SCUOLA DI DOTTORATO
UNIVERSITÀ DEGLI STUDI DI MILANO-BICOCCA

Department of Physics

Ph. D. program in Physics and Astronomy, XXXVIII cycle
Curriculum of Applied Physics and Electronics

Studies towards a sub-100 ps time-resolved LHCb RICH detector and a novel method for the primary vertex time estimation during LHC Run 4

Lorenzo Malentacca
Registration number: 896622

Tutor: Gianluigi Pessina

Supervisor: Marta Calvi

CERN Supervisor: Floris Keizer

Coordinator: Stefano Ragazzi

Academic Year 2024/2025

Abstract

The increase in instantaneous luminosity expected for the High-Luminosity LHC (HL-LHC) will significantly rise track multiplicity and hit occupancy in the LHCb detector. To preserve the excellent particle-identification (PID) performance of the LHCb Ring-Imaging Cherenkov (RICH) detector, the integration of sub-nanosecond photon hit-time information has been identified as a key strategy for the detector upgrade. This thesis presents studies supporting the RICH detector upgrade programme, covering both the Long Shutdown 3 (LS3) Enhancements and the Upgrade II phase. The work focuses on implementing and validating sub-nanosecond timing capability in the RICH system through detailed photodetector characterisation, hardware design and software development. In the first part of the thesis, analogue and digital measurements were performed on Multi-anode Photomultiplier Tubes (MAPMTs), the photodetectors used in the current RICH detector, to characterise their single-photon response. A single-photon time resolution (SPTR) of 145 ps was measured for the 1-inch MAPMT and 125 ps for the 2-inch MAPMT. These results define the intrinsic timing performance of the sensors and serve as a reference for the development of a novel time-resolved optoelectronic readout chain for the RICH detector. Based on these measurements, a novel method for the estimation of the primary vertex time (PV t_0) was developed using only RICH reconstruction information. The PV t_0 provides the event time reference required to apply a software time gate on RICH photon hit times, which has been demonstrated to improve the detector PID performance during LHC Run 4. A precision of 85 ps (FWHM) was achieved for 94 % of reconstructed PV in a sample of ten thousand simulated bunch-crossing events, marking a first-time measurement for the LHCb experiment. The final part of the thesis describes the design and testing of a flexible PCB (flex-PCB) for the readout of Silicon Photomultiplier (SiPM) arrays, the main photodetector candidate for the RICH Upgrade II detector. The main challenge with SiPMs is their high dark-count rate (DCR), which increases with radiation. DCR can be mitigated by operating SiPMs at cryogenic temperatures (~100 K). The flex-PCB features ~15 cm analogue traces, enabling SiPM operation at cryogenic temperatures while keeping the readout electronics at room temperature. Laboratory measurements and particle-beam tests at room temperature have demonstrated stable signal integrity and sub-100 ps single-photon time resolution, confirming the suitability of this flex-PCB design for future RICH detector upgrades. These studies are important for the ongoing and future developments of the LHCb RICH detectors as well as the integration of sub-nanosecond timing information into the detector readout system and the event reconstruction.

Contents

Introduction

In this thesis, I present studies supporting the upgrade of the Ring-Imaging Cherenkov (RICH) detectors of the LHCb experiment at the Large Hadron Collider (LHC). The work has been carried out in preparation for the LHCb Long Shutdown 3 (LS3) Enhancements (2026–2030) and the long-term Upgrade II programme (2034–2035). It focuses on the integration of sub-100 ps time information in the RICH detector, an essential element to address the increased detector occupancy expected during High-Luminosity LHC (Run 5, 2036 onward).

The timing information will be integrated in the RICH detector during the LS3 Enhancements by coupling the current photodetectors, the Multi-anode Photomultiplier Tubes (MAPMTs), with novel readout electronics based on the 25 ps time-resolved FastRICH ASIC. During LHC Run 4, this timing information will be used to improve the detector particle identification performance by including the photon hit time in the RICH reconstruction algorithms. The integration of such precise time information requires a detailed understanding of the intrinsic timing performance of the MAPMTs and their coupling with the new time-resolved electronics. The first part of this thesis addresses these aspects through analogue and digital characterisation of the MAPMT single-photon response.

The RICH detector will be the only LHCb sub-detector with sub-nanosecond timing precision during LHC Run 4, making it essential to develop reconstruction algorithms to fully exploit this capability. It has been demonstrated that a software time gate around the predicted time of arrival of the photons at the RICH detector reduces the combinatorial background in the reconstruction, improving the detector PID performance. The primary vertex time (PV t_0) is a key element in the prediction of the photon time of arrival. I therefore developed a novel method for the estimation of the PV t_0 using the RICH reconstruction only in LHC Run 4, which represents a first-time measurements for the LHCb experiment.

For the Upgrade II phase, the RICH detectors will rely on new photodetector technologies with improved timing performance and spatial granularity to cope with the increased photon occupancy expected for LHC Run 5. Silicon Photomultipliers (SiPMs) are a promising candidate due to their excellent single-photon detection efficiency, time resolution and fine granularity. Their main challenge is the high dark-count (DCR) rate, which increased with radiation damage. The SiPM DCR can be mitigated by operating the sensors at cryogenic temperatures. I designed a flexible PCB (flex-PCB) to transmit analogue signals over centimetre-scale distances between an SiPM array at cryogenic temperatures and the readout electronics at room temperatures. I also characterised the flex-PCB under pulsed-laser illumination and in beam test at the CERN SPS North Area facility to measure the SiPM array signal integrity and timing performance.

The studies and the results presented in this thesis serves as a reference for the current and future development of the LHCb RICH detectors, particularly for the integration of sub-nanosecond timing information into the detector readout system and reconstruction.

Chapter 1

The upgrades of the LHCb experiment at the LHC

This chapter presents the Large Hadron Collider beauty (LHCb) experiment at the Large Hadron Collider (LHC), with emphasis on its upgrade programme. Section 1.1 introduces the European Organisation for Nuclear Research (CERN) and its main facility, the LHC. Section 1.2 summarises the key physics concepts studied by LHCb and provides an overview of the detector. Section 1.3 outlines the upgrades of the LHCb sub-detectors, where the use of a precise hit-time information and single-photon detection technologies are of particular interest to the Ring-Imaging Cherenkov (RICH). Section 1.4 describes the main ongoing R&D activities for the RICH detector upgrade.

1.1 CERN and the Large Hadron Collider

The Large Hadron Collider (LHC) is the world's largest and most powerful particle accelerator [1]. It is located at CERN (the European Organization for Nuclear Research), beneath the Franco–Swiss border near Geneva, and represents a landmark achievement in the field of high-energy physics. The machine spans 26.7 km in circumference and lies between 50 and 175 metres underground. Housed in the tunnel originally constructed for the LEP (Large Electron–Positron Collider), the LHC is designed to probe the fundamental constituents of matter by colliding protons or heavy ions at unprecedented energies.

Over the past seven decades, CERN has led the field of high-energy physics, producing some of the most significant experimental discoveries and technological innovations in the discipline. Among its many scientific milestones are the first observation of anti-nuclei in 1965 [2], the detection of neutral currents in 1974 [3], and the Nobel-Prize-winning discovery of the W and Z bosons in 1983 [4, 5]. The LEP collider [6] allowed precision electroweak measurements, including the determination of the number of light neutrino families in 1989 [7]. Further landmark results include the synthesis of antihydrogen atoms in 1995 [8], the discovery of pentaquark states in 2015 [9], and the observation of the Higgs boson in 2012, which confirmed mechanism of electroweak symmetry breaking predicted by the Standard Model and earned the Nobel Prize for Peter Higgs and François Englert [10, 11].

CERN has also been a driver of technological innovation, with several technologies originally developed for particle physics now widely used across other fields. One of the most notable examples is the World Wide Web, first proposed and implemented at CERN to facilitate information sharing among scientists [12]. Moreover, advances in accelerator, magnet and detector technology have had a profound impact on medical applications, including cancer therapy with particle beams [13] and medical imaging techniques such as PET [14] and MRI [15]. These cross-disciplinary contributions underscore CERN's broader role as a hub for cutting-edge research and innovation.

To maintain CERN's leadership in particle physics well into the next decades and to continue pushing the frontiers of fundamental research, an ambitious upgrade programme has been launched for the LHC and its experiments.

1.1.1 The LHC and its main experiments

The LHC is the final stage of a complex accelerator chain, illustrated in Figure 1.1. The injection chain begins with protons extracted from hydrogen atoms or lead ions from heated samples. These particles are initially accelerated by LINAC4, a linear accelerator. The beam then passes through a succession of synchrotrons: the Proton Synchrotron Booster (PSB), the Proton Synchrotron (PS), and the Super Proton Synchrotron (SPS), where energies reach up to 450 GeV. The final stage of acceleration occurs in the LHC, where particles are accelerated up to multi-TeV energies. Looking further ahead, CERN is developing plans for the Future Circular Collider (FCC), a proposed 90 km tunnel that would accommodate both a high-luminosity electron–positron collider (FCC-ee) and, eventually, a 100 TeV proton–proton collider (FCC-hh). The FCC is currently in the feasibility study phase, with the aim of delivering a comprehensive strategy in time for the next update of the European Strategy for Particle Physics [16, 17].

In the LHC, the particles are boosted by radiofrequency (RF) cavities operating at 400 MHz. A total of 1232 superconducting dipole magnets bend the particle beam in a circular trajectory, while 392 quadrupole magnets focus the beam. Each magnet is cooled with liquid helium and produces fields up to 8 T [1]. The particle beams are structured into bunches circulating in opposite directions within two separate beam pipes. These bunches, each containing over 10^{11} protons, are spaced by 25 ns and brought into collision at four interaction points. These collisions occur at a design rate of up to 40 MHz and at a peak centre-of-mass energy of 13.6 TeV, as achieved during 2024 operation.

At the heart of the LHC's scientific mission are four main experiments:

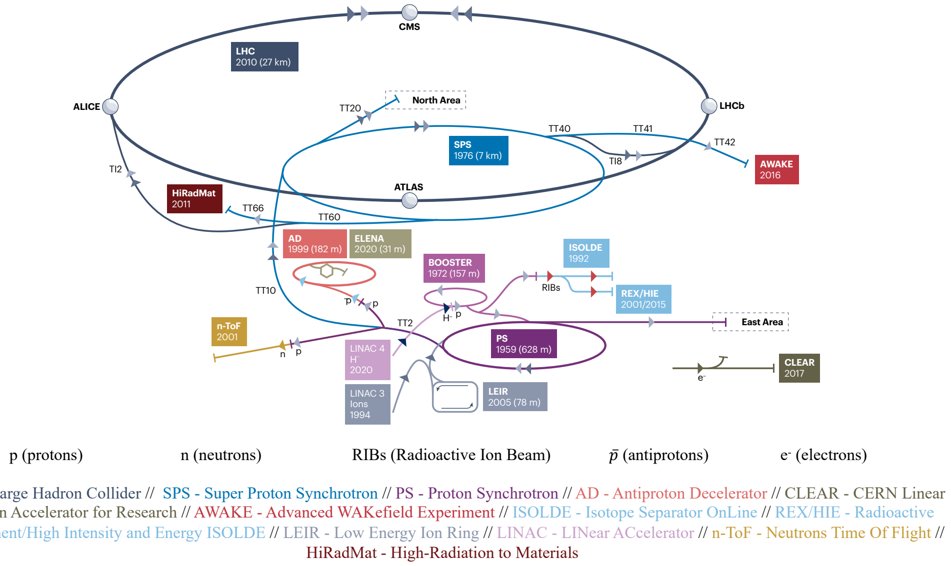

Figure 1.1: Schematic view of the CERN accelerator complex, showing the individual stages of acceleration and main experiments. Edited from [22].

- **ATLAS** [18] and **CMS** [19], two general-purpose detectors located at diametrically opposite points of the ring, were designed to explore a broad range of phenomena, including the discovery of the Higgs boson and searches for new physics beyond the Standard Model. ATLAS is the largest detector ever built, while CMS is the heaviest due to its dense superconducting solenoid.
- **ALICE** [20] focuses on the study of strongly interacting matter under extreme conditions, particularly the quark–gluon plasma created in lead–lead collisions.
- **LHCb** [21], the experiment most relevant to this thesis, is a forward spectrometer, optimised for the study of beauty and charm hadron decays. LHCb investigates the matter–antimatter asymmetry in the universe by studying rare charge-parity (CP) symmetry-violating decays.

1.1.2 The LHC and LHCb future programme

This section provides a summary of the future programme of the LHC and the LHCb experiments, using the timeline in Figure 1.2. The LHC began delivering data in 2010 with an initial proton–proton collision energy of 7 TeV (LHC Run 1), later increased to 8 TeV, collecting over 25 fb^{-1}

of data across multiple experiments. LHC Run 2 (2015–2018) operated at 13 TeV and saw the LHC reach its nominal luminosity design goals. The subsequent long shutdown (LS2, 2019–2021) prepared the accelerator and its experiments for higher data rates, higher pile-up, and extended physics capabilities.

The current operation phase, LHC Run 3, began in 2022 and features proton–proton collisions at a centre-of-mass energy of 13.6 TeV. After LHC Run 3, the accelerator will enter the Long Shutdown 3 (LS3) phase, scheduled between 2026 and 2030, during which the High-Luminosity LHC (HL–LHC) will be installed. The HL–LHC aims to reach peak luminosities of up to $5 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ and to deliver an integrated luminosity of approximately $3000\,\mathrm{fb}^{-1}$ by the end of its operational lifetime [23]. This requires extensive upgrades to the injector chain, new beam optics and manipulation technologies, and substantial detector upgrades to cope with increased radiation levels, pile-up, and data throughput [24].

During LS3, the ATLAS and CMS experiments will undergo major upgrades to cope with the increased event multiplicities and detector occupancies expected in the HL–LHC. In order to spread the available resources to experimental updates, major upgrades for LHCb and ALICE are planned for Long-Shutdown 4 (LS4, 2034-2035). During LHC Run 4, the luminosity at LHCb will be actively levelled to match LHC Run 3 conditions, limiting event pile-up in the detector and radiation damage near the interaction point. Luminosity levelling is achieved by dynamically adjusting beam parameters such as transverse separation, crossing angle, β^{*}, and crab-cavity voltage [25].

As described in this chapter, the LHCb experiment will be completely revised (LHCb Upgrade II) in preparation for the luminosity increase of a factor of 7.5 from LHC Run 5 onwards. A total integrated luminosity of around $300\,\mathrm{fb}^{-1}$ is expected by the end of LHC Run 5. The increase in luminosity leads to a substantial rise in primary vertex multiplicity, track density, detector occupancy, and radiation fluence. To cope with these challenges, LHCb sub-detectors will implement detector hit time information and higher detector granularity in Upgrade II to reduce combinatorial background and improve reconstruction in high-occupancy regions. LHCb Upgrade II is a major task. The limited duration of LS4 poses a challenge for installation and commissioning. Therefore, some LHCb sub-detectors will profit of the three years of LS3 to anticipate part of their upgrade programmes.

1.2 The LHCb experiment

This section introduces the LHCb experiment. Section 1.2.1 provides a brief overview of the physics case behind the LHCb experiment and its upgrades. A general overview of the LHCb detector and its structural changes for LHCb Upgrade II is given Section 1.2.2.

Figure 1.2: Timeline of the LHC and LHCb experiment. Although the HL-LHC will start operation in 2030, the luminosity in the LHCb experiment will not be increased until LHC Run 5. LHCb Upgrade II will be installed during LS4, while some sub-detectors will introduce some enhancements already during LS3. Edited from [26].

1.2.1 The LHCb physics case

One of the most profound open questions in modern physics is the observed imbalance between matter and antimatter in the universe. This disparity has been the focus of several generations of physicists since the first observation of antimatter in 1932 [27]. This imbalance, known as the baryon asymmetry, implies that matter and antimatter do not behave identically under all interactions. The Standard Model (SM) allows for violations of charge–parity (CP) symmetry through complex phases in the weak interaction [28]. However, the magnitude of CP violation predicted by the SM is too small to explain the observed asymmetry. Additional sources of CP violation beyond the SM are needed. Understanding the origin of this asymmetry remains one of the primary goals of flavour physics, and a central motivation for the LHCb experiment.

The unitary Cabibbo–Kobayashi–Maskawa (CKM) matrix describes quark flavour transitions in weak interactions [29]:

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

The CKM matrix connects the mass eigenstates for the quarks $(d,\, s,\, b)$ with their weak eigenstates $(d',\, s',\, b')$. A widely used parametrisation is the Wolfenstein [30], which expands each element as a power series in the small parameter $\lambda \approx |V_{us}| \simeq 0.22$, thereby reflecting the hierarchical structure of the CKM matrix elements:

$$V_{\mathrm{CKM}} = \begin{pmatrix} 1 - \frac{\lambda^2}{2} & \lambda & A\lambda^3(\rho - i\eta) \\ -\lambda & 1 - \frac{\lambda^2}{2} & A\lambda^2 \\ A\lambda^3(1 - \rho - i\eta) & -A\lambda^2 & 1 \end{pmatrix} + \mathcal{O}(\lambda^4) \tag{1.2}$$

The four Wolfenstein parameters A, ρ, λ, and η determine all flavour-changing transitions in the quark sector. A non-zero value of η introduces CP violation. The CKM matrix elements are not defined by the SM but are constrained by the unitarity condition. The unitarity condition requires that $\sum_i V_{ij} V_{ik}^* = 0$ for $j \neq k$, from which six constraints arise. Each constraint can be represented by a closed triangle in the complex plane. Of particular interest is the triangle resulting from $j = d$ and $k = b$, often referred to as the Unitarity Triangle:

$$V_{ud} V_{ub}^* + V_{cd} V_{cb}^* + V_{td} V_{tb}^* = 0 \tag{1.3}$$

The Unitarity Triangle applies directly to beauty decays and is shown in Figure 1.3. The three angles of the Unitarity Triangle are defined as:

$$\alpha = \arg\left(-\frac{V_{td} V_{tb}^*}{V_{ud} V_{ub}^*}\right), \quad \beta = \arg\left(-\frac{V_{cd} V_{cb}^*}{V_{td} V_{tb}^*}\right), \quad \gamma = \arg\left(-\frac{V_{ud} V_{ub}^*}{V_{cd} V_{cb}^*}\right) \tag{1.4}$$

and the sides can be written as:

$$R_u = \left|\frac{V_{ud} V_{ub}^*}{V_{cd} V_{cb}^*}\right| = \sqrt{\rho^2 + \eta^2}, \quad R_t = \left|\frac{V_{td} V_{tb}^*}{V_{cd} V_{cb}^*}\right| = \sqrt{(1-\rho)^2 + \eta^2} \tag{1.5}$$

Measurements of the Unitarity Triangle are therefore directly related to the CKM matrix and CP violation. Any deviation from the expected values of the Unitarity Triangle parameters would indicate physics beyond the SM. Decays of beauty hadrons, in particular B mesons, provide a sensitive probe for potential deviations from the SM. The relatively long lifetime of B mesons offers a clear experimental signature, enabling efficient selection of signal decays and suppression of background.

The LHCb detector is specifically optimised for the study of the decays of heavy beauty and charm hadrons produced in the LHC high-energy collisions. It builds on results from previous B-physics experiments, such as BaBar and Belle [31, 32]. Using the data collected during LHC Run 1, Run 2 and Run 3, the LHCb experiments has already placed some of the most stringent constraints on the parameters of the CKM matrix. However, the obtained results remains statistically limited, demonstrating that further reduction in uncertainty can be obtained with the LHCb Upgrade II detector.

Beyond CKM studies, the LHCb experiment plays a crucial role in testing lepton flavour universality, a fundamental prediction of the SM. The prediction states that all leptons have identical couplings to gauge bosons (photon and W and Z bosons) [34]. The LHCb experiment can test the lepton universality by measuring ratios of purely leptonic or semi-leptonic decays that involve the same quark transition but different lepton flavours (e.g. R_K, R_{K^*}, $R_{D^{(*)}}$). Given the wide spectrum of beauty hadrons produced in the LHC collision, a large variety of lepton universality tests can be performed. Any deviation from the unity in the ratio is a clear sign of new physics beyond the SM. Another key strength of the LHCb experiment is the study of rare and very rare decays.

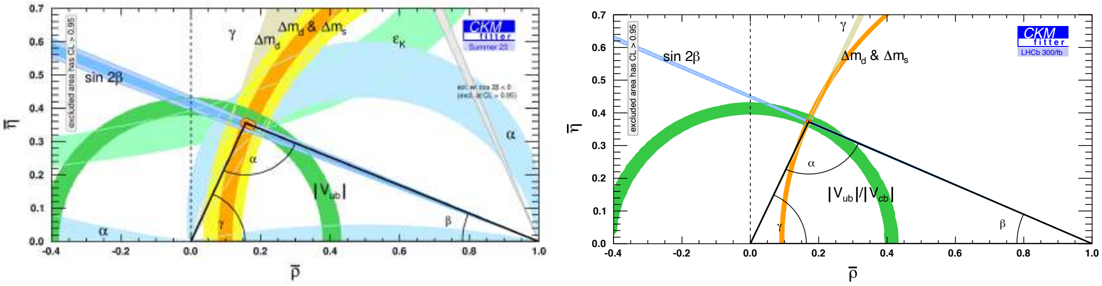

Figure 1.3: Constraints from the dominant CKM observables to the apex of the Unitarity Triangle ($\bar{\rho}$, $\bar{\eta}$) with global inputs as of 2023 (left) and the projection after LHCb Upgrade II measurements in LHC Run 5 with 300 fb^{-1} of data (right) [33].

These processes are mediated by flavour-changing neutral currents, which are forbidden at tree level in the SM and proceed only via higher-order loops. Because of this, even small contributions from new particles can affect their rates or angular distributions. Therefore, the LHCb precision measurements of these rare decays can reveal indirect signs of physics beyond the SM. The LHCb experiment has also contributed to the discovery and study of exotic hadrons. Observations of states such as tetraquarks and pentaquarks have provided evidence for hadronic structures beyond the conventional meson and baryon classifications [9].

Despite the increasing experimental precision, all the results remain consistent with a sole CP violation predicted by the SM, which is insufficient to explain the matter–antimatter asymmetry of the universe [35]. This discrepancy motivates LHCb long-term programme to continue probing the SM with increased precision. The LHCb flavour physics programme will greatly benefit from the increase of statistics obtained during LHC Run 5 [36]. Figure 1.3 (right) shows the expected gain on the CKM parameters estimation precision. The precision bands at 300 fb^{-1} set tight constraints on the SM predictions. In addition, similar improvements are expected in lepton universality tests and searches for rare decays and exotic hadrons. These predictions are based on the assumption that the LHCb Upgrade II detector will maintain the excellent detector performance of LHCb under HL-LHC conditions. The increased luminosity results in a higher track multiplicity and detector occupancy, which is a major challenge for the detector performance. The detector upgrades foreseen for LHCb Upgrade II to cope with these harsh conditions are described in Section 1.3.

1.2.2 Overview of the Detector

The LHCb detector is a single-arm forward spectrometer with a length of approximately 20 m and a height of 7 m. It covers an angular region between approximately 10 and 300 mrad in the horizontal plane and 10 to 250 mrad in the vertical plane, corresponding to a pseudorapidity range of $2 < \eta < 5$ [37], defined as

$$\eta = -\ln\left[\tan\left(\frac{\theta}{2}\right)\right], \tag{1.6}$$

The LHCb detector is designed to detect beauty hadrons boosted in the forward direction due to their relatively large mass.

A schematic of the current LHCb detector is shown in Figure 1.4. The LHCb detector is composed of several sub-detectors, each with a specific function, such as tracking, particle identification and energy measurement. The tracking system is composed by the VErtex LOcator (VELO), the Upstream Tracker (UT) and the Scintillating Fibre tracker (SciFi). The VELO is the closest detector to the particle interaction region. It is a silicon pixel tracker that reconstructs the primary vertices with micrometre precision and the initial segment of charged particle tracks. VELO tracks are combined with the segments from the Upstream Tracker (UT) and the Scintillating Fibre tracker (SciFi) to form the charged particle track vector. The momentum of the charged particles is measured using the LHCb dipole magnet, located between the UT and the SciFi detectors. The magnet creates a field with an average bending power of 4 Tm, allowing momentum measurement with a resolution of about 0.5 %–1 %. Charged hadron identification is provided by two Ring Imaging Cherenkov (RICH) detectors (RICH 1 and 2). RICH 1 is located upstream, between the VELO and the magnet, and operates in a high track-density environment. RICH 2 is positioned downstream of the SciFi tracker. Electromagnetic and hadronic calorimeters (ECAL and HCAL) are placed downstream of RICH 2. They measure the energy of photons, electrons, and hadrons. Behind the calorimeters are the muon stations. Most particles are absorbed before reaching this stage, making the muon system highly effective at identifying muons from signal decays.

The LHCb Upgrade II detector layout is shown in Figure 1.5. The VELO will introduce precise hit time information, with a resolution of 50 ps. This capability will reduce combinatorial background from multiple vertices by assigning tracks to the correct interaction in time. The upstream tracking systems will be redesigned. The inner region of the SciFi tracker will be replaced by a silicon tracker. Together, the new silicon region and the outer SciFi form the Mighty Tracker (MT). The current UT will be replaced by the Upstream Tracker (UP), with higher granularity. UP will play a critical role in reducing incorrect matches of the VELO and MT track segments. In addition, scintillating fibre trackers are foreseen on the side of the magnets, hence Magnet Stations, to track low-momentum ($p < 5\,\mathrm{GeV}/c$) charged particles exiting the detector. The Time of Internally Reflected CHerenkov Light (TORCH) detector may be installed between the MT and RICH 2. TORCH extends charged hadron identification in the low momentum region (1–10 GeV/c), where standard RICH reconstruction is less effective. The electromagnetic calorimeter will be replaced by the PicoCal, a system with finer granularity and picosecond-order time resolution. The hadronic calorimeter will be removed and replaced by an iron wall thicker than 1 m, positioned in front of the muon stations.

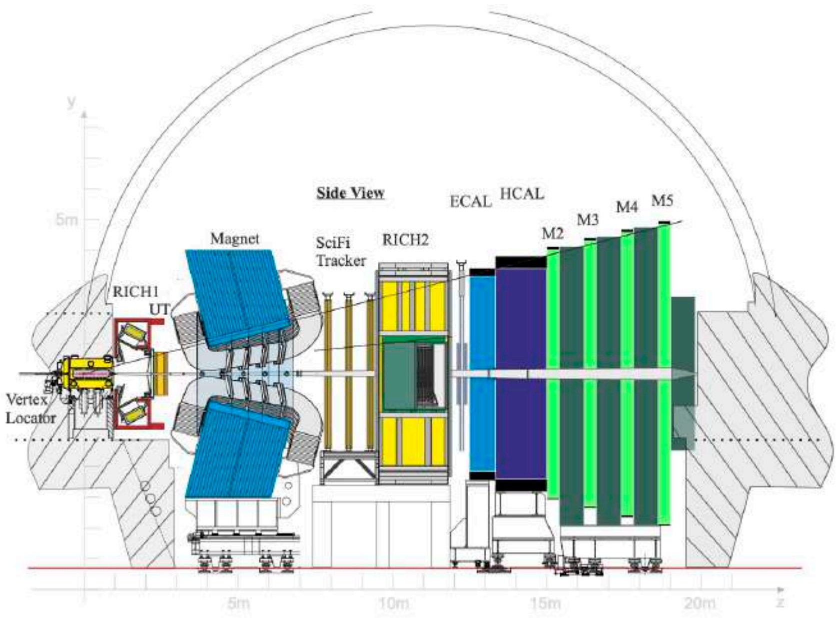

Figure 1.4: Schematic view of the current LHCb detector. Source: [38].

The LHCb sub-detectors are described in detail in Section 1.3, with focus on their main design characteristics and the modifications for Upgrade II

1.3 The LHCb detector upgrades

This section describes the LHCb sub-detectors together with their upgrade plans for Upgrade II. Section 1.3.1 presents the tracking detectors, Section 1.3.2 the particle identification system, and Section 1.3.3 the trigger and data acquisition. For each sub-detector, the current design is summarised, followed by the improvements required to cope with the HL–LHC challenges.

At the LHC Run 5 instantaneous luminosity of $1.5 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$, the multiplicity of primary vertices, track density, detector occupancy and radiation fluence will increase substantially. To address these challenges, two common strategies are adopted across the sub-detectors: the introduction of timing information to reduce combinatorial background, and higher detector granularity to mitigate the increase of detector occupancy. These measures will maintain (and improve) the overall reconstruction performance in the harsh environment of the HL-LHC. A baseline design of each sub-detector has been proposed for LHCb Upgrade II. In addition, alternative scenarios have been evaluated that assume a reduced instantaneous luminosity in LHC Run 5 or accept compromises in detector performance in order to lower cost [33].

