generate a density map of the material within the angular acceptance above the detector.

Applications of muography include the search for a hidden chamber in the Second Pyramid of Giza using a spark chamber [116], the mapping of archaeological structures using a scintillating fibre tracker [117], the observation of cosmic ray shadowing by the Moon using the Baksan underground scintillation telescope [118], the imaging of the internal structure of volcanos using emulsion detectors [119], the imaging of the Fukushima reactor core using a scintillator telescope [120] and the measurement of the overburden of an underground tunnel using a drift-tube detector [121]. These have in common that the detectors need to be installed in confined spaces underneath large structures. The portable, compact and high-resolution semiconductor tracker, which is designed for muon scattering tomography, would therefore also be suitable for muography.

Muon scattering tomography

For muon scattering tomography, detectors are located above and below an imaging region wherein objects of interest are located, as illustrated in figure 4.2. Individual muons are tracked into and out of the imaging region and the scattering angle is measured, which depends on the amount of material in the path of the muon. The requirement to track both the incoming and outgoing muons limits the imaging region.

MST was first demonstrated at the Los Alamos National Laboratory in 2003 [122] for the passive detection of nuclear contraband in commerce at ports and borders. Although nuclear material is radioactive, its detection via radiation detectors can be thwarted by a small amount of lead shielding [123]. To be able to still detect the nuclear material, conventional X-rays would need to be of extremely high energy [124]. Instead, cosmic-ray muons provide a high sensitivity to high-Z materials and can easily penetrate through

containers or lorries from all angles, without requiring the driver to leave the vehicle. The original MST demonstrator consisted of an aluminium drift-tube array. Plates of known thickness and composition were placed between the tracking planes below the imaging region, in order to get an estimate of the outgoing muon momentum [125]. This drift-tube array has been developed into a drive-through scanning portal at Freeport in the Bahamas [126].

Subsequent projects have used resistive-plate chamber detectors [127] and gas-electron multipliers [128], as well as extruded scintillator bars and wavelength-shifting fibres read out by PMTs and SiPMs for the CRIPT (cosmic-ray inspection and passive tomography) project [129] and the muon portal project [130] respectively. A wide range of applications has been explored, such as the imaging of sealed nuclear-waste drums [131, 132], the search for radioactive material in scrap metal containers to avoid accidental melting and mixing [133] and the mapping of the material composition inside blast furnaces [134].

The angular resolution of MST systems is determined by the granularity, internal muon scattering and separation of the detector planes. The previously reported detectors have a relatively large granularity of around 1 mm, a large separation of the detector planes of typically 10 cm, a total target volume of the order of several cubic-metres and an angular resolution of the order of 10 mrad. The semiconductor tracker presented here is much more compact, with a separation between the planes of 23.8 mm and a separation of the outer modules of less than 40 cm, resulting in a target volume of 1.27 l. Despite these small dimensions, the 80 μm silicon strip width allows an angular resolution of approximately 1 mrad to be achieved.

4.2 ATLAS Semiconductor Tracker module

The SCT module shown in figure 4.3 consists of two layers of single-sided p-in-n silicon sensors, each containing 768 strips with 80 μm pitch and mounted back-to-back on a thermal conductor. The layers are at a stereo angle of ± 20 mrad, which gives position information in the y -direction. The active area of the sensors is $128 \times 64 \text{ mm}^2$. The module is designed to have low mass and radiation length, which reduces muon scattering in the tracker to a minimum.

The silicon strips are read out by ABCD [135, 136] application-specific integrated circuits mounted on a hybrid [137]. These ABCDs implement signal amplification, analogue to digital conversion, data compression and masking of noisy strips. As the module was designed for the ATLAS experiment, the ABCDs are clocked at the LHC bunch crossing frequency of 40 MHz. Although the arrival of a cosmic ray muon is asynchronous to the 25 ns clock, the ability to stretch the signal to more than 25 ns using the ABCD shaping amplifier, combined with the three-consecutive-bin readout mode, still results in a high detection efficiency.

Due to a power consumption of 6.0 W per module, cooling is required to prevent thermal runaway and to reduce thermal noise. The unirradiated sensors are fully depleted at a 150 V reverse bias voltage.




Figure 4.3: ATLAS Semiconductor Tracker module mounted on its aluminium handling frame with the adapter card.

4.3 Electronic readout design

The decision to use ATLAS SCT modules avoids the need for the design, construction and testing of silicon sensors and analogue readout electronics. A simplified trigger and readout design, as shown in figure 4.4, is used instead of the readout in the ATLAS experiment. The adapter card, which is mounted on each module, is introduced in section 4.3.1. A trigger is provided by the coincidence of scintillator paddle detectors positioned above and below the tracker. The trigger controller was described in section 2.2.2, and the logic implemented in the readout boards is presented in section 4.3.2.




Figure 4.4: Schematic diagram of the electronic readout architecture.

4.3.1 Module adapter card

The adapter card provides the interface for the clock, control and data lines between each silicon module and the readout board, as well as the ABCD bias voltages, the silicon bias voltage and the temperature readout of the modules, as illustrated in figure 4.5. The card is mounted on the aluminium handling frame of a module, as shown in figure 4.3, and connected to the hybrid using a 36-pin connector with 1.27-mm pins.

Four Maxim Integrated MAX9180-EXT-T [138] low-voltage differential signal (LVDS) line drivers and receivers on the card ensure the signal integrity between the readout board and silicon module, and help to reduce the risk of electrical damage to the ABCDs through the external connections. Five twisted-pair cables between the readout board and adapter card carry two LVDS serial data lines, an LVDS serial command line, an LVDS clock line and a 3.3 V line to power the LVDS buffers. Two pairs of pins on the card connect the thermistors on the module hybrid to an external temperature interlock.

The compact design of the adapter card benefits its integration with the SCT module, as shown in figure 4.3. In order to allow the card to be mounted flat on the aluminium plate of the module, the 36-pin connector to the hybrid is surface mount and the electrical components are located on the top surface. However, this also results in a high density of signal traces in the top layer of the PCB, which requires careful routing of the signals, in particular the differential lines.




Figure 4.5: The adapter card contains a filter for the silicon bias voltage, LVDS buffers for the clock, command and data lines between the readout board and the SCT module, power regulator circuits for the ABCD bias voltages and the connection to the temperature interlock. The analogue and digital grounds are joined on the SCT module.

Two versions of the card were produced: a two-layer PCB ('V1') with switching power regulator circuits and a four-layer PCB ('V2') without power regulator circuits.

Version 1 adapter card

The V1 card and its PCB layout are shown in figure 4.6. This card was used for the presented MST tracker. Two Analog Devices ADP2303 [139] switching power regulator circuits provide $VCC=3.5\text{ V}$ at 900 mA and $VDD=4.0\text{ V}$ at 400 mA to the module hybrid. These regulator circuits allow the use of a single-channel high-power low-cost input power supply for all SCT modules, with a voltage between 5 and 20 V. Since the circuits produce a stable output voltage over a range of input voltages, they reduce the risk of electrical damage to the ABCDs.

The high density of signals and components in the two-layer V1 design caused switching noise from the regulator circuits to couple to other sensitive parts of the circuit, such as the silicon bias voltage line on the opposite side of the PCB. Figure 4.7 was recorded using an oscilloscope that was AC-coupled to the silicon bias voltage on the 36-pin connector, while the VCC regulator circuit was terminated using a 4Ω resistor to emulate the power consumption by the SCT module. Noise peaks of an amplitude around 70 mV are observed




Figure 4.6: The V1 adapter card and the top (left) and bottom (right) layout of the two-layer PCB. The power regulator switching nodes are shown in yellow and the silicon bias voltage in red.

at the 700 kHz switching rate of the power regulator. Although the SCT modules contain additional filters for the bias voltage supply, this switching noise increases the overall noise behaviour of the tracker.

Version 2 adapter card

In order to achieve better noise performance, a second version of the adapter card was developed. Although not used for the scattering data, this board was used during the acquisition of the system characterisation data described in section 4.5. For this version, the power regulators are removed, and instead two voltage levels, 3.5 and 4.0 V, are required from an external power supply. Additionally, the PCB contains four layers for additional electrical ground planes, as shown in figure 4.9. These modifications result in a higher stability of the silicon bias voltage, as shown in figure 4.8.




Figure 4.7: The voltage on the silicon bias output to the SCT module, measured using an AC-coupled oscilloscope while the VCC power regulator was switching at 700 kHz.




Figure 4.8: Identical to figure 4.7, but for the V2 adapter card without power regulators.

Electrical grounding scheme

The silicon bias voltage supply passes through an LC filter before reaching the SCT module, which contains a further double RC filter [107]. The electrical grounds of the silicon bias voltage and the ABCD bias voltages are floating, and are joined on the module hybrid, close to the ABCDs, as illustrated in figure 4.5. A second connector on the adapter card allows the silicon bias voltage to be daisy chained between modules, which reduces the required number of power supply channels. Similarly, the ABCD bias voltage was provided using two power supply channels, one for each set of four modules.

The daisy chaining of the modules introduces ground loops because the grounds are connected at the SCT module hybrid and at the power supply. An additional ground loop is introduced by the temperature interlock because the channels of the interlock are internally connected to a common ground [140], while the thermistors on the SCT module hybrids are also connected to ground [137]. In order to reduce these ground loops, single power leads were used between the supplies and the light-tight box and the cable lengths in the daisy chain were kept as short as possible. No significant difference in the recorded noise hits was observed when running with or without the temperature interlock connected.

The ground at the readout board, which is part of the LVDS buffer circuit, is connected to the ground of the ABCD bias voltages using a switch on the adapter card. To avoid a ground loop via the Ethernet connection between the readout boards and the DAQ PC, the switch should only be set for one adapter card per readout board or per daisy chain of SCT modules, whichever covers a larger set of SCT modules.




Figure 4.9: Four-layer PCB layout of the V2 adapter card. The switching power regulators have been removed and two electrical ground planes are added. A set of surface-mount resistors on the top layer can be used to choose the electrical ground configuration.

4.3.2 Trigger and readout

This section presents the logic for the configuration and readout of the SCT modules, which was implemented in the FPGA of the readout board, as shown in figures 4.10 and 4.11. When the two scintillator detectors result in a coincidence (A in figure 4.10), the trigger controller output to the readout board toggles (B). While the readout boards are processing, a gate is returned to the trigger controller (C), which prohibits any further triggers to be sent. The trigger controller output is registered by the readout board (D) on the next clock edge of the 40 MHz clock. On receipt of the trigger, the FPGA logic counts down a user-defined number of clock cycles before it issues a trigger command (E) to all ABCDs. This delay synchronises the trigger command with the arrival of the silicon signal at the end of the ABCD pipeline (F). The ABCDs read three consecutive 25 ns time bins at the end of the pipeline.

The signals D, E and C are controlled by the **ABCD trigger and configuration encoder** logic shown in figure 4.11. Additionally, this logic receives ABCD configuration commands from the **RGMII RX and configuration buffer** logic, and forwards these commands to the SCT modules via the LVDS command line, which is also used for trigger commands. For testing purposes, it is therefore possible to issue a readout trigger command through the software and Ethernet interface, rather than the trigger controller interface used during normal data acquisition. When the **trigger and configuration encoder** logic issues a trigger command to the modules, it simultaneously activates the **ABCD decoder and buffer** logic for the receipt of data.

The ABCDs respond to the trigger command by sending out data (G) on the serial data lines connected to a readout board. There are up to eight serial data lines, two per module,




Figure 4.10: Timing diagram for signals in the trigger controller (TC), readout board (RB) and modules (MOD).

There is one 16 kbit buffer per ABCD configured as dual-port RAM with 1-bit write port width and 8-bit read port width. Each buffer is logically subdivided into 32 slots of 512 bits. The upper bits of the buffer write pointer are generated from an event counter and select the slot. The lower write pointer bits are generated from a counter that increments for each incoming serial data bit and wraps around to zero at 511. The buffering allows the bandwidth of the downstream electronics and Ethernet to be used efficiently, while preserving the ability to maintain a higher instantaneous trigger rate. As the ABCD data are zero suppressed, the length of the data stream depends on the occupancy in the silicon. The chosen buffer size guarantees that at least 28 hits per silicon layer can be stored per event. In the rare case of a buffer overwrite, the buffer address pointer wraps around and the event gets rejected in software, as the preamble pattern is no longer present.




Figure 4.11: FPGA logic diagram for the configuration and readout of SCT modules. The ABCD trigger and configuration encoder and ABCD decoder and buffer logic was adapted to the SCT modules, while the Ethernet interface shaded in grey has not been modified for these studies.

The Ethernet packets from the two readout boards are routed through an Ethernet switch to the control PC. When the CPU usage of the PC is high, Ethernet packages can occasionally get lost. Although a small rate of event loss is tolerable, the synchronisation between the tracker stations is lost. The event IDs in the Ethernet headers from both readout boards are therefore compared in software after data-taking has finished. If these headers are not the same, packets are discarded until synchronisation is restored.

The trigger rate is ultimately limited by the time taken by the ABCDs to transmit the hit information, which is done at one bit per 25 ns clock period. The shortest ABCD pattern is 53 bits at zero occupancy, giving a maximum trigger rate of around 750 kHz. The longest ABCD pattern, at 50% occupancy with alternating hit strips, is 6563 bits, corresponding to a trigger rate of 6 kHz. For the chosen readout buffer size, the 100 MBit s⁻¹ Ethernet link limits the trigger rate to 195 kHz. In practice, the downstream performance of the Ethernet switch and control PC also have to be taken into account. The tracker was successfully operated at an event rate of 30 kHz when it was used as a beam telescope during CERN LHCb upgrade beam tests, as described in appendix A. During cosmic ray data taking, the scintillator detectors trigger the tracker at approximately 2.5 Hz, which is well below the limit.

4.4 Software

The detector control software aims to provide an intuitive way to operate the tracker and is described in section 4.4.1. The decoding of the data returned by the ABCDs is presented in section 4.4.2.

4.4.1 Detector configuration interface

The Java graphical user interface (GUI) is shown in figure 4.12. In response to an action in this high-level interface, the software generates the corresponding sequence of bits to configure the SCT modules, following the ABCD manual in [135]. This bit-sequence is sent to the readout board as an Ethernet packet with custom protocol and containing a unique target identifier in the packet payload. As discussed in the previous section, the FPGA on the readout board subsequently sends the configuration data as a stream of serial commands to the ABCDs. In this section, the basic functionality of the GUI is outlined. The GUI software was written in a modular way, in order to provide a high flexibility to extend the panel with new configuration procedures.

The **Get ready for start** button initialises the SCT modules. It sends the following sequence of configuration commands:

- *Soft reset*: clears all buffers inside the ABCDs.
- *BC reset*: sets the beam-crossing counter to zero.
- *Configuration*: writes to the configuration register. For example, this sets the readout mode and calibration mode and configures the ABCD to be a master or slave.

- *Threshold/Calibration DAC*: sets the threshold value during operation and the calibration pulse amplitude in test mode, which are stored in the same 16-bit register.
- *Bias DAC*: sets the bias current value in the input transistor and shaper bias current value in the ABCD.
- *Strobe delay*: loads the delay of the calibration pulse in test mode with respect to the system clock phase.
- *Trim DAC*: applies a channel-to-channel threshold correction.
- *Mask*: masks individual silicon strips.
- *Enable data*: enables the ABCDs to receive triggers and return hit data.

The initialisation sequence is performed for each ABCD and module in turn. The configuration for each ABCD can be changed by navigating to the relevant settings under the **ABCD** tab in figure 4.12. The GUI stores these settings, which can be exported and imported as XML files to simplify the initialisation of the GUI for new installations.

The **Start run** button activates the external trigger, which is used to record cosmic-ray data. The **Hardware trigger delay** sets the number of clock periods between signals D and E in figure 4.10. The **Stop** button finishes the run, closes the data file and disables the external trigger, until a new run is started again.

Several other modes of data taking are implemented for the characterisation of the SCT modules and front-end readout, for example the **Software trigger** and **Calibration pulse** data runs, as well as the parameter scans in the **Noise scan**, **Strobe delay scan**




Figure 4.12: Java GUI of the detector control software. The **Main** menu contains several parameter scan procedures, the **Settings** menu provides the readout board options and the **ABCD** menu stores the configuration of the ABCDs.

and **Threshold/Calibration scan** panels. These procedures are described in more detail in section 4.5.

4.4.2 Data decoding

The Ethernet packets containing the ABCD data are written to a file, which is read and decoded using a C++ algorithm. Each event starts with a 5-bit preamble followed by a 4-bit L1 counter, an 8-bit beam-crossing number and the zero-suppressed data. The algorithm checks the preamble and verifies whether the L1 counter is incremented by one with respect to the previous event. The decoding follows the decision tree in table 4.2. Reading the table from left to right, the algorithm uses the fixed bit-patterns that prefix each type of data in order to branch to the corresponding decoding step.

For each isolated or non-isolated hit, the chip-ID and channel number of the hit are used to calculate the strip number between 1 and 768 (corresponding to six ABCDs per layer with 128 strips per ABCD). These numbers are stored in a ROOT tree with one branch per silicon layer. This tree is subsequently used for hit clustering and tracking, which will be presented in chapter 5.

In addition to the hit information, several data types exist for the debugging of the ABCDs. The send-ID mode returns the ABCD configuration data, and is the default response to a trigger when the data-taking mode is disabled. The buffer error messages during operation are defined in [135]. The FPGA error message is generated by the decoding software in the rare case of an inconsistent pattern from the ABCDs, which is generally caused by an overwritten buffer in the FPGA RAM, as was described in section 4.3.2.

Bit sequence				Decoded data
0 1	14-bit: chip-ID + channel + hit pattern	1	0	Isolated hit (store and return)
			1 0 0 0	Isolated hit (store) + trailer
			1 + hit pattern	Non-isolated hit (store all and return)
		0		FPGA Error: buffer overwritten
0 0 1				No-hit data packet
0 0 0	4-bit: chip-ID	1 1 1		ABCD Send-ID (output configuration)
		0 0 1		ABCD Error: no data available
		0 1 0		ABCD Error: buffer overflow
		1 0 0		ABCD Error: buffer soft reset required

Table 4.2: Decision tree of the C++ algorithm for the decoding of zero-suppressed ABCD data. The bit sequence is preceded by the preamble, L1 counter and beam-crossing number, and followed by a trailer of fifteen zeros.

4.5 Characterisation of the sensor and front-end electronics

This section presents the noise and gain characteristic of the sensor coupled to the new electronic readout system. The test protocols are implemented in the Java GUI introduced in section 4.4.1. The noise scan is outlined in section 4.5.1 and the calibration scan for the three-point gain test in section 4.5.2.

The presented characteristics are for a typical SCT module out of the 80 modules that were tested for the Forward Search Experiment (FASER) at the LHC using the electronic readout system presented in this thesis. A test system for FASER was set up at CERN by the author, as described in the technical proposal [141] of the experiment. The dataset analysed in this study was recorded during the characterisation by the FASER group, where the modules were tested one-by-one on a dedicated cooling plate. The temperature measured by the thermistors on the module hybrids was around 15° C during operation.

4.5.1 Noise scan

During the noise scan, which is started from the GUI panel shown in figure 4.13, the threshold is incremented and a user-defined number of software triggers is sent at a rate of approximately 1 kHz at each threshold step. The threshold setting of the ABCD discriminator covers a range from 0 to 640 mV with steps of 2.5 mV, which corresponds to input charges from 0 to 12.8 fC with a step size of 0.05 fC at the nominal front-end gain of 50 mV fC⁻¹ [135]. No hardware triggers are sent and no calibration pulse is injected. The threshold DAC value in the configuration is stored by adding a tag in the FPGA to the data for each trigger. The data are recorded at 150 V bias voltage resulting in the full depletion of the silicon.

The noise scan data for five ABCD channels are shown in figure 4.14. A complementary error function is fitted to the data of each channel:

$$fit(x) = A \left[1 - erf\left(\frac{x - \mu}{\sqrt{2}\sigma}\right) \right] \quad (4.3)$$

Vt50 is defined to be the threshold at which a hit is recorded 50 % of the time on average and is equal to the value of the fitted parameter μ . The contribution of noise (ENC) can




Figure 4.13: The noise scan panel is used to record a single data file containing a scan of the threshold parameter. At each threshold step, a *Threshold/Calibration* configuration command and an *Enable data* command are sent to the ABCDs, followed by the trigger commands.




Figure 4.14: Noise scan distributions for five channels, including the complementary error function fit. The threshold at the 50% hit level is V_{t50} , and the transition width of the error function provides the ENC.




Figure 4.15: The V_{t50} distribution obtained from the error function fit in the noise scan for all channels of a module. The spread corresponds to a threshold DAC range of approximately 10 to 30.

be obtained from the value of the fitted parameter σ , using the conversion factors given in section 4.5.2. The noise is described for the three-point gain test in order to ensure that input charges are calculated in the linear-gain regime of the ABCD amplifier. As shown in figure 4.15, the ABCD discriminator offsets for the individual channels range from 20 to 80 mV. Because the same threshold setting is programmed for all channels in the ABCD, this offset is corrected in the ATLAS experiment using the channel-to-channel threshold trim circuitry.

4.5.2 Three-point gain

The three-point gain measurement characterises the gain and equivalent noise charge of the ABCD channels. The measurement utilises the internal 100 fF capacitor charged by a chopper circuit to inject test charges into each ABCD input, triggered by commands generated in software. The gain is inferred from calibration scans with charge amplitudes of 1.5 fC, 2.0 fC and 2.5 fC.

Every fourth channel in the ABCDs is tested simultaneously with charge injection, which is selected by a 2-bit calibration address in the configuration command. The calibration scan, which is started from the GUI panel in figure 4.16, performs the threshold parameter scan four times with a different calibration address each time, in order to record data for all strips with the injected charge. The duration of the strobe signal is 200 ns and the delay with respect to the clock phase is tuned using a 5-bit strobe delay over the range of 50 ns. At each threshold step, the software generates a custom command to the ABCDs, consisting




Figure 4.16: The calibration scan measures the response of all ABCD channels to injected test charge pulses. For the three-point gain test, three calibration scans were recorded using different settings of the 8-bit test charge amplitude.




Figure 4.17: The V_{t50} of a typical ABCD channel as a function of the injected test charge amplitude. The gain is the slope of the linear fit.

of a calibration pulse command plus a number of zeros specified by the **Calibration delay** setting in figure 4.12 plus a readout trigger command. This calibration delay synchronises the readout trigger with the arrival of the test charge pulse at the end of the ABCD pipeline.

The fit in equation 4.3 is used to obtain the V_{t50} threshold values, which are plotted as a function of the calibration pulse amplitude in figure 4.17. The gain of the channel is then obtained as the slope of the linear fit, with an error of approximately 20%. This error is dominated by the 5% error in the amplitude of the test charge, which includes all variations between the ABCDs, as well as switching pickup noise during the test charge injection. Within one ABCD, the skew is smaller, at $<2\%$ for test charge amplitudes between 0.8 and 4.0 fC. The distribution of the gain for all channels of a module is shown in figure 4.18. The mean gain of 51 ± 4 mV/fC agrees with the nominal value of 50 mV/fC [136].

The Equivalent Noise Charge (ENC) at the front-end including the silicon strip detector is calculated from the transition width σ in equation 4.3 for a threshold scan with a 2.0 fC test charge amplitude. The measured gain of each channel is used to convert the threshold DAC units into the equivalent input charge. Since the noise is measured in addition to the test charge, the ABCD amplifier is in the linear regime. The ENC distribution shown in figure 4.19 has a mean of $1573 \pm 136 e^-$, which agrees with the measured ATLAS reference value of $1550 \pm 100 e^-$ [136]. Figures 4.20 and 4.21 show the variation in ENC for the channels in the top and bottom silicon layer of a module. The periodicity of the noise reflects the division of channels over the twelve ABCDs, six for each silicon layer, on the module hybrid.




Figure 4.18: The gain distribution of all ABCD channels in a typical module.




Figure 4.19: The Equivalent Noise Charge obtained from a calibration scan with a 2 fC test charge amplitude.




Figure 4.20: ENC as a function of the channel number of the top silicon layer of the module.




Figure 4.21: ENC for the bottom silicon layer. The boundaries between the six ABCDs are drawn.

4.5.3 Noise characteristic and the semiconductor tracker

The results in this section confirm preliminary measurements done before MST data-taking, showing that the quality of the modules has not been affected by their long-term storage.

The MST tracker is operated at a higher temperature of around 35° C compared to 15° C during the noise characterisation. Since the ENC typically increases by $6 e^-$ per degree Celsius [142], the tracker is expected to have 5 to 10 % more noise. Furthermore, the tracker was built using V1 adapter cards, and contains multiple, daisy-chained SCT modules, which increases the noise as discussed in section 4.3.1. The tracker is therefore operated at a higher threshold. A global threshold of 55 DAC, corresponding to 137.5 mV, was found to yield a negligible noise level with only a small trade-off in efficiency. This efficiency was measured by fitting a track, as described in section 5.1, to all layers of silicon except the layer of interest and extrapolating the track to the layer of interest, to look for a hit in a window of 10 silicon strips around the predicted position. The obtained efficiency ranged from 86 % to 93 %, which can be improved in the future using threshold trimming and V2 adapter cards.




Figure 4.22: Measurements on the alignment of the sensor within the handling frame, in millimetres or degrees using a precision microscope.




Figure 4.23: The handling frames are mounted on a bespoke support for cooling and alignment.

4.6 Mechanics

The support for the semiconductor tracker was designed to meet a number of critical requirements. The support should align up to ten modules to within approximately the width of a silicon strip and should allow the tracker to be positioned horizontally (in a particle beam) or vertically (for cosmic-ray data taking). The design should also be rigid and compact, with a length of ~ 40 cm and adjustable spacing between the modules. The mechanical design is presented in section 4.6.1, and the technical drawings are included in appendix C. The modules require active cooling during operation, which should be integrated into the support with a short thermal path to the hybrid. The cooling is discussed in section 4.6.2. The light-tight environment is important to reduce sensor noise and is described in section 4.6.3.