1.3.1 Tracking system upgrades

The LHCb tracking system provides precise measurements of charged particle trajectories, enabling the reconstruction of primary vertices (PVs), displaced secondary vertices from long-lived hadrons, and momentum determination through deflection in the dipole magnet. In the current LHCb detector, the system consists of the Vertex Locator (VELO) surrounding the interaction point, the Upstream Tracker (UT) placed before the magnet, and the Scintillating Fibre (SciFi) tracker downstream of the magnet. The schematic of the LHCb tracking system is shown in Figure 1.6. Tracks in LHCb are classified according to the detectors they traverse. Long tracks traverse the full spectrometer and generate hits in the VELO, UT and SciFi. They are the most important for physics analysis. Downstream tracks originate from the decays of long-lived particles, such as K^0_S and Λ hadrons, downstream of the VELO, and have hits in the UT and SciFi only. Finally, upstream tracks traverse the VELO and UT, and enhance the detector acceptance at low momentum.

For LHCb Upgrade II, the increased luminosity and harsher radiation environment of HL-LHC will require significant modifications, including improved radiation hardness, enhanced granularity, and the introduction of timing capabilities. The VELO will introduce hit time information with 50 ps precision, while UT and SciFi detectors will be replaced by detectors with higher granularity: the Upstream Precision Tracker (UP) and the Mighty Tracker (MT), respectively. An additional tracking element, the magnet side stations, will extend coverage to low-momentum tracks. The

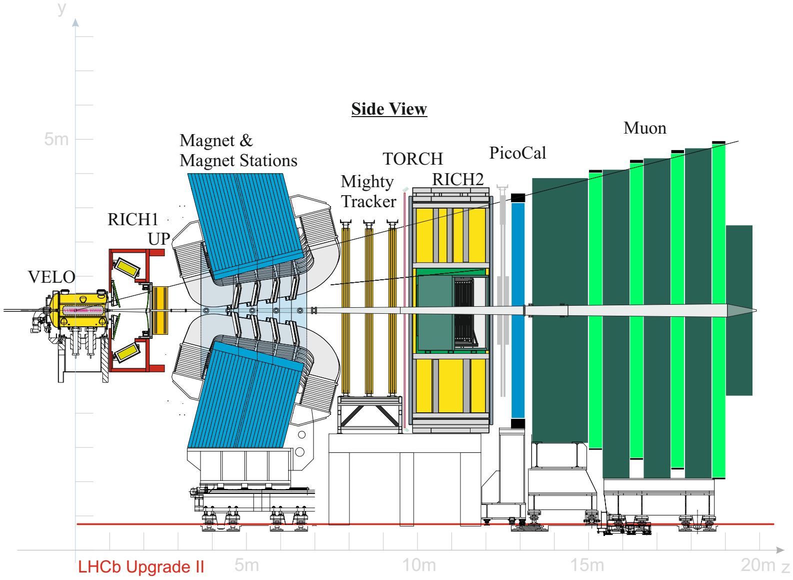

Figure 1.5: Schematic view of the proposed LHCb Upgrade II detector. Source: [33].

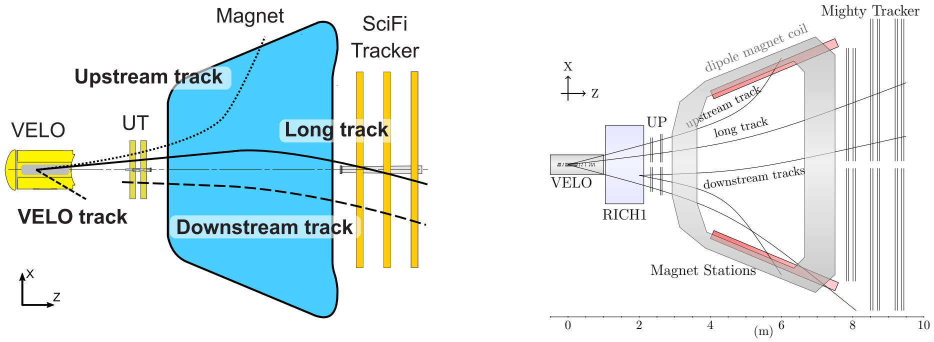

Figure 1.6: Schematic of the tracking system in the current LHCb detector (left) [39] and in the LHCb Upgrade II (right) [33]. A sketch of the different types of track is shown in both figures.

schematic of the LHCb Upgrade II tracking system is shown in Figure 1.6 (right).

Vertex Locator (VELO)

The VELO is the innermost sub-detector of LHCb, providing precise measurements of charged particle trajectories close to the interaction point. Its main functions are the reconstruction of primary vertices, the identification of secondary decay vertices, and the measurement of track impact parameters with a resolution of a few micrometres. The detector consists of two retractable halves, each equipped with silicon pixel sensors arranged to form a diamond-shaped acceptance around the beam [40, 41], as shown in Figure 1.7. A 150 µm aluminium RF foil separates the sensors from the beam vacuum. During operation, the halves close to a distance of 5.11 mm from the beam. The sensors have a pixel pitch of $55 \times 55\,\mu\mathrm{m}^2$ and are read out by the VeloPix ASIC, derived from the Timepix3 ASIC design [42]. The system operates at the full LHC bunch crossing rate of 40 MHz without pixel hit time information. The non-uniform radiation field at the VELO location results in an integrated fluence of up to $8 \times 10^{15}\,\mathrm{n_{eq}/cm^2}$ during its lifetime [40].

The HL-LHC running conditions will increase the instantaneous luminosity and pile-up by an order of magnitude compared to LHC Run 3. For the VELO detector this implies unprecedented data rates, pattern recognition challenges, and a radiation fluence up to $5 \times 10^{16}\,\mathrm{n_{eq}/cm^2}$. To cope with these conditions while maintaining current tracking performance, the VELO will introduce hit time information, forming a 4D tracking system, and possibly reduce the pixel size. The fundamental role of timing in suppressing pile-up is illustrated in Figure 1.8 (left), where an LHC Run 5 event reconstructed in the VELO is shown, with and without a time selection of tracks. By adding the time dimension, vertices can be separated in both space and time. This reduces combinatorial background and preserves the PV reconstruction efficiency of LHC Run 3 under HL-LHC conditions, as shown in Figure 1.8 (right). Timing information will also reduce the CPU usage in the online selection, improve track matching with UP and MT and provide a PV timestamp to other sub-detectors that use time information.

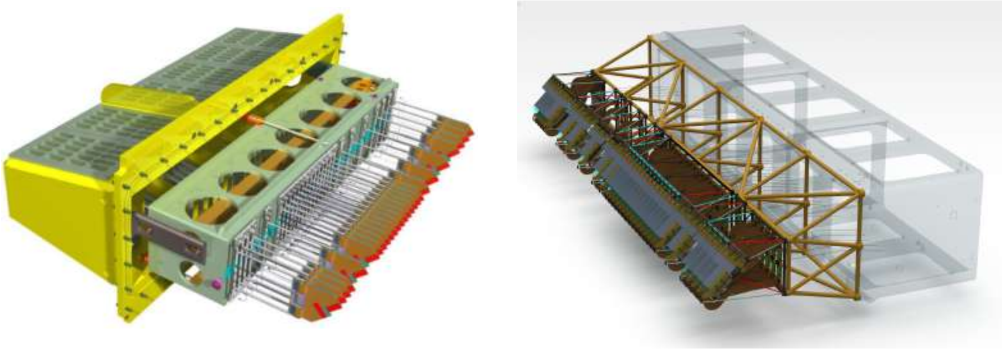

Figure 1.7: (Left) The current VELO design [41]. (Right) Illustration of the VELO stations for the Upgrade II baseline detector [33]. In both figures only one half of the detector is shown.

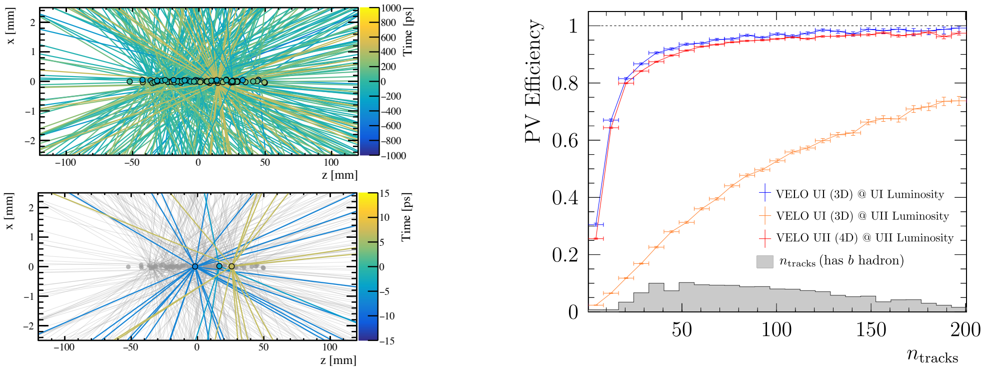

Figure 1.8: (Left, top) Tracks produced in a bunch crossing with 42 pp collisions, as seen from a detector with no timing capability. (Left, bottom) Tracks selected in a 30 ps time window, showing a drastic reduction of the vertex multiplicity to O(1). Tracks are coloured according to time of production. (Right) PV reconstruction efficiency as a function of track multiplicity for the VELO Upgrade II detector, and for the current VELO in both LHC Run 3 and HL-LHC conditions. Source: [33].

The baseline specification requires a track timestamp precision of 20 ps, corresponding to a per-hit time resolution of about 50 ps when combining multiple measurements, and a track spatial resolution of 10 μm. A new detector technology is needed to meet such requirements. The chosen baseline sensor technology is 3D silicon, which combines high spatial and temporal resolution with radiation tolerance for the full detector lifetime. Associated with a dedicated ASIC, this technology is expected to reach a hit efficiency of 90–95 % [33]. Candidate ASIC developments include designs based on 28 nm CMOS technology, a highly-radiation hard technology which offers sufficient integration density for pixel-level digital processing. Ongoing R&D includes the PicoPix ASIC [43], based on the Timepix4 [44], targeting sub-35 ps time-stamping with a 48–55 μm pitch, and the TimeSPOT ASIC [45], developed within the IGNITE project, aiming for a 45 μm pitch and similar timing performance. Both approaches are being qualified for high-rate operation and high-pile-up environments, with future developments exploring scalable multi-wafer prototypes and high-rate readout architectures [33]. A possible combination with silicon photonics is also being investigated.

The VELO Upgrade II design is shown in Figure 1.7 (right). The detector will have 32 stations covering $2.0 < \eta < 4.8$, each equipped with 150 µm thick sensors. A cylindrical RF shield with 75 μm thickness will reduce the material budget in front of the sensors. Only minimal descoping of the VELO is possible without a significant loss of performance. A small cost reduction can be obtained at lower instantaneous luminosity by reducing the number of optical links and readout boards. More substantial savings would require removing the first two and last two stations, but this comes at the expense of reduced acceptance and lower tracking efficiency.

Upstream Tracker (UT), future Upstream Precision Tracker (UP) and magnet stations

The UT is located between the RICH 1 detector and the LHCb dipole magnet. It provides an early momentum estimate and charge determination by combining hits with VELO tracks and exploiting the magnetic field between the two detectors. This information is used in the software trigger to speed up matching with the downstream SciFi Tracker and to suppress fake tracks formed by mismatched VELO and SciFi segments. The UT is also important for the reconstruction of downstream tracks. The UT detector consists of four planes of silicon microstrip detectors arranged in two stations, with vertical and stereo strip orientations to provide 3D space points, as shown in Figure 1.9. The geometry ensures full coverage within the LHCb acceptance, using finer segmentation near the beam pipe to limit track occupancy. The sensors are mounted on lightweight, actively cooled carbon-fibre staves to minimise the material budget. The detector is designed to maintain high efficiency and low noise under the radiation conditions for an integrated luminosity of at least $50, \mathrm{fb}^{-1}$.

In the HL-LHC environment, the higher luminosity and pile-up will lead to unacceptably high occupancies in the UT if it remains based on silicon strip technology. It will therefore be replaced by the Upstream Pixel (UP) detector, composed of four planes of fine-pitch silicon pixels based on

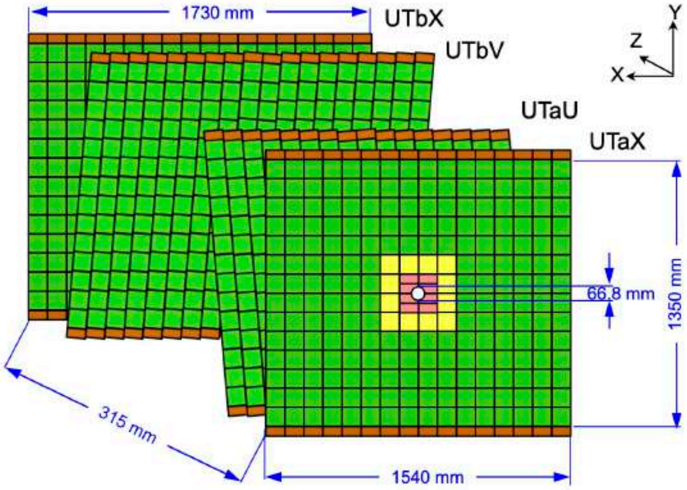

Figure 1.9: (Left) Drawing of the four UT silicon planes with indicative dimensions. Different colours (green, yellow and pink) designate different sensor granularities [38].

DMAPS technology [33]. The pixel layout reduces occupancy and provides precise, unambiguous y-coordinate measurements, enabling stand-alone tracking in this region. The UP fulfils the role of the UT in linking VELO and downstream tracks, with improved greater robustness against ghost tracks and improved momentum resolution for long tracks. In combination with the Mighty Tracker and the Magnet Stations, it also extends the acceptance for reconstructing downstream and upstream tracks.

The Upgrade II tracking system includes magnet stations to extend the physics reach for low-momentum tracks ($p < 5\,\mathrm{GeV}/c$) by instrumenting the sides of the dipole magnet with scintillating fibre detectors [33]. In the current LHCb detector, most of these tracks are absorbed into the magnet iron and are reconstructed with low efficiency and poor momentum resolution ($\delta p/p \sim$10–20 %). The magnet stations will improve the resolution to about 1 % and significantly increase the efficiency, yielding important gains for decays with soft tracks, such as $D^{*+} \to D^0\pi^+$ [33]. Each magnet station consists of four quadrants, and each quadrant into three panels. A quadrant measures $350 \times 100\,\mathrm{cm}^2$, covering $y \in [10, 95]\,\mathrm{cm}$ and $z \in [400, 750]\,\mathrm{cm}$, optimised for tracks with $0.5 < p < 5\,\mathrm{GeV}/c$ and $1.5 < \eta < 5.0$. Each panel contains four layers of scintillating bars with a parallelogram cross-section of 5 mm side and wrapped in 100 μm Tyvek. Light is collected via 1 mm wavelength-shifting fibres embedded in the bars and transported outside the magnet with clear fibres.

SciFi tracker and Mighty Tracker (MT)

The SciFi Tracker is located downstream of the LHCb dipole magnet and provides charged-particle tracking and momentum measurement in conjunction with the upstream tracking detectors. It consists of twelve detection planes arranged in three stations (T1–T3), each with a particular orientation to reconstruct the horizontal and vertical position of track interactions [38, 46]. The active medium is 250 μm diameter plastic scintillating fibres, arranged in six staggered layers to form fibre mats approximately 2.4 m long, as shown in Figure 1.10 (left). Scintillation light is detected by 128-channel Silicon Photomultipliers (SiPM) arrays at the end of each module. SiPM technology is described in detail in Section 1.4.2.

The SciFi achieves a single-hit resolution better than 100 μm and > 99 % tracking efficiency. The detector is designed to operated in the radiation and occupancy conditions expected for at least 50 fb^{-1} of integrated luminosity [47]. Radiation damage to fibres and SiPMs is mitigated through shielding and cooling. The SiPMs are mounted on a cooling bar at -40°C and electrically connected to the readout electronics via a custom flexible PCB, shown in Figure 1.11. The flex-PCB-based readout is a key feature of the SciFi design, allowing high-density channel routing from the cooled SiPM to the readout electronics at room temperature [46]. This solution is directly relevant to this thesis. SiPMs are a promising candidate as photodetectors for the RICH Upgrade II detector, but they require cryogenic cooling to operate in a high-radiation environment. As part of this work, a flex-PCB prototype for the readout of SiPM arrays at cryogenic temperature was designed and tested. Design and results are presented in Chapter 4.

In the HL-LHC environment, the SciFi technology alone cannot sustain the required performance due to the higher particle flux, integrated radiation dose, and lack of y-segmentation, which increases

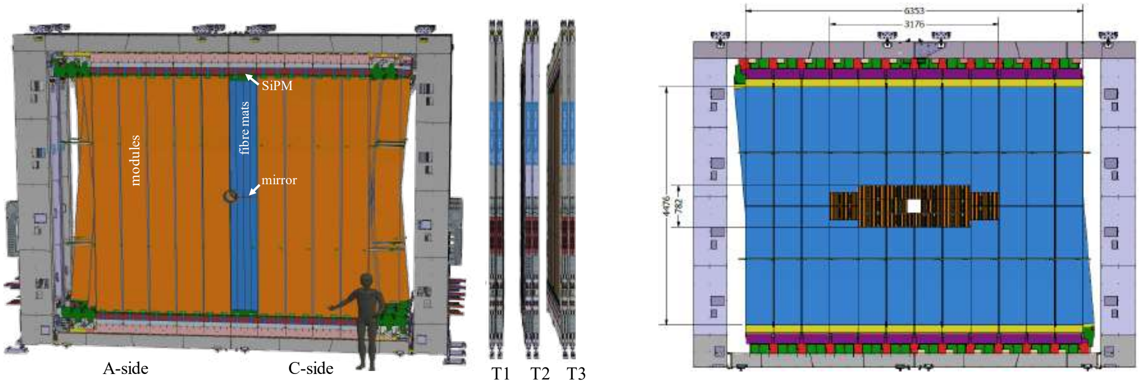

Figure 1.10: (Left) Front and side views of the 3D model of the SciFi Tracker detector [38]. (Right) Front view of the Mighty-Tracker detector modules [33].

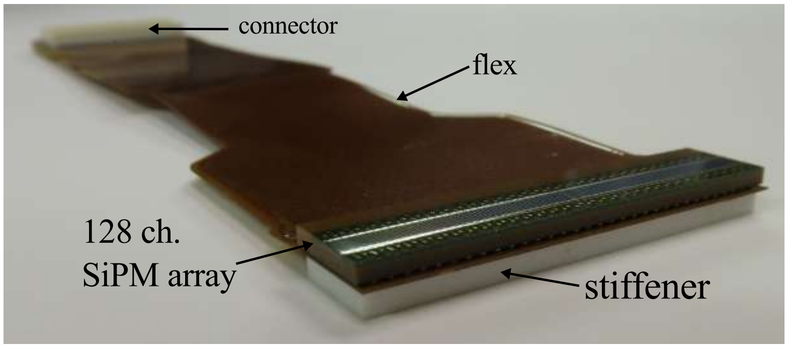

Figure 1.11: A 128-channels SiPM array bonded to a flex-PCB used in the SciFi tracker detector [38].

the probability of hit overlaps and fake tracks. The Mighty Tracker (MT) will replace the SciFi with a hybrid system: a high-granularity silicon pixel detector around the beam pipe and scintillating-fibre modules in the outer regions, as shown in Figure 1.10 (right). As for the UP, the MT uses DMAPS silicon pixels to cover the high-occupancy central region, surrounded by "Mighty-SciFi" modules for the outer area. Radiation damage in the fibres and the tracking performance define the region that must be covered by silicon, and hence the boundary between the two systems [33]. Originally, the MT detector design assumed six silicon pixel layers of $3\,\mathrm{m}^2$ each, corresponding to $18\,\mathrm{m}^2$ in total. However, the expected performance of the Mighty-SciFi has improved, driven by cryogenic operation of the SiPMs and improvements from the latest generation of SiPMs. These advances, together with a better understanding of track reconstruction, reduce the radiation damage requirements. As a result, the pixel coverage in the baseline Upgrade II scenario has been lowered to $2.1\,\mathrm{m}^2$ per layer, $12.6\,\mathrm{m}^2$ in total. The cost and complexity of the MT can be reduced by limiting the pixel coverage to about $1.3\,\mathrm{m}^2$ per layer. Additional savings are possible by decreasing the detector height or removing outer modules, although at the expense of physics performance.

1.3.2 Particle identification system upgrades

Particle identification (PID) 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. At present, PID relies on three subsystems: the RICH detectors, the calorimeter system, and the muon stations. The HL-LHC environment, with its higher detector occupancies and harsher radiation conditions, will challenge the performance of all these detectors. To preserve, and where possible extend, their capability, each subsystem will undergo a dedicated upgrade. These upgrades share the common goal of improving resolution, efficiency, and robustness, ensuring that LHCb can continue to deliver high-precision flavour-physics measurements during LHC Run 5.

RICH detectors

The RICH system provide charged hadron identification at LHCb. In particular, the RICH detectors are able to distinguish pions (π), kaons (K) and protons (p) across a momentum range of 10 to 100 GeV/c, and even down to 2.6 GeV/c for π–K separation. In beauty and charm decays, most of the products are pions, making hadron identification essential to reduce background and distinguish between different final states. For instance, the RICH detectors play a critical role in measuring CP asymmetry in the two-body hadronic decays $B \to h^+h^-$, where h is a charged hadron [48].

The operation of the RICH detectors is based on the Cherenkov effect. When a charged particle travels through a medium with a speed higher than the speed of light in the medium, it polarises the surrounding atoms. The constructive interference of electromagnetic waves from the polarised atoms results in the Cherenkov radiation, emitted at a characteristic angle relative to the particle direction. This angle, known as the Cherenkov angle θ_c, is given by:

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

where $n(\epsilon)$ is the refractive index of the radiator medium for a photon of energy ϵ, and $\beta = v/c$ is the particle velocity in units of the speed of light in vacuum. The minimum particle velocity required to emit Cherenkov light is obtained for $\theta_c = 0$, which defines the Cherenkov threshold:

$$\beta_{th} = \frac{1}{n(\epsilon)} \tag{1.8}$$

Figure 1.12 shows the reconstructed Cherenkov angle as a function of the particle momentum for different radiators. Different particle types cross the Cherenkov threshold at different momenta, while at high momenta ($\beta \to 1$) the Cherenkov angle saturates. The principle of the RICH detector is to measure the Cherenkov angle, and hence the particle velocity, and to combine it with the particle momentum from the LHCb tracking system to determine the particle mass. The RICH reconstruction algorithm is described in detail in Chapter 3.

The RICH system consists of two detectors. The RICH 1 detector is placed upstream of the LHCb magnet while the RICH 2 detector is downstream. The Cherenkov radiator used in the RICH 1 detector is C_4F_{10} gas for a particle momentum range of 10 to 65 GeV/c. The RICH 2 detector contains CF_4 gas for a momentum range of 15 to 100 GeV/c. Cherenkov photons emitted in the gas radiator are focused onto a plane of photodetectors by the RICH mirror system, composed of a tilted spherical mirror and a planar mirror. The resulting pattern from a single charged track is a characteristic Cherenkov ring on the photodetector plane. The photodetector plane of RICH1 and the inner region of RICH2 is equipped with 1-inch multi-anode photomultiplier tubes (MAPMTs), while the outer region of RICH2 uses 2-inch MAPMTs. The schematic layout of the two RICH detectors is shown in Figure 1.13. Although their designs are very similar, the RICH 1 detector is

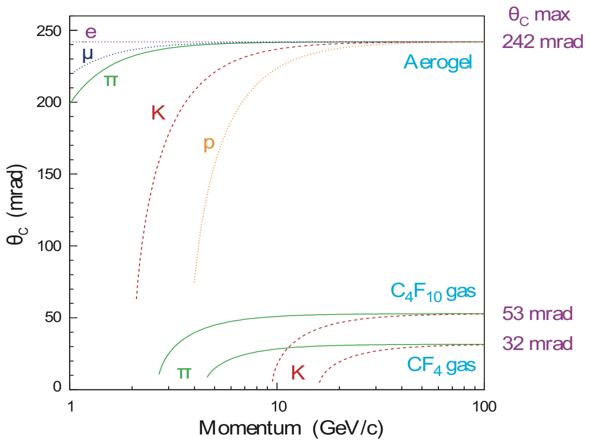

Figure 1.12: Cherenkov angle dependence on momentum for different kind of particles in different Cherenkov radiators. Source: [21].

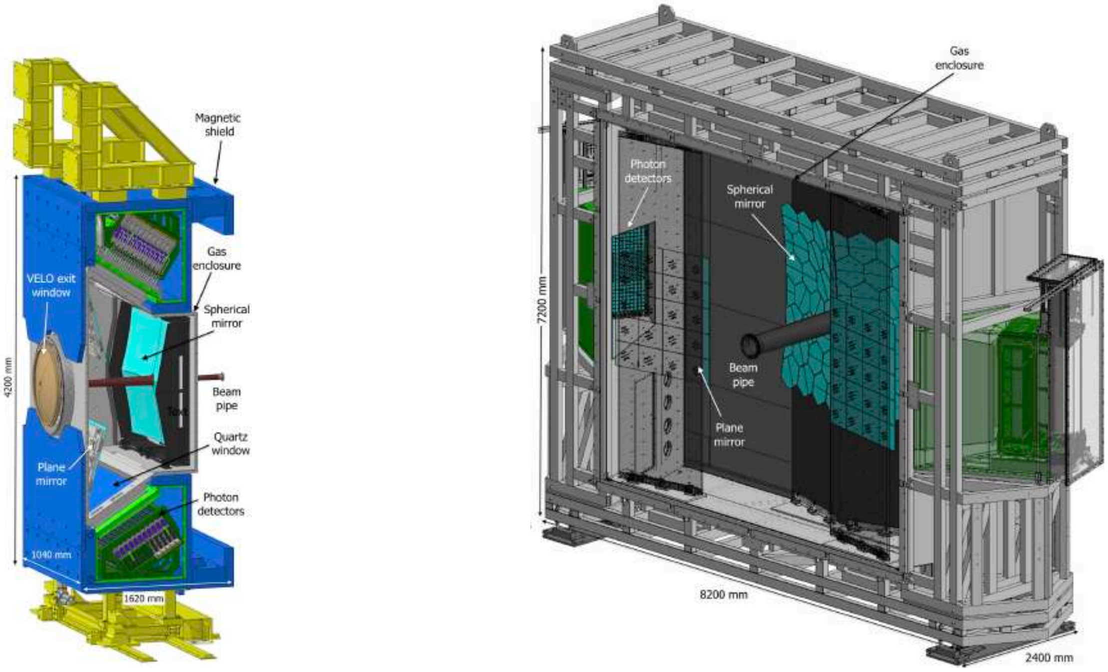

Figure 1.13: Schematic view of the RICH 1 (left) and RICH 2 (right) detectors. Source: [38].

more compact and operates with higher track multiplicity and detector occupancy.

With the increased luminosity expected for LHC Run 5, the performance of the RICH reconstruction will significantly degrade given the ten times larger track and photon densities. In the high occupancy regions, information is lost when more than one photon hit the same photodetector pixel. The increased rates also put strong demands on the readout electronics. To cope with these conditions, the RICH system requires finer photodetector granularity and the addition of precise photon timing. These improvements reduce occupancy and remove background, thereby preserving PID performance at the LHC Run5 instantaneous luminosity. In particular, the time information on the photon hits can be used to reduce event complexity by removing combinatorial background from track-photon misassociations. This strategy and its implementation in the RICH reconstruction are described in Chapter 3.

To meet these requirements, the RICH detector will introduce a complete-new optoelectronic readout system, including a novel front-end readout electronics, based on the FastRICH ASIC [49] and capable of timestamping a photon hit with 25 ps precision, and new photodetectors, with higher granularity and sub-100 ps time resolution. The characteristics of the FastRICH ASIC, the photodetector candidates and the R&D in preparation for the novel optoelectronic readout chain are described in Section 1.4. The benefits of the photon hit time information on the RICH detector performance are described in Section 3.1.2.

Other key parameters to improve the RICH detectors performance are the photon yield, i.e. the number of Cherenkov photons emitted by a track detected on a Cherenkov ring by the photodetectors, and the Cherenkov angle resolution. The latter is affected by several contributions:

- Emission point uncertainty: optical aberrations of the mirrors introduce an error on the Cherenkov angle reconstruction, as the photons emitted along the particle trajectory follow different paths to the photodetector plane.

- Chromatic dispersion: the unknown energy of the Cherenkov photons results in an error due to the variation of the refractive index of the gas radiator and the gas enclosure quartz window as a function of the photon wavelength.