4.6.1 Mechanical alignment of the SCT modules

The SCT modules are contained in aluminium handling frames, which provide protection of sensitive parts during testing and storage. The alignment of the handling frames, rather than the sensors, would make the tracker design more robust and would allow the modules to be relatively easily interchanged or displaced. Since this approach relies on the alignment of the sensor within the module handling frame, the measurements in figure 4.22 were made for seven modules. For example, the distances from the edges of the sensor to the outer




Figure 4.24: Aluminium support for cooling and alignment of the modules. To show the sensor, the aluminium cover plate has been removed from one of the modules. A sample is clamped in the plastic holder in the imaging region, which is marked in red. Each tracking station comprises four modules plus a spare.

surfaces of the handling frame are $A = 87.771 \pm 0.032 \text{ mm}$ and $B = 23.352 \pm 0.054 \text{ mm}$, where the error in the mean of the seven modules is given. The rotation of the sensor with respect to the frame was measured as $C = 21.14 \pm 0.38 \text{ mrad}$. This error results in a translational error of $24 \text{ }\mu\text{m}$ at the edges of the sensor. Since these errors are less than the $80 \text{ }\mu\text{m}$ width of the silicon strips, the strategy to align the handling frames can be adopted. The residual misalignments can be compensated later in software.

The support is shown in figures 4.23 and 4.24. As shown in appendix C, it consists of a slotted plate and a cold bar with a wide cooling channel, which is sealed with an O-ring by attaching the bar to the slotted plate. The aluminium surfaces of the cold bar and the slotted plate are precision manufactured with a flatness of approximately $20 \text{ }\mu\text{m}$. In order to constrain the modules to be aligned and parallel, the handling frames are clamped into the right angle between the cold bar and the slotted plate, as shown in figure 4.25. Two threaded holes were drilled in the handling frames, and using M5 bolts with an aluminium bar across, the module is pulled towards the slotted plate, while a third bolt fixed by a square nut is used to push the module onto the cold bar. The 1.0 mm-thick aluminium




Figure 4.25: The handling frame of the SCT module is aligned using two bolts to pull the module onto the slotted plate and a third bolt to push the module onto the cold bar. The corner with the shortest thermal path to the module hybrid is attached to the support.

cover plates often protruded from the sides of the handling frames, thereby preventing good alignment, and were replaced by smaller plates. The thickness of the new plates is 0.5 mm, in order to reduce the multiple scattering in the modules whilst maintaining mechanical integrity.

The spacing between the modules on the support can be adjusted, allowing the tracker to be optimised for high resolution at large spacing or greater muon acceptance at small spacing. Two tracking stations, each containing four modules plus a spare, are mounted on the support. In this study, the module spacing is 23.8 ± 0.1 mm with a 155.2 ± 0.1 mm separation between the tracking stations. The imaging volume of 1.27l is indicated in red in figure 4.24.

4.6.2 Cooling of the tracker

Water coolant at 10°C at the inlet flows through the support at around 5 to 10 l min^{-1} to dissipate heat from the modules. The modules are orientated such that the corner with the shortest thermal path to the sensors is in contact with the cold reservoir. The flat surfaces provide a good thermal contact even without thermal paste.

The ATLAS upgrade temperature interlock [140] is used to switch off the ABCD bias voltage and the cooling circulation when the temperature at the module hybrid reaches below 12°C or above 41°C . This avoids potential damage to the sensors due to thermal runaway, in case of failure of the cooling system, or condensation on cold parts, in case of failure of the power supply. The light-tight box is flushed with nitrogen to further avoid condensation. The cooling and alignment provided by the support was found to be stable over weeks of data taking, with temperature fluctuations within 0.3°C and with an approximately constant trigger rate, efficiency of the silicon sensors and scattering angle resolution.

4.6.3 Light-tight environment

The silicon sensors are operated in a light-tight environment in order to reduce noise. The tracker and scintillator detectors were placed inside a black lightweight expanded-polypropylene box [143], as shown in appendix A. The box has a front-loading door and a width of 40 cm, depth of 50 cm and height of 60 cm. A patch-panel was built into the rear wall, providing light-tight connections for the silicon bias voltage, ABCD bias voltage and readout board bias voltage, as well as a cable conduit for the Ethernet cables to the readout boards and the temperature, trigger and gate signal cables. Additionally, the patch-panel provided the inlet and outlet of the coolant and an inlet with flow-control for the nitrogen gas. The light-tight box was originally designed to be able to operate the SCT modules without their handling frames. The aluminium cover plates of the handling frames are expected to significantly block light to the sensors and relax the light-tight requirements for the box.

4.7 Chapter conclusion

Cosmic-ray muon scattering tomography exploits highly penetrating muons as a probe to passively image objects. While a variety of detectors and applications exist, these are generally large-scale systems with an angular resolution of the order of 10 mrad. In this chapter, a high-resolution silicon strip tracker was presented, which is the first published use of semiconductor technology in the field of MST. The tracker is compact, with a height of less than 40 cm, whilst achieving an angular resolution of approximately 1 mrad, as will be presented in the next chapter. It is therefore an important development towards portable MST detectors.

The apparatus contains ATLAS SCT modules and a novel, low-power, scalable electronic readout system. A mechanical support was designed for the precision-alignment and active cooling of the modules. In the future, some cooling requirements could be avoided if the same silicon sensors were read out by custom electronics that were not tied to the relatively high LHC clock frequency.

The noise characteristic of the sensors was measured, in order to confirm both the high performance of the SCT modules after long-term storage and the low noise injection by the novel readout system. The analysis in this chapter used data recorded for the characterisation of SCT modules for the FASER experiment, which was performed using the readout system from this thesis including the V2 adapter card. The measured mean channel gain of $51 \pm 4 \text{ mV/fC}$ and mean ENC of $1573 \pm 136 \text{ e}^-$ agree with the original module specifications. The high efficiency and low noise performance in the ATLAS experiment could therefore be met, which is an important conclusion for the FASER experiment. For the MST tracker, however, the noise is expected to be higher due to ground loops from the power supply to the modules, the relatively high operating temperature of around 35 °C and the V1 adapter card. The threshold setting was not tuned individually per channel but, instead, the same value was programmed for all channels and chosen to yield a negligible noise level at the expense of a small loss in efficiency, which ranges from 87 to 93 % for the silicon layers. These thresholds can be trimmed to reach an expected efficiency close to 100 % in the future.

Chapter 5

Performance of the MST tracker

In this chapter, the angular resolution of the semiconductor tracker and the cosmic-ray imaging results for plastic, iron and lead samples are presented, in order to demonstrate the high performance of the apparatus.

Section 5.1 describes the detector hit clustering and the association of clusters with muon tracks. Data from 8000 muons recorded in the absence of a sample were subsequently used to fine-tune the module alignment in software by a translation and rotation, as described in section 5.2. A Geant4 simulation of the tracker, which is used to better understand the scattering angle resolution of the tracker, is introduced in section 5.3. The image reconstruction software and cosmic-ray imaging results are presented in section 5.4.




Figure 5.1: Number of adjacent silicon strips that make up a cluster.




Figure 5.2: Number of clusters of hits per silicon layer in the tracker.

5.1 Track identification

Starting from the detector hits in the ROOT file obtained from the C++ decoding algorithm described in section 4.4.2, the first step in the analysis is to cluster adjacent hit strips in the same silicon layer. The cluster position is assigned to the boundary between the strips when two strips are hit. The number of adjacent strips that form a cluster is shown in figure 5.1. The analysis discards 1.3 % of the clusters on the basis that these contain more than two strips. The number of clusters per silicon layer in figure 5.2 is typically low, with a mean occupancy of 0.5 clusters per layer per event in the absence of cosmic-ray tracks.

The following algorithm was developed to identify muon tracks in the relatively high-multiplicity events when the tracker was operated in LHCb beam tests described in appendix A. Cosmic-ray events typically contain one muon track and are less challenging to identify. Given the approximately 90 % efficiency of the silicon layers during operation, as discussed in section 4.5.3, only around 19 % of the muon tracks is expected to cause a cluster in each of the sixteen layers in the tracker. Therefore, the algorithm has been designed to also be able to track muons that are detected in three out of four layers, which is anticipated to increase the overall track efficiency to approximately 81 %.

The even and odd silicon layers are treated separately in the fitting procedure due to the stereo angle. A pair of lines is found for each tracking station using linear regression in the xz -plane (as was defined in figure 4.3), and combined using the stereo angle to obtain the gradient and intercept in the yz -plane. The first step in the association of clusters to




Figure 5.3: Example of track identification in software. Different initial tracks are chosen from the set of five potential tracks, and the associated clusters removed. Tracks are assigned to the remaining clusters, with preference given to tracks with four clusters and tracks with lowest mean squared deviation. Configuration B contains three tracks and no unassigned clusters.

tracks is to draw all possible tracks connecting three or four clusters in different modules within the acceptance, with a loose cut on the residuals to remove non-physical tracks. In the example in figure 5.3, this step is shown on the left hand side, where five tracks are created. However, some of these tracks share the same cluster, which is unlikely given the low muon track multiplicity and $80\text{ }\mu\text{m}$ silicon strip pitch. The algorithm therefore steps through the tracks and removes the associated clusters, as shown in the second column in figure 5.3. Each choice of initial track leaves a different set of remaining clusters. To assess which initial track gives the most likely set of tracks, additional tracks are assigned to the remaining clusters, with preference given to tracks containing four clusters followed by tracks with the smallest mean squared deviation. The associated clusters are removed, until it is no longer possible to draw tracks.

The most likely configuration of muon tracks is defined as the one that (i) leaves the minimum number of clusters unassigned. In configurations where the number of unassigned clusters is the same, preference is given to (ii) the one with the lowest mean squared deviation. It is also favoured to have (iii) one less track if this only increases the number of unassigned clusters by one. In the example of figure 5.3, the best configuration is B, based on criterion (i), as this configuration assigns all clusters to three muon tracks. Based on criterion (iii), configuration A is preferred over C and D.

5.2 Software alignment of modules

The alignment of the modules in the x -direction is corrected in software using data taken in the absence of a sample. Track identification is performed as described in section 5.1. However, when a muon track with a cluster in each silicon layer of both tracking stations is identified, the predicted position of the cluster on each layer is calculated from a linear fit through the fifteen other layers. The deviation from this predicted position on a typical layer is plotted for 8000 muons in figure 5.4, and shows a double-peak structure with a non-zero mean. The modules are shifted by half of the deviation from the predicted position in the x -direction, which is iterated 20 times, at which point the mean deviation from zero is typically $0.1\text{ }\mu\text{m}$. After alignment, the angle between the line connecting the centres of the layers and the normal to the layers is less than 0.5 mrad and no further correction is applied. A rotation of modules around the centre in the xy -plane results in deviations in the x -direction that are different for positive and negative y -coordinates of the clusters. This creates the two-peak structure in figure 5.4, which is corrected by the same iterative approach as the translation.

The deviation from the predicted position after alignment is shown in figure 5.5. The standard deviation of the distribution is $43\text{ }\mu\text{m}$ due to internal scattering of the relatively low energy cosmic-ray muons used for the alignment. Owing to the precise alignment of the modules on the mechanical support, the correction to the mean in the x -direction is typically below $100\text{ }\mu\text{m}$. Because a similar misalignment in the x and y -direction is expected from




Figure 5.4: Deviation of clusters from the predicted positions in the x -direction, based on linear regression through all silicon layers excluding the layer of interest.




Figure 5.5: Identical plot to the one shown in figure 5.4, after a correction of the module alignment in software.

the mechanics, it is decided not to correct the alignment in the y -direction, as the spatial resolution in this direction is around 2 mm. The moderate angular acceptance of the tracker limits the error in the x -direction due to misalignment in the z -direction, which is therefore also not corrected in the analysis.

5.3 Detector simulation and angular resolution

The angular resolution of the tracker is limited by internal scattering, which depends on the muon momentum, the angle of incidence, the material budget of the module and the number of modules in the tracker. Additionally, the strip width, any remaining misalignment after calibration, inefficiencies of the modules and cross talk or electronic noise reduce the fit quality of the tracks. To better understand some of these contributions, a description of the tracker is created in Geant4 in section 5.3.1. In the absence of a sample, the scattering angle resolution from simulation and from cosmic-ray data are compared in section 5.3.2.

5.3.1 Geant4 detector description

The Geant4 description of the MST tracker is shown in figure 5.6. The tracker geometry is simulated by modules in air with a random offset in the x -direction sampled from a Gaussian with standard deviation of $80\ \mu\text{m}$ and a rotation up to 3 mrad in the xy -plane. The sensors are designed to have a low radiation length, and contribute $1.17\% X_0$ averaged over each sensor [107]. Additionally, the protective covers on either side of the silicon introduce 1.0 mm of aluminium with a radiation length of $1.12\% X_0$. To simulate the spatial resolution of the sensors, each hit in the sensor is assigned to the centre of an $80\ \mu\text{m}$ strip. The module hybrid and the electronic readout characteristics are not included. The output from the strip




Figure 5.6: Geant4 simulation of the tracker, including the silicon sensors at a stereo angle in red, the aluminium cover plates in blue and a sample in black. The sample was removed for the angular resolution studies in this chapter.

assignment has the same format as the experimental data file, which means that identical track identification, calibration and image reconstruction can be used in both cases.

The cosmic-ray muon source is modelled as a planar source, centred with the tracker and positioned 2.9 mm above the top silicon sensor. The source has an area of $66.6 \times 132.6 \text{ mm}^2$, as muons generated outside this area cannot go through the tracker. For 100 % efficient silicon sensors, $3.7 \pm 0.2 \%$ of the source particles result in a track through the tracker due to the acceptance. At 90 % silicon efficiency, $2.6 \pm 0.2 \%$ of the source particles result in a reconstructed track. The cosmic-ray angular and energy distributions from [109] are implemented. The energy distribution was shown in figure 4.1.

5.3.2 Scattering angle resolution results

Firstly, the tracker simulation is run for mono-energetic muons with infinite energy, with 4 GeV average cosmic-ray energy and with 0.5 GeV energy, and the cosmic-ray angular distribution. The angle between the incoming and outgoing muon tracks is shown in figure 5.7. The width of the Gaussian distribution for infinite energy is 0.9 mrad and arises from the silicon strip width and module alignment errors. This width increases to 1.5 mrad for 4 GeV muons and 7.3 mrad for 0.5 GeV muons due to internal scattering.

Secondly, the simulation is run with the cosmic-ray energy distribution and 40000 events in the absence of a sample, in order to compare the scattering angle distribution between simulation and experiment, as shown in figure 5.8. The scattering angle distribution is a

convolution of a Gaussian distribution, which has a width dependent on the cosmic-ray energy as shown in equation 4.1, and the broad energy spectrum of cosmic-ray muons. The benefit of reducing pixel dimensions and alignment errors is only significant for the narrowest Gaussians associated with the highest energy muons.

The error in the scattering angle is estimated using the error in the gradient m of the straight line fit to the hits in the silicon sensors, which is calculated following the approach in [144]:

$$(\Delta m)^2 = \frac{1}{n-2} \frac{\sum d_i^2}{\sum (z_i - \bar{z})^2} \quad (5.1)$$

Here, $d_i = x_i - mz_i - c$ are the residuals, n is the number of points on the track and z_i is the position of the silicon plane i . This equation is individually evaluated for the even and odd silicon layers and combined in quadrature to get the overall track fitting error. The results shown in figure 5.9 show a mode of less than 1 mrad. The error from simulation with mono-energetic particles with infinite energy and 100% efficient silicon sensors shows the lower limit in the error from the strip width and small deviations in the module position. As expected, the error distribution broadens when the simulation is run with the cosmic-ray energy distribution, and broadens further when inefficiencies are introduced in the silicon sensors.

The experimental results show a slightly wider distribution in figure 5.8 and a larger error in the gradient in figure 5.9 than in simulation. These results can be attributed to misalignments that are not corrected in the tracker calibration or to uncertainty in cluster assignment due to detector noise or cross talk for example.




Figure 5.7: Simulated angle between incoming and outgoing muon tracks in the absence of a sample, for infinite, 4 GeV and 0.5 GeV energy.




Figure 5.8: Angle between incoming and outgoing muon tracks for 40000 scattering events in simulation and in experiment.




Figure 5.9: Error in the gradient of the straight line fit to the hit silicon strips in the tracker. The curves show experiment, simulation with infinite energy and simulation with the cosmic-ray muon energy spectrum at two different silicon sensor efficiencies.

5.4 Image reconstruction

This section presents the imaging results for plastic, iron and lead samples, using cosmic-ray muons in experiment. The Angle Statistics Reconstruction (ASR) algorithm is used for the image reconstruction, as described in section 5.4.1. Section 5.4.2 contains images for the different samples and projections, as well as the combination of three objects within the sampling volume. Section 5.4.3 reflects on the image recording time.

5.4.1 Angle Statistics Reconstruction

ASR is chosen for the image reconstruction in this study, because it is robust to changes in the system setup, easy to implement and has been shown to perform better than Point of Closest Approach (PoCA) algorithms [145]. For a given volume element, the algorithm stores a score for each event that has the shortest distance between the volume element and the incoming and outgoing muon tracks below a threshold distance. After all data are processed, the events in the volume element are listed in ascending order by their score. The significance of each event in the list is weighted by the path length of the track through the volume element. Furthermore, to smooth the image, events in nearby volume elements are added to the list with a low weight of 30 % for adjacent and 15 % for diagonal volume elements. The event at a fixed fraction of the total weighted number of events is chosen, and the score stored in that event is used as the final score to create the ASR image. The fixed fraction is referred to as the ‘quantile’. For a detailed description of the algorithm, its motivation and performance see [145].

The score for ASR in this study is the absolute value of the scattering angle in the xz -plane. The 2 mm spatial resolution in the yz -plane means that the scattering angle in this




Figure 5.10: ASR reconstruction of a lead sample with 45000 muons from experiment. The black dotted lines show the sample contours. From left to right: smoothed and non-smoothed ASR, for a 0.5 quantile and for a 0.8 quantile.

direction is not measured accurately enough for tomography, but the information is still used to localise the scattering event. Figure 5.10 shows the effect of smoothing and quantile choice in the xy -projection of a $15 \times 25 \times 55 \text{ mm}^3$ lead sample imaged using 45000 muons. As discussed in [145], the algorithm performs well over a large range of quantiles, and in this thesis a 0.8 quantile and $4 \times 4 \times 4 \text{ mm}^3$ volume elements are found to give best contrast between the sample and background.

5.4.2 Cosmic-ray imaging results

The $15 \times 25 \times 50 \text{ mm}^3$ iron and lead samples are clamped in the holder shown in figure 5.13. The holder is made of black acetal and white laminated plastic in order to minimise background scattering and is mounted on a translation stage, which is outside the acceptance. The results in figure 5.11 demonstrate the ability to distinguish between air and the plastic sample holder. Given the small muon scattering angle from plastic, as was shown in table 4.1, this result confirms the high sensitivity of the semiconductor tracker. When introducing iron and lead samples, the multiple scattering significantly increases and a new colour scale is adopted. On this scale, the lead shows better visual contrast, as expected from its higher atomic number and density than iron.

Projections of the lead sample in the xz -plane and yz -plane are shown in figure 5.12. As a result of the cosmic-ray angular spectrum and the acceptance of the tracker, most muons tracks travel vertically through the tracker, resulting in the observed reduced definition in the z -direction. This effect is asymmetric in the yz -projection, where the side closest to the centre of the sampled volume gives a better representation of the sample size, which is superimposed as a black dotted line. Non-vertical tracks that pass the sample on the




Figure 5.11: ASR images with 45000 muons from experiment. The black dotted lines show the sample contours. From left to right: air, plastic holder, iron and lead samples.




Figure 5.12: ASR images of a lead sample with 45000 muons from experiment, projected in the xz -plane and the yz -plane. The colour scales are chosen to present the best contrast.




Figure 5.13: The sample holder consists of a 6 mm-thick acetal head and a laminated plastic (tufnol) pole.




Figure 5.14: ASR image of three lead samples with 74000 muons from experiment. The cross sections are indicated on the top, and the sample with the lowest peak ASR value has a thickness of 15 mm compared to 50 mm for the other samples.

side away from the centre are more likely to go out of acceptance, resulting in the conical reconstructed image of the sample.

Figure 5.14 shows the image of three lead samples of $15 \times 25 \times 50 \text{ mm}^3$, $15 \times 15 \times 50 \text{ mm}^3$ and $15 \times 15 \times 15 \text{ mm}^3$, recorded using 74000 muons. The contours of the samples are overlaid. The 15 mm-thick sample has a lower score in the ASR algorithm than the 50 mm-thick samples. Further tests of the resolving power for samples of different size, orientation and composition are highly interesting for the various applications discussed in section 4.1.3.

5.4.3 Image recording time and tracker improvements

The scanning time is a limiting factor of MST with cosmic-ray muons, which have an approximate flux of $1 \text{ cm}^{-2} \text{ min}^{-1}$ at sea level. The rate from the simulated source of area 88.3 cm^2 in section 5.3.1 is therefore approximately 1.47 s^{-1} . For 90 % efficient silicon sensors, this results in an estimated track rate of $3.89 \times 10^{-2} \text{ s}^{-1}$, which is an upper limit as it assumes a 100 % efficiency of the scintillator trigger detectors. In experiment, 45841 ± 241 tracks were recorded over a time of 340.4 ± 0.5 hours, which gives a track rate of $(3.74 \pm 0.02) \times 10^{-2} \text{ s}^{-1}$.

Cosmic-ray muon imaging has found many applications, some that can sacrifice high resolution imaging for short acquisition and object detection times [127, 146], and some that require high resolution images with less emphasis on the acquisition time [131, 116]. An example of early detection of a lead sample is given in figure 5.15 for 100, 1000, 5000 and 25000 muons. The larger volume elements for 100 muons give an average of 1.4 muons per volume element. The relatively high score in the volume element at $x = -25 \text{ mm}$ and $y = 45 \text{ mm}$ is therefore the effect of a scattering outlier. The lead sample stands out at




Figure 5.15: ASR images of a lead sample. From left to right: 100, 1000, 5000 and 25000 muons from experiment. The volume element size is increased at low statistics, in order to keep a sufficient number of muons in each element. The colour scales are chosen to present the best contrast.

less than 5000 muons, which is likely to be reduced by improvements in image reconstruction and recognition algorithms. Heavier objects such as uranium are also expected to give higher contrast.

The rate from the tracker could be improved by optimising the threshold setting to achieve a silicon sensor efficiency close to 100 %, which increases the track rate to $5.44 \times 10^{-2} \text{ s}^{-1}$ estimated from simulation. Additionally, an array of sensors could be used, or the module position and the number of modules optimised, to increase the angular acceptance of the tracker. Although an array of sensors would require new mechanics and cooling, the readout electronics are designed to be able to scale to such a larger system.

5.5 Chapter conclusion

In this chapter, the high performance of the semiconductor tracker for cosmic-ray muon scattering tomography was presented. After track identification and software alignment of the modules, the muon scattering angle distribution obtained in experiment agreed with a Geant4 simulation of the detector. The simulation provides a scattering angle resolution of 0.9 mrad for high-momentum muons and 1.5 mrad at the 4 GeV average energy of cosmic rays, which is limited by internal muon scattering. Images have been recorded using 45000 muons and Angle Statistics Reconstruction, and show good contrast between air, plastic, iron and lead materials. A lead sample can be visually localised with less than 5000 muons.

The semiconductor tracker and corresponding Geant4 simulation provide an important test-bed for further studies. In particular, the development of reconstruction and image processing algorithms is a key area for improvement in the image quality [147, 148], where for example the use of machine learning could be explored. Additionally, the module positioning and the number of modules can be optimised in the future, in order to maximise sensitivity in the presence of internal scattering and to achieve the best compromise between muon acceptance (modules closer together) and angular resolution (modules further apart). The mechanics of the tracker were designed to easily accommodate changes in the module arrangement. The amount of information per muon event would improve with a measurement of the muon momentum, or with a more accurate measurement of the scattering angle in the yz -plane, for example using pixel rather than strip sensors or moving to a crossed module geometry.

Appendix A

The semiconductor tracker as a beam telescope

The semiconductor tracker was used as a beam telescope during two test periods of the LHCb RICH UpgradeIa photon detector system in a charged particle beam at the SPS facility at the CERN North Area in 2017. The tracking information is useful for the data analysis and the online monitoring software of the tracker benefitted the alignment of the system with the beam.

Test setup at the beam area

The RICH UpgradeIa setup (see section 2.3.1) and the tracker were operated in light-tight boxes (see section 4.6.3) as shown in figures A.1 and A.2. These boxes were fixed rigidly on a custom metal alignment plate. The alignment plate was mounted on a translation table, with a pivot point directly below the lens between the plate and the table. The two mechanically coupled boxes could therefore be optimally aligned with the beam using a translation in the plane perpendicular to the particle beam and a rotation around the pivot point.

Prior to the installation of the tracker, the SCT modules were mounted one by one on the aluminium support described in section 4.6. A 24 mm-thick spacer was used to maintain a constant separation between the silicon sensors. The angular alignment of the sensors was verified using reflected laser light with a lever arm of 5 m, resulting in an error of less than 0.1 mrad. The tracker was subsequently attached to a translation stage as shown in figure A.3 and placed inside the light-tight box. Using the alignment plate as a reference, the translation stage of the tracker was moved until the centre of the upper half of the sensors was aligned with the centre of the lens of the RICH UpgradeIa system. This nominal position was chosen to avoid the beam passing through the SCT module hybrid containing the ABCD ASICs.




Figure A.1: The tracker and the RICH system are mechanically fixed onto a single metal alignment plate, which contains translational and rotational freedom for the alignment of the apparatus in the charged particle beam. A side view is given on the left and a top view on the right.




Figure A.2: The semiconductor tracker (left) and the RICH system (right) in light-tight boxes at the test area of the SPS facility at CERN. The boxes are mounted on the alignment plate and the translation table is visible in the background.