- Pixel size: a commonly used rule for a good pattern reconstruction is that the Cherenkov angle resolution σ_θ should not smear the ring over more than one pixel, i.e. $\sigma_\theta \cdot f \leq A_p$, where f is the focal length and A_p the pixel area.

- Track-direction measurement: uncertainties introduced by the tracking system and the non-linearity of the particle trajectories through the radiator.

The reduced photodetector pixel size foreseen for RICH Upgrade II detector improves the pixel error on the Cherenkov angle resolution. In the baseline scenario for Upgrade II, the photodetectors in the high occupancy regions of RICH 1 and RICH 2 have a pixel area of $1.4 \times 1.4\,\text{mm}^2$, compared to the $2.8 \times 2.8\,\text{mm}^2$ of the MAPMTs currently used. In addition, choosing a photodetector with

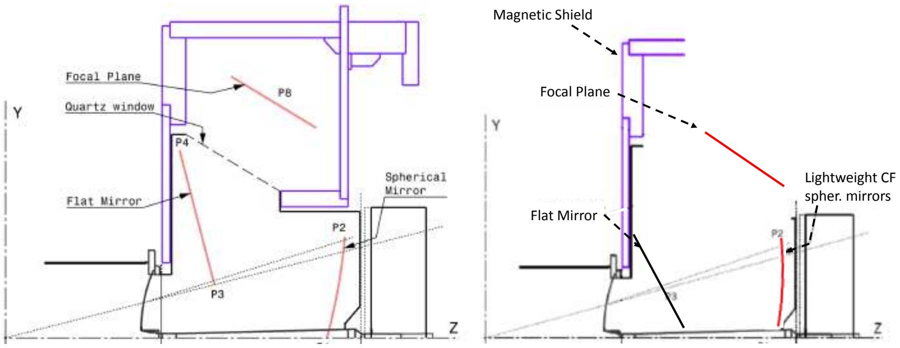

Figure 1.14: The RICH 1 optical system in the current (left) and Upgrade II (right) configuration. The flat mirror is brought into the detector acceptance to reduce optical aberrations during Upgrade II. Source: [50].

higher sensitivity in the green region reduces the chromatic error. The main photodetector candidates to replace the MAPMTs are SiPMs and MicrochannelPlate Photomultipliers (MCP-PMTs). Both technologies offer sub-100 ps single-photon time resolution (SPTR) and higher pixel density, but present different challenges in terms of radiation hardness, dark count rate, and rate capability. A dedicated description of these photodetector technologies and their possible integration into the RICH Upgrade II detector is discussed in Section 1.4.2. For RICH Upgrade II, the planar mirror will be replaced with a lightweight composite mirror and placed inside the acceptance, as shown in Figure 1.14. This configuration reduces optical aberrations in the RICH system and lowers the emission point error. Table 1.1 summarises the improvements expected on the Cherenkov angle resolution with the RICH Upgrade II detector. Different scenarios are considered for Upgrade II, varying in photodetector granularity and detector geometry. The table reports the resolution contributions only for the highest-performance scenario, corresponding to the most advanced and cost-intensive option [33].

As discussed in Section 1.1.2, the limited duration of LS4 poses a challenge for the installation and commissioning of the RICH Upgrade II detector. For this reason, the RICH collaboration will anticipate the installation of the novel FastRICH-based readout electronics during LS3, while the photodetectors and the mirrors will be replaced in LS4. With the so-called LS3 Enhancements, the photon hit time information can be integrated in the RICH reconstruction and used during LHC Run 4 to improve the detector PID performance. An additional benefit is the opportunity to develop and validate timing-based reconstruction algorithms ahead of LHC Run 5. The photon hit time information in LHC Run 4 will also enable the estimation of the primary vertex time with

Table 1.1: Expected contributions to the Cherenkov angle resolution for RICH 1 and RICH 2 in the current and Upgrade II configurations. Source: [33].

	RICH 1		**RICH 2**	
	Current	Upgrade II	Current	Upgrade II
N_{pe}	62	52	39	25
Chromatic	0.59	0.28	0.33	0.19
Pixel	0.46	0.22	0.20	0.10
Emission	0.38	0.13	0.37	0.05
Total [mrad]	**0.82**	**0.22**	**0.45**	**0.13**

the RICH detectors. This represent a first-time measurement for the LHCb experiment and it is one of the central topics of this thesis. Chapter 3 describes the integration of the photon hit time in the RICH reconstruction and presents an algorithm to estimate the primary vertex time based from RICH reconstruction information only.

TORCH

The Time Of internally Reflected CHerenkov light (TORCH) detector has been proposed as an additional PID system for LHCb Upgrade II. Its purpose is to extend π/K and K/p separation into the low-momentum region (2–15 GeV/c), where the standard RICH reconstruction becomes less effective because the kaon threshold in the C_4F_{10} radiator is 9.3 GeV/c and in CF_4 is 15.6 GeV/c. TORCH combines time-of-flight and Cherenkov imaging, providing hadron identification below these thresholds [33]. The detector is located upstream of the RICH 2 detector and is made of 18 highly-polished 1 cm-thick quartz plates, covering the full LHCb acceptance. Charged particles crossing the quartz generate Cherenkov photons, which propagate by total internal reflection to the edges of the plate, as shown in Figure 1.15. A focusing block with a cylindrical mirror projects the photons onto an array of photodetectors. The photon hit time and position, combined with the reconstructed track vector, allow reconstruction of both the photon path and the Cherenkov angle. To achieve π/K and K/p separation in the 2–20 GeV/c momentum, a time-of-flight resolution of around 15 ps is required [51]. This translates into a SPTR below 70 ps for an expected yield of 30 photons per track. To reach this performance in the HL-LHC environment, the detector occupancy, which peaks in the central region of the detector near the beam line, is required to be below 20 %. This sets the granularity of the photodetectors and the electronics, resulting in about 225,000 readout channels.

Baseline photodetectors for TORCH are the Micro-Channel Plate Photomultiplier Tubes (MCP-PMTs), developed in collaboration with Photek [1]. MCP-PMTs offer excellent intrinsic SPTR of 20–30 ps and low dark-count rate (DCR). Beam tests with prototype modules achieved single-photon

---

[1] Photek (`https://www.photek.com`)

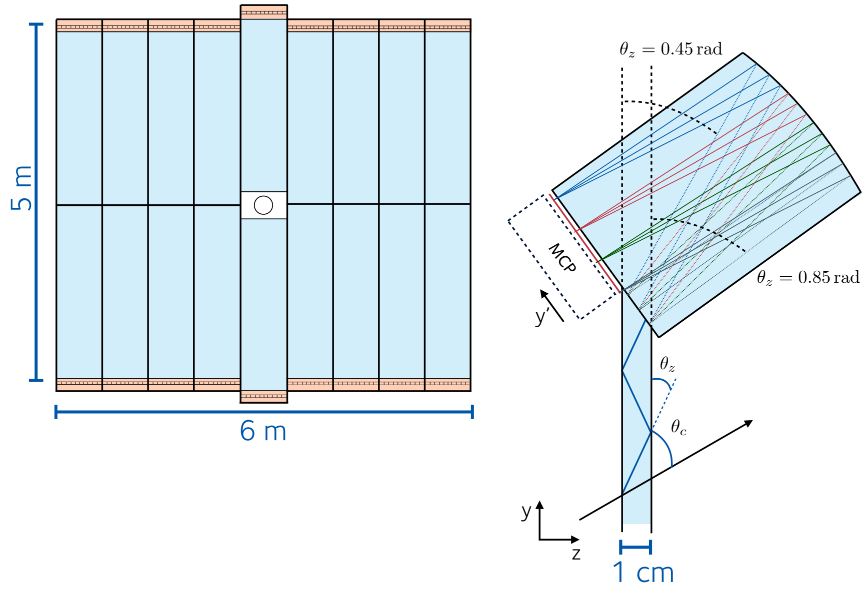

Figure 1.15: (Left) Layout of the 18-module (quartz bar) TORCH detector. (Right) A single TORCH module with the focusing block. Source: [33].

resolutions in the range 80–115 ps, approaching the target performance [52, 53]. The main challenge for MCP-PMTs is the limited lifetime with high integrated anode charge. Therefore, alternative technologies, such as SiPMs, are also under investigation. These two technologies are described in detail in Section 1.4.2 as they are also candidate photodetectors for the RICH Upgrade II detector. Many considerations made for the RICH detector and the Upgrade II environment are valid for TORCH as well. Different scenarios for the integration of TORCH in the LHCb Upgrade II detector are under evaluation. The baseline design maintains full acceptance, while reduced-coverage options remove several central and outer bars to limit cost and complexity. However, even with 75 % coverage, significant impact on the physics reach is expected [33]. Beyond low-momentum PID, TORCH also serves as a downstream timing layer. By comparing its measurements with the VELO prediction, ghost tracks can be suppressed and the reconstruction of long-lived particle decays improved [54].

Calorimeters and muon stations

The LHCb calorimeter system follows a classical layout of an electromagnetic calorimeter (ECAL) [55] followed by a hadronic calorimeter (HCAL) [56]. Its main purpose is to measure the energies and the positions of clusters deposited by electrons, photons and hadrons [57]. The ECAL is 25 radiation lengths thick to ensure full containment of electromagnetic showers, while the HCAL is 5.6 interaction lengths thick, limited by the available space [38]. Both detectors are sampling calorimeters based on plastic scintillator tiles, with wavelength-shifting (WLS) fibres transporting the light to the readout system. To cope with the different particle fluxes across the detector surface, the ECAL is segmented into three regions (inner, middle, outer) with increasing cell sizes moving outwards from the beam pipe, for a total 6016 cells. The ECAL cells have a shashlik structure, as shown in Figure 1.16, with alternated scintillator and lead layers. The scintillation light transported by the WLS fibres is readout by dedicated PMTs. Radiation damage studies indicate that the innermost ECAL cells, exposed to doses up to 25 kGy at the shower maximum, will need replacement during LS3 [38]. The HCAL is segmented in two regions with a larger granularity with respect to the ECAL, given the typical spread of hadronic showers. The HCAL is composed of 1488 with interleaved iron and scintillator tiles. The same PMT type as in ECAL is used for the readout.

At the Upgrade II luminosity of $1.5 \times 10^{34}\,\mathrm{cm^{-2}s^{-1}}$, the expected particle flux and accumulated radiation dose will exceed the capabilities of the present calorimeter modules. A new electromagnetic calorimeter, named PicoCal, has been proposed [54, 58]. Its design aims to preserve or exceed the performance of the current ECAL by introducing both higher granularity and timing information at the level of tens of picoseconds. The detector modules are arranged in rhombic-like regions, as shown in Figure 1.17, which follow the profile of the radiation dose and occupancy downstream of the LHCb magnet. The inner region, subject to the highest occupancy and radia-

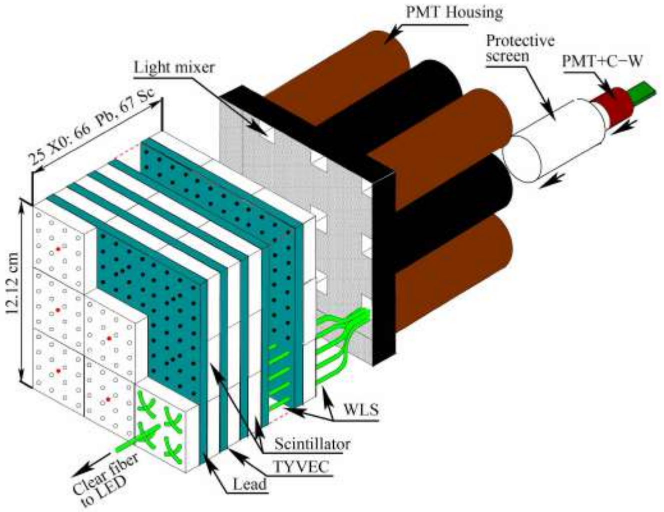

Figure 1.16: Schematic of an ECAL cell. Source: [38].

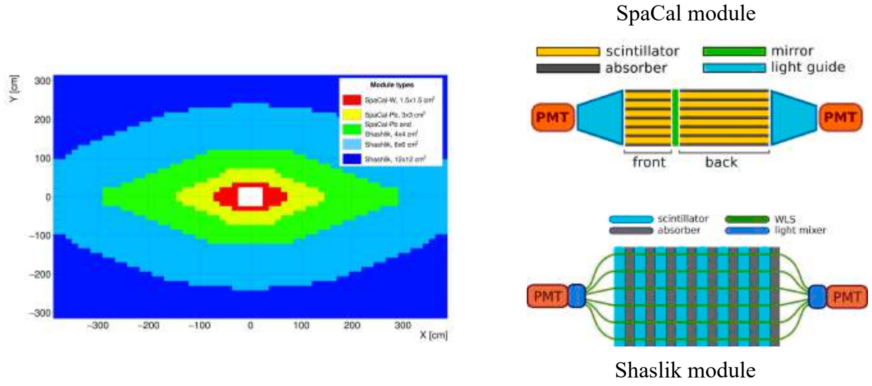

Figure 1.17: (Left) Boundaries between regions with different module segmentation for the PicoCal design. Technology used in each region is also indicated: SpaCal-W and SpaCal-Pb for the inner regions, and Shashlik for the outer regions. Source: [33]. (Right) Schematics of a SpaCal and a shashlik module in double-sided readout configuration. Source: [58].

tion, will be equipped with high granularity tungsten–GAGG SpaCal modules, with small Molière radius. The surrounding region will employ lead–plastic SpaCal modules with relaxed granularity requirements [59]. The modules are rotated by a few degrees relative to the beam axis to optimise energy resolution. A mirror is placed near the depth of the shower maximum to divide the modules, which are read on both sides by fast photodetectors. This provides access to the longitudinal shower profile, improving both timing resolution and cluster reconstruction. The outer region will use a combination of refurbished Shashlik modules, upgraded with new WLS fibres, and newly built SpaCal modules. The total number of channels increases to about 31,000, compared with the 6100 in the current ECAL, making the PicoCal a major cost driver. In addition to the baseline configuration, a potential down scoping option has been evaluated. In this scenario, most of the detector modules would use single-sided readout, except for the 176 innermost modules. This reduces the number of channels by about 30 %, but at the cost of losing longitudinal shower profile information and degrading the timing resolution over most of the detector area [33].

As mentioned in Section 1.2.2, the hadronic calorimeter will be removed and replaced by an iron wall thicker than 1 m, positioned in front of the muon stations. In the Upgrade II conditions, the HCAL does not provide sufficient resolution for hadronic energy measurements, while its role as an absorber becomes more critical to suppress background in the downstream muon detectors. The new shielding improves muon identification performance by reducing punch-through from hadrons and provides a more uniform response across the acceptance [54].

Muon system

The LHCb muon detector provides muon identification and contributes to the global event reconstruction. In LHC Runs 1 and 2 it consisted of five stations (M1–M5) [60]. Station M1, located upstream of the calorimeters, was used in the hardware Level-0 trigger but is no longer needed in the trigger-less architecture, and has been removed for LHC Run 3 [61]. The current system therefore consists of four stations (M2–M5), located downstream the calorimeters and interleaved with 80 cm thick iron absorbers to filter low energy particles [38]. Each station is divided into four regions (R1–R4) of increasing area from the beam pipe outward. The stations instruments with multi-wire proportional chambers (MWPCs), for a total of 1104 MWPCs covering about 385 m^2. The MWPCs are made up of four independent layers (or gaps), each consisting of anode wires between two cathode planes, to achieve a high efficiency and a high redundancy. The MWPCs are read out by radiation-tolerant front-end ASICs (CARIOCA and DIALOG), digitised and transmitted at the 40 MHz bunch crossing rate. The major Run 3 upgrade consisted of new readout electronics (nODE boards) to make the system compatible with the trigger-less DAQ (see Section 1.3.3), while keeping the detector modules themselves unchanged [38].

In Upgrade II conditions, particle fluxes in the innermost regions (R1, R2) will approach 1 MHz/cm^2, far beyond the sustainable rate of the present MWPCs [54]. The baseline scenario for the muon system upgrade consists of a substantial increase in granularity and the replacement of MWPCs in R1 and R2 with micro-Resistive WELL (μRWELL) chambers, which can withstand MHz/cm^2 fluxes. In R3 of stations M2 and M3, the MWPCs will be replaced by higher granularity versions, while in the outer regions (R4 and large parts of R3) the existing MWPCs are expected to perform adequately. Additional shielding will be installed upstream of M2, replacing the present hadronic calorimeter. This adds about four interaction lengths to reduce the flux in the most critical regions. The readout scheme will also change: instead of combining analogue signals from two consecutive gaps into a single channel, each gap will be digitised independently and a coincidence logic applied in the back-end, suppressing noise from uncorrelated low-energy particles. This requires new front-end ASICs for both μRWELLs and MWPCs. The number of readout channels will increase to about 720,000, compared with 120,000 in the current system [33].

1.3.3 Data processing upgrades

From LHC Run 3 onwards, the LHCb detector operates in a trigger-less configuration. All subdetectors are read out synchronously at the LHC bunch-crossing rate of 40 MHz, corresponding to about 30 MHz of visible *pp* collisions. This represents a major departure from the previous design, where a hardware Level-0 trigger reduced the event rate to 1 MHz before software processing [61]. The trigger-less architecture allows the full event information to be available for software reconstruction and selection at the earliest stage [62, 63]. The LHCb readout architecture in the current

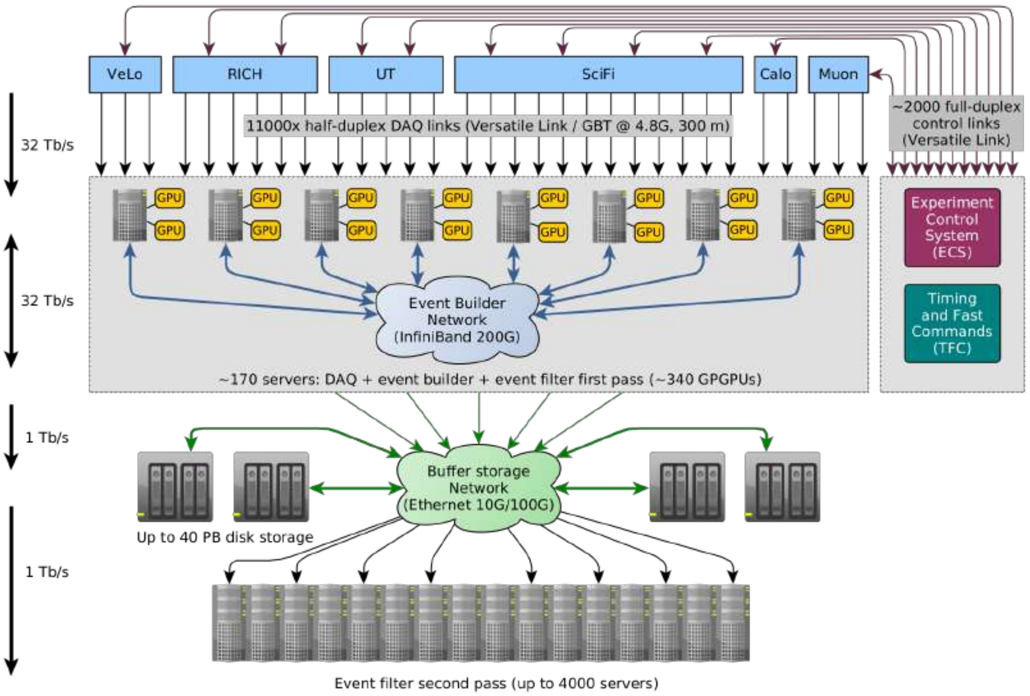

Figure 1.18: Scheme of the upgraded online system. The data received from each sub-detector is combined in the event builder farm before being processed by the trigger system. Source: [38].

LHC Run 3 detector is shown in Figure 1.18, while the dataflow scheme is presented in Figure 1.19

The front-end electronics of the sub-detectors transmit data continuously through radiation-hard optical links (the Versatile Links) using the GBT protocol at 100 Gbit/s. FPGA-based PCIe40 boards [65], configured as TELL40 for data acquisition, receive the fragments, align them by bunch-crossing ID, and forward them into the event-builder network. PCIe40 boards, in a different configuration, are also used to provides the bi-directional fast- and slow-controls communication channel to the front-end electronics through the Versatile Link. Complete events are assembled at the full 30 MHz rate by a farm of event-builder servers, each hosting both CPUs and GPUs. The LHCb software trigger operates in two software stages. The High-Level Trigger 1 (HLT1) runs on the GPUs and performs a fast, partial reconstruction. It uses information from the VELO, UT, SciFi, and muon detectors to reconstruct tracks and reduce the 30 MHz input rate to about 1 MHz [66]. The accepted events are written to a disk buffer while real-time alignment and calibration of the VELO, UT, muon system, and RICH optics are performed. The second stage (HLT2) then processes the buffered events on a CPU farm. At this stage a full, offline-quality reconstruction is

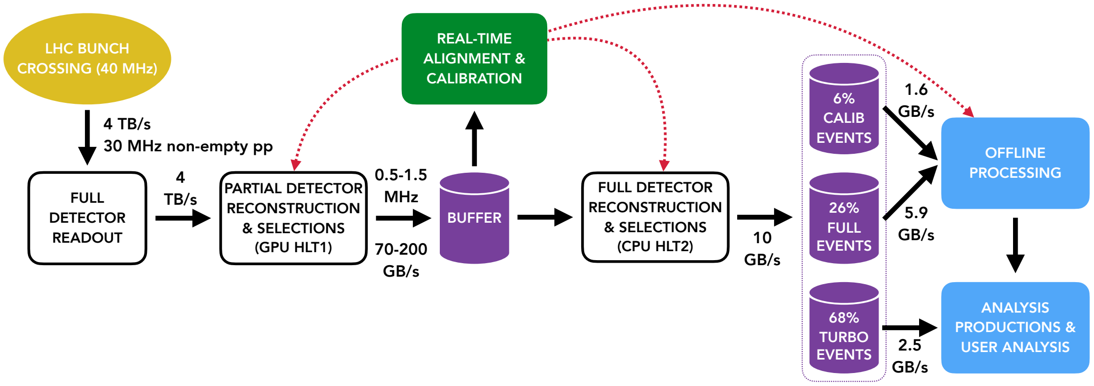

Figure 1.19: LHCb dataflow, from the LHC bunch crossing of 40 MHz down to the offline processing and analysis production. Source: [64].

carried out using the updated alignment constants. Unlike HLT1, HLT2 includes all sub-detectors, such as the RICH, allowing both inclusive and exclusive selections. The final output is written to permanent storage at a rate of 2–5 GB/s, either in full or in reduced “Turbo” format, which stores only the physics candidates and their associated primary vertices [66]. In both the online and offline chains, VELO tracks are reconstructed first. In the offline reconstruction they are extrapolated to the SciFi as part of the third forward-tracking stage, while online the intermediate UT requirement reduces CPU time. PID in the RICH detectors relies on the momentum measurement from tracking, obtained with a Kalman fit.

At the Upgrade II luminosity of $1.5 \times 10^{34}\,\text{cm}^{-2}\text{s}^{-1}$, LHCb detector is expected to produce up to 200 Tb of data per second. This volume must be processed in real time and reduced by about four orders of magnitude before being sent to permanent storage [54]. The trigger system will remain two-staged, but its algorithms must evolve to exploit time information from the sub-detectors and to cope with the higher occupancies. GPU acceleration will be extended to large fractions of the HLT2 reconstruction. The event-filter farm is foreseen to expand to several thousand heterogeneous nodes, including GPUs and potentially other architectures such as FPGAs or ARM-based processors, depending on their cost and energy efficiency [33]. The software frameworks (Allen for GPUs and Gaudi for CPUs) are being unified to allow for flexible scheduling of reconstruction and selection across different processors. This flexibility will be essential to balance throughput, energy consumption and performance.

The trigger-less model introduced in Run 3 has proven to be robust and scalable. For Upgrade II, the addition of timing information, extended GPU utilisation, and flexible heterogeneous computing resources will be crucial to sustain the physics programme in HL-LHC. With these developments,

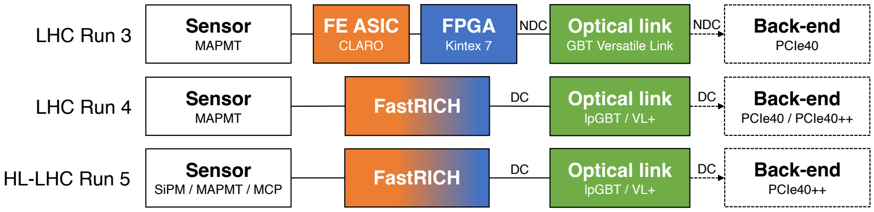

Figure 1.20: Schematic of the optoelectronic readout chain evolving from LHC Run 3 to Run 5 Source: [67].

LHCb will continue to record a broad spectrum of beauty, charm, and exotic signatures with offline-quality reconstruction available directly in real time.

1.4 The pathway for the RICH Upgrade II detector

The increased luminosity foreseen for the LHC Run 5 poses significant challenges for the RICH detectors. Higher event rate and track multiplicity result in a greater detector occupancy. To maintain, and possibly improve, the excellent PID performance of the RICH detectors, two key improvements are required: the introduction of the photon hit time information and an increase of the photodetector granularity. As discussed in Section 1.3.2, these improvements necessitate an entirely new optoelectronic readout system, including novel readout electronics and new photodetectors. Figure 1.20 shows the evolution of the optoelectronic readout chain for the RICH detector from the current LHC Run 3 to Run 5. The new readout electronics will be introduced already during LS3, while the photodetectors will be replaced in LS4.

The readout electronics of the current RICH detector is based on the CLARO ASIC [68] at the front-end and on an Kintex 7 FPGA for the hit digitalisation. Their functionalities have been integrated into the novel FastRICH ASIC [49], designed to meet the requirements for the RICH detector upgrades. The FastRICH is a 65 nm CMOS technology chip built on the existing FastIC ASIC [69], with an integrated TDC of 25 ps bins and a radiation hardness up to 10^{13} $\mathrm{n_{eq}/cm^2}$. Among other specifics, the FastRICH features advanced data compression schemes and bandwidth optimization capabilities with limited power consumption. The FastRICH main characteristics and the design choices behind them are described in Section 1.4.1.

For LHC Run 5, new photodetectors will replace the MAPMTs used in the current RICH detector. The MAPMTs offer limited granularity and time resolution, which restrict their suitability for Upgrade II. Although no single replacement technology has yet met all requirements, several promising candidates are under active evaluation. Section 1.4.2 reviews the characteristics of the current MAPMTs and presents the main alternatives under investigation: SiPMs and MCP-PMTs.

Both technologies offer sub-100 ps time resolution and higher pixel density, but each presents specific challenges in terms of radiation hardness, DCR and rate capabilities.

In parallel, the RICH collaboration has developed several prototype optoelectronic readout chains with picosecond-timing capabilities over the past few years. These chains are based on the FastIC ASIC coupled with an external TDC and include the photodetector candidates for the RICH upgrades. The chains were extensively tested in beam test campaigns at the CERN SPS facility over the past few years, providing valuable inputs for the design of the FastRICH ASIC and the final RICH Upgrade II optoelectronic readout. Section 1.4.3 documents the evolution of the prototype chains.

1.4.1 The FastRICH ASIC

The FastRICH is an ASIC developed in 65 nm CMOS technology and specifically designed for the requirements of the LHCb RICH detectors upgrades. The FastRICH integrates the functionalities of both the CLARO ASIC and the Kintex 7 FPGA used in the current RICH detector into a single ASIC, with ~ 25 ps time resolution and additional features to cope with higher detector occupancies and data rate. The chip is an evolution of the FastIC ASIC [70, 69], with its architecture adapted to tolerate the radiation levels expected in LHC Run 5 ($2 \cdot 10^{13}$ 1 MeV $n_{\mathrm{eq}}/\mathrm{cm}^2$ and 12 kGy) [67]. The FastRICH has 16 analogue input channels with a input current range from 50 μA to 2 mA. The doubled number of channels compared to the CLARO chip supports the increased of granularity foreseen for the RICH Upgrade II detector. A configurable analogue front-end stage allows operation with appropriate speed, stability, and input current range for different photodetector characteristics and impedances, making the ASIC compatible with MAPMTs, SiPMs and MCP-PMTs. This flexibility is required because the FastRICH-based readout will operate with MAPMTs during LHC Run 4 and with new photodetectors from Run 5 onward.