Beam alignment and online monitoring

Owing to its careful alignment with respect to the RICH system, the tracker could be used to monitor the alignment of the apparatus in the charged particle beam. The online monitor shown in figure A.5 was developed for this purpose. The monitor is based on an existing `Java` programme, which reads the tracker data that are also written to file. The decoding of the ABCD data in `Java` follows the same procedure as the offline `C++` decoding described in section 4.4.2. The monitor displays a histogram of the number of hits as a function of the silicon strip number for each layer of silicon. Since the tracker was operated with eight SCT modules containing two silicon layers per module, this produces 16 histograms. A 2-dimensional plot with one row per layer of silicon and one bin per group of 16 silicon strips was added in order to provide an overview of the beam pattern. The histograms show a staggered pattern, which is due to the stereo angle between the alternating layers of silicon.

The readout trigger for the tracker is generated by the scintillator detectors positioned around the lens in the RICH system as described in section 2.3.1. While the beam profile is approximately Gaussian distributed in the horizontal plane, the observed hit pattern in figure A.5 has sharp edges corresponding to the sides of the scintillator detectors. Outside this scintillator window, the hits are predominantly caused by other tracks in the beam that happen to arrive at approximately the same time as the track that caused the trigger. The shape of the hit distribution within the scintillator window provides information on the position and alignment of the beam with respect to the scintillator trigger detectors and the centre of the lens in the RICH system. The apparatus is aligned if the distribution peaks in the centre of the scintillator window, as shown in the example of figure A.5. However, the apparatus is misaligned if the distribution is skewed to one side, and the translation table should be moved until the peak is found.

The angular alignment with the beam was achieved using a combination of the hit position on the first tracker module and the angle between the particle tracks and the line connecting the centres of the SCT modules. The lever arm of 1.5 m between the trigger scintillator detectors and the tracker results in a sensitivity of approximately 1 mrad, based on the mechanical alignment between the lens and the tracker with an error of about 1 mm. Although the resolution of the sensors is lower in the vertical direction, the error is dominated by the mechanical alignment, and is therefore similar to the horizontal direction.




Figure A.3: Front view of the tracker mounted on the translation stage, which allows the tracker to be aligned with the centre of the RICH system in the plane perpendicular to the beam.




Figure A.4: Offline event display showing the reconstructed particle tracks for a hadronic interaction near the lens of the RICH system. Hits are shown in red, and form displaced pairs due to the stereo shift between the silicon layers. The centre of the tracker is at $z = 200$ mm.




Figure A.5: Screenshot of the online monitor for the semiconductor tracker in a charged particle beam. A histogram of the number of hits per silicon strip is produced for each of the 16 silicon layers, as well as a 2-dimensional colour-scale plot with one row per silicon layer and one column per group of 16 strips.

Tracking information for analyses

Since the tracker is located downstream of the RICH setup, hadronic interactions between the 180 GeV protons and pions in the beam and the borosilicate lens can be detected. A typical example, where four tracks converge to the interaction point, is shown in figure A.4. The relatively large surface area of $64 \times 128 \text{ mm}^2$ of the silicon modules is beneficial for the detection of such events. This information was used to understand events with a large number of hits in the MAPMTs and to filter data for photon detection efficiency studies.

The tracking information is important for the correlation studies between multi-track events and the number of photoelectron hits in the MAPMTs, as was reported in [80] for the 2015 beam tests. The number of hits per event in the RICH Upgrade Ia MAPMTs is shown in figure A.6, for a selection of events with one track (red curve) and two tracks (black curve) using the semiconductor tracker. Although a clear correlation is observed, the number of hits is not proportional to the number of tracks, because each MAPMT pixel is read out in binary mode and the probability that two photons will hit the same pixel increases with the number of tracks. In figure A.6, the most likely number of hits per MAPMT on the Cherenkov ring is 4.5 ± 1.0 and 7.8 ± 1.5 for one and two tracks respectively. The same factor of increase of approximately 1.7 was obtained in the 2015 beam tests and in simulation when the binary readout effect is included.




Figure A.6: Total number of MAPMT hits per event for a selection of events with one track or two tracks in the semiconductor tracker.

Appendix B

4-dimensional primary vertex distribution

This appendix contains the algorithm that generates the 4-dimensional primary vertex coordinates. The calculation that resulted in the `gauss4D` function and the principle of the Markov chain are described in section 3.1.1. Recurring parameters in the 4-dimensional function are cached for speed and clarity. The `smearVertex` function is called once per primary vertex and contains the Markov chain with a number of iterations that can be defined through the `NMarkovChainSamples` parameter. The `gauss4D` function is called twice per iteration in the Markov chain.

The number of Markov chain iterations was set to the default `NMarkovChainSamples = 1000` in this study. The tool is expected to be computationally demanding, and the number of Markov chain iterations could be optimised in the future. Since the RICH reconstruction algorithms can be run standalone, all studies presented in this thesis could be run on a single Monte-Carlo sample of ten thousand events. The speed of the event generator was therefore not critical.

The size of the 4-dimensional Gaussian perturbation in the Markov chain is robust to changes in the beam parameters, provided these do not change the order of magnitude of the spread of the PV distribution function. The standard deviations for the perturbation were fixed at $\sigma_{x,y} = 10 \mu\text{m}$, $\sigma_z = 5 \text{ mm}$ and $\sigma_t = 100 \text{ ps}$.

BeamSpotMarkovChainSampleVertex.cpp**1/3**

~/LHCbCMake/RichTTFUStudies-Sim09/Gauss/Gen/Generators/src/component/

```

1 // local
2 #include "BeamSpotMarkovChainSampleVertex.h"
3 //-----
4 // Implementation file for class : LHCbAcceptance
5 // 2019 : Floris Keizer and Chris Jones
6 //-----
7 //=====
8 // Standard constructor, initializes variables
9 //=====
10 BeamSpotMarkovChainSampleVertex::
11 BeamSpotMarkovChainSampleVertex( const std::string& type,
12                                   const std::string& name,
13                                   const IInterface* parent )
14 : GaudiTool ( type, name , parent )
15 {
16   declareInterface< IVertexSmearingTool >( this ) ;
17   declareProperty( "Xcut" , m_xcut = 4. ) ; // times SigmaX
18   declareProperty( "Ycut" , m_ycut = 4. ) ; // times SigmaY
19   declareProperty( "Zcut" , m_zcut = 4. ) ; // times SigmaZ
20   declareProperty( "BeamParameters" ,
21                     m_beamParameters = LHCb::BeamParametersLocation::Default ) ;
22   declareProperty( "NMarkovChainSamples", m_nMCSamples = 1000 );
23 }
24
25 //=====
26 // Initialize
27 //=====
28 StatusCode BeamSpotMarkovChainSampleVertex::initialize( )
29 {
30   StatusCode sc = GaudiTool::initialize( ) ;
31   if ( sc.isFailure() ) return sc ;
32
33   // Random number generators to perturb the four-vector in the Markov chain
34   IRndmGenSvc * randSvc = svc< IRndmGenSvc >( "RndmGenSvc" , true ) ;
35   sc = ( sc
36         && m_gaussDistX.initialize( randSvc , Rndm::Gauss( 0. , 0.010 ) ) //mm
37         && m_gaussDistY.initialize( randSvc , Rndm::Gauss( 0. , 0.010 ) ) //mm
38         && m_gaussDistZ.initialize( randSvc , Rndm::Gauss( 0. , 5. ) ) //mm
39         && m_gaussDistT.initialize( randSvc , Rndm::Gauss( 0. , 0.1 ) ) //ns
40         && m_flatDist.initialize ( randSvc , Rndm::Flat ( 0. , 1. ) ) );
41   if ( sc.isFailure() )
42   { return Error( "Could not initialize random number generators" ); }
43
44   release( randSvc ) ;
45
46   return sc ;
47 }
48
49 //=====
50 // Function representing the product of two 4D Gaussian PDFs
51 //=====
52 double BeamSpotMarkovChainSampleVertex::gauss4D( LHCb::BeamParameters * beamp ,
53                                                    const HepMC::FourVector & vec ) const
54 {
55   // Import beam parameters
56   const auto emittance = beamp -> emittance() ;
57   const auto betastar = beamp -> betaStar() ;
58   // RMS bunch length in mm:
59   const auto sz = beamp -> sigmaS() ;
60   const auto aX = beamp -> verticalCrossingAngle()/2. ;
61   // aY is the half-crossing angle around the y-axis, in the horizontal (xz) plane:
62   const auto aY = beamp -> horizontalCrossingAngle()/2. ;
63
64   const auto c = Gaudi::Units::c_light ;
65   const auto Pi = Gaudi::Units::pi;
66
67   const auto sx = std::sqrt( emittance*betastar ) ;
68   const auto sy = sx ;
69   const auto x = vec.x() ;
70   const auto y = vec.y() ;
71   const auto z = vec.z() ;
72   const auto t = vec.t() ;
73
74   const auto sX2 = std::pow( sx , 2 ) ;
75   const auto sY2 = std::pow( sy , 2 ) ;
76   const auto sZ2 = std::pow( sz , 2 ) ;
77   const auto x2 = std::pow( x , 2 ) ;
78   const auto y2 = std::pow( y , 2 ) ;

```

BeamSpotMarkovChainSampleVertex.cpp**2/3**

~/LHCbCMake/RichTTFUStudies-Sim09/Gauss/Gen/Generators/src/component/

```

79     const auto z2 = std::pow( z , 2 ) ;
80     const auto t2 = std::pow( t , 2 ) ;
81     const auto c2 = std::pow( c , 2 ) ;
82
83     const auto c20 = std::cos( 2*aX ) ;
84     const auto c40 = std::cos( 4*aX ) ;
85     const auto c01 = std::cos( aY ) ;
86     const auto c02 = std::cos( 2*aY ) ;
87     const auto c2p2 = std::cos( 2*(aX + aY) ) ;
88     const auto c2p1 = std::cos( 2*aX + aY ) ;
89     const auto c2m1 = std::cos( 2*aX - aY ) ;
90     const auto c2m2 = std::cos( 2*(aX - aY) ) ;
91     const auto c4p2 = std::cos( 2*(2*aX + aY) ) ;
92     const auto c4p1 = std::cos( 4*aX + aY ) ;
93     const auto c4m1 = std::cos( 4*aX - aY ) ;
94     const auto c4m2 = std::cos( 4*aX - 2*aY ) ;
95     const auto s10 = std::sin( aX ) ;
96     const auto s30 = std::sin( 3*aX ) ;
97     const auto s1p2 = std::sin( aX + 2*aY ) ;
98     const auto s1p1 = std::sin( aX + aY ) ;
99     const auto s1m1 = std::sin( aX - aY ) ;
100    const auto s1m2 = std::sin( aX - 2*aY ) ;
101    const auto s3p2 = std::sin( 3*aX + 2*aY ) ;
102    const auto s3p1 = std::sin( 3*aX + aY ) ;
103    const auto s3m1 = std::sin( 3*aX - aY ) ;
104    const auto s3m2 = std::sin( 3*aX - 2*aY ) ;
105
106    // Return the function
107    return std::exp( ( (-6*c2*t2)/sY2 - (18*c2*t2)/sZ2 - (16*x2)/sX2 - (8*x2)/sY2 - (8*x2)/sZ2 -
(16*y2)/sY2 - (16*y2)/sZ2 - (16*z2)/sX2 - (8*z2)/sY2 - (8*z2)/sZ2 - (8*(c2*sY2*t2+(sY2-sZ2)*(
x2-2*y2+z2))*c20)/(sY2*sZ2) - (6*c2*(sY2 - sZ2)*t2*c40)/(sY2*sZ2) + (c2*t2*c4m2)/sY2 - (c2*t2*
c4m2)/sZ2 + (4*c2*t2*c2m2)/sZ2 - (4*x2*c2m2)/sY2 + (4*x2*c2m2)/sZ2 + (4*z2*c2m2)/sY2 - (4*z2*c
2m2)/sZ2 + (16*x*y*c2m1)/sY2 - (16*x*y*c2m1)/sZ2 - (4*c2*t2*c4m1)/sY2 + (4*c2*t2*c4m1)/sZ2 + (
8*c2*t2*c01)/sY2 - (8*c2*t2*c01)/sZ2 - (2*c2*t2*c02)/sY2 - (6*c2*t2*c02)/sZ2 - (16*x2*c02)/sX2
+ (8*x2*c02)/sY2 + (8*x2*c02)/sZ2 + (16*z2*c02)/sX2 - (8*z2*c02)/sY2 - (8*z2*c02)/sZ2 + (4*c2
*t2*c2p2)/sZ2 - (4*x2*c2p2)/sY2 + (4*x2*c2p2)/sZ2 + (4*z2*c2p2)/sY2 - (4*z2*c2p2)/sZ2 - (16*x*
y*c2p1)/sY2 + (16*x*y*c2p1)/sZ2 + (c2*t2*c4p2)/sY2 - (c2*t2*c4p2)/sZ2 - (4*c2*t2*c4p1)/sY2 + (
4*c2*t2*c4p1)/sZ2 + (16*c*t*y*s10)/sY2 - (16*c*t*y*s10)/sZ2 + (16*c*t*y*s30)/sY2 - (16*c*t*y*s
30)/sZ2 - (4*c*t*x*s1m2)/sY2 + (4*c*t*x*s1m2)/sZ2 + (4*c*t*x*s3m2)/sY2 - (4*c*t*x*s3m2)/sZ2 +
(8*c*t*x*s1m1)/sY2 + (24*c*t*x*s1m1)/sZ2 - (8*c*t*y*s1m1)/sY2 - (24*c*t*y*s1m1)/sZ2 - (8*c*t*x
*s3m1)/sY2 + (8*c*t*x*s3m1)/sZ2 - (8*c*t*y*s3m1)/sY2 + (8*c*t*y*s3m1)/sZ2 - (8*c*t*x*s1p1)/sY2
- (24*c*t*x*s1p1)/sZ2 - (8*c*t*y*s1p1)/sY2 - (24*c*t*y*s1p1)/sZ2 + (8*c*t*x*s3p1)/sY2 - (8*c*
t*x*s3p1)/sZ2 - (8*c*t*y*s3p1)/sY2 + (8*c*t*y*s3p1)/sZ2 + (4*c*t*x*s1p2)/sY2 - (4*c*t*x*s1p2)/
sZ2 - (4*c*t*x*s3p2)/sY2 + (4*c*t*x*s3p2)/sZ2 )/32. ) / (8.*std::pow(Pi,3)*sX2*sY2*sZ2);
108 }
109
110 //=====
111 // Markov chain sampler : performed once per primary vertex
112 //=====
113 StatusCode BeamSpotMarkovChainSampleVertex::smearVertex( LHCb::HepMCEvent * theEvent )
114 {
115
116     // Load the beam parameters
117     LHCb::BeamParameters * beamp = get< LHCb::BeamParameters >( m_beamParameters ) ;
118     if ( ! beamp ) Exception( "No beam parameters registered" ) ;
119
120     // The sampled point, starting at the origin.
121     HepMC::FourVector x( 0 , 0 , 0 , 0 ) ;
122
123     // Repeat for n-loops, provided a sampled point is obtained within the (x,y,z) limits
124     unsigned int iLoop = 0; // sanity check to prevent infinite loops...
125     bool OK = false;
126     while ( iLoop++ < m_nMCSamples && !OK )
127     {
128         for ( unsigned int repeat = 0; repeat < m_nMCSamples; ++repeat )
129         {
130             // Compute the function for this point
131             const auto f = gauss4D( beamp , x ) ;
132
133             // Smear the point (random walk).
134             const HepMC::FourVector y( x.x() + m_gaussDistX(),
135                                         x.y() + m_gaussDistY(),
136                                         x.z() + m_gaussDistZ(),
137                                         x.t() + m_gaussDistT() );
138
139             // Compute the function for the new point
140             const auto g = gauss4D( beamp , y ) ;
141

```


BeamSpotMarkovChainSampleVertex.cpp**3/3**

~/LHCbCMake/RichTTFUStudies-Sim09/Gauss/Gen/Generators/src/component/

```

142     if ( f < g )
143     {
144         // The new point is better: always keep it
145         x = y;
146     }
147     else
148     {
149         // The new point is worse: randomly keep the point a fraction of the times
150         // depending on the probabilities f and g
151         const auto ratio = ( fabs(f)>0 ? g/f : 0.0 );
152         const auto r = m_flatDist();
153         if ( r < ratio ) { x = y; }
154     }
155 }
156
157 // Check if the spatial part of x is within the defined limits
158 OK = ( ( fabs(x.x()) < ( m_xcut * beamp->sigmaX() ) ) &&
159        ( fabs(x.y()) < ( m_ycut * beamp->sigmaY() ) ) &&
160        ( fabs(x.z()) < ( m_zcut * beamp->sigmaZ() ) ) );
161
162 // Reset and repeat
163 if ( !OK ) { x = HepMC::FourVector(0,0,0,0); }
164
165 }
166 if ( UNLIKELY(!OK) )
167 {
168     Warning( "Markov Chain sampling for PV (x,y,z,t) failed" ).ignore();
169 }
170 else
171 {
172     // Shift the sampled point to the average beam spot position
173     x.setX( x.x() + beamp -> beamSpot().x() ); // Offset the centre of the beamspot
174     x.setY( x.y() + beamp -> beamSpot().y() );
175     x.setZ( x.z() + beamp -> beamSpot().z() );
176
177     // update the values for all vertices belonging to the PV
178     auto * pEvt = theEvent -> pGenEvt();
179     if ( pEvt )
180     {
181         for ( auto vit = pEvt -> vertices_begin(); vit != pEvt -> vertices_end(); ++vit )
182         {
183             const auto pos = (*vit) -> position();
184             (*vit) -> set_position( HepMC::FourVector( pos.x() + x.x(),
185                                                         pos.y() + x.y(),
186                                                         pos.z() + x.z(),
187                                                         pos.t() + x.t() ) );
188         }
189     }
190 }
191 return StatusCode::SUCCESS;
192 }
193
194 // Declaration of the Tool Factory
195 DECLARE_TOOL_FACTORY( BeamSpotMarkovChainSampleVertex )
196 end

```



Appendix C


Mechanics of the semiconductor tracker


This appendix contains the technical drawings of the slotted plate, cold bar and mounting bar. When assembled, these components form the mechanical support presented in section 4.6.

The rubber seal between the cold bar and slotted plate avoids leakage of the coolant, which could damage the SCT modules. Two threaded, quarter-inch BSPT, brass straight hose couplings are fixed onto the cold bar as the coolant inlet and outlet. When the tracker was positioned upright for cosmic-ray data taking, the coupling near the top was used as the inlet, in order to allow gravity to aid the flow.

The mounting bar shows the dimensions of the slots. The separation between the slots was chosen to be the same as the 16 mm width of the SCT module handling frames, which provides the possibility to mount the modules without spacing. For example, four modules could be mounted as a stack with cover plates only on the outer sides, which would reduce the total thickness of the aluminium cover plates by 6 mm for the incoming and outgoing tracking stations.







Glossary

ADC Analogue-to-Digital Converter.

ALD Atomic Layer Deposition.

ASIC Application-Specific Integrated Circuit.

ASR Angle Statistics Reconstruction.

CP Charge-Parity.

DAC Digital-to-Analogue Converter.

DAQ Data Acquisition.

DCR Dark-Count Rate.

DLL Delta Log-Likelihood.

DTM Data Transmission Module.

EC Elementary Cell.

ENC Equivalent Noise Charge.

FASER Forward Search Experiment.

FCC Future Circular Collider.

FEB Front-End Board.

FPGA Field-Programmable Gate Array.

GBT Gigabit Transceiver.

GUI Graphical User Interface.

HL-LHC High-Luminosity Large Hadron Collider.

HLT High-Level Trigger.

IO Input-Output.

LGAD Low-Gain Avalanche Diode.

LVDS Low-Voltage Differential Signal.

MAPMT Multi-Anode Photomultiplier Tube.

MC Monte-Carlo.

MCP Micro-Channel Plate.

MST Muon Scattering Tomography.

PCB Printed Circuit Board.

PDF Probability Density Function.

PDMDB Photon Detector Module Digital Board.

PMT Photon Multiplier Tube.

PV Primary Vertex.

QE Quantum Efficiency.

RICH Ring-Imaging Cherenkov.

SCT Semiconductor Tracker.

SEU Single-Event Upset.

SIN Signal-Induced Noise.

SiPM Silicon Photomultiplier.

SPAD Single-Photon Avalanche Diode.

SPS Super Proton Synchrotron.

TCM Trigger and Control Module.

TDC Time-to-Digital Converter.

ToT Time-over-threshold.

Bibliography

- [1] T. Massam et al. “Experimental observation of antideuteron production”. *Nuov Cim A*, 63:10, 1965. doi:10.1007/BF02898804.
- [2] F.J. Hasert et al. “Observation of neutrino-like interactions without muon or electron in the gargamelle neutrino experiment”. *Phys. Lett. B*, 46(1):138–140, 1973. doi:10.1016/0370-2693(73)90499-1.
- [3] UA1 Collaboration. “Experimental observation of isolated large transverse energy electrons with associated missing energy at $s=540$ GeV”. *Phys. Lett. B*, 122(1):103–116, 1983. doi:10.1016/0370-2693(83)91177-2.
- [4] UA2 Collaboration. “Observation of single isolated electrons of high transverse momentum in events with missing transverse energy at the CERN pp collider”. *Phys. Lett. B*, 122(5-6):476–485, 1983. doi:10.1016/0370-2693(83)91605-2.
- [5] ALEPH Collaboration, DELPHI Collaboration, L3 Collaboration, OPAL Collaboration, SLD Collaboration, LEP Electroweak Working Group, SLD Electroweak and Heavy Flavour Groups. “Precision electroweak measurements on the Z resonance”. *Physics Reports*, 427(5-6):257–454, 2006. doi:10.1016/j.physrep.2005.12.006.
- [6] ALEPH Collaboration. “Determination of the number of light neutrino species”. *Phys. Lett. B*, 231(4):519–529, 1989. doi:10.1016/0370-2693(89)90704-1.
- [7] G. Baur et al. “Production of antihydrogen”. *Phys. Lett. B*, 368(3):251–258, 1996. doi:10.1016/0370-2693(96)00005-6.
- [8] LHCb Collaboration. “Observation of $J/\psi p$ Resonances Consistent with Pentaquark States in $\Lambda_b^0 \rightarrow J/\psi K^- p$ Decays”. *Phys. Rev. Lett.*, 115(7):072001, 2015. doi:10.1103/PhysRevLett.115.072001.
- [9] ATLAS Collaboration. “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”. *Phys. Lett. B*, 716(1):1–29, 2012. doi:10.1016/j.physletb.2012.08.020.
- [10] CMS Collaboration. “Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC”. *Phys. Lett. B*, 716(1):30–61, 2012. doi:10.1016/j.physletb.2012.08.021.
- [11] G. Charpak et al. “The use of multiwire proportional counters to select and localize charged particles”. *Nucl. Instrum. Meth.*, 62(3):262–268, 1968. doi:10.1016/0029-554X(68)90371-6.
- [12] ATLAS Collaboration. “ATLAS barrel toroid: Technical Design Report”. *CERN-LHCC-97-019*, 1997. <https://cds.cern.ch/record/331065>.
- [13] E. Mobs. “The CERN accelerator complex”. *OPEN-PHO-ACCEL-2018-005*, 2018. <http://cds.cern.ch/record/2636343>.

- [14] M. Benedikt et al. “Future Circular Collider - European Strategy Update Documents”. *CERN-ACC-2019-0005*, 2019. <https://fcc-cdr.web.cern.ch/>.
- [15] FASER Collaboration. “FASER: ForwArD Search ExpeRiment at the LHC”. *UCI-TR-2019-01*, 2019. <https://arxiv.org/abs/1901.04468>.
- [16] ATLAS Collaboration. “The ATLAS Experiment at the CERN Large Hadron Collider”. *JINST*, 3: S08003, 2008. doi:10.1088/1748-0221/3/08/S08003.
- [17] CMS Collaboration. “The CMS experiment at the CERN LHC”. *JINST*, 3:S08004, 2008. doi:10.1088/1748-0221/3/08/S08004.
- [18] ALICE Collaboration. “The ALICE experiment at the CERN LHC”. *JINST*, 3:S08002, 2008. doi:10.1088/1748-0221/3/08/S08002.
- [19] LHCb Collaboration. “The LHCb Detector at the LHC”. *JINST*, 3:S08005, 2008. doi:10.1088/1748-0221/3/08/S08005.
- [20] Edited by G. Apollinari et al. “High-Luminosity Large Hadron Collider (HL-LHC) Technical Design Report V.0.1”. *CERN-2017-007-M*, 2017. doi:10.23731/CYRM-2017-004.
- [21] LHCb collaboration. “Physics case for an LHCb Upgrade II: Opportunities in flavour physics, and beyond, in the HL-LHC era”. *CERN-LHCC-2018-027*, 2018. <https://arxiv.org/abs/1808.08865>.
- [22] P.A.M. Dirac. “The quantum theory of the electron”. *Proceedings of the Royal Society A*, 117(778): 610–624, 1928. doi:10.1098/rspa.1928.0023.
- [23] C.D. Anderson. “The Positive Electron”. *Phys. Rev.*, 43(6):491–494, 1933. doi:10.1103/PhysRev.43.491.
- [24] A.D. Sakharov. “Violation of CP invariance, C asymmetry, and baryon asymmetry of the universe”. *Sov. Phys. Usp.*, 34(5):392–393, 1991. doi:10.1070/pu1991v034n05abeh002497.
- [25] M. Kobayashi and T. Maskawa. “CP-Violation in the Renormalizable Theory of Weak Interaction”. *Progress of Theoretical Physics*, 49(2):652–657, 1973. doi:10.1143/PTP.49.652.
- [26] L. Wolfenstein. “Parametrization of the Kobayashi-Maskawa Matrix”. *Phys. Rev. Lett.*, 51(21): 1945–1947, 1983. doi:10.1103/PhysRevLett.51.1945.
- [27] A.J. Bevan et al. “The Physics of the B Factories”. *Eur. Phys. J. C*, 74(11):3026, 2014. doi:10.1140/epjc/s10052-014-3026-9.
- [28] LHCb collaboration. “Measurement of CP violation in $B^0 \rightarrow J/\psi K_S^0$ and $B^0 \rightarrow \psi(2S) K_S^0$ decays”. *Journal of High Energy Physics*, 11, 2017. doi:10.1007/JHEP11(2017)170.
- [29] M. Thomson. *Modern Particle Physics*. Cambridge University Press, 2013. doi:10.1017/CBO9781139525367.
- [30] S. Bifani et al. “Review of lepton universality tests in B decays”. *J. Phys. G: Nucl. Part. Phys.*, 46: 023001, 2019. doi:10.1088/1361-6471/aaf5de.
- [31] LHCb collaboration. “Observation of $J/\psi p$ resonances consistent with pentaquark states in $\Lambda_b^0 \rightarrow J/\psi K^- p$ decays”. *Phys. Rev. Lett.*, 115(7):072001, 2015. doi:10.1103/PhysRevLett.115.072001.
- [32] L. Canetti et al. “Matter and antimatter in the universe”. *New J. Phys.*, 14:095012, 2012. doi:10.1088/1367-2630/14/9/095012.
- [33] LHCb collaboration. “LHCb Tracker Upgrade Technical Design Report”. *CERN-LHCC-2014-001*, 2014. <https://cds.cern.ch/record/1647400>.