Each analogue input channel includes a constant-fraction discriminator (CFD) to minimise time-walk effects and eliminate the need to transmit time-over-threshold (ToT) information to the back-end, thereby reducing the output data volume [49]. In addition, the ASIC can provide a non-linear ToT energy signal using a Leading-Edge Discriminator (LED). Analogue studies on MAPMT single-photon response described in Chapter 2 provided key input on the expected analogue pulse shape. CFD or LED pulses are digitised on-chip using a time-to-digital converter (TDC) with a bin size of approximately 25 ps. The digital output therefore includes the timestamp of each photon hit. This additional information is a powerful tool to improve RICH reconstruction performance. However, it increases the volume of data to be transmitted by the detector. To address this, the FastRICH ASIC includes several features to optimise the data throughput. A configurable hardware shutter is implemented to discard out-of-time photons within the 25 ns of an LHC Bunch Crossing (BX), as Cherenkov photons produced in a BX are expected in a window of only a few nanoseconds (see Section 3.1.2). Data are zero-suppressed and encoded into variable-length packets, which are transmitted through one of four independent SLVS output links, compatible with lpGBT/VTRx+ optical link chipset. This gives the chip a total maximum output bandwidth of 5.12 Gbps. In addition, the number of active output links is configurable, allowing the optimisation of the chip for different photon-occupancy regions of the detector. The digital circuitry is implemented using triple-modular redundancy (TMR) to improve radiation tolerance, making the chip robust by design against single-event upsets. Despite these features, the power consumption of the chip is kept below 16 mW per channel, easing the cooling requirements for the operation in the experiment.

As introduced in Section 1.3.2, the integration of the FastRICH-based readout electronics during LS3 allows to exploit the photon hit time information in the RICH reconstruction during LHC Run 4. Simulation studies have demonstrated that the photon hit time and the hardware shutter are essential to reduce the combinatorial background and improve the RICH detector performance [67]. This is described in more detail in Chapter 3, which also presents a novel method to use the RICH photon hit time for the reconstruction of the primary vertex time in LHC Run 4.

The first FastRICH chips have been produced and tested, showing encouraging results in terms of operability and stability. A first batch of ASICs has been assembled onto front-end boards, to be coupled with different photodetectors. At the moment of the writing, the first particle-beam test with FastRICH-based readout electronics is ongoing.

1.4.2 Photodetector candidates for Upgrade II

The photodetectors currently used in the RICH detectors are MAPMTs. These will remain in use during LHC Run 4, coupled with the novel FastRICH-based readout electronics. However, due to their intrinsic limitations in granularity and timing, MAPMTs will be replaced in the RICH Upgrade II detector. The main candidates to replace them are SiPMs and Microchannel Plate Photomultiplier Tubes (MCP-PMTs). Both technologies offer sub-100 ps SPTR and higher pixel density, but present different challenges.

This section describes in detail the MAPMT technology, as well as the SiPM and MCP-PMT options under consideration for the RICH Upgrade II detector.

Multi-Anode Photomultiplier Tubes (MAPMTs)

The photodetectors currently used in the RICH system are MAPMTs produced by Hamamatsu[2]. Two custom models are used: the 1-inch R13742 and the 2-inch R13743, both with an 8×8 pixel array [38]. The two models are shown in Figure 1.21 (left). The 1-inch MAPMTs have a pixel size of $2.88\times2.88\,\mathrm{mm}^2$ and are installed in RICH 1 and in the central region of RICH 2. The

---

[2] Hamamatsu (`https://www.hamamatsu.com/`)

2-inch MAPMTs, with 6×6 mm^2 pixels, are used in the outer regions of RICH 2, where the occupancy is lower. Both models use a UV-glass entrance window and a super-bialkali photocathode, which provides a quantum efficiency (QE) of 30–35 % around a 350 nm wavelength, as shown in Figure 1.22 (left).

Figure 1.21 (right) shows the photoelectron multiplication scheme of an MAPMT. Photoelectrons emitted from the photocathode are accelerated towards one of the 64 independent dynode chains, each corresponding to a single MAPMT pixel. The electron strikes the first dynode and produces multiple secondary electrons through impact ionisation. These are then accelerated to the next dynode stage, repeating the multiplication process. The signal is amplified over 12 stages in the 1-inch MAPMTs and 10 stages in the 2-inch versions, reaching a nominal gain of 10^6 at 1 kV. This process leads to a relatively long transit time of 5.1 to 5.3 ns for the 1-inch and 2-inch MAPMTs respectively [71]. Having per-pixel multiplication chains results in gain spread between pixels within the same array. In addition, statistical fluctuations in the multiplication process cause gain variation even within a single pixel, as shown in Figure1.22(right). This gain variation, typical of MAPMTs, motivated the design of the on-chip CFD in the FastRICH ASIC, as described in Section1.4.1. The CFD minimises time-walk effects induced by gain fluctuations, eliminating the need to transmit ToT information for this purpose.

Despite their high QE, high gain, and low dark count rate, MAPMTs are not suitable for the RICH Upgrade II detector. Their granularity is limited by the geometry of the internal photoelectron multiplication chain, which prevents scaling to smaller pixel sizes. The most advanced Upgrade II scenario foresees pixel sizes of 1.4×1.4 mm^2 in the high-occupancy regions, a dimension incompatible with MAPMT technology. In addition, their time resolution is limited by the transit-time spread (TTS), which is expected to be approximately 150 ps [72]. MAPMTs will remain in

Figure 1.21: Left: the MAPMTs used in LHC Run 3 RICH detectors with the 2-inch model on the left and the 1-inch model on the right. Right: scheme of the internal structure of the MAPMT [38].

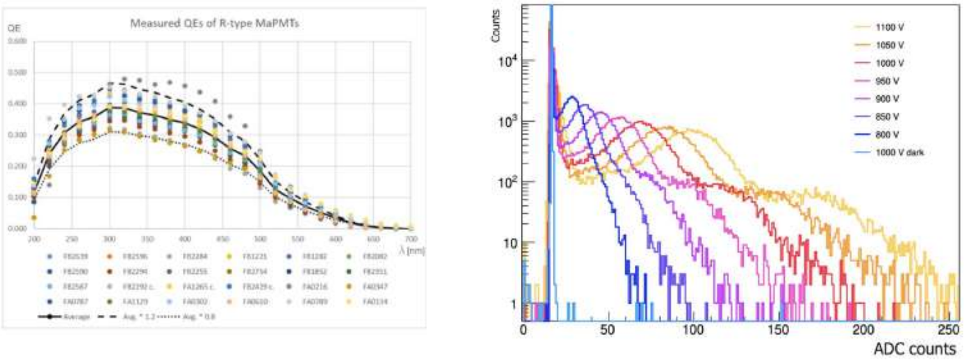

Figure 1.22: (Left) QE curves for a batch of 1-inch MAPMTs from the production: the ultra bi-alkali photocathode allows to reach excellent QE values. (Right) Typical signal amplitude spectra for an MAPMT channel as a function of the HV value. The first peak correspond to the noise pedestal while the second to single-photon signals. The width of the single-photon peak is due to the gain variation. Source [38]

use during LHC Run 4. For this reason, understanding their single-photon response and timing behaviour is important both for their coupling with the FastRICH ASIC and for defining the expected SPTR of the RICH detector in LHC Run 4. This thesis includes a dedicated study of their single-photon response, with particular focus on the TTS, presented in Chapter2.

Silicon Photomultipliers (SiPMs)

SiPMs are currently used in the LHCb SciFi detector as described in Section 1.3.1. These detectors offer several attractive features for the RICH Upgrade II detector, including small pixel size, high and green-shifted QE and adequate SPTR. They also represent a cost-effective solution for the RICH Upgrade II detector.

Silicon Photomultipliers (SiPMs) are solid-state photodetectors composed of an array of Single-Photon Avalanche Diodes (SPADs) operating in Geiger mode [74, 75]. To illustrate the functioning of the SiPMs, it is useful to first describe the functioning of Avalanche Photodiodes (APDs) and SPADs. As shown in Figure 1.23, an APD is a p-n photodiode biased in reverse direction below the breakdown voltage, i.e. the voltage beyond which the current in the diode increases exponentially. In this configuration, a photon hitting the diode in the depletion region can generate an electron–hole pair. These charge carriers are then separated and accelerated in opposite directions due to the strong electric field applied. Their velocity increases to the point that, upon colliding with the crystal lattice, they generate additional electron–hole pairs, initiating a chain reaction known as the avalanche effect. In APDs, the resulting current signal is proportional to the number of incident photons, amplified by a multiplication factor typically ranging from 50 to 300, depending on the applied voltage. SPADs are essentially APDs biased above the breakdown voltage. In these conditions, the absorption of even a single photon in the depletion region can trigger a self-sustained avalanche with an internal gain of the order of 10^6 [75]. In this regime, known as Geiger mode, the SPAD output signal saturates, regardless of the number of incoming photons. To reset the SPAD after the saturation and restore normal operation, a quenching circuit is required [73]. The SPAD dimensions range from 10 to $100\,\mu m^2$. An SiPM is commonly composed of thousands of SPADs, each able to detect individual photons and trigger a local avalanche [76]. The summed response from multiple SPADs gives the SiPM photon-counting capability, while the saturating avalanche process of the SPADs ensures high gain with limited fluctuations and low sensitivity to magnetic field. Furthermore, SiPMs require a much lower bias voltage than MAPMTs and MCP-PMTs, which simplifies the powering requirements in the RICH detectors.

SiPMs can be manufactured with dimensions below the millimetre scale, meeting the granularity requirements for the RICH Upgrade II detectors. In addition, the QE of SiPMs peaks in the visible spectrum, especially in the 450–500 nm region, and can reach above 45 % [76]. As discussed in Section 1.3.2, the reduced pixel size, combined with the Upgrade II optics, and the high, green-shifted photon sensitivity improve the Cherenkov angle resolution. In terms of timing performance, intrinsic SPTR between 70 and 200 ps full width half maximum (FWHM) have been reported for commercially available devices [76]. Simulation studies have demonstrated that such time resolution, combined with the aforementioned detector upgrades would preserve the excellent RICH PID performance in the high-luminosity environment of Upgrade II [33].

The main challenge for the use of SiPMs in the RICH detectors is radiation-induced DCR. Neutron fluence damages the silicon lattice, creating generation-recombination centres that trigger spontaneous avalanches, increasing the DCR [77]. The expected fluence is ~ $3 \times 10^{12}\,\mathrm{n_{eq}/cm^2}$ for

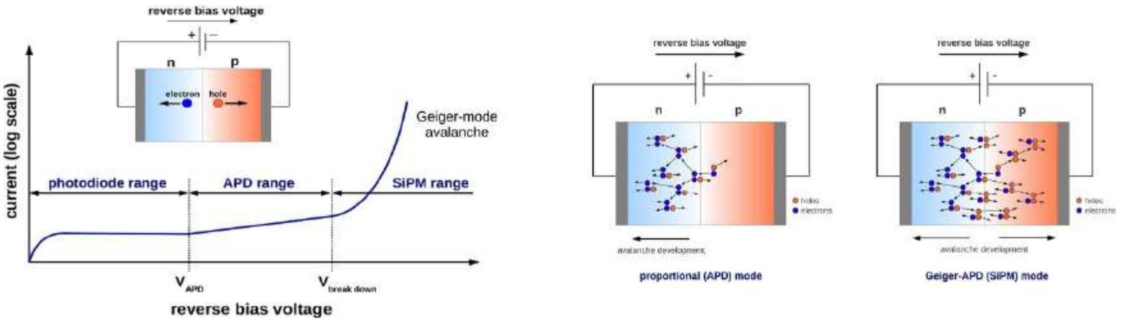

Figure 1.23: (Left) Operation regimes of solid state p-n junction, i.e. photodiode, avalanche photodiode (APD) and SPAD or SiPM range. (Right) In the APD only electrons can sustain the avalanche, whereas in a SPAD holes perform impact ionization as well [73].

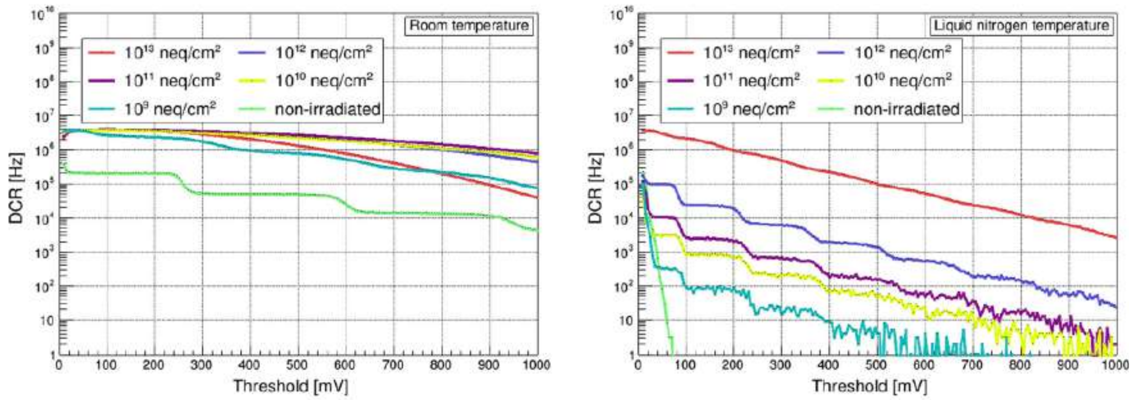

Figure 1.24: (Left) DCR measurements of FBK SiPMs at room temperature after the irradiation with different fluences. (Right) DCR measurements on the same SiPMs repeated at liquid nitrogen temperature [81]. Note: in the plot on the left the DCR is limited by the rate capability of the system.

RICH 1 detector. Studies have shown that SiPM irradiated at such a fluence can reach DCR value of the order of 100 MHz/mm^2 [78, 79]. At this rate, it becomes impossible to distinguish in-time Cherenkov photons from dark-counts. The DCR increases linearly with fluence but decreases exponentially with temperature, making cryogenic cooling the primary strategy to mitigate DCR [80]. Operation at cryogenic temperatures (down to approximately 100 K) has been shown to highly suppress the DCR even after radiation damage [81]. Figure 1.24 (left) shows the DCR of five FBK[3] $1 \times 1\,\mathrm{mm}^2$ SiPMs before and after irradiation with fluences ranging from 10^9 to 10^{13} 1 MeV neutron equivalent per cm^2 ($\mathrm{n_{eq}/cm^2}$). For fluences $> 10^{10}\,\mathrm{n_{eq}/cm^2}$, the DCR saturates at few MHz, reaching the limit of the acquisition system [81]. Figure 1.24 (right) shows the DCR from the same irradiated FBK SiPMs measured at liquid-nitrogen temperature. For fluences up to $10^{12}\,\mathrm{n_{eq}/cm^2}$, the DCR is restored to the non-irradiated level (or even lower), indicating that cryogenic cooling is an effective solution for DCR mitigation. However, implementing a cryostat solution for large photodetector areas presents several challenges. One of the main challenges is the transmission of signals from cryogenic-cooled SiPMs to room-temperature readout electronics. To address this, a dedicated flexible PCB (flex-PCB) for the readout of SiPM array at liquid-nitrogen temperatures has been designed and tested as part of this thesis. Details are presented in Chapter 4, including the flex-PCB design, signal integrity simulations and measurements, and the SPTR of the SiPM array. Additional strategies to mitigate the DCR in SiPMs are also under evaluation. For example, annealing reduces radiation damage [77]. Therefore, periodic annealing cycles at 40 °C during long shutdowns may help restore performance [79]. Micro-lenses installed on SiPMs to focus the

---

[3] Fondazione Bruno Kessler (`https://www.fbk.eu/`)

light onto a smaller SiPM area have demonstrated improvements in the SiPM detection efficiency and time resolution [82]. This technique allows to operate the SiPMs at lower voltages, reducing the DCR. Neutron shielding using polyethylene is another option to reduce exposure to secondary particles, hence limiting the radiation damage.

Another possible challenge for SiPMs is their signal shape. Due to the large number of SPADs in parallel, the SiPM has a large input capacitance, typically of the order of hundreds of picofarads. This results in a recovery time between 20 and 100 ns, which can lead to inefficiencies in the high occupancy regions of the RICH detector. While the input capacitance reduces with smaller pixel dimensions, further R&D is needed. In parallel, simulations studies are ongoing to quantify the effect of the recovery time on the detector performance.

Micro-Channel Plate PMTs (MCP–PMTs)

MCP-PMTs are valid candidates for the photon detection system in the RICH Upgrade II, particularly due to their excellent SPTR, high gain and granularity, and low dark count rate. These characteristics match well with the needs of a time-resolved RICH detector operating in a high-occupancy environment. An MCP-PMT consists of a photocathode, one or more microchannel plates, and a segmented anode. The principle of an MCP-PMT is shown in Figure 1.25. A photoelectron emitted at the photocathode is accelerated by an electric field towards the MCP and enters a microchannel – a microscopic, tilted glass capillary coated with a secondary emissive layer. As the photoelectron traverses the channel, it generates secondary electrons through multiple collisions with the channel walls. This cascade process results in a gain of up to 10^6 Me, analogous to the amplification in the dynode chain of MAPMTs. To increase gain and suppress ion feedback, two MCPs are typically stacked in a Chevron configuration. The time resolution improves with reduced spacing between the MCPs and increased inter-plate bias voltage. The amplified signal is collected at the anode, which can be segmented to provide spatial resolution [83]. In addition, MCP-PMTs are radiation hard and have low sensitivity to the magnetic field. Flexibility in the choice of photocathode, MCP structure, and anode layout has led to several MCP-PMT designs, tailored for different applications. In the context of LHCb Upgrade II, MCP-PMTs developed by Photek have been evaluated extensively for the TORCH detector, as described in Section 1.3.2. Many of the considerations made by the TORCH collaboration are compatible with the requirements of the RICH Upgrade II detector. These devices have demonstrated a time resolution of 23–30 ps [84, 85] and quantum efficiencies in the range of 20–25 % in the visible spectrum [86]. In addition, the MCP-PMTs are radiation hard and have a low sensitivity to the magnetic field due to the MCP design.

Two other interesting families of the MCP-based detectors are the hybrid detectors [87] and the Large-Area Picosecond photodetectors (LAPPD)[88]. The charge collection in the hybrid detector is performed by the TimePix 4 ASIC, developed for the VELO Upgrade II. This solution provides

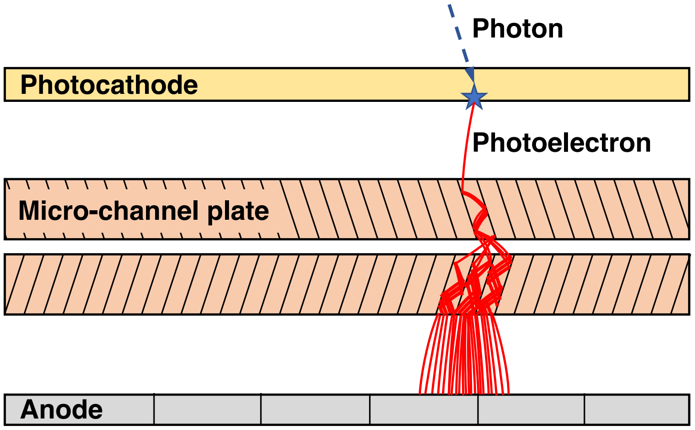

Figure 1.25: Schematic of the cross section of an MCP detector, showing the photo- cathode, two micro-channel plates in chevron orientation and the anodes. The photoelectron current is illustrated in red. Similar to the MAPMT, a resistor network is used to (negatively) bias the photocathode and MCPs. Source [26].

a spatial resolution of 10 μm and has a TDC with a bin size of approximately 200 ps. The time resolution can be further improved by collecting the charge from multiple anodes. LAPPDs are developed by Incom and are characterised by their large MCP area with atomic layer deposition (ALD) coating. They have an active area of 195×195 mm^2, but R&D is still ongoing to reach anode granularity of ~1 mm^2 over the full area. A new version, the High-Rate Picosecond photodetector (HRPPD), has been recently produced. The HRPPDs have a pixel size of 3.25 mm^2 over an active area of 104×104 mm^2. The anodes are directly coupled to the MCPs, leading to a time resolution of ~20 ps and improved charge sharing [89].

Despite these advantages, MCP-PMTs face challenges in terms of lifetime and rate capability. Their ageing is driven by the accumulated anode charge, which degrades the gain and quantum efficiency over time. Although ALD coating applied to the microchannels reduces this effect and increases the lifetime of the MCP-based detectors, this is still work in progress. Studies have shown undesired side effects of ALD coating, including lower saturation current and increased gain recovery time after saturation [90]. In addition, MCPs exhibit gain saturation at high instantaneous currents, which leads to recovery times and potential dead-time effects under continuous high-rate operation [72]. For these reasons, while MCP-PMTs offer outstanding timing performance, further improvements in lifetime and rate tolerance are required before large-scale deployment in the RICH Upgrade II can be confirmed.

1.4.3 A prototype optoelectronic readout chain with sub-nanosecond time capabilities and its evolutions

The development of timing capabilities for the RICH detectors required an intermediate step before the availability of the FastRICH ASIC. The RICH collaboration developed a prototype optoelectronic chain to study the timing performance of different photodetector technologies and to validate multi-channel readout electronics with sub-nanosecond precision. The scheme of the prototype optoelectronic chain and its evolutions are shown in Figure 1.26.

The first prototype and the 2022 LHCb RICH beam test

The first prototype of the readout chain [91] is built around the FastIC ASIC [69], an 8-channel front-end chip designed in 65 nm CMOS technology, predecessor of the FastRICH. FastIC processes signals in current mode with a low input impedance of about 16 Ω and includes a fast discriminator stage with nominal time jitter of 25 ps. The leading edge of the comparator output encodes the time-of-arrival (ToA), while the pulse width retains non-linear ToT information [70]. This design allows off-chip correction for time-walk and provides flexibility to couple the ASIC to different sensor types, such as MAPMTs, SiPMs or MCP-based photodetectors. The ASICs are mounted on custom front-end boards (FEBs), each hosting up to eight chips. The FEB layout is compatible with the CLARO-based FEB of the current RICH detector and optimised for short analogue trace lengths between the sensor and the FastIC inputs. Digital outputs of the FastIC are routed to a dedicated digital board carrying a Xilinx Kintex-7 FPGA. The FPGA is used for configuring the FastIC registers, transmitting data to the back-end, receiving a read-out trigger signal and an external clock, and providing a gate signal that can be used for flow control. To provide the hit timestamp, a custom multi-channel time-to-digital converter (TDC) is implemented in the FPGA using a multi-phase clock sampling architecture [92]. This TDC uses 16 sampling flip-flops (FFs). Each FF is clocked by one of eight phase-shifted copies of a 420 MHz clock, giving a theoretical bin

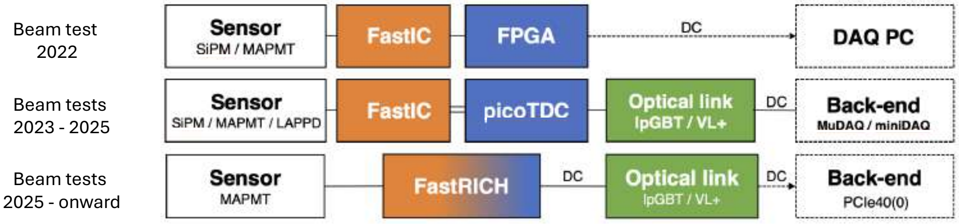

Figure 1.26: Schematic of the prototype optoelectronic readout chain and its evolutions. Source [91].

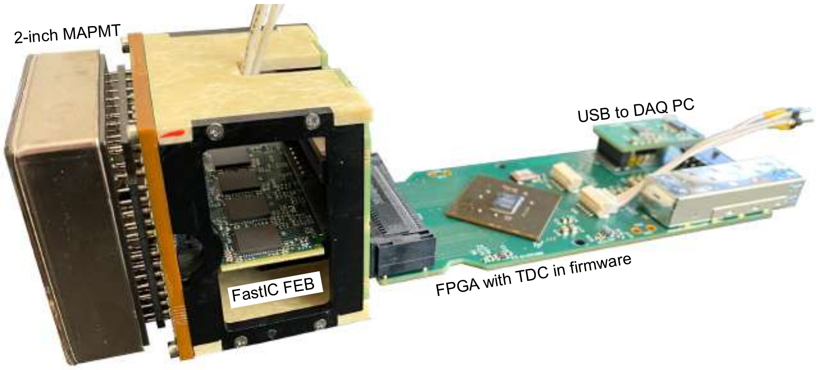

Figure 1.27: First prototype of optoelectronic readout chain with a 2-inch MAPMT. Source [91].

width of 148.8 ps. In practice, the bin widths within each 16-bin block are non-uniform and span in the range 60–250 ps [92]. The number of input channels of the TDC-in-FPGA is limited to 34. Figure 1.27 shows the first prototype optoelectronics chain assembled with a 2-inch MAPMT.

The first prototype was tested at the CERN SPS in 2022 with a beam of 180 GeV/c charged-hadrons, mainly pions and protons. The optoelectronic chains were placed inside a light-tight box in the beam line, as shown in Figure 1.28 (A). MAPMT and SiPM photodetectors were used during the beam test campaign. A borosilicate lens was used both as Cherenkov radiator and to focus the Cherenkov photons onto the photodetectors. Figure 1.28 (B) shows the photodetector with the nominal Cherenkov-ring position indicated by the orange dotted-line and the hitmaps from data. The schematic of the setup is illustrated in Figure 1.28 (C). A crossed-finger scintillator detector and a MCP-PMT placed on the beam line were used as trigger source. The coincidence of the two was performed by the trigger-logic unit in the FPGA and distributed to the readout boards, marking the start of the TDC readout window [91]. A SPTR of 182 ± 24 ps was obtained for the 1-inch MAPMTs, comparing the MAPMT signals to the MCP-PMT reference, and correcting for time-walk and non-uniformities in the TDC bin width. The result is consistent with the expected 150 ps and was validated with analogue and digital measurements presented in Chapter 2. This result represented the first demonstration of sub-nanosecond timing with a full prototype optoelectronic chain in the LHCb RICH beam tests. The study concentrated on the 1-inch MAPMTs, which will be the main photodetector in the RICH detector during LHC Run 4, while later campaigns have broadened the programme to include 2-inch MAPMTs, SiPM arrays and MCP-based photodetectors.

Evolution of the prototype and a novel data acquisition system

The experience gained with the first prototype chain in 2022 highlighted some limitations. The optoelectronic chain was therefore redesigned for the 2023 beam test campaign onward, providing

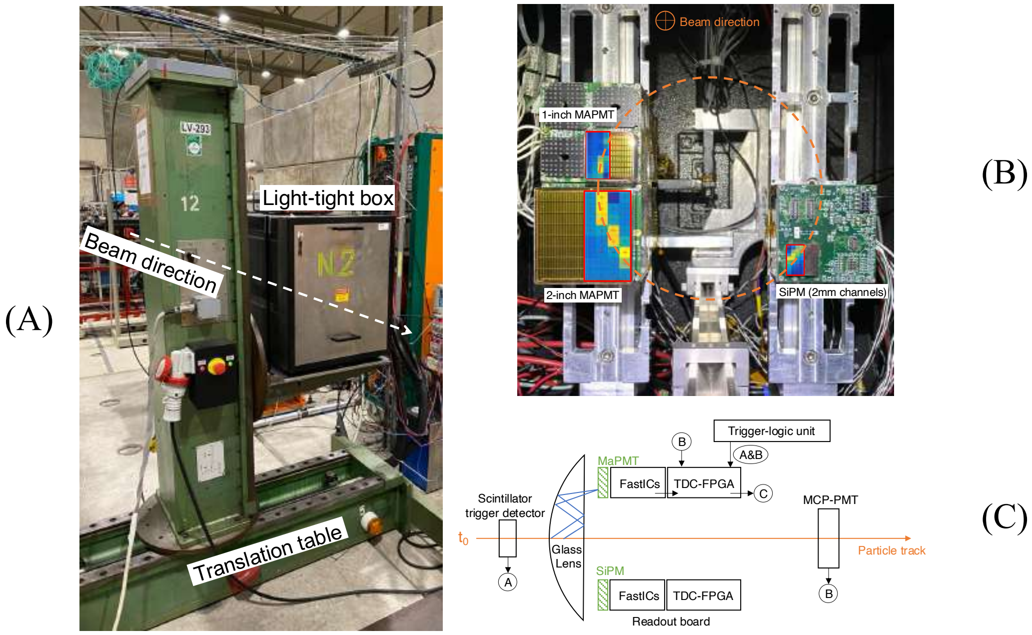

Figure 1.28: (A) Image of the light-tight box mounted on a remotely-operated translation table for alignment with the beam. (B) Image of part of the beam test setup viewed from the beam direction. The orange-dotted line represents the position of the Cherenkov photon ring and the integrated hit maps are overlaid on the MAPMT and SiPM sensors. Source [91].

improved modularity, finer TDC resolution, higher number of channels and optical link readout. These features were achieved with an ASIC-only solution and a new modular carrier board, as shown in Figure 1.29. The carrier board hosts dedicated plug-ins for each key component of the readout chain. Two FastIC plug-ins, with four FastICs each, allow the readout of 64 channels of the sensor connected to the prototype. The picoTDC plug-in hosts the picoTDC, a 64-channel TDC ASIC in 65 nm CMOS technology [93]. It enables high-precision timing measurements with selectable bin sizes of 3 ps or 12 ps and an intrinsic jitter below 1 ps. The data transmission is handled by the lpGBT plug-in, based on the Low-Power Gigabit Transceiver ASIC, a radiation tolerant serializer/deserializer device designed in 65 nm CMOS technology [94]. The lpGBT is the baseline chip adopted for data transmission in the LHCb Upgrade II. Its integration in this prototype allows thorough testing of its operation and compatibility with the readout electronics. Its output signals are sent to the back-end via optical links, using either a Versatile Link Plus transceiver (VTRx+) [95] or a small form-factor pluggable (SFP) module. This plug-in meets the high standard of future HL-LHC LHCb detector operation with 10 Gbps bandwidth. This data transmission configuration will already be deployed in LHC Run 4 together with the FastRICH readout. Its use in this system provides important integration advantages and long-term benefits

Figure 1.29: Second prototype of optoelectronic readout chain consisting of a carrier board with (from left to right) the photon sensor board, FastIC ASIC plugins, picoTDC plugin and optical link plugin.

for the experiment. The successful implementation of this communication chain can serve as a model for other LHCb sub-detectors in view of Upgrade II. At the back-end, the data streams are received by the MuDAQ system. The MuDAQ consists of a lightweight environment developed for beam tests that provides clock distribution, trigger synchronisation, and multiplexed data readout for up to seven carrier boards, with event building and monitoring functionalities similar to the LHCb MiniDAQ [96]. Data transmission to the acquisition computer is performed via a Gigabit Ethernet link. The prototype chain thus combines the precise photodetector pulse discrimination of the FastIC with a stable, high-resolution TDC ASIC and reliable optical data transmission to the back-end. Figure 1.30 shows a picture of the described data acquisition system.