- [34] K. Hennessy. “LHCb VELO upgrade”. *Nucl. Instrum. Meth. A*, 845:97–100, 2017. doi:10.1016/j.nima.2016.04.077.
- [35] Mark Williams. “VELO Upgrade II: Overview”. *3rd workshop on LHCb Upgrade II*, March 2018. <https://indico.in2p3.fr/event/16795>.
- [36] LHCb Collaboration. “LHCb VELO Upgrade Technical Design Report”. *CERN-LHCC-2013-021*, 2013. <https://cds.cern.ch/record/1624070>.
- [37] R. Ballabriga et al. “Asic developments for radiation imaging applications: The medipix and timepix family”. *Nucl. Instrum. Meth. A*, 878:10–23, 2018. doi:10.1016/j.nima.2017.07.029.
- [38] G. Aglieri Rinella et al. “The NA62 GigaTracker”. *Nucl. Instrum. Meth. A*, 845:147–149, 2017. doi:10.1016/j.nima.2016.06.045.
- [39] A. Lai. “Sensors, electronics and algorithms for tracking at the next generation of colliders”. *Nucl. Instrum. Meth. A*, 927:306–312, 2019. doi:10.1016/j.nima.2019.02.050.
- [40] P. Hopchev. “SciFi: A large Scintillating Fibre Tracker for LHCb”. *Proceedings of the Fifth Annual LHCP*, 2017. <https://arxiv.org/abs/1710.08325>.
- [41] B. Leverington. “SciFi R&D for Upgrade II”. *4th Workshop on LHCb Upgrade II*, April 2019. <https://indico.cern.ch/event/790856/>.
- [42] M. Calvo Gomez et al. “Studies of a neutron shielding for the upgraded LHCb detector”. *LHCb-INT-2017-014*, 2017. <https://cds.cern.ch/record/2268537>.
- [43] M. Deckenhoff. “Scintillating Fibre and Silicon Photomultiplier Studies for the LHCb Upgrade”. *CERN-THESIS-2015-318*, 2015. <https://cds.cern.ch/record/2140068>.
- [44] C. Allaire et al. “Beam test measurements of Low Gain Avalanche Detector single pads and arrays for the ATLAS High Granularity Timing Detector”. *JINST*, 13:P06017, 2018. doi:10.1088/1748-0221/13/06/P06017.
- [45] T. Ackernley et al. “Mighty Tracker: Design studies for the downstream silicon tracker in Upgrade Ib and II”. *LHCb-INT-2019-007*, 2019. https://lphe.epfl.ch/tell1/SciFi_Upgrade_II/Mighty_Tracker.pdf.
- [46] R. Schimassek et al. “HVC MOS Pixel Detectors - Methods for Enhancement of Time Resolution”. *IEEE Nuclear Science Symposium*, pages 1–5, 2016. doi:10.1109/NSSMIC.2016.8069903.
- [47] C.L. Da Silva. “Magnet side station”. *4th Workshop on LHCb Upgrade II*, April 2019. <https://indico.cern.ch/event/790856/>.
- [48] LHCb collaboration. “Expression of Interest for a Phase-II LHCb Upgrade: Opportunities in flavour physics, and beyond, in the HL-LHC era”. *CERN-LHCC-2017-003*, 2017. <https://cds.cern.ch/record/2244311/>.
- [49] LHCb RICH Collaboration. “Performance of the LHCb RICH detector at the LHC”. *Eur. Phys. J., C*(73):2431, 2013. doi:10.1140/epjc/s10052-013-2431-9.
- [50] M. Aleksa et al. “Strategic R&D Programme on Technologies for Future Experiments”. *CERN-OPEN-2018-006*, 2018. <https://cds.cern.ch/record/2649646>.
- [51] C. D’Ambrosio. “The future of RICH detectors through the light of the LHCb RICH”. *Nucl. Instrum. Meth. A*, 876:194–197, 2017. doi:10.1016/j.nima.2017.02.076.
- [52] M.J. Charles and R. Forty. “TORCH: Time of flight identification with Cherenkov radiation”. *Nucl. Instrum. Meth. A*, 639:173–176, 2011. doi:10.1016/j.nima.2010.09.021.

- [53] N. Harnew et al. “Status of the TORCH time-of-flight project”. *Nucl. Instrum. Meth. A*, 2019. doi:10.1016/j.nima.2018.12.007.
- [54] T. Gys et al. “Latest results from the TORCH R&D Project”. *Nucl. Instrum. Meth. A*, 912:53–56, 2018. doi:10.1016/j.nima.2017.10.034.
- [55] N.H. Brook et al. “Testbeam studies of a TORCH prototype detector”. *Nucl. Instrum. Meth. A*, 908:256–268, 2018. doi:10.1016/j.nima.2018.07.023.
- [56] L. Castillo Garcia. “Testing micro-channel plate detectors for the particle identification upgrade of LHCb”. *Nucl. Instrum. Meth. A*, 695:398–402, 2012. doi:10.1016/j.nima.2011.12.021.
- [57] T. Gys et al. “Performance and lifetime of micro-channel plate tubes for the TORCH detector”. *Nucl. Instrum. Meth. A*, 766:171–172, 2014. doi:10.1016/j.nima.2014.04.020.
- [58] T. M. Conneely et al. “Extended lifetime MCP-PMTs: Characterisation and lifetime measurements of ALD coated microchannel plates, in a sealed photomultiplier tube”. *Nucl. Instrum. Meth. A*, 732:388–391, 2013. doi:10.1016/j.nima.2013.07.023.
- [59] N. Akchurin et al. “First beam tests of prototype silicon modules for the CMS High Granularity Endcap Calorimeter”. *JINST*, 13:P10023, 2018. doi:10.1088/1748-0221/13/10/P10023.
- [60] Loris Martinazzoli. “ECAL - Testbeam Results and Simulations”. *4th Workshop on LHCb Upgrade II*, April 2019. <https://indico.cern.ch/event/790856/>.
- [61] Andreas Schopper. “The LHCb ECAL upgrade(s) and ongoing R&D”. *4th Workshop on LHCb Upgrade II*, April 2019. <https://indico.cern.ch/event/790856/>.
- [62] J. Albrecht et al. “The upgrade of the LHCb trigger system”. *JINST*, 9:C10026, 2014. doi:10.1088/1748-0221/9/10/C10026.
- [63] J. P. Cachemiche et al. “The PCIe-based readout system for the LHCb experiment”. *Topical Workshop on Electronics for Particle Physics*, 2015. doi:10.1088/1748-0221/11/02/P02013.
- [64] LHCb collaboration. “LHCb Trigger and Online Upgrade Technical Design Report”. *CERN-LHCC-2014-016*, 2014. <https://cds.cern.ch/record/1701361>.
- [65] LHCb collaboration. “Trigger Schemes”. <http://lhcb.web.cern.ch/lhcb/speakersbureau/html/TriggerScheme.html>, accessed: 23-06-2019.
- [66] N. Neri et al. “4D fast tracking for experiments at high luminosity LHC”. *JINST*, 11:C11040, 2016. doi:10.1088/1748-0221/11/11/C11040.
- [67] LHCb Collaboration. “LHCb PID Upgrade Technical Design Report”. *LHCB-TDR-014*, 2013. <https://cds.cern.ch/record/1624074>.
- [68] Baszczyk et al. “LHCb RICH upgrade - Production Readiness Review report for the Elementary Cell of the Photo-Detector”. *LHCb-2016-TBD*, 2016. http://lhcb-project-rich-upgrade.web.cern.ch/PRR_ofTheEC/Soi2ECka6KB1u0RKd/EC-RR-v2r5.pdf.
- [69] M. Fiorini. “The upgrade of the LHCb RICH detectors”. *Nucl. Instrum. Meth. A*, 2018. doi:10.1016/j.nima.2018.12.003.
- [70] Hamamatsu Photonics. “Multi-anode photomultiplier tube R13742 Tentative Data Sheet”. 2015. Private communication.
- [71] C. Gotti. “An ASIC for fast single photon counting in the LHCb RICH upgrade”. *JINST*, 12:C03016, 2017. doi:10.1088/1748-0221/12/03/C03016.

- [72] M. Andreotti et al. “Radiation hardness assurance of the CLARO8 front-end chip for the LHCb RICH detector upgrade”. *Nucl. Instrum. Meth. A*, 876:126–128, 2017. doi:10.1016/j.nima.2017.02.035.
- [73] M. Baszczyk et al. “CLARO: an ASIC for high rate single photon counting with multi-anode photomultipliers”. *JINST*, 12:P08019, 2017. doi:10.1088/1748-0221/12/08/P08019.
- [74] P. Carniti et al. “CLARO-CMOS, a very low power ASIC for fast photon counting with pixellated photodetectors”. *JINST*, 7:P11026, 2012. doi:10.1088/1748-0221/7/11/P11026.
- [75] P.J. Garsed and S.A. Wotton. *PDMDB Motherboard*. <https://www.hep.phy.cam.ac.uk/lhcb/PDMDB/PDMDB-Motherboard-v1.2.pdf>, accessed: 20-06-2019.
- [76] P. Moreira et al. “The GBT-SerDes ASIC prototype”. *JINST*, 5(11):C11022, 2010. doi:10.1088/1748-0221/5/11/C11022.
- [77] V.M. Placinta and L.N. Cojocariu. “Radiation Hardness Studies and Evaluation of SRAM-Based FPGAs for High Energy Physics Experiments”. *Proceedings of Science*, 313(TWEPP-17), 2018. doi:10.22323/1.313.0085.
- [78] M.P. Blago and F. Keizer. “High rate tests of the photon detection system for the LHCb RICH Upgrade”. *Nucl. Instrum. Meth. A*, 876:101–103, 2017. doi:10.1016/j.nima.2017.02.014.
- [79] Hamamatsu Photonics. *Picosecond light pulser C10196 with M10306-30 laser head*. <http://www.hamamatsu.com.cn/UserFiles/DownFile/Product/20130915210506017.pdf>, accessed: 20-06-2019.
- [80] Baszczyk et al. “Test of the photon detection system for the LHCb RICH Upgrade in a charged particle beam”. *JINST*, 12:P01012, 2017. doi:10.1088/1748-0221/12/01/P01012.
- [81] L.W. Nagel and D.O. Pederson. “SPICE (Simulation Program with Integrated Circuit Emphasis)”. *Memorandum No. ERL-M382*, University of California, Berkeley, 1973. <http://www2.eecs.berkeley.edu/Pubs/TechRpts/1973/22871.html>.
- [82] M.P. Blago. “Characterisation of the Photon Detection System for the LHCb RICH Detector Upgrade”. *CERN-THESIS-2017-121*, 2017. <http://cds.cern.ch/record/2280735>.
- [83] G. Cavallero et al. “High-rate, SIN and pile-up”. *RICH upgrade Integration and Commissioning meeting*, December 2018. <https://indico.cern.ch/event/781194/>.
- [84] S. Gambetta. “Status of MaPMT testing at CERN”. *LHCb RICH meeting*, May 2019. <https://indico.cern.ch/event/817180/>.
- [85] Xilinx inc. “7 Series FPGAs Data Sheet: Overview”. 2018. https://www.xilinx.com/support/documentation/data_sheets/ds180_7Series_Overview.pdf.
- [86] L. Cadamuro et al. “Characterization of the Hamamatsu R11265-103-M64 multi-anode photomultiplier tube”. *JINST*, 9:P06021, 2014. doi:10.1088/1748-0221/9/06/P06021.
- [87] The Gaudi Project. <http://gaudi.web.cern.ch/gaudi/>, accessed: 20-06-2019.
- [88] H. Braun. “Emittance Diagnostics”. *CERN Accelerator School*, Dourdan, France, 2008. <http://cas.web.cern.ch/sites/cas.web.cern.ch/files/lectures/dourdan-2008/braun-emittance.pdf>.
- [89] Z. Xu. “Simulation study of the ECAL with precision timing for the LHCb Upgrade II”. *4th Workshop on LHCb Upgrade II*, April 2019. <https://indico.cern.ch/event/790856/>.
- [90] Geant4: A simulation toolkit. <https://geant4.web.cern.ch>, accessed: 20-06-2019.