This upgraded prototype was used in the beam test campaigns of 2023, 2024 and 2025. The 2023 campaign represented the first successful coupling of the FastIC to the PicoTDC and lpGBT/VTRX+ operation, establishing the foundations for stable multi-channel operation of the chain. The subsequent campaigns built on this basis, consolidating the setup and extending the characterisation of different photodetectors. The 2024 campaign focused on the MAPMTs, which will continue to be the photodetector for the RICH detector during LHC Run4. In 2025, the emphasis shifted to Upgrade II photodetector candidates, in particular SiPM arrays and MCP-based photodetectors. For these campaigns the setup was similar to that used in 2022 and is described in detail in Chapter 4, with particular focus on the 2025 beam test where the custom flex-PCB developed for SiPM arrays was also tested.

The design and operation of this second-generation prototype chain placed the collaboration in a strong position for the integration of the FastRICH ASIC in 2025. At the time of writing, work is ongoing on a third prototype in which the FastRICH replaces both the FastIC and the PicoTDC, integrating their functions into a single ASIC, as illustrated in Figure 1.26. With this development, the optoelectronic chain takes the form that will be used in the RICH detectors during LHC Run 4. The forthcoming beam tests in late 2025 and 2026 will therefore be of central importance: the first

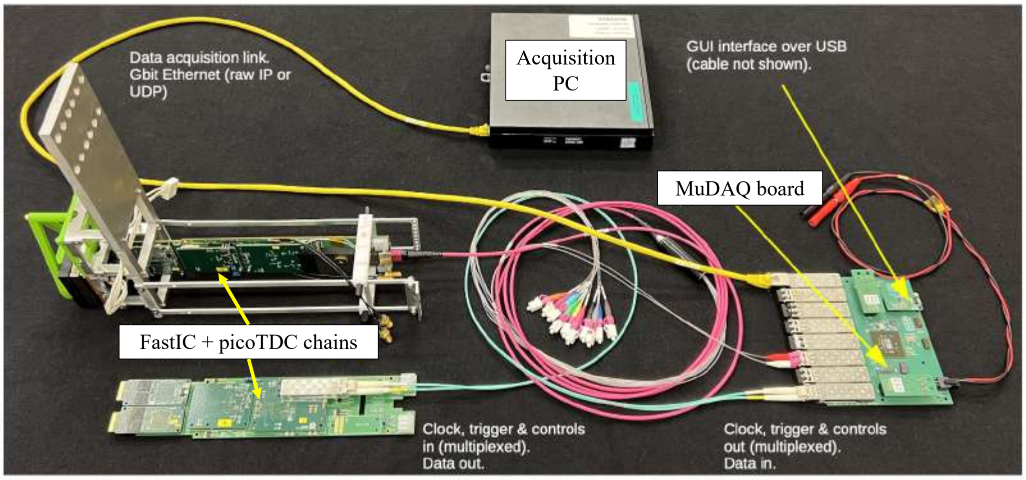

Figure 1.30: Connection in between the different system components reproduced in the laboratory.

will test the coupling of the FastRICH to MAPMTs, while the second will focus on the integration and stable operation of multiple boards.

Chapter 2

Analogue and digital studies on MAPMT single-photon response

Multi-anode photomultiplier tubes (MAPMTs) are the photodetectors used in the current LHCb RICH detectors. They will remain in use for LHC Run 4, coupled to the new FastRICH-based readout electronics. This system provides a timestamp for each detected photon, adding timing information that can improve particle identification (PID) performance. A full characterisation of the MAPMT single-photon response is therefore required to quantify their intrinsic timing performance and to optimise the coupling to current-mode front-end electronics such as the FastIC and FastRICH ASICs. This chapter presents such a characterisation with analogue and digital readout systems under pulsed-laser illumination. Section 2.1 reports analogue measurements of single-photon pulse shape and time resolution in individual channels of the two MAPMT types described in Section 1.4.2. Section 2.2 describes multi-channel digital measurements on MAPMT time resolution performed with the time-resolved readout electronics prototype developed for the LHCb RICH beam tests.

2.1 MAPMT analogue single-photon response

This section presents the analogue measurements performed on the MAPMTs to study their single-photon response. The measurements were performed on a single channel of both the 1-inch and a 2-inch MAPMT. These measurements provided key input for coupling MAPMTs to the FastIC ASIC, used in the time-resolved readout prototype for the LHCb RICH beam tests. The pulse shape also served as an input for the design of the FastRICH analogue chain. In addition, they define the intrinsic single-photon time resolution (SPTR) from a fully illuminated MAPMT channel, setting a reference for multi-channel digital readout with both laser illumination and in beam tests. The setup for single-photon illumination, its electrical simulation and the waveform analysis method are presented in Section 2.1.1. Section 2.1.2 describes the analysis and results on the pulse shape characterisation. The focus was on key parameters such as amplitude, rise time, and gain, which were then used to estimate the expected current at the front-end electronics. Section 2.1.3 presents the measurement of the transit-time spread (TTS), which quantifies the fluctuations in the photon arrival time at the anode defining the MAPMT time resolution.

2.1.1 Experimental setup and waveform analysis

This section describes the setup used to measure the single-photon analogue response of an individual channel of the two MAPMT types used in the RICH detector. The setup was used to record single-photon waveforms and extract pulse parameters such as amplitude, rise time, charge, and TTS. Figures 2.1 and 2.2 show the setup and the corresponding connection scheme overview. The illumination was provided by a picosecond pulsed laser[1], which delivers light pulses at a wavelength of 405 nm. The laser optical power was set to approximately 200 nW with an illumination rate of 10 kHz. This configuration ensured a time jitter of ~50 ps (FWHM) for each pulse, corresponding to the best timing performance of the laser. Laser pulses were guided by a single-mode optical fibre to a to a microscope objective[2]. The lens was used to spread the light over the MAPMT surface. A neutral-density filter with optical density 3.00 was placed before the MAPMT for single-photon illumination. The MAPMTs were covered by a mask with 2 mm and 4 mm holes aligned to a single channel for the 1-inch and 2-inch sensors, respectively, to minimise crosstalk effects. The MAPMTs were mounted on the same baseboards used in the current RICH system [97]. Each MAPMT anode was routed to a multi-pin output connector through the custom fanout board shown in Figure 2.3 (left). Each output line included a termination resistor between an output pin and a ground pin, allowing the analogue signal to be probed directly across the two. The fanout is shown in Figure 2.3, with the multi-pins connectors on the front and the termination resistors on the back. Analogue signals were acquired with a differential probe[3] equipped with a solder-in tip[4]. As shown in Figure 2.3 (right), the probe tip was soldered into a custom two-pin plug-in connector. This configuration simplified the connection to the fanout board, although it may have introduced additional noise in the analogue readout. Analogue signals were then sampled with a 12 GHz oscilloscope[5]. The synchronisation output (`Sync-out`) signal from the laser driver was sent

---

[1] PicoQuant TAIKO pulsed laser (`https://www.picoquant.com/products/category/picosecond-pulsed-driver/taiko-pdl-m1#specification.`).

[2] Edmund Optics, 40× DIN 0.65 Achromatic Commercial Grade Objective (`https://www.edmundoptics.com/p/40x-din-achromatic-commercial-grade-objective/5383/`).

[3] Agilent 1169A 12 GHz probe (`https://www.keysight.com/pd-1169A`).

[4] Agilent N5426A ZIF tip (`https://www.keysight.com/pd-N5426A`).

[5] Agilent DSA91204A 12 GHz oscilloscope (`https://www.keysight.com/pd-DSA91204A`).

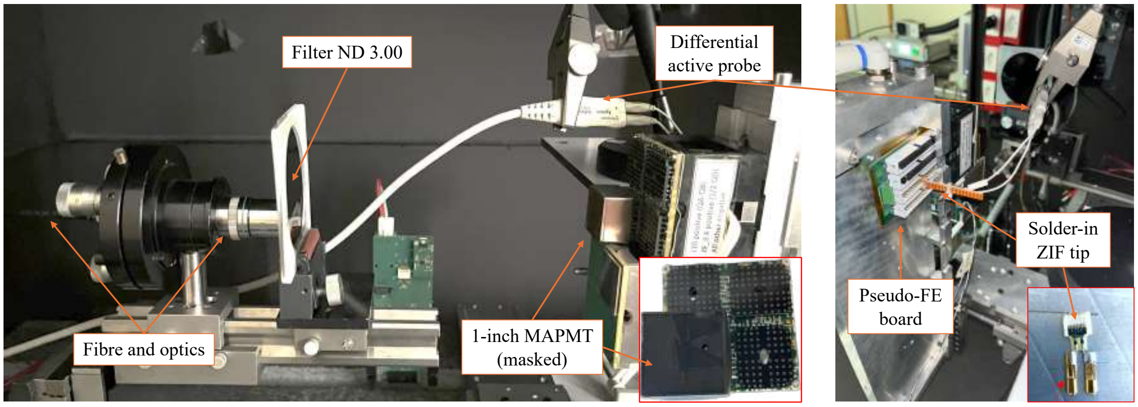

Figure 2.1: Experimental setup for single-photon illumination and analogue readout of the 1-inch MAPMT. A cover is placed on the MAPMT with a 2 mm hole to illuminate a single channel and reduce crosstalk effects. Same setup is used for the for the 2-inch MAPMT, applying a mask with a 4 mm hole.

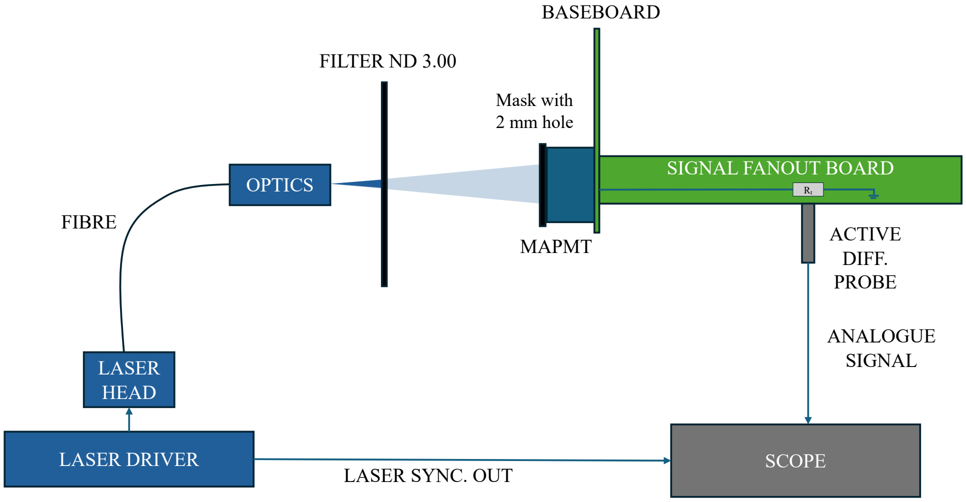

Figure 2.2: Overview of the experimental setup connections for the 1-inch MAPMT.

to the oscilloscope. The `Sync-out` signal was used as the oscilloscope trigger and as a time reference for the time resolution measurements, as described in Section 2.1.3.

The termination resistor value was chosen based on circuit simulations. The setup equivalent

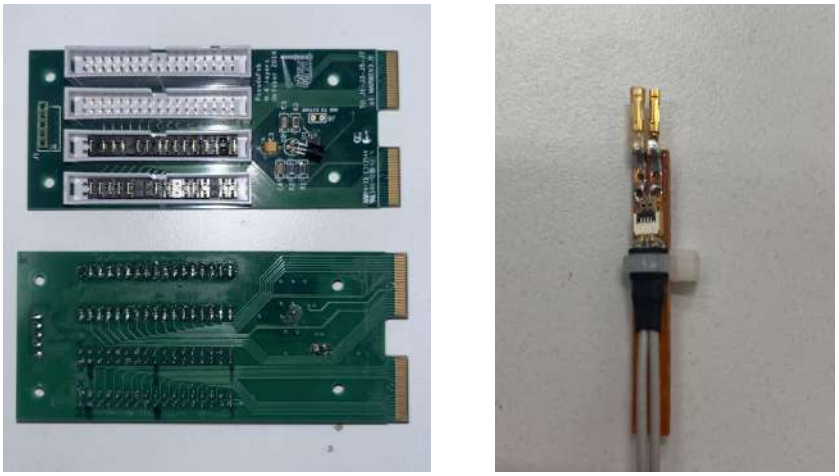

Figure 2.3: (Left) The signal fanout board routes analogue signals from individual MAPMT channels through dedicated trace lines. The top side (upper image) includes high-density connectors to allow individual channel readout. The bottom side (lower image) provides solder pads for termination resistors. (Right) Agilent N5426A solder-in ZIF tip soldered in a custom two-pin plug-in connector.

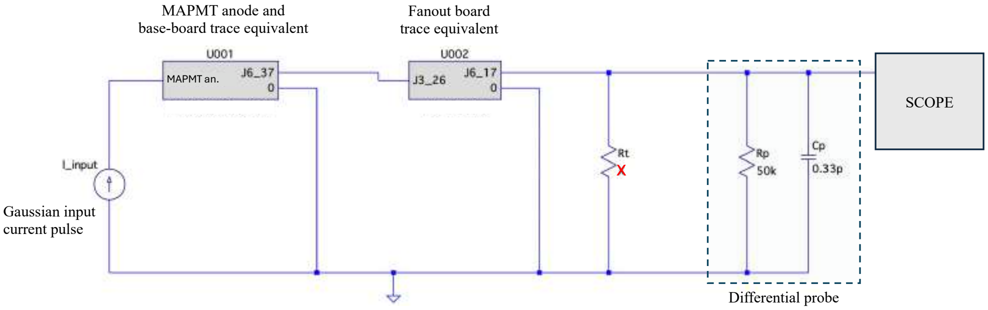

Figure 2.4: LTspice simulation schematic of the readout system.

circuit was implemented in the LTSpice[6] and it is shown in Figure 2.4. The equivalent circuits of the MAPMT anode, the baseboard trace, and the fanout board trace were provided by the board designers. These were modelled as RLC circuits reproducing the impedance and parasitic effects of the connections and vias from the MAPMT anode to the termination resistor. The current collected at the anode for a single-photon pulse was modelled as a Gaussian pulse with a standard deviation of 350 ps. The amplitude was selected to obtain a total charge of 2 Me. The simulated waveforms for three termination resistor values (30, 100, and 510 Ω) are shown in Figure 2.5. They

---

[6] LTSpice: `https://www.analog.com/en/resources/design-tools-and-calculators/ltspice-simulator.html`

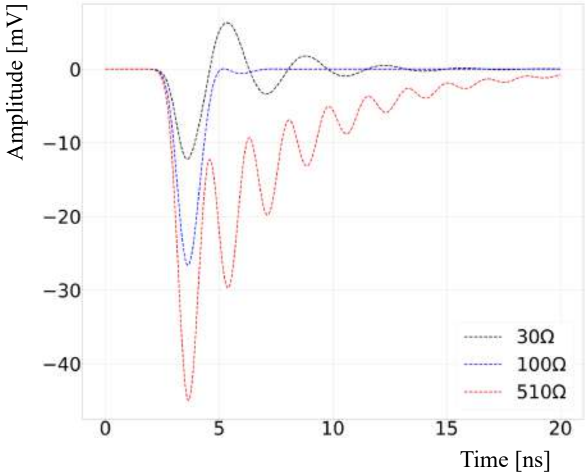

Figure 2.5: Voltage response of the simulated readout system to a Gaussian current input (σ = 350 ps and 2 Me charge) for 30, 100, and 510 Ω termination resistor.

illustrate the effect of the termination resistor on the pulse shape. Table 2.1 summarises the results for the three termination values. The rise time is measured as the 10–90 % interval of the pulse leading edge, while the charge is obtained from the integrated area of the voltage pulse divided by the termination value. The simulations show that the termination resistor had no measurable effect on the pulse leading edge: the rise time remained consistent within a few picoseconds for all resistor values. This behaviour is expected, since the rising edge is determined by the MAPMT capacitance and the shape of the input current pulse, rather than by the external load. In contrast, the termination value strongly affected the pulse amplitude and decay time. Higher resistances produced larger amplitudes and longer tails, as the same current flowed across a higher impedance. Nevertheless, the integrated charge, normalised to the resistance, remained constant within 3 %, confirming that the MAPMT gain is independent of the termination choice. For the 30 and 510 Ω terminations, the simulated waveforms show additional peaks on the trailing edge of the pulses. These peaks originate from reflections due to the impedance mismatch between the termination resistor and the trace impedance, estimated to be 98 Ω. This interpretation is confirmed by the absence of such peaks in the 100 Ω case, where the termination matches the trace impedance and reflections are suppressed. Importantly, these reflections appear only on the trailing edge and do not affect the leading edge of the pulse. For this reason, although 100 Ω provides better matching, the final measurements were performed with a 510 Ω termination. The higher termination value increases the pulse amplitude for the same input current, improving the signal-to-noise ratio and facilitating waveform fitting. The ringing from the impedance mismatch on the trailing edge was considered acceptable, as it did not affect the timing performance of the MAPMT response.

Table 2.1: Simulated rise time and charge for different termination resistor values. The input current pulse was a Gaussian with 350 ps standard deviation and a total charge of about 2 Me.

Termination R [Ω]	Rise time (10–90 %) [ps]	Charge [Me]
30	732 ± 3	1.99 ±0.01
100	733 ± 3	1.99 ±0.01
510	735 ± 3	1.95 ±0.01

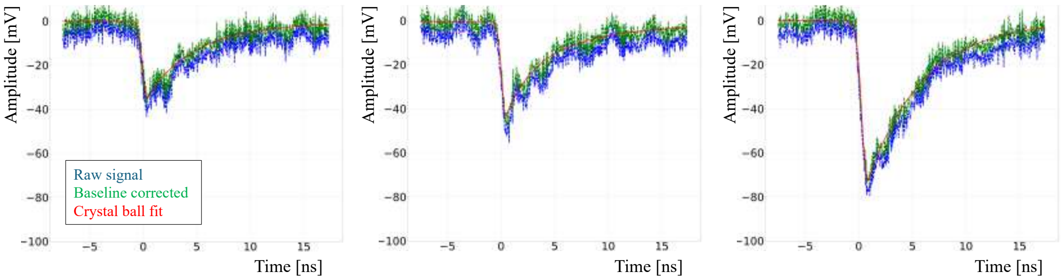

Figure 2.6: Examples of crystal-ball fit on the MAPMT single-photon pulses after baseline correction. The pulses were recorded with the MAPMT biased at −1000 V.

Fifty thousand waveforms were recorded for both MAPMT types at bias voltages of −900 V and −1000 V. Typical waveforms recorded from the 1-inch MAPMT channel at −1000 V are shown in Figure 2.6. As expected, the noise level is relatively high, with a peak-to-peak amplitude of about 10 mV, which most likely originates from the probing system described above. The recorded waveforms agree with the simulated ones, including the reflections on the trailing edge observed for the mismatched termination values. Each waveform was first baseline-corrected and then fitted with a Crystal Ball function, consisting of a Gaussian leading edge and an exponential tail. The fit was iterated, and the best result was selected based on the reduced χ^2 value. The fit primarily suppresses high-frequency noise and allows extraction of the relevant parameters from a smooth analytical function instead of raw data points. The amplitude, rise time, total charge, and time-of-arrival (ToA) used for the TTS analysis were all extracted from the fit parameters, ensuring consistency across measurements.

2.1.2 Results on single-photon pulse shape

This section presents the results on the single-photon pulse shape obtained from the waveform analysis introduced in Section 2.1.1. The discussion focuses on the pulse amplitude, rise time, gain, and expected input current at the front-end electronics. Figure 2.7 shows the pulse-height spectrum for the 1-inch MAPMT channel at −900 V. The spectrum agrees with the studies conducted in the

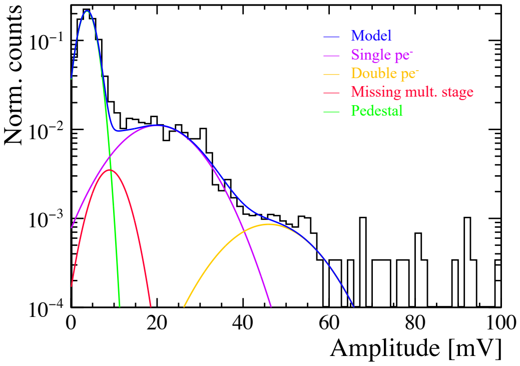

Figure 2.7: pulse-height spectrum of a 1-inch MAPMT channel at −900 V. The distribution is fitted with four Gaussian components describing the noise pedestal (green), the single-photon peak (purple), the double-photon peak (yellow), and events with one missing multiplication stage in the dynode chain (magenta). The local minimum between the noise and single-photon peaks defines the amplitude cut for single-photon selection.

Quality Assurance commissioning phase for the current RICH detector [98], shown in Figure 1.22. Since the oscilloscope was triggered by the laser `Sync-out` signal and the light intensity was highly attenuated to ensure single-photon operation, only ~3 % of the triggers contain a photon pulse. The most populated peak at low amplitude corresponds indeed to the noise pedestal. Its width of a few millivolts reflects the high-frequency noise observed in Figure 2.6. The second peak, or plateau, originates from single-photon pulses. Its spread results from the gain fluctuations in the MAPMT dynode chain. The valley between the noise pedestal and the single-photon peak is caused by events in which one multiplication stage is skipped in the dynode chain [98]. A smaller contribution at higher amplitudes arises from double-photon pulses, where two photons hit the same pixel within a few picoseconds. All four contributions were modelled with Gaussian functions to fit the pulse-height spectrum, as shown in Figure 2.7. This is a simplified approach compared to the procedure adopted in the MAPMT Quality Assurance [98]. The fit simplifies the pulse-height spectrum and provides a straightforward way to discard noise pulses applying an amplitude cut. The cut value is defined as the local minimum between the noise peak and the single-photon peak. Figure 2.8 shows the pulse-height spectra for the two MAPMT types at −900 V and −1000 V. The same fit strategy is applied to all spectra. The average single-photon amplitude increases with the applied voltage, reflecting the higher gain achieved at larger bias values. Additionally, the 1-inch MAPMT shows a higher average single-photon amplitude than the 2-inch MAPMT, due to its typically larger gain compared to the 2-inch device. For the 2-inch MAPMT, the ratio between the single-photon peak and the noise pedestal is higher than for the 1-inch. This is because the same laser illumination settings were used for both MAPMT types to preserve the laser timing performance. With identical

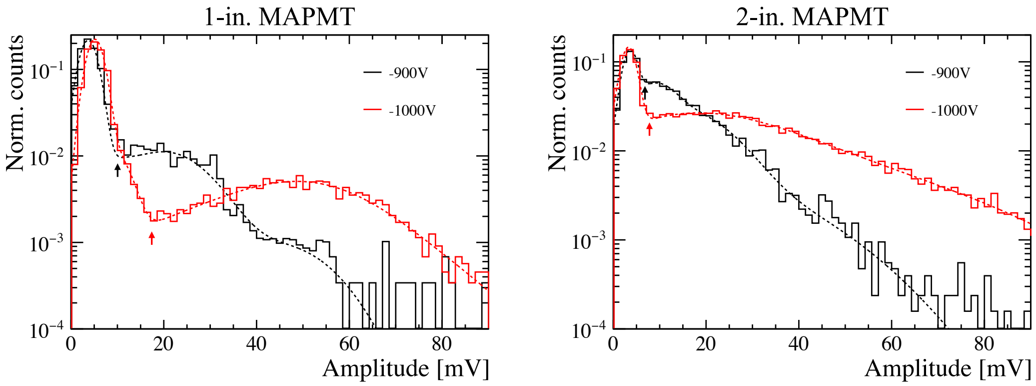

Figure 2.8: Pulse-height spectra of a single channel for the 1-inch (left) and 2-inch (right) MAPMTs at −900 V and −1000 V with the fit superimposed. The arrows indicate the amplitude cut, defined as the local minimum between the noise and single-photon peaks in the fit.

illumination, the larger 2-inch channel collected more photons, resulting in a higher occupancy. Nevertheless, the occupancy remained ~10 %, ensuring single-photon operation with double-photon probability below 10 %. The arrows in Figure 2.8 indicate the amplitude cuts applied to the fitted waveforms for selecting photon pulses.

For each selected pulse, the rise time was extracted as the time interval between 10 % and 90 % of the leading edge of the fitted waveform. Figure 2.9 shows the rise-time distributions for the 1-inch and 2-inch MAPMT channels at −900 V and −1000 V. The distributions were fitted with a double-Gaussian function: one Gaussian for the core and another to account for the tails. The average rise time was extracted from the peak position of the fit. The widths of the distributions reflect event-to-event fluctuations in the multiplication process and in the initial photoelectron trajectory. No significant dependence on the bias voltage was observed for the rise time. The average rise time is approximately 640 ps for the 1-inch MAPMT and 300 ps for the 2-inch. The faster rise time in the 2-inch MAPMT is likely due to the reduced number of dynodes (ten compared to twelve in the 1-inch device), resulting in a shorter overall electron transit time across the multiplication chain. The numerical values extracted from the fits are summarised in Table 2.2.

In the MAPMTs, the single-photon gain is defined as the total charge collected at the anode when a photon hitting the photocathode is converted into a photoelectron and amplified through the multiplication chain. In this study, the total charge of each pulse was obtained by integrating the fitted waveform and dividing the area by the termination resistance. This value was then expressed in Me for comparison with the gain values from Quality Assurance campaign [98] and Hamamatsu datasheet. Figure 2.10 shows the charge distribution for the 1-inch MAPMT channel at −900 V. Following the fit strategy used for the pulse-height spectra, the distribution is fitted with

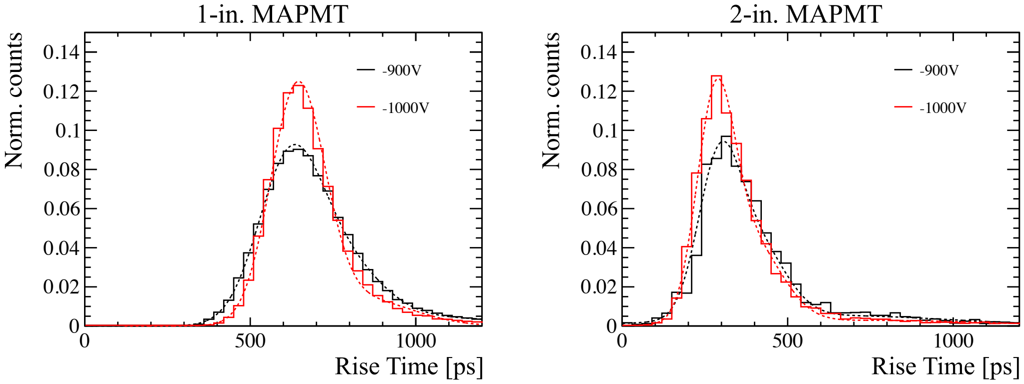

Figure 2.9: MAPMT single-photon rise time for −900 V and −1000 V.

Table 2.2: Average single-photon rise time for the 1-inch and 2-inch MAPMT channels at −900 V and −1000 V. The uncertainties represent the standard deviation of the distributions.

MAPMT type	Bias [V]	Rise time [ps]
1-inch	−900	635 ± 60
	−1000	640 ± 65
2-inch	−900	315 ± 50
	−1000	290 ± 45

Gaussian functions. A main Gaussian corresponds to single-photon events, and its peak position is taken as the average single-photon gain, while a second Gaussian accounts for double-photon pulses. The third component, describing events in which one multiplication stage of the dynode chain is skipped, was neglected in this gain analysis. Because of the large electronic noise in the readout system, most of these low-amplitude events fall below the amplitude selection threshold and are therefore excluded. The omission of this component has no impact on the results, as its expected contribution is an order of magnitude lower than the single-photon distribution.

Figure 2.11 show all four charge distributions, where the double-Gaussian fit is applied. As expected, the measured gain increases with the applied bias voltage. For the 2-inch MAPMTs, the distributions exhibit a more pronounced tail at high charge values, caused by the higher photon occupancy and resulting double-photon probability. The average single-photon gain values are summarised in Table 2.3, along with the results from the Quality Assurance campaign [98] and the values provided by Hamamatsu. The three sets of measurements are broadly consistent: the present results show a slight overestimation compared with the QA values in some cases, but remain compatible within uncertainties with the Hamamatsu ones. It should be noted that gain measurements are particularly sensitive to the experimental setup and to the analysis method. The three values

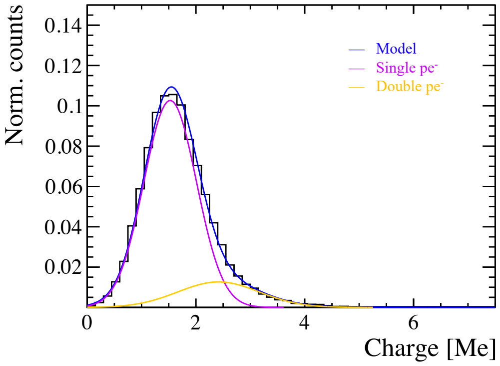

Figure 2.10: Charge distribution for the 1-inch MAPMT at −900 V with the fit model superimposed. The two Gaussian components describe single-photon and double-photon pulses. The peak position of the single-photon Gaussian defines the average single-photon gain.

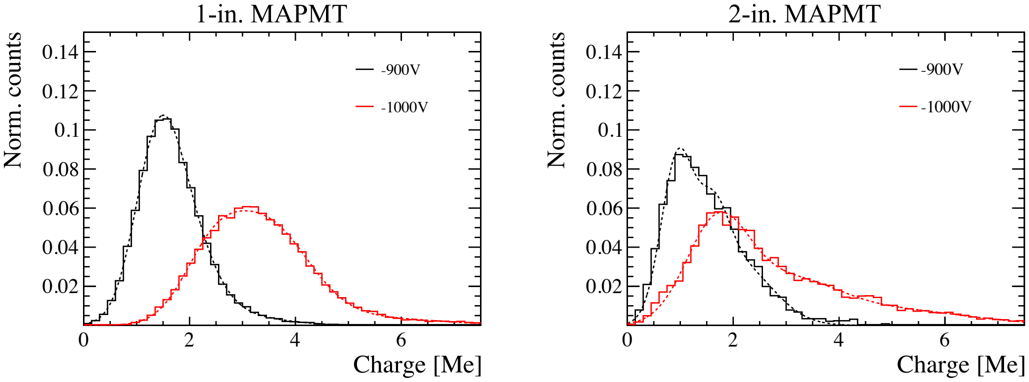

Figure 2.11: Charge distributions for single-photon pulses recorded from the 1-inch (left) and 2-inch (right) MAPMT channels at −900 V and −1000 V. The distributions are fitted with a double Gaussian function to account for the tail at high charge values from occasional double-photon events. The single-photon gain is extracted from the peak position of the main Gaussian component.

reported in Table 2.3 were obtained using different acquisition chains and analysis procedures. For instance, during the gain calibration campaigns on the current LHCb RICH detectors, aimed at the gain uniformity across the full detector plane, gain values were found to be slightly lower than the Hamamatsu ones [99].

As an independent validation of the analysis chain, a known input pulse was injected into the system, and the output charge was measured using the same procedure described in this section.

Table 2.3: Average single-photon gain for the 1-inch and 2-inch MAPMT channels at −900 V and −1000 V. Results are compared with the values from the Quality Assurance (QA) campaign and those provided by Hamamatsu. Uncertainties represent the standard deviation of the distributions.

MAPMT type	Bias [V]	Measured [Me]	QA [Me]	Hamamatsu [Me]
1-inch	−900	1.4 ± 0.4	1.1 ± 0.2	—
	−1000	3.0 ± 0.6	2.4 ± 0.5	2.6 ± 0.5
2-inch	−900	1.2 ± 0.3	0.7 ± 0.2	—
	−1000	1.8 ± 0.5	1.6 ± 0.4	1.9 ± 0.5

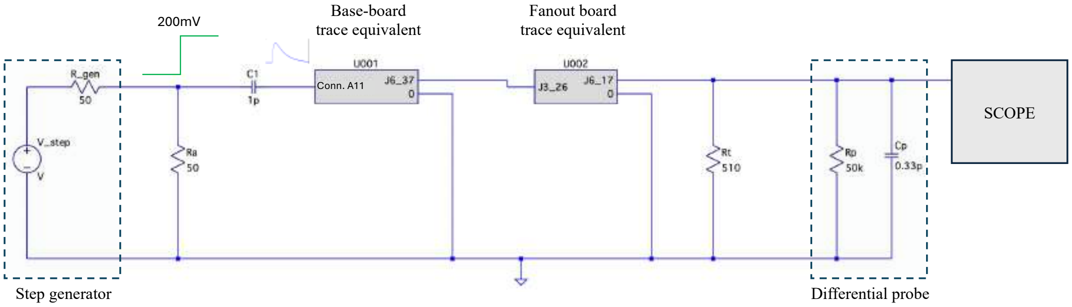

Figure 2.12: Schematic of the setup used for injecting a 1.25 Me pulse into the system.

The schematic of the setup for this test pulse measurement is shown in Figure 2.12. A 200 mV step pulse was generated with an Active Technology Pulse RIDER PG-1072 pulse generator [100] and sent through a 1 pF capacitor to generate a 1.25 Me MAPMT-like pulse. The pulse was then injected into the baseboard anode pin and measured across the termination resistor on the fanout board, following the same path as for the MAPMT measurements. Waveforms were also recorded and analysed with the same procedure. The resulting charge distribution is shown in Figure 2.13. The distribution mean of 1.22 Me agrees with the injected charge. This measurement demonstrated that the full acquisition and analysis chain accurately reconstructs the injected charge and validated the gain extraction method.

The measured gain and rise time values can be combined to estimate the peak current expected at the input of current-mode readout ASICs, such as FastIC and FastRICH. The nominal input impedance of such ASICs is approximately 30–50 Ω depending on the input frequency [69]. Using Ohm's law $I(t) = V(t)/R$, the expected current pulse can be obtained by converting the simulated voltage pulse across the 30 Ω termination resistor, shown in Figure 2.5. As a first approximation, the pulse across the 30 Ω termination can be modelled by a Gaussian function:

$$I(t) = I_{\text{peak}} \exp\left(-\frac{t^2}{2\sigma^2}\right), \tag{2.1}$$

where I_{peak} is the peak current. The total charge of the pulse, i.e. the gain, is obtained by integrating the current:

$$Q = \int I(t)\, dt = I_{\text{peak}}\, \sigma \sqrt{2\pi}. \tag{2.2}$$

The rise time (10–90 %) is obtained by solving $I(t) = 0.9\, I_{\text{peak}}$ and $I(t) = 0.1\, I_{\text{peak}}$ for t_{90} and t_{10} respectively, yielding

$$\Delta t_{90-10} = \sqrt{2}\,\sigma \left(\sqrt{\ln 10} - \sqrt{\ln \frac{10}{9}} \right) \approx 1.687\, \sigma. \tag{2.3}$$

From Equation 2.2, σ can be expressed in terms of the total charge and peak current as

$$\sigma = \frac{Q}{I_{\text{peak}} \sqrt{2\pi}}. \tag{2.4}$$

Substituting Equation 2.4 into Equation 2.3 gives

$$\Delta t_{90-10} = \frac{1.687\, Q}{I_{\text{peak}} \sqrt{2\pi}}, \tag{2.5}$$

and therefore the peak current can be written as

$$I_{\text{peak}} = \frac{1.687}{\sqrt{2\pi}} \frac{Q}{\Delta t_{90-10}} \approx 0.673\, \frac{Q}{\Delta t_{90-10}}. \tag{2.6}$$

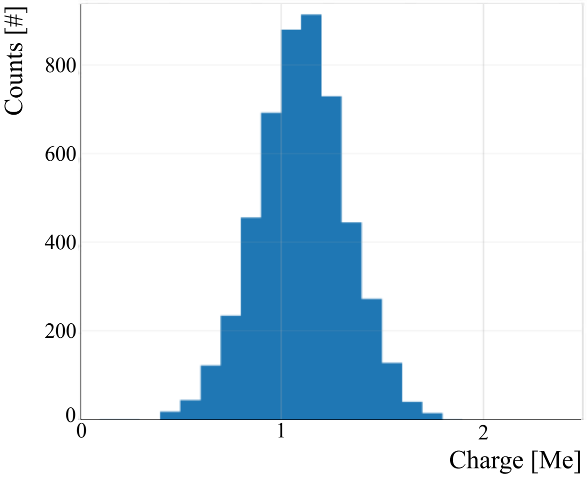

Figure 2.13: Distribution of the measured charge from the charge injection measurement.

This relation provides a simple estimate of the expected input current based solely on the measured gain and rise time. The expected input currents listed in Table 2.4 are obtained from the measured gain and rise time values in Tables 2.2 and 2.3. As expected, the shorter rise time of the 2-inch MAPMT results in a higher peak current for the same gain. In addition, this analytical estimate is consistent with the circuit simulations presented in Section 2.1.1. The simulated $30\,\Omega$ output pulse has a peak voltage of about $11\,\mathrm{mV}$, corresponding to a current of approximately $366\,\mu\mathrm{A}$. Using the Gaussian approximation with the simulated pulse gain of 2 Me and rise time of 732 ps, the expected peak current is about $300\,\mu\mathrm{A}$. The two values are in good agreement, showing that Equations 2.6 provides a reliable first-order estimate of the input current.

This estimate provided a useful reference for the FastRICH development, supporting the definition of the lower bound of the input-current dynamic range. In the current RICH detector, dedicated gain-calibration scans are performed to ensure uniform response across the photon-detector plane. For each group of 16 MAPMTs of 1-inch (or 4 of 2-inch) sharing the same bias line, the average anode gain is adjusted to about 1.2 Me. This value was used as a reference value to estimate the lower bound for the FastRICH input-current dynamic range. The peak current for a gain of 1.2 Me corresponds to ~ $200\,\mu\mathrm{A}$ for a 1-inch MAPMT channel. However, the gain can vary significantly between different MAPMTs, between anodes within the same device, and even from event to event within a single channel. To account for gain fluctuations, a safety margin was therefore included in the FastRICH design. From the gain distribution measured in this study, the standard deviation was found to be approximately 0.3 Me for a gain of 1.2 Me. This spread was used to define a conservative lower limit on the expected current. A 3σ variation below the nominal gain corresponds to $1.2-3\times0.3 = 0.3\,\mathrm{Me}$. For this lower-limit gain value, the estimated peak current is about $50\,\mu\mathrm{A}$. This estimated current value supported the definition of the lower bound of the input-current dynamic range of the FastRICH. In addition, it provided an important reference for the coupling between the MAPMTs and the FastIC ASIC, employed in the time-resolved readout chain developed for LHCb RICH beam tests.

Table 2.4: Summary of rise time, measured single-photon gain, and estimated peak current for the 1-inch and 2-inch MAPMT channels at different bias voltages.

MAPMT type	Voltage [V]	Rise time [ps]	Gain [Me]	I_{peak} [μA]
1-inch	−900	635 ± 60	1.5 ± 0.4	255 ± 65
	−1000	640 ± 65	3.0 ± 0.6	510 ± 100
2-inch	−900	315 ± 50	1.2 ± 0.3	430 ± 100
	−1000	290 ± 45	1.8 ± 0.5	650 ± 170

2.1.3 Results on the single-photon transit-time spread

The intrinsic time resolution of the MAPMTs will be the dominant contribution to the SPTR of the RICH detectors during LHC Run 4. Quantifying this resolution requires measuring the TTS, which reflects the statistical fluctuations in the photoelectrons transit time from the photocathode to the anode. Experimentally, the TTS is obtained by measuring the spread of the single-photon ToA with respect to a reference time. The analysis presented in this section was performed on the data sets described in Section 2.1.1. The ToA was extracted from the fitted waveforms passing the amplitude selection, while the laser `Sync-out` signal was used as the time reference. To accurately measure the pulse ToA, an in-software constant-fraction discriminator (CFD) was applied. For each fitted waveform, the pulse height was determined, and the ToA was defined as the time at which the signal reached a fixed fraction of this height. Using a CFD rather than an absolute threshold removes the dependence on pulse amplitude and suppresses the time-walk effect. This effect is caused by gain fluctuations in the electron multiplication chain. Pulses with different amplitudes would otherwise cross a fixed threshold at different times, introducing an artificial spread in the measured ToA, as shown in Figure 2.14. Figure 2.15 (left) illustrates the procedure on a representative waveform, showing the ToA extracted for CFD fractions of 10 % and 90 %. The corresponding laser `Sync-out` signal and its ToA are shown in Figure 2.15 (right). For the MAPMT pulses, the ToA values are obtained from the Crystal-Ball fit to minimise the impact of noise. The time difference between the MAPMT channel and the laser reference, $\text{ToA}_{\text{ch}} - \text{ToA}_{\text{Sync.out}}$, is computed for each event to build the time-difference distributions used for the TTS extraction. Figure 2.16 shows the resulting distributions for the 1-inch MAPMT at −900 V and for CFD fractions of 20 %, 40 %, and 80 %.

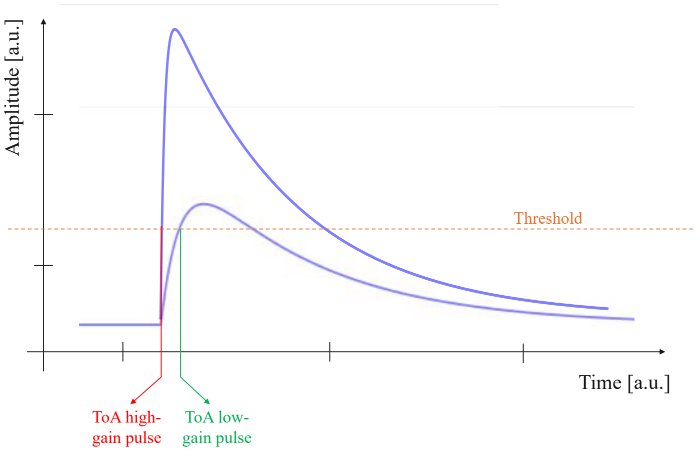

Figure 2.14: Time-walk effect on the MAPMT signal ToA.

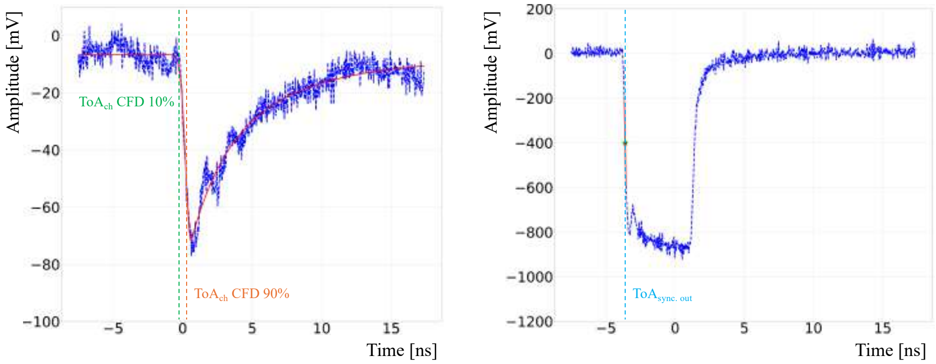

Figure 2.15: (Left) In-software CFD method used to extract the ToA on a single-photon MAPMT pulse. The CFD is applied to the fit of the waveform. (Right) Laser `Sync-out` waveform and its ToA used as time reference.

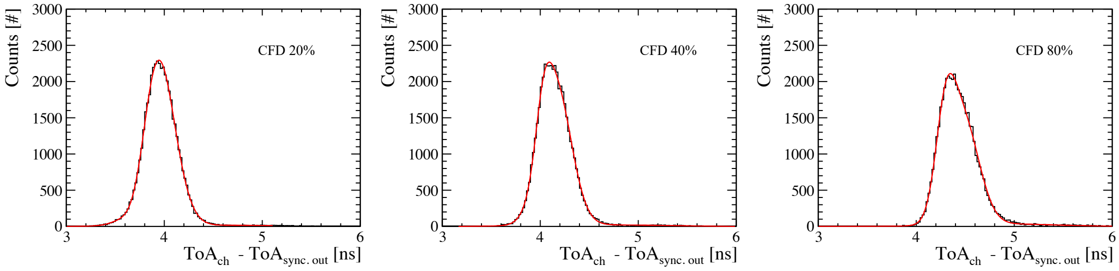

Figure 2.16: Time difference distributions for the 1-inch MAPMT at −900 V with different CFD fractions.

As expected, the mean of the distribution shifts to later times for higher CFD fractions, since the threshold is crossed closer to the pulse peak. For the CFD at 80 %, the right-hand tail is slightly more pronounced because the fit does not always reproduce the waveform accurately near the pulse peak, while the leading edge is fitted more reliably. Each distribution is then fitted with a double Gaussian function. The main Gaussian component describes the core distribution, while the secondary one accounts for the tails. The TTS is extracted as the FWHM of the fit.

The same analysis was performed for both MAPMT types at bias voltages of −900 V and −1000 V. The resulting TTS values as a function of the CFD fraction are shown in Figure 2.17. Operating at −1000 V yields better time resolution than at −900 V, with optimal performance for CFD fractions between 10–20 %. The 2-inch MAPMT shows better timing performance than the 1-inch device, in agreement with the Hamamatsu specifications. This is likely due to the

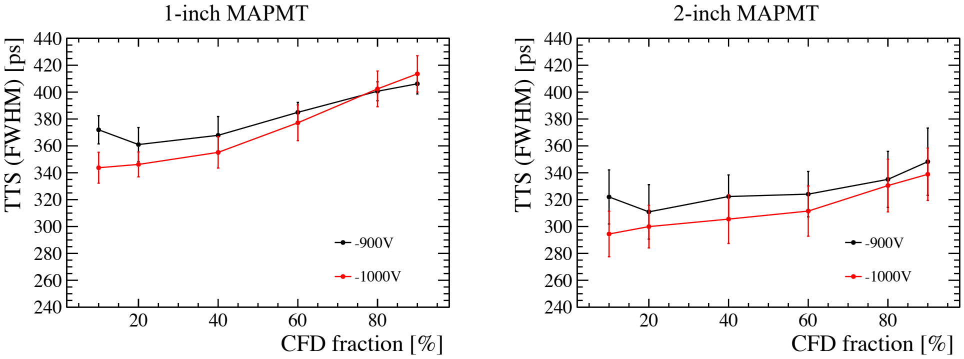

Figure 2.17: TTS as a function of the CFD fraction for the two MAPMT types at −900 V and −1000 V.

Table 2.5: Best TTS values for the two MAPMT types at −900 V and −1000 V. The values correspond to the optimal CFD fractions of 20 % for the −900 V and 10 % for the −1000 V.

MAPMT type	−900 V (20 % CFD)	−1000 V (10 % CFD)
1-inch	155 ± 6 ps	145 ± 5 ps
2-inch	132 ± 5 ps	125 ± 5 ps

smaller number of multiplication stages in the 2-inch MAPMT. For comparison with the nominal values, the FWHM from the fits was converted into the equivalent Gaussian standard deviation ($\sigma_{meas} = \mathrm{FWHM}_{\mathrm{meas}}/2.355$). Then , the contribution from the laser time jitter was subtracted in quadrature:

$$\sigma_{\mathrm{ch}} = \sqrt{\sigma_{\mathrm{meas}}^2 - \sigma_{\mathrm{laser}}^2}\,. \tag{2.7}$$

The best time resolutions were obtained at −1000 V with a CFD fraction of 10 %, yielding 145 ± 5 ps for the 1-inch MAPMT and 125 ± 5 ps for the 2-inch device. The best results for both MAPMT types and bias voltages are summarised in Table 2.5.

Finally, the stability of the timing performance was investigated by monitoring the TTS of the 2-inch MAPMT channel over more than one month, including periods of continuous operation lasting multiple days. The measurements were performed at a bias voltage of −900 V, and the TTS was extracted using a CFD fraction of 20 %. Figure 2.18 shows the measured TTS values as a function of time. The results demonstrate that the TTS remains stable within uncertainties over the entire measurement period, with no indication of degradation even after extended operation. This confirms the reliability of the MAPMT time response, a key requirement for stable performance of the RICH detectors throughout LHC Run 4, where continuous high-rate operation is expected.

The analogue measurements presented here define the intrinsic single-photon timing performance

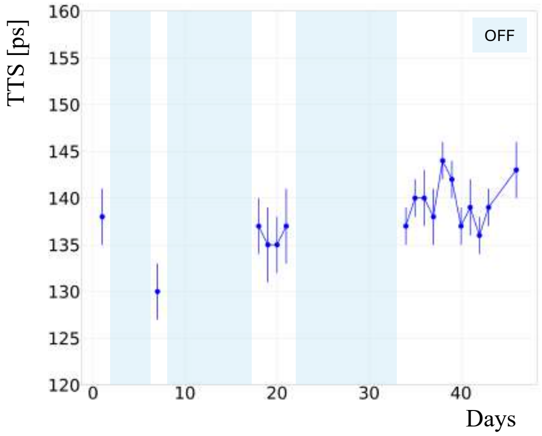

Figure 2.18: TTS of a 2-inch MAPMT channel at −900 V, measured over more than one month time period.

of the MAPMTs, obtained under controlled illumination and single-channel readout conditions. These results provide reference values for the multi-channel TTS measurements performed with the time-resolved readout electronics developed for the LHCb RICH beam tests.

2.2 Digital measurements on MAPMT single-photon time resolution

This section describes the digital measurements for the SPTR of a 1-inch MAPMT. The measurements were performed using the multi-channel readout electronics developed for the LHCb RICH beam test and presented in Section 1.4.3. They provided the TTS across multiple channels to assess the uniformity of the MAPMT timing performance and served as a direct validation of the results obtained in the LHCb RICH 2022 beam-test campaign [91]. The measurements validated the working point used during the 2022 beam test campaign and delivered a reference value for the TTS under controlled laser illumination conditions.

2.2.1 Experimental setup

This section describes the setup used for single-photon illumination and multi-channel readout of the 1-inch MAPMT. The setup is shown in Figure 2.19 (left). Same laser system as for the analogue measurements was used. The MAPMT signals were read out using the multi-channel chain developed for the 2022 LHCb RICH beam test shown in Section 1.4.3. The chain includes a front-end board (FEB) with eight FastIC ASICs and a digital board with a Kintex-7 FPGA. Each MAPMT anode was connected to one FastIC channel, operated in leading-edge time-over-threshold (ToT) mode. Due to the 34 available readout channels, only half of the MAPMT was read out, corresponding to the same half illuminated in the 2022 beam test, as shown in Figure 1.28 (B). The comparator output from each FastIC channel was sent to a time-to-digital converter (TDC) implemented in the FPGA. The TDC provided timestamps for each photon event with a nominal bin size of about 150 ps. Each event contained the ToA and ToT of the pulse, enabling off-chip time-walk correction. Data were transmitted via USB2 to the DAQ PC for storage and offline analysis. A graphical user interface (GUI) running on the DAQ PC provided the operator with a flexible configuration and run control environment. In this laser setup, the FPGA was also used to generate the laser trigger at 20 kHz rate. A copy of this trigger signal is stored in a TDC-in-FPGA channel to include the laser reference time in data.

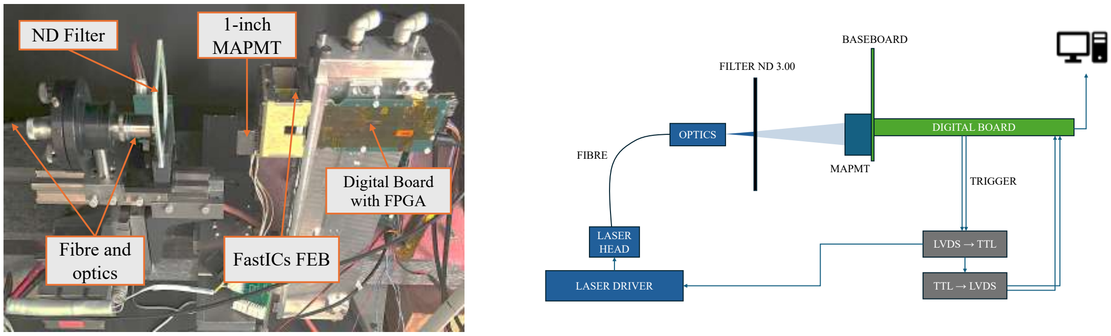

Figure 2.19: Setup for the digital measurements of a 1-inch MAPMT SPTR with picosecond-pulsed laser illumination. (Left) Photograph showing the optical chain with neutral-density filters, fibre coupling, and the 1-inch MAPMT connected to the FastIC front-end and the digital board with FPGA. (Right) Schematic of the setup with all the connections. The FPGA provides the trigger signal to the laser driver and stores a copy of it in the TDC as time reference.

2.2.2 Single-photon time resolution analysis method

The SPTR was extracted following the procedure developed in Refs. [91, 99]. The analysis accounts for the main effects that contribute to the ToA measurement: variations of the TDC bin width and time-walk due to signal amplitude fluctuations. The non-uniform bin widths of the TDC in FPGA impacts both the time reference channel and the MAPMT channel under study. To control this contribution, the data are first grouped according to the time reference TDC bin. In this way, each subset contains events with the same TDC uncertainty arising from the time reference ToA estimation. To mitigate the time-walk effect, the data in each reference-bin subset are further divided according to the ToT of the MAPMT signal. Within a ToT bin the pulse amplitude is approximately constant, and the residual time-walk is negligible. For each pair of (reference bin, ToT bin), the distribution of the time difference between the MAPMT channel and the time reference is computed. These distributions are fitted with a CrystalBall function, where the time resolution is extracted as the Gaussian sigma parameter. Several fits are performed changing the fit range in order to account for parameter fluctuations, as shown in Figure 2.20. A subset of Gaussian sigmas is obtained selecting only the fits with parameters satisfying some minimum requirements to assess the good quality of the fit. Once corrected for the reference ToA bin uncertainty, the Gaussian σ closest to the average σ of the subset defines the time resolution for that bin. The results from all reference ToA bins are recombined using a weighted average, with the individual uncertainties as weights, to obtain the SPTR for each MAPMT channel ToT bin. This procedure is repeated for a range of ±3 ToT bins around the most populated one. The final SPTR for the MAPMT channel results from the average sigma of the ToT bins, weighted by their statistics. This two-step binning procedure ensures that the extracted resolution reflects the intrinsic timing performance

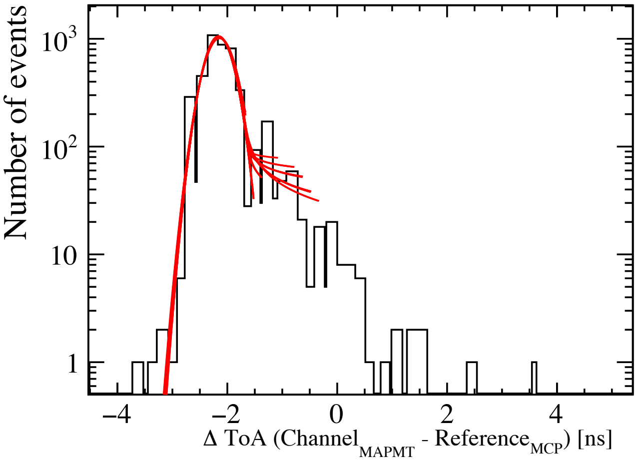

Figure 2.20: Example distribution of the time difference between a MAPMT channel and the time reference. These data are a subset for one reference ToA bin and one ToT bin. The CrystalBall-based iterative fit procedure (red lines) is used to extract the Gaussian σ, which defines the SPTR for the selected reference ToA and ToT bin. The distribution is obtained for beam test data, where an MCP-PMT was used as time reference, as described in Section 1.4.3. It is shown for the purpose of illustrating the analysis method, applied also to the laser measurements. Source [91].

of the system, minimising the effects of time-walk and digital readout. Nevertheless, the measured Gaussian σ_{meas} still includes contributions from the laser illumination and the readout chain. The intrinsic SPTR of the MAPMT channel is obtained as

$$\sigma_{\mathrm{MAPMT}} = \sqrt{\sigma^2_{\mathrm{meas}} - \sigma^2_{\mathrm{ref}} - \sigma^2_{\mathrm{TDC}} - \sigma^2_{\mathrm{FE}}}\,, \tag{2.8}$$

- $\sigma_{\mathrm{ref}} \sim 21\,\mathrm{ps}$: laser jitter;
- $\sigma_{\mathrm{TDC}} \sim 150/\sqrt{12}\,\mathrm{ps}$: contribution from the TDC, estimated from the nominal TDC bin width;
- $\sigma_{\mathrm{FE}} \sim 25\,\mathrm{ps}$: intrinsic jitter of the FastIC discriminator.