- [91] The GAUSS project. <http://lhcbdoc.web.cern.ch/lhcbdoc/gauss/>, accessed: 20-06-2019.
- [92] Fabio Sauli. *Instrumentation in High Energy Physics*, volume 9. World Scientific Publishing, 1992. doi:10.1142/1356.
- [93] M. Adinolfi et al. “Performance of the LHCb RICH photo-detectors and readout in a system test using charged particles from a 25ns-structured beam”. *Nucl. Instrum. Meth. A*, 603(3):287–293, 2009. doi:10.1016/j.nima.2009.02.011.
- [94] I. H. Malitson. Interspecimen comparison of the refractive index of fused silica. *J. Opt. Soc. Am.*, 55 (10):1205–1209, 1965. doi:10.1364/JOSA.55.001205.
- [95] T. Blake et al. “Quenching the scintillation in CF4 Cherenkov gas radiator”. *Nucl. Instrum. Meth. A*, 791:27–31, 2015. doi:10.1016/j.nima.2015.04.020.
- [96] I. Frank and Ig. Tamm. “Coherent Visible Radiation of Fast Electrons Passing Through Matter”. *Selected Papers*, B.M. Bolotovskii, V.Y. Frenkel, R. Peierls (eds), Springer Berlin:29–35, 1991. doi:10.1007/978-3-642-74626-0_2.
- [97] The BOOLE project. <http://lhcbdoc.web.cern.ch/lhcbdoc/boole/>, accessed: 20-06-2019.
- [98] J. Albrecht et al. “Upgrade trigger and reconstruction strategy: 2017 milestone”. *LHCb-PUB-2018-005*, 2018. <https://cds.cern.ch/record/2310579>.
- [99] R. Forty and O. Schneider. “RICH pattern recognition”. *LHCb/98-040*, 1998. <http://cds.cern.ch/record/684714/files/lhcb-98-040.pdf>.
- [100] LHCb RICH group. “Performance of the LHCb RICH detector at the LHC”. *CERN-LHCb-DP-2012-003*, 2012. <https://cds.cern.ch/record/1495721>.
- [101] The ‘OTMeasurementProvider’ tool. http://lhcb-doxygen.web.cern.ch/lhcb-doxygen/lhcb/latest/d0/d27/_o_t_measurement_provider_8cpp_source.html, accessed: 20-06-2019.
- [102] M.V. Nemallapudi et al. “Single photon time resolution of state of the art SiPMs”. *JINST*, 11: P10016, 2016. doi:10.1088/1748-0221/11/10/P10016.
- [103] E.V. Antamanova et al. “Anode current saturation of ALD-coated Planacon MCP-PMTs”. *JINST*, 13:T09001, 2018. doi:10.1088/1748-0221/13/09/T09001.
- [104] T. Gys. “Micro-channel plates and vacuum detectors”. *Nucl. Instrum. Meth. A*, 787:254–260, 2015. doi:10.1016/j.nima.2014.12.044.
- [105] M. Fiorini et al. “Single-photon imaging detector with $\mathcal{O}(10)$ ps timing and sub-10 μm position resolution”. *JINST*, 13:C12005, 2018. doi:10.1088/1748-0221/13/12/C12005.
- [106] M.J. Minot et al. “Large Area Picosecond Photodetector (LAPPD) - Pilot production and development status”. *Nucl. Instrum. Meth. A*, 2018. doi:10.1016/j.nima.2018.11.137.
- [107] A. Abdesselam et al. “The Barrel Modules of the ATLAS SemiConductor Tracker”. *Nucl. Instrum. Meth. A*, 568:642–671, 2006. doi:10.1016/j.nima.2006.08.036.
- [108] F. Keizer et al. “A compact, high resolution tracker for cosmic ray muon scattering tomography using semiconductor sensors”. *JINST*, 13:P10028, 2018. doi:10.1088/1748-0221/13/10/P10028.
- [109] S. Chatzidakis et al. “Developing a cosmic ray muon sampling capability for muon tomography and monitoring applications”. *Nucl. Instrum. Meth. A*, 804:33–42, 2015. doi:10.1016/j.nima.2015.09.033.
- [110] C. Patrignani et al (Particle Data Group). “Cosmic Rays”. *Chin. Phys. C*, 40:100001, 2017. <http://pdg.lbl.gov/2017/reviews/rpp2017-rev-cosmic-rays.pdf>.

-
- [111] P.K.F. Grieder. *Cosmic Rays at Earth*. Elsevier Science, 2001. doi:10.1016/B978-0-444-50710-5.X5000-3.
- [112] H.A. Bethe. “Molière’s Theory of Multiple Scattering”. *Phys. Rev.*, 89(6):1256–1266, 1953. doi:10.1103/PhysRev.89.1256.
- [113] G.R. Lynch and O.I. Dahl. “Approximations to multiple Coulomb scattering”. *Nucl. Instrum. Meth. B*, 58(1):6–10, 1991. doi:10.1016/0168-583X(91)95671-Y.
- [114] S. Eidelman et al. “Passage of particles through matter”. *Review of particle physics*, B(592), 2004. <http://pdg.lbl.gov/2004/reviews/passagerpp.ps>.
- [115] P. La Rocca et al. “Cosmic Ray Muons as Penetrating Probes to Explore the World around Us”. *IntechOpen*, Cosmic Rays, 2018. doi:10.5772/intechopen.75426.
- [116] L.W. Alvarez et al. “Search for Hidden Chambers in the Pyramids”. *Science*, 167(3919):832–839, 1970. doi:www.jstor.org/stable/1728402.
- [117] M. Menichelli et al. “A scintillating fibres tracker detector for archaeological applications”. *Nucl. Instrum. Methods A*, 572:262–265, 2007. doi:10.1016/j.nima.2006.10.317.
- [118] Y.M. Andreyev et al. “Observation of the Moon Shadow in Cosmic Ray Muons”. *Cosmic Research*, 40:559, 2002. doi:10.1023/A:1021553713199.
- [119] H.K.M. Tanaka et al. “Development of an emulsion imaging system for cosmic-ray muon radiography to explore the internal structure of a volcano”. *Nucl. Instrum. Meth. A*, 575:489–497, 2007. doi:10.1016/j.nima.2007.02.104.
- [120] K. Takamatsu et al. “Cosmic-ray muon radiography for reactor core observation”. *Annals of Nuclear Energy*, 78:166–175, 2015. doi:10.1016/j.anucene.2014.12.017.
- [121] E. Guardincerri et al. “3D Cosmic Ray Muon Tomography from an Underground Tunnel”. *Pure and Applied Geophysics*, 174(5):2133–2141, 2017. doi:10.1007/s00024-017-1526-x.
- [122] K.N. Borozdin et al. “Radiographic imaging with cosmic-ray muons”. *Nature*, 422:277, 2003. doi:10.1038/422277a.
- [123] B. Schwarzschild. “Cosmic-Ray Muons Might Help Thwart Transport of Concealed Fissile Material”. *Physics Today*, 56(19), 2003. doi:10.1063/1.1583524.
- [124] W.C. Friedhorsky et al. “Detection of high-Z objects using multiple scattering of cosmic ray muons”. *Review of Scientific Instruments*, 74:4294, 2003. doi:10.1063/1.1606536.
- [125] L.J. Schultz et al. “Image reconstruction and material Z discrimination via cosmic ray muon radiography”. *Nucl. Instrum. Meth. A*, 519:687–694, 2004. doi:10.1016/j.nima.2003.11.035.
- [126] Decision Sciences Corp. *Multi-Mode Passive Detection System*. <https://decisionssciences.com>, accessed: 20-06-2019.
- [127] P. Baesso et al. “Towards a RPC-based muon tomography system for cargo containers”. *JINST*, 9: C10041, 2014. doi:10.1088/1748-0221/9/10/C10041.
- [128] K. Gnanvo et al. “Imaging of high-Z material for nuclear contraband detection with a minimal prototype of a muon tomography station based on GEM detectors”. *Nucl. Instrum. Meth. A*, 652: 16–21, 2011. doi:10.1016/j.nima.2011.01.163.
- [129] V. Anghel et al. “A plastic scintillator-based muon tomography system with an integrated muon spectrometer”. *Nucl. Instrum. Meth. A*, 798:P10041, 2015. doi:10.1016/j.nima.2015.06.054.

- [130] F. Riggi et al. “The Muon Portal Project: Commissioning of the full detector and first results”. *Nucl. Instrum. Meth. A*, 912:16–19, 2018. doi:10.1016/j.nima.2017.10.006.
- [131] C. Thomay et al. “Passive 3D imaging of nuclear waste containers with Muon Scattering Imaging”. *JINST*, 11:P03008, 2016. doi:10.1088/1748-0221/11/03/P03008.
- [132] D. Poulson et al. “Application of muon tomography to fuel cask monitoring”. *Philosophical transactions of the Royal Society A*, 377:2137, 2018. doi:10.1098/rsta.2018.0052.
- [133] M. Furlan et al. “Application of Muon Tomography to Detect Radioactive Sources Hidden in Scrap Metal Containers”. *IEEE Transactions on Nuclear Science*, 61(4):2204–2209, 2014. doi:10.1109/TNS.2014.2321116.
- [134] P. Zanuttigh et al. “Study of the capability of muon tomography to map the material composition inside a blast furnace (Mu-Blast)”. *EU publications*, 2019. doi:10.2777/24858.
- [135] ATLAS collaboration. “ABCD3T ASIC V1.2”. 2000. https://twiki.cern.ch/twiki/pub/Main/AtlasSiliconRodGroup/abcd3t_spec_v1.2.pdf.
- [136] F. Campabadal et al. “Design and performance of the ABCD3TA ASIC for readout of silicon strip detectors in the ATLAS semiconductor tracker”. *Nucl. Instrum. Meth. A*, 552:292–328, 2005. doi:10.1016/j.nima.2005.07.002.
- [137] Y. Unno et al. “Cu/Polyimide barrel hybrids”. FDR 5.3, ATLAS project, 2001. <http://jsdhp1.kek.jp/~unno/notes/SCT/FDR/k4bhspec.pdf>.
- [138] Maxim Integrated. *400Mbps, Low-Jitter, Low-Noise LVDS Repeater in an SC70 Package*. <https://datasheets.maximintegrated.com/en/ds/MAX9180.pdf>, accessed: 20-06-2019.
- [139] Analog Devices Inc. *ADP2303, 3A, 700KHz NonSynchronous Step-Down DC-DC Switching Regulator*. <https://www.analog.com/en/products/adp2303.html>, accessed: 20-06-2019.
- [140] B. Hommels and S. Sigurdsson. *Temperature Interlock Unit*. <https://www.hep.phy.cam.ac.uk/~hommels/interlock>, accessed: 20-06-2019.
- [141] FASER collaboration. “Technical Proposal for FASER: ForwArd Search ExpeRiment at the LHC”. *CERN-LHCC-2018-036*, 2018. <https://arxiv.org/abs/1812.09139>.
- [142] V.A. Mitsou. “Production Performance of the ATLAS Semiconductor Tracker Readout System”. *IEEE Transactions on Nuclear Science*, 53(3):729–734, 2006. doi:10.1109/TNS.2006.874727.
- [143] Kangabox Tower (black). <https://www.kangaboxuk.com>, accessed: 20-06-2019.
- [144] G.L. Squires. *Practical Physics*. Cambridge University Press, 4th edition, 2012. doi:10.1017/CBO9781139164498.
- [145] M. Stapleton et al. “Angle Statistics Reconstruction: a robust reconstruction algorithm for Muon Scattering Tomography”. *JINST*, 9:P11019, 2014. doi:10.1088/1748-0221/9/11/P11019.
- [146] J. Burns et al. “A drift chamber tracking system for muon scattering tomography applications”. *JINST*, 10:P10041, 2015. doi:10.1088/1748-0221/10/10/P10041.
- [147] G. Wang et al. “Bayesian Image Reconstruction for Improving Detection Performance of Muon Tomography”. *IEEE Transactions on Image Processing*, 18(5):1080–1089, 2009. doi:10.1109/TIP.2009.2014423.
- [148] G. Wang et al. “Statistical Image Reconstruction for Muon Tomography Using a Gaussian Scale Mixture Model”. *IEEE Transactions on Nuclear Science*, 56(4):2480–2486, 2009. doi:10.1109/TNS.2009.2023518.



Illustration: Astrid Keizer-Huijsing

Copyright © 2019 Floris Keizer
All rights reserved

The work in part II of this dissertation is funded, in part,
by AWE, Aldermaston, Reading RG7 4PR, UK.