2.2.3 Data sets and working point selection

The MAPMT working points were established using reference datasets at a low laser intensity. For each channel, the noise pedestal at the FastIC discriminator was found with a dedicated threshold scan procedure [91]. The relative single-photon detection efficiency (PDE) and the SPTR were measured as a function of the FastIC discriminator threshold offset from the noise pedestal. Figure 2.21 (left) shows the fraction of detected photons per laser trigger event, which is used as a measure of the relative PDE, for a typical 1-inch MAPMT channel. The low laser intensity results in approximately 3 % of the events containing a photon, which defines the occupancy for the channel under study. The PDE follows a typical S-curve due to the typical wide signal amplitude spectrum of the MAPMT. The working point was chosen at an offset of $20.3\,\mu\mathrm{A}$ above pedestal, consistent with the beam test settings. Figure 2.21 (right) shows the SPTR a function of the threshold offset. The SPTR is obtained with the method described in Section 2.2.2 and shows good timing performance of the MAPMT for a broad range of threshold. The red mark in Figure 2.21 highlights the optimal working point at which the following data sets were recorded. This working point is sufficiently close to the noise pedestal, allowing to have the best time resolution without being affected by the noise. This measurements also validated the threshold used in the 2022 LHCb RICH beam test.

At the selected working point, data sets with ten million events each were recorded for three different photon occupancies, defined as:

$$\mathrm{Occupancy} = \frac{N_{\mathrm{photons}}}{N_{\mathrm{triggers}}}. \tag{2.9}$$

Here, N_{photon} is the number of photons detected in one channel and N_{triggers} the number of generated triggers. The photon occupancies were adjusted by inserting neutral-density filters of optical densities 3.00, 2.70, and 2.30, yielding $2 \pm 1\,\%$, $10 \pm 5\,\%$, and $20 \pm 10\,\%$, respectively. The uncertainty represents the minimum and maximum average occupancies on the sensor due to non-uniform illu-

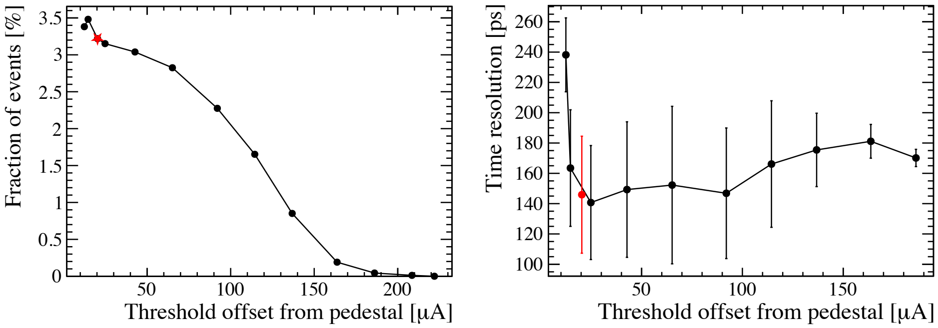

Figure 2.21: Fraction of events with a recorded hit (left) and SPTR (right) of a single MAPMT channel as a function of the threshold offset. The red mark represents the optimal operational threshold. Source [91].

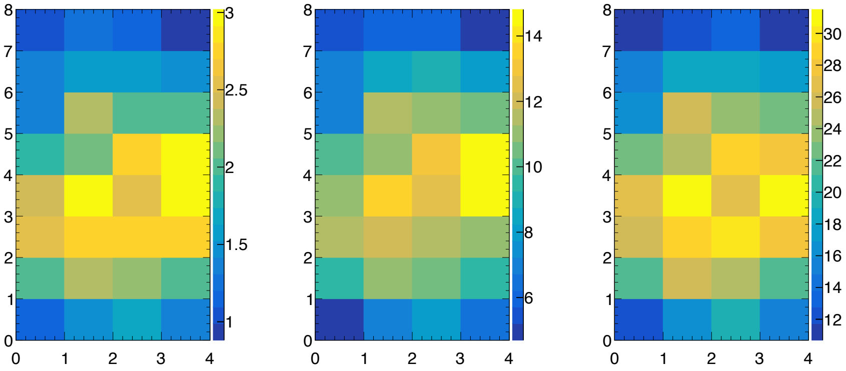

Figure 2.22: MAPMT hitmap obtained with the picosecond pulsed laser setup in condition of peak occupancy approximately 2 %, 10 %, 20 %. The occupancy is defined as the percentage of trigger events for which the channel registers a hit. Source [91].

mination of the laser light spot, as shown in Figure 2.22. Laser intensity and trigger frequency were maintained to ensure identical laser timing performance. A MAPMT bias of −1000 V was used in all three datasets. Figure 2.23 shows the number of channel hit per event in the 1-inch MAPMT for the three photon occupancy conditions. As the occupancy increases, more channels record a photon hit per event, adding workload on the electronics and potentially affecting performance. The SPTR results obtained with these data sets are described in Section 2.2.4.

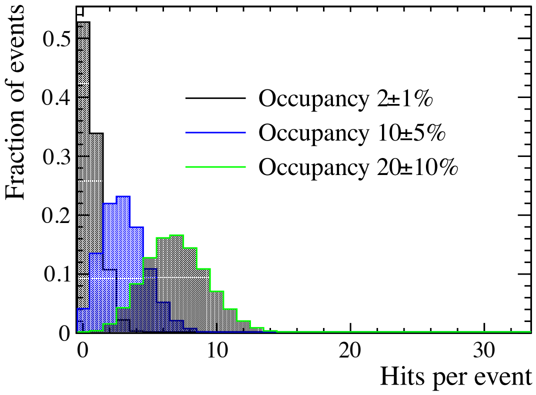

Figure 2.23: Number of hits per event in the 1-inch MAPMT for the three photon occupancy conditions. Source [91].

2.2.4 Results on single-photon time resolution

The SPTR results are shown in Table 2.6 and Figure 2.24. The SPTR results are generally consistent within uncertainty with the analogue TTS value of 145±5 ps for the 1-inch MAPMT at −1000 V shown in Section 2.1.3. These results are also in good agreement with the beam test SPTR, with 182 ± 24 ps during beam test and 202 ± 24 ps with the laser dataset with 20 ± 10 % photon occupancy. While this photon occupancy reproduces most closely the beam test conditions, the illumination pattern differs. In the beam test the MAPMT was exposed to a focused Cherenkov ring, while in the laser setup the light from the single-mode fibre and optics produces a diffuse spot over several channels. At higher occupancies, the increased hit rate enhances electronic switching noise from the CMOS outputs of the FastICs, as well as cross-talk between adjacent channels. This effect leads to a degradation in the measured time resolution, from 162±19 ps at 2% occupancy to 202±24 ps at 20%, as reported in Table 2.6. When a mask is applied to illuminate only one channel, the contribution of switching noise from neighbouring channels is removed. With the mask, the resolution at 20 %

Table 2.6: SPTR of the 1-inch MAPMT measured with laser illumination at different occupancies. Results are reported as the average of all the illuminated channels and for a single channel for masked measurements where only a single channel is exposed. Source [91].

Photon Occupancy [%]	Illuminated channels	σ_t [ps]
2 ± 1	All	162 ± 19
10 ± 5	All	178 ± 27
	Single channel (mask)	173 ± 25
20 ± 10	All	202 ± 24
	Single channel (mask)	170 ± 24

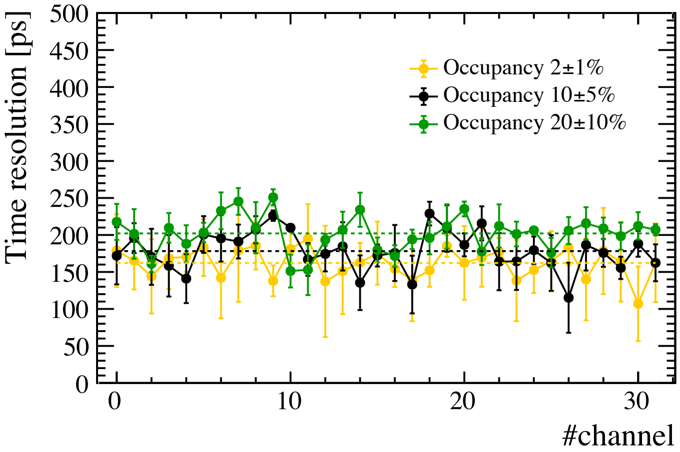

Figure 2.24: SPTR of the illuminated MAPMT channels as a function of the peak occupancy of the corresponding data set. Source [91].

occupancy recovers to 170 ± 24 ps, close to the value obtained at lower occupancy. This suggests that the observed degradation originates from cross-talk and/or increased electronic noise in high-multiplicity events rather than intrinsic limitations of the MAPMT. In addition, when the mask is applied, most photons hit the photocathode near the centre of the illuminated channel rather than the edges. The uncertainty in the photon impact position within a channel can worsen the SPTR for two main reasons. First, the average transit time of the photoelectron depends on the emission point: electrons produced near the centre follow a shorter path to the dynodes than those generated at the edges. Second, the electric field is less uniform near the channel boundaries, causing larger fluctuations in the electron transit time. By concentrating the illumination at the channel centre, the mask therefore reduces both the systematic offset and the spread in transit times, improving the timing performance of the MAPMT channel. In order to study this, a setup for high-precision characterisation of single-photon sensors was built and it's described in Section 2.2.5.

2.2.5 High-precision illumination setup for single-photon characterisation of photodetectors

The setup described in this section was developed in collaboration at CERN for the high-resolution characterisation of single-photodetectors. The setup enables illumination of a photodetector with a sub-100-μm laser spot that can be scanned across the sensor area with motorised XY translation stages This allows measurements of PDE and SPTR uniformity within a single channel. Figure 2.25 shows the main components of the setup. The light source is a picosecond pulsed laser[7], which

---

[7] Photonics PILAS pulsed laser (`https://www.nktphotonics.com/products/pulsed-diode-lasers/pilas/`).

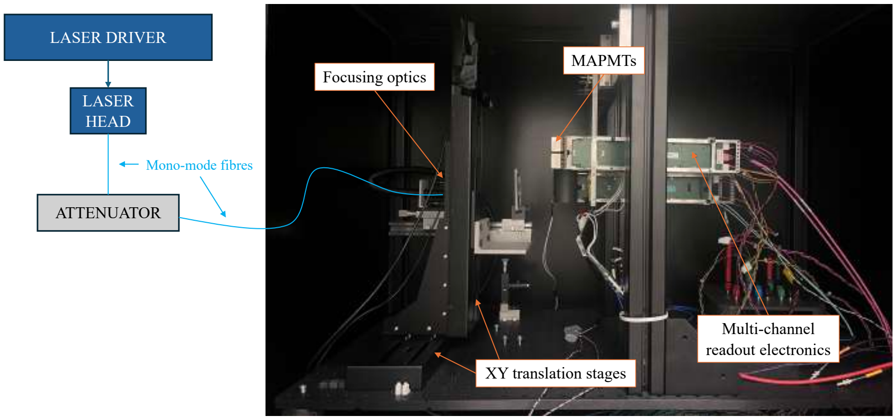

Figure 2.25: Setup for high-precision, single-photon characterisation of photodetectors.

delivers 20 ps (FWHM) pulses at 403±2 nm. Laser pulses are delivered through single-mode fibres to the focusing optics, after passing through a digital attenuator to ensure single-photon illumination. The optics are mounted on motorised XY stages that can scan the focused spot over an area of about 45 cm^2. The stages are controlled remotely via USB connection. Python scripts have been integrated into the acquisition software to automate scan&data-taking runs. Photodetectors are mounted on a dedicated mechanical support in front of the optics. Readout is provided by the second version of the prototype electronics chain for the LHCb RICH upgrade, based on FastIC and picoTDC (see Section 1.4.3).

Preliminary scans performed with this setup on MAPMTs have revealed characteristic features of their response. Variations in detection efficiency and timing performance were observed within individual channels, with effects more pronounced near pixel edges. These first studies confirm the sensitivity of the system and its suitability for systematic characterisation of photodetectors at the sub-pixel level. A summary of the initial results is available in the summer student project report [101]. A more detailed investigation, including full scans of both 1-inch and 2-inch MAPMTs, has since been carried out and a publication is in preparation.. The flexible design of the setup also allows its adaptation to other photodetector technologies, such as SiPMs and MCP-PMTs. The measurement campaign will therefore be extended to these devices, providing essential input for the photodetector R&D programme of the LHCb RICH community in view of the Upgrade II detector. In addition, the entire setup is placed inside a cold room that can operate at a stable temperature of $-20\,^{\circ}$C. This is particularly important for SiPMs characterisation at cold temperature, which reduces the dark-count rate and improve the performance.

2.3 Chapter conclusion

In this chapter, analogue and digital measurement of the MAPMT single-photon response were performed under picosecond-pulsed laser illumination. The analogue measurements provided a detailed characterisation of the single-photon pulse shape and time resolution for the 1-inch and 2-inch MAPMTs. The TTS was found to be 145 ps for the 1-inch MAPMT and 125 ps for the 2-inch sensor, defining the intrinsic timing capability of the two MAPMT types. Digital measurements on the 1-inch MAPMT were performed with the first time-resolved readout electronics developed by the LHCb RICH group as a prototype for the LHC Run 4 optoelectronics system of the RICH detector. These tests yielded a SPTR of 162 ps, representing the first time resolution result obtained with this optoelectronics prototype under controlled laser illumination. The results presented in this chapter provided a reference for subsequent beam-test campaigns with the same sensors and contributed to the collaboration upgrade strategy for the LHCb RICH detectors in view of the LHCb LS3 enhancements and the Upgrade II programme.

Chapter 3

Simulation studies on primary vertex time reconstruction with the LHC Run 4 RICH detectors

This chapter presents simulation studies of a novel method to estimate the primary vertex time (often referred to as PV t_0) using the LHCb RICH detectors during LHC Run 4. The estimation is possible due to the introduction of picosecond time information in the LHCb RICH detector LS3 Enhancements, as discussed in Chapter 1. The method is integrated into the RICH reconstruction algorithm and uses the intrinsic picosecond-timing capabilities of RICH detectors to estimate the PV t_0. The PV t_0 plays a key role for applying a time gate around the predicted photon hit time for each track. The time gate reduces combinatorial background from other tracks, improving the particle identification (PID) performance of the RICH detector. This PV t_0 estimation will be a first-time measurement for the LHCb experiment, setting a milestone for the use of time information ahead of LHCb Upgrade II. The proposed method uses the photon-objects (POs) produced in the RICH reconstruction for each track-hit association. Each PO is, in turn, associated with a PV, based on the track–PV association from tracking information. The hit time information contained in the PO can be used to estimate its PV t_0. However, only a fraction of POs result from a correct PO–PV association. The resolution of the estimated PV t_0 strongly depends on this fraction, which can be improved with high-quality PO selection. The results in this chapter demonstrate that the PV t_0 can be estimated with a precision of 85 ± 3 ps FWHM (66 ± 2 ps RMS) for 94 % of the PVs in LHC Run 4 conditions.

Section 3.1 introduces the RICH reconstruction. It describes how time information can be integrated into the RICH reconstruction to improve the detector PID performance. The method developed for the PV t_0 estimation is described in Section 3.2. The section also demonstrate the importance of high-quality POs for improving the PV t_0 resolution. Section 3.3 introduces different PO selection criteria, compares their impact on the PV t_0 resolution, and discusses how event topology influences the result. The possible implementation of this algorithm in the LHCb experiment and the chapter conclusions are described in Section 3.4 and Section 3.5, respectively.

3.1 Fast-timing information in the RICH reconstruction for LHC Run 4

Simulation studies have shown that photon hit time information combined with the excellent intrinsic time resolution of the RICH detector can improve the PID performance [67]. In high-multiplicity events, overlapping tracks generate photons in the same region of the detector. This produces a high combinatorial background and the RICH reconstruction may fail to separate Cherenkov rings using only the spatial information on the photon hits. The photon hit time is a powerful tool to remove combinatorial background from out-of-time photons and improve the detector PID performance. In this context, the PV t_0 is a key element, providing the time reference to predict the photon hit time for each track. This enables an efficient reduction of combinatorial background by applying a time gate around the predicted time.

Section 3.1.1 describes the RICH reconstruction currently used in the LHC Run 3, which relies only on the spatial information of the photon hits. Section 3.1.2 introduces the improvement on the PID performance achievable with the use of time information, highlighting the role of the PV t_0. It also describes how the hit time information can be integrated into the LHC Run 4 RICH reconstruction, including the PV t_0 estimation algorithm.

3.1.1 LHC Run 3 RICH reconstruction

The RICH PID reconstruction algorithm is implemented in the LHCb High-Level Trigger [102]. The algorithm exploits the reconstructed trajectories of charged particles to predict the expected photon hit positions on the photodetector plane for a given particle mass hypothesis. Based on this prediction, the algorithm calculates the likelihood of the observed photodetector hits and looks for the set of mass hypotheses that maximise the overall likelihood.

Figure 3.1 illustrates the schematic of the LHC Run3 RICH reconstruction. Photodetector hits are stored in the RICH SmartID, which provides the interface between the photodetector and the RICH reconstruction algorithms. The SmartID contains any event information required by the reconstruction, such as the RICH detector, photodetector panel, column, row and MAPMT pixel involved in the event [26]. The RICH SmartID and reconstruction algorithms are identical for simulation and experiment data. In the first stage of the reconstruction, hits are clustered and the

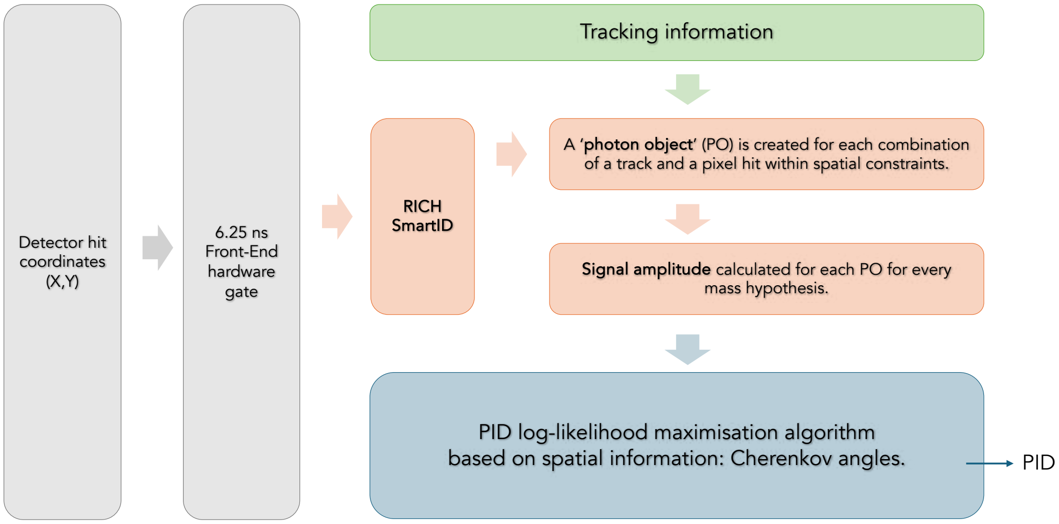

Figure 3.1: Schematic of the workflow of the RICH reconstruction algorithm in LHC Run 3.

cluster coordinates are calculated, assuming hits occur at the centre of the MAPMT pixel. Using tracking information, the algorithm reconstructs the segments of particle tracks within the RICH radiators. For each track segment, the algorithm calculates the expected Cherenkov-ring radius for every track mass hypothesis. A spatial acceptance in the photodetector plane is defined from these radii, including a small tolerance to account for angular resolution. Hits outside this spatial acceptance are discarded to limit the number of possible photon associations in later stages [103]. The algorithm then combines track segments and detector hits into Cherenkov POs (POs) for each possible hit–track association within spatial constraints. Each PO is a candidate for a true combination and contains all the required information about the track and the hit, including the parameters needed to reconstruct the corresponding Cherenkov angle. For each PO, the algorithm calculates the Cherenkov angle by solving the equation of the photon trajectory through the RICH optical system [104]. Because the exact emission point is unknown, and can be anywhere along the particle trajectory in the radiator, the segment mid-point is taken as the emission point. The reconstruction algorithm assumes that there are N sources of hits in the photodetector and defines the event hypothesis $\mathbf{h} = (h_1, \ldots, h_N)$, where h_j is either the mass-hypothesis for track j or a parameter associated with background j. For each h_j, the algorithm defines the signal amplitude

$a_{ij}(h_j)$, which describes the probability that source j produced a photon hit in pixel i:

$$a_{ij}(h_j) = \epsilon_i \, b_{ij}(h_j) \tag{3.1}$$

where ϵ_i is the photodetector efficiency and $b_{ij}(h_j)$ is the expected number of Cherenkov photons arriving at pixel i that are emitted by track j:

$$b_{ij}(h_j) = \lambda_j(h_j) \iint_{\text{pixel } i} f_{h_j}(\theta, \phi) \, \mathrm{d}\theta \, \mathrm{d}\phi \tag{3.2}$$

Here $\lambda_j(h_j)$ is the expected total number of photons emitted by track j and the integral gives the fraction of them that arrives at pixel i. The probability density function $f_{h_j}(\theta, \phi)$ of the emitted photons is taken to be a Gaussian θ, with expected Cherenkov angle $\theta_c(h_j)$ and Cherenkov angle resolution $\sigma(\theta)$, and uniform in the angle ϕ:

$$f_{h_j}(\theta, \phi) = \frac{1}{(2\pi)^{3/2}\sigma(\theta)} \exp\left[-\frac{1}{2}\left(\frac{\theta - \theta_c(h_j)}{\sigma(\theta)}\right)^2\right] \tag{3.3}$$

For each PO, the signal amplitude is calculated for all possible mass hypotheses of track j and and stored for the likelihood maximisation step. Once the signal amplitude is computed for all POs, the expected signal in pixel i is the sum, over all the sources, of the signal amplitudes for that pixel:

$$\nu_i(\vec{h}) = \sum_{j=1}^{N} a_{ij}(h_j) \tag{3.4}$$

Finally, the event likelihood is the product, over all pixels M^{tot}, of the Poisson probability to to observe n_i hits given the expected signal $\nu_i(\vec{h})$:

$$\mathcal{L}(\vec{h}) = \prod_{i=1}^{M^{\text{tot}}} \mathcal{P}_{\nu_i(\vec{h})}(n_i) \,, \tag{3.5}$$

where n_i can be either 0 or >=1, corresponding to the absence or presence of a hit, respectively. The algorithm maximises the likelihood by determining the optimal set of mass hypotheses for the event tracks. Initially, all tracks are assumed to be pions, as they are the most abundant particle type in LHCb. The algorithm then evaluates the change in log-likelihood, $\Delta \ln L$, by modifying the mass hypothesis of one track at a time while keeping all others fixed. The change that gives the largest increase in likelihood is selected, and the track mass hypothesis is updated if $\Delta \ln L$ exceeds a predefined threshold. This iterative process continues until no further improvement is possible, meaning all tracks have been assigned their optimal mass hypotheses.

3.1.2 Time information in the RICH detector

As discussed in Section 1.4, the LHCb RICH detector will be equipped with the novel FastRICH ASIC, which will add a timestamp to the photodetector hits during LHC Run 4. While tracking resolution, RICH spatial resolution and luminosity remain the same from LHC Run 3 to Run 4, the addition of time information offers a new dimension to the RICH detector. According to simulation studies, the time information on the photodetector hits and the intrinsic time resolution of the RICH detector provide a powerful tool to enhance the PID performance of the RICH detector, especially in high-multiplicity processes [38]. In addition, the time information can also be used to obtain an estimate of the PV t_0 using the RICH reconstruction. A novel method for the PV t_0 is presented in Section 3.2, marking a first-time measurement for the LHCb experiment. The development of algorithms including time information is an important step for LHC Run 5. The use of timing in LHC Run 5 will be a necessity rather than an enhancement and the experience gathered during LHC Run 4 will greatly help in the transition.

Because of the nature of Cherenkov photons and the RICH detector optics, photons from tracks produced in a given PV arrive at the photodetectors within a picosecond-scale time interval [54]. It has been demonstrated that the hit time of Cherenkov photons produced by a given track can be predicted with a precision of 10 ps [26]. The prediction uses the LHCb tracking information, the reconstructed photon paths in the RICH detector and the PV time. The hit time prediction is defined as

$$t_{\text{pred}} = t_{\text{pv}} + \frac{d_{\text{pv},A}}{c}\sqrt{1+\left(\frac{mc}{p}\right)^2} + \frac{d_{A,E}}{c}\,n\cos\theta_c + \left[d_{E,M1} + d_{M1,M2} + d_{M2,HIT}\right]\frac{n}{c} \tag{3.6}$$

where t_{pv} is the PV time, $d_{\text{pv},A}$ 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 photodetector hit. In addition, p is the particle momentum from the LHCb tracking detectors, n is the refractive index at the average photon emission energy and m is the mass hypothesis for the particle. For RICH 2, a track curvature correction is applied to this formula to compensate for the curved trajectory of the particles through the magnetic field [26]. The PV timestamp t_{pv} and the PV z-position (contained in $d_{\text{pv},A}$) are often combine into the PV t_0, defined as:

$$t_0 = t_{\text{pv}} - \frac{z_{\text{pv}}}{c} \tag{3.7}$$

where $\frac{z_{\text{pv}}}{c}$ is the time displacement given by the PV z-coordinate. The time dependence on the x and y coordinates of the PV is not included, since the PV spread on x and y axes is negligible compared to $\sigma_z \sim 50$ mm. The PV t_0 uniquely defines at which time the particles in the event are assumed to originate, and is therefore the main quantity of interest in this chapter. PV t_0

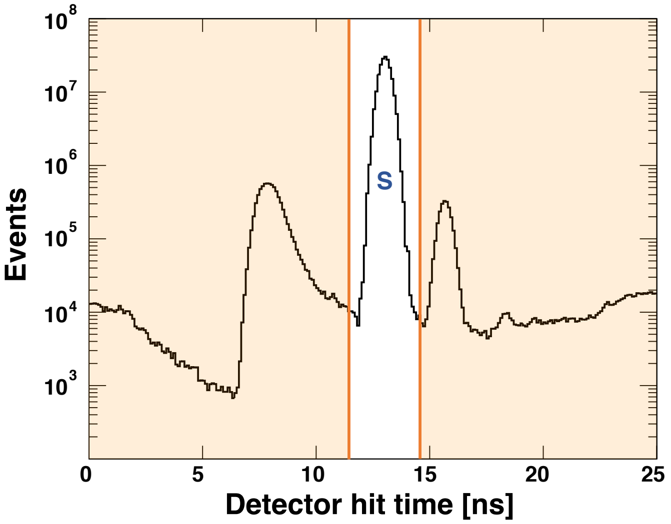

Figure 3.2: Simulated hit time distribution in the RICH 1 detector, showing the signal peak ('S') and a front-end hardware gate to exclude background hits [38].

uncertainties arise from both the intrinsic PV time spread and the PV Z distribution:

$$\sigma_{t_0} = \sqrt{\sigma_{t_{\mathrm{PV}}}^2 + \left(\frac{\sigma_z}{c}\right)^2}. \tag{3.8}$$

For nominal LHC Run 4 conditions, $\sigma_{t_{\mathrm{pv}}} \simeq 200\,\mathrm{ps}$ and $\sigma_z \simeq 50\,\mathrm{mm}$, yielding a total spread of $\sigma_{t_0} \approx 260\,\mathrm{ps}$ for the PV t_0.

Figure 3.2 shows the simulated photon hit time distribution for the RICH 1 detector from several LHC bunch crossings. Given the bunch crossing at $t = 0$, the Cherenkov signal peak for RICH 1 ("S") is at ~13 ns. The spread of a few nanoseconds observed in the Cherenkov signal peak is dominated by the PV t_0. Additional contributions come from variations in track and photon path lengths through the detector. Background hits originate from various sources, including beam interactions, sensor noise, and certain photon paths. For example, particles and photons travelling directly towards the photodetectors cause early hits, while photon trajectories with extra reflections in the mirror system produce late hits [38]. To exclude out-of-time background hits, a hardware time shutter can be implemented in the front-end electronics, as shown in Figure 3.2. This shutter, centred on the signal peak, does not rely on reconstruction information. It provides a first-level suppression of out-of-time background hits, lowering the CPU load required for the reconstruction. Taking into account the PV t_0 spread and the range of possible track and photon paths through the experiment, an optimal shutter width is about 2 ns [38]. A similar shutter has been already implemented in

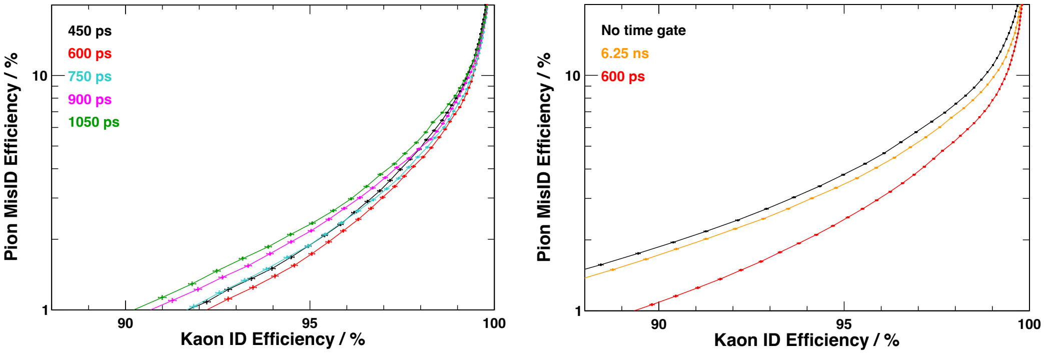

Figure 3.3: (Left) PID curves at $L = 2.0 \times 10^{33}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ for different software time gate applied to the photon hits, considering the MAPMT TTS of ~150 ps [38]. (Right) PID curves at $L = 3.0 \times 10^{33}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ and $\sigma_{\mathrm{MaPMT}} \sim 150\,\mathrm{ps}$, comparing 6.25 ns (Run 3) and 600 ps (Run 4) time gates [67].

FPGA during LHC Run 3, but limited to 6.25 ns by the CLARO ASIC constraints [68]. The FastRICH ASIC will allow implementation down to 2 ns around the signal peak, as described in Section 1.4.1.

The photon hits selected by the shutter are timestamped for finer software selection. In the RICH reconstruction, the photon hits are combined with tracks to form POs. Given the PV t_0, these objects contain all the information needed to calculate the time-of-flight of the tracks and the time-of-propagation of the photon candidates. This allows the prediction of the expected hit time for each PO using Equation3.6. A software time gate can then be applied around the predicted hit time to discard out-of-time photons. POs within the gate are more likely to be correct hit–track associations, while the remaining objects are considered fake and excluded from the log-likelihood maximisation.

The impact of different software time gates on the detector performance has been evaluated using the LHCb Simulation Framework. The LHC Run 3 RICH detector was simulated with added timing information under different luminosity scenarios. Figure 3.3 shows the PID curves obtained from these simulations with different software time gates. The PID curves are a standardised and sensitive probe of RICH performance, where the bottom right corner represents the ideal case of high kaon identification and low pion misidentification [102]. Figure 3.3 (left) shows the PID curve at a luminosity of $2.0 \times 10^{33}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. The MAPMT transit-time spread (TTS) (σ_{MAPMT}) of ~150 ps is included as a spread to the photon hit time. The curves show that the optimal software time gate is $\pm 2\sigma \sim 600\,\mathrm{ps}$. Narrower time gates cause loss of signal photons, while wider gates increase the combinatorial background. In both cases, the result is a worse PID performance than the 600 ps

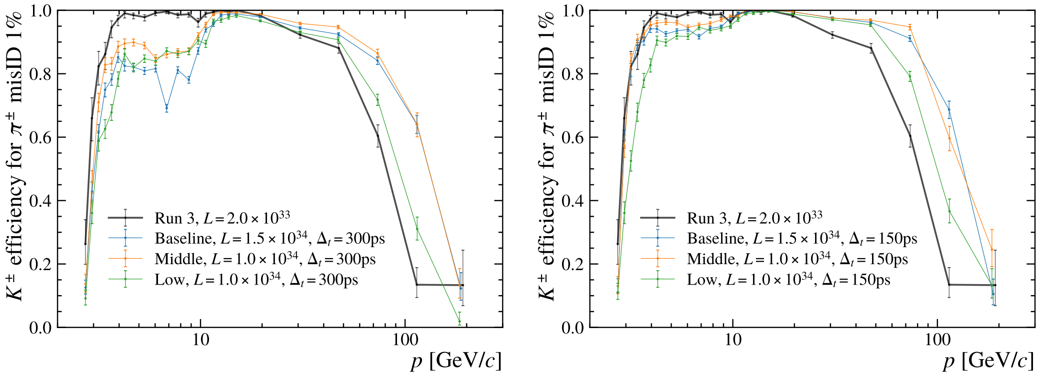

Figure 3.4: Kaon efficiency for 1 % pion misidentification versus momentum, for the LHC Run 3 and the three RICH Upgrade II options at different luminosities, for (left) 300 ps and (right) 150 ps time windows [33].

time gate [38]. Figure 3.3 (right) compares the expected PID performance for LHC Run 3 and Run 4. Here, the simulations were perform at a luminosity of $3.0 \times 10^{33}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ in order to provide a safety margin for additional tracks and the corresponding increase in photon occupancies in the RICH detectors [38]. The PID curves show that a software time gate of 600 ps applied for LHC Run 4 results in better PID performance than Run 3.

During Upgrade II, the increased luminosity will pose significant challenges for the RICH reconstruction and its performance. As discussed in Section1.4, new photodetectors with improved timing performance and higher granularity will be installed to maintain PID performance under these conditions. Figure 3.4 shows simulation results for the kaon identification efficiency requiring a pion misidentification probability below 1 %. The results were obtained for three Upgrade II scenarios with different combinations of detector granularity and luminosity [33]. Two software time gates are considered in each case: 300 ps, corresponding to a nominal single photoelectron time resolution of 75 ps, and 150 ps, corresponding to a time resolution slightly below 40 ps. The results demonstrate that reduced sensor time resolution and optimal time gating mitigate the performance degradation from higher occupancy. At high momenta, all configurations deliver better performance than in LHC Run 3, with the largest gains seen for finer granularity and better timing. At low momenta, the performance falls short of the LHC Run 3 level in all cases, but increases steadily as the photodetector time resolution improves. These results justify the strong focus of the RICH R&D programme on improving the timing performance of the photodetectors.

The PID performance results presented above assume a perfect knowledge of the PV t_0 from simulation truth. In real data, the PV t_0 must be measured, and any uncertainty on its value will reduce the effectiveness of the software time gate. The next subsection describes how the RICH reconstruction in LHC Run 4 can be adapted to estimate the PV t_0 directly from the detector information.

3.1.3 The LHC Run 4 RICH reconstruction

The PID performance results in Section 3.1.2 were obtained using the PV t_0 from truth-matched simulation events. As shown in Equation 3.6, the PV t_0 is a key ingredient in the prediction of the photon hit time. Any uncertainty on the PV t_0 propagates into the predicted photon hit time, ultimately reducing the efficiency of the software time gate. The LHCb tracking system will provide the PV t_0 only in Run5, when track hit time information becomes available in the VELO. During LHC Run 4, the PV t_0 can be estimated using only RICH reconstruction information. A novel method for this estimate is introduced in Section 3.2. The method uses 3D spatial reconstruction to find the correct associations between POs and PVs. Intrinsic PO properties are used to identify high-quality PO candidates, allowing the PV t_0 estimate from a selected subset of POs, which are more likely to originate from correct PO–PV associations.

Figure 3.5 summarises the changes introduced by the integration of photon hit time information in the LHC Run 4 RICH reconstruction. Compared to the LHC Run 3 workflow (Figure 3.1.1), the LHC Run 4 reconstruction includes three additional stages before the likelihood calculation. A hardware time gate at the front-end level rejects a large fraction of out-of-time background hits

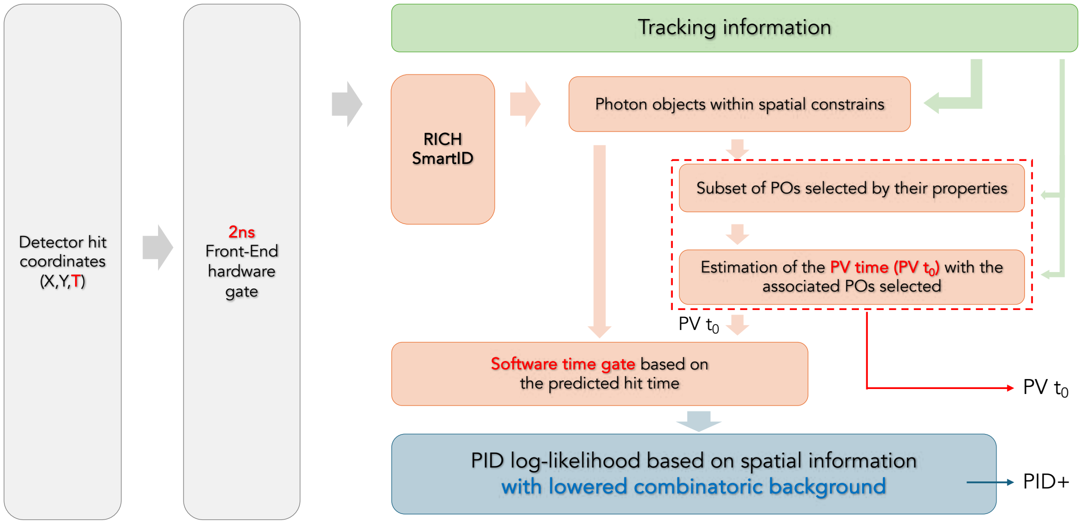

Figure 3.5: Workflow of the RICH reconstruction algorithm in LHC Run 4, integrating the time gates and the PV t_0 estimation algorithm.

before they enter the reconstruction chain, reducing the CPU load of the subsequent reconstruction steps. The accepted photon hits are combined with tracks to form POs. The RICH PV t_0 estimation algorithm is integrated at this point. A subset of the POs is used to estimate the PV t_0, which is then applied as input to the software time gate on the full PO set. The same log-likelihood algorithm described in Section 3.1.1 is then run on a PO sample with reduced combinatorial background, improving PID performance. Although developed for the RICH reconstruction, the estimated PV t_0 can also benefit other sub-detectors, such as the LHCb tracking system.

3.2 Primary vertex time with the RICH detectors

As discussed in Section 3.1.2, the RICH reconstruction during LHC Run 4 will apply a software time gate to the reconstructed photon hits to reduce combinatorial background and improve PID performance. The software time gate is centred around the predicted predicted photon hit time. A key element of the prediction is the PV t_0, described in Equation 3.7. This section introduces a novel method to estimate the PV t_0 using only the RICH reconstruction information during LHC Run 4. The method was implemented in the LHCb simulation framework. It relies on the picosecond timing of photon hits provided by the novel FastRICH-based readout and the association between POs and PVs. This represents a first-time measurement for the LHCb experiment and sets a milestone for future time-based applications in the LHCb Upgrade II detector. Section 3.2.1 describes the simulation setup used for developing and testing the PV t_0 estimation method. Section 3.2.2 explains the algorithm used to reconstruct the PV t_0 from the time information in the POs, using a representative LHC Run 4-like event. Finally, Section 3.2.3 investigates the role of PO purity in determining the PV t_0 resolution.

3.2.1 Simulation setup

For the studies presented in this chapter, the LHCb Simulation Framework was used. A sample of ten thousand bunch-crossing events was generated using the LHCb GAUSS framework [105], which interfaces with the Geant4 simulation toolkit [106]. The GAUSS framework is responsible for simulating particle trajectories and momenta, energy loss, and interactions with the sub-detectors and magnetic field. The Gauss version used is v56r5, which includes a detailed model of all LHCb Run 3 sub-detector geometries and materials. The sample was produced at the expected LHC Run 4 instantaneous luminosity of $2 \times 10^{33}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ with an average of 5.3 PVs generated per bunch crossing. The events are 'minimum bias', meaning that no bias or filter was applied to favour specific physics processes of interest. The sample was then processed through Boole, the LHCb digitisation framework, which simulates the detector responses to the hits produced in Gauss [107]. Boole version v45r0 was used. The output of Boole closely resembles the real experimental data and is used as input to the RICH reconstruction algorithms for PID. The RICH SmartID and reconstruction algorithms are identical for simulation and experiment data. The PV t_0 estimation algorithm is integrated into the RICH reconstruction framework. In addition to the reconstructed track parameters and detector hit information, the RICH reconstruction has access to the full Monte Carlo (MC) event history stored by the Gauss framework, which allows for performance evaluation. In this study, the RICH reconstruction was performed for both RICH detectors using 'long tracks', that are reconstructed in all LHCb tracking stations. Given that the FastRICH characteristics were not yet implemented in the Boole framework, an emulation tool was used to propagate the MC information on the hit time from Gauss to the reconstruction. Uncertainties on the hit time, such as the FastRICH TDC bin size of 25 ps and the sensor time jitter, were implemented a posteriori in the reconstruction.

3.2.2 PV t_0 estimation method

This section describes how the RICH detectors alone can provide a PV t_0 estimate during LHC Run 4, using the RICH reconstruction. The proposed method exploits the intrinsic time capabilities of the RICH detectors. With the novel FastRICH ASIC, the photodetector hits will be timestamped with a precision of ~25 ps. The time information will be stored for each photodetector hit and used in the RICH reconstruction. As discussed in Section 3.1.1, a PO is created in the reconstruction for each combination of a photodetector hit and a particle track within some geometrical constraints. Using the LHCb tracking system, each track can be associated with a PV, resulting in hundreds of signal photons per PV. This creates a multitude of POs per PV that can be used to estimate the PV t_0 with high precision.

Figure 3.6 (left) shows the PO hit time distribution in the RICH 1 detector for an event with five PVs, which corresponds to the average number of PVs expected during LHC Run 4. A bin size of 25 ps has been chosen to simulate the FastRICH TDC sampling. Each peak in the distribution corresponds to one PV in the event. In this case, the photon hit time could be directly used to determine the PV t_0. Figure 3.6 (right) shows the PO hit time distribution in the RICH 2 detector for the same event. RICH 2 is located further downstream and is physically larger than RICH 1, resulting in poorer separation between the PVs. For each PO, the PV time t_0 can be calculated by inverting the photon hit time prediction (Eq. 3.6), where the PV time was taken from Monte Carlo (MC) information, and replacing the predicted hit time by the detected hit time t_{hit}:

$$t_0 = t_{hit} - \frac{|r_A|}{c}\sqrt{1+\left(\frac{mc}{p}\right)^2} - \frac{d_{A,E}}{c}\, n\cos\theta_c - \left[d_{E,M_1} + d_{M_1,M_2} + d_{M_2,hit}\right]\frac{n}{c} \tag{3.9}$$

Here, $|r_A|$ is the reconstructed track path from the PV to the RICH entrance window, $d_{A,E}$ is the reconstructed track path inside the RICH gas radiator from the entrance window to the Cherenkov

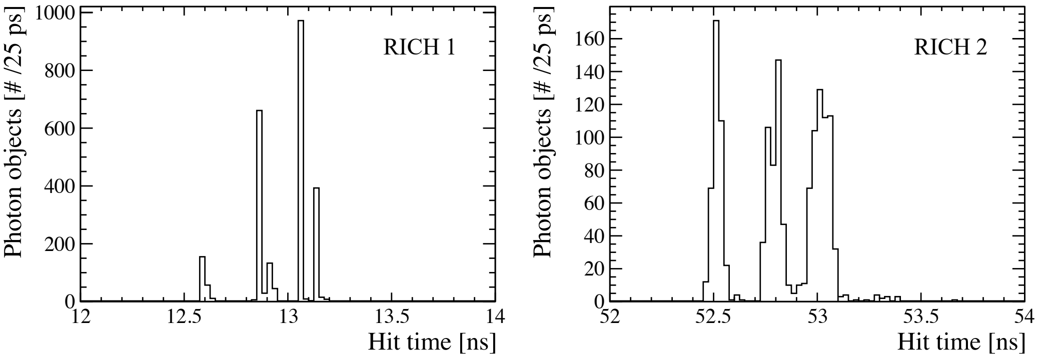

Figure 3.6: Distributions of the PO hit times in RICH 1 (left) and RICH 2 (right) for an event with five PVs .

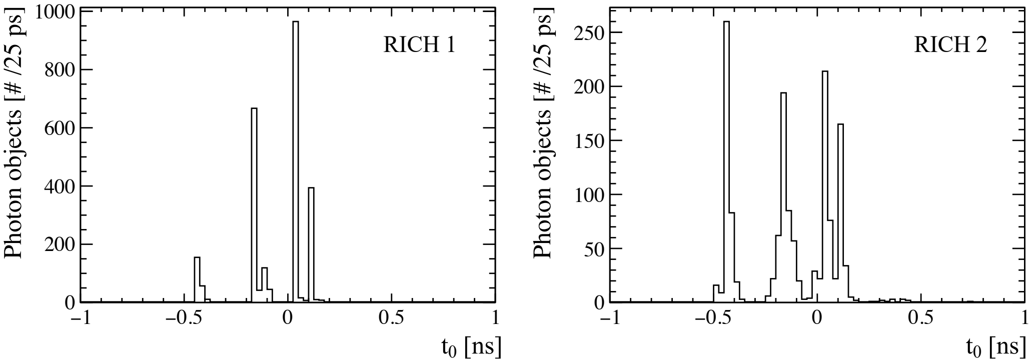

Figure 3.7: Distributions of the t_0 calculated for each PO in RICH 1 (left) and RICH 2 (right) for an event with five PVs.

emission point, and d_{E,M_1}, d_{M_1,M_2} and $d_{M_2,hit}$ are the reconstructed photon paths from the emission point to the photodetector plane through the mirror reflection points [26]. Since the Cherenkov emission point cannot be precisely determined, it is approximated as the midpoint of the reconstructed track segment within the radiator. In addition, p is the particle momentum from the LHCb tracking detectors, n is the refractive index at the average photon emission energy and m is the mass hypothesis for the particle. Since the RICH PV t_0 is calculated prior to the likelihood maximisation, the mass hypothesis is set to the pion mass. Figure 3.7 shows the distribution of the PV t_0 calculated for each PO in RICH 1 and RICH 2. The combined information provides a direct estimation of the t_0 of the PVs in the event. These figures, however, do not yet reflect real experimental conditions, as the photodetector time jitter has not been included.

During LHC Run 4, the RICH detectors will still use MAPMTs as photodetectors. The MAPMT

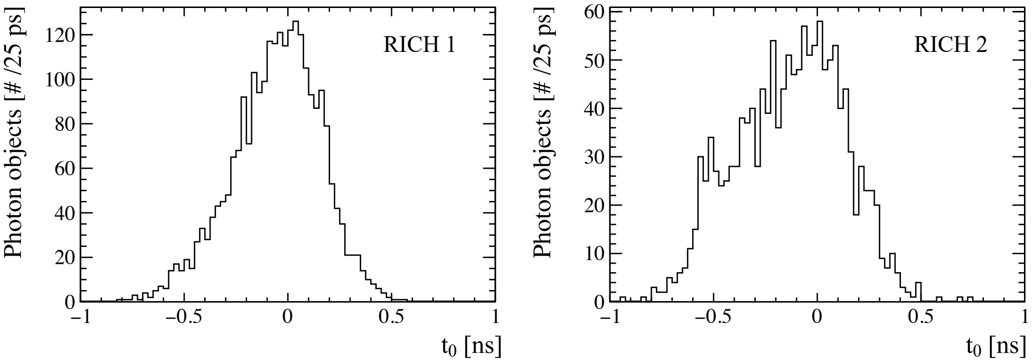

Figure 3.8: Distributions of the t_0 calculated for each PO in RICH 1 (left) and RICH 2 (right) for an event with five PVs with the addition of the MAPMT TTS to the photon hit time. The TTS is modelled as a Gaussian distribution with a σ of ~145 ps for the 1-inch MAPMT and ~125 ps for the 2-inch MAPMT.

TTS was measured to be ~145 ps for the 1-inch MAPMT and ~125 ps for the 2-inch MAPMT, as shown in Chapter 2. Figure 3.8 shows the t_0 for RICH 1 and RICH 2 with the addition of the MAPMT TTS to the photon hit time. The TTS is modelled as a Gaussian distribution with a standard deviation corresponding to the measured spread for the 1-inch and 2-inch MAPMTs. The PV peaks overlap, making it difficult to identify the RICH PV t_0 directly from the photon hit time alone. A dedicated approach is therefore required for the PV t_0 reconstruction.

The dedicated approach is based on the association between POs and PVs. The RICH PV t_0 can be obtained by averaging the reconstructed PV time ($t_{0,i}$) over the N POs associated with the PV:

$$RICH\,PV\,t_0 = \frac{1}{N}\sum_{i=1}^{N} t_{0,i} \tag{3.10}$$

A key consideration for this summation is that the number N does not need to include all the POs produced for the PV, but can instead contain only those POs that are considered highly likely to be correctly associated with the PV. This selection strongly improves the resolution of the RICH PV t_0. Especially in high pile-up events, a significant number of POs is incorrect, with the pixel hit not associated with the right track and PV. Including them in the summation would increase the spread of the RICH PV t_0.

Figure 3.9 (left) shows the hits produced by three tracks in a low occupancy region of the RICH 1 photodetector plane. The hits arrives in the characteristic Cherenkov ring pattern, as described in Section 1.3.2. The hits in green are associated with a single track by the RICH reconstruction. The hits in yellow are associated with two different tracks, because they are located at the intersection of two Cherenkov rings, and the reconstruction cannot geometrically distinguish which track generated them. As a consequence, two POs are created for these hits, but only one corresponds to a correct hit-track association. For the incorrect POs at the intersection of *track 1* and *track 2* rings, the association is still with the correct PV because the two tracks originate from the same PV. Since such tracks reach the RICH detector within a few picoseconds of each other [26], the t_0 estimated from the incorrect POs is close to the correct PV t_0 value. On the contrary, the incorrect POs generated at the intersection of *track 2* and *track 3* rings are associated with a wrong PV. The t_0 estimated from these incorrect POs can differ from the correct PV t_0 by hundreds of picoseconds, increasing the spread of the RICH PV t_0. This remains close to an ideal scenario, as the rings are located in a low-occupancy region of the detector and their overlap is minimal. Therefore, the RICH PVt_0 estimates for the two PVs in this example are expected to be close to the correct PVt_0.

In the high occupancy region of the detector, the situation is significantly more complex. Figure 3.9 (right) shows the photon hit distribution in the high-occupancy region of RICH 1 for an event with five PVs. Cherenkov rings from multiple tracks overlap and a single photon hit can be associated with more than ten tracks. This generates many incorrect POs per PV which can significantly worsen the RICH PV t_0 estimate. Figures 3.10 and 3.11 show the PV t_0 distributions for the five primary vertices in the event, for RICH 1 and RICH 2 respectively. In both figures, the plots on the left are obtained using all POs associated with each PV, while those on the right

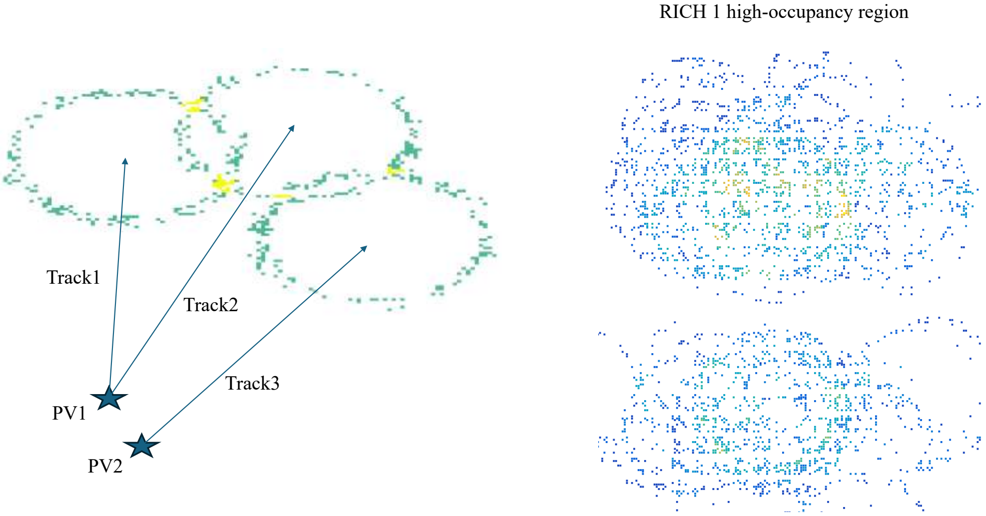

Figure 3.9: Left: spatial distribution of the photon hit produced by three tracks in the RICH 1 low occupancy region. The hits in green are associated with a single track generating a PO. The hits in yellow are associated with two tracks producing two POs. Right: photon hit distribution in the high-occupancy region of RICH 1 for an event with five PVs, where a photon hit can be associated with more than ten tracks.

use only the correct POs, identified with MC truth information. In each plot, the reconstructed RICH PV t_0 (given by the mean of the distribution) is indicated by a vertical dotted line and the MC PV t_0 by a plain vertical line. When all POs are used, the estimated RICH PV t_0 can deviate from the MC value by tens of picoseconds. In contrast, when only the correct POs are used, the difference between the RICH and the MC PV t_0 is reduced to just a few picoseconds. These examples illustrate that PO mis-associations are the main limitation in the RICH PV t_0 estimation and underlines the need for effective PO selection criteria. In the next section, the method is extended to the full sample of ten thousand events and the RICH PV t_0 resolution is evaluated with focus on the effect of the quality of POs.

3.2.3 Purity of photon-objects and PV t_0 resolution

The RICH PV t_0 reconstruction method was studied using the simulation sample of ten thousand bunch-crossing events at LHC Run 4 beam conditions, described in Section 3.2.1. Figure 3.12 shows the distribution of the difference between the RICH PV t_0 and MC PV t_0 for all the PVs in the sample. The RICH PV t_0 resolution is extracted as the FWHM of a double-Gaussian fit to the distribution. The RMS of the distribution is also considered as an indicator of the spread from the tails. In Figure 3.12 (left), the RICH PV t_0 is calculated using all POs associated with each PV. The resulting FWHM is ~400 ps with an RMS of approximately 200 ps, consistent with the intrinsic spread of the PV t_0. Figure 3.12 (right) shows the result using instead only the correct POs from MC information. The spread is ~25 ps, with a fitted FWHM of approximately 30 ps, representing the best RICH PV t_0 resolution achievable with ideal PO selection. The strong improvement observed between the two distributions emphasises the critical role of the PO selection. Each PV produces, on average, more than 15 tracks, yielding hundreds of POs per PV in each RICH detector. This provides sufficient PO statistics to enable for such a selection.

A figure of merit, PV purity, has been introduced to quantify the impact of PO selections on the RICH PV t_0 resolution. It is defined as:

$$\text{PV purity} = \frac{N\text{correct POs associated with PV}}{N\text{all POs associated with PV}} \tag{3.11}$$

Figure 3.13 shows the RICH PV t_0 resolution as a function of the PV purity and the number of POs associated per PV. The purpose of this plot is not to measure an absolute PV t_0 resolution, but it illustrates the trend between PV purity, the number of POs per PV, and resolution. The RICH PV t_0 resolution improves as PV purity increases. A higher number of POs per PV also improves the resolution by reducing statistical fluctuations. However, at low PV purity, incorrectly associated POs degrade the resolution and limit the benefit of having more POs. Increasing PV purity typically reduces the number of POs per PV. Under LHC Run 4 conditions, the average PV purity without selection is approximately 25 % for RICH 1 and 29 % for RICH 2. High-quality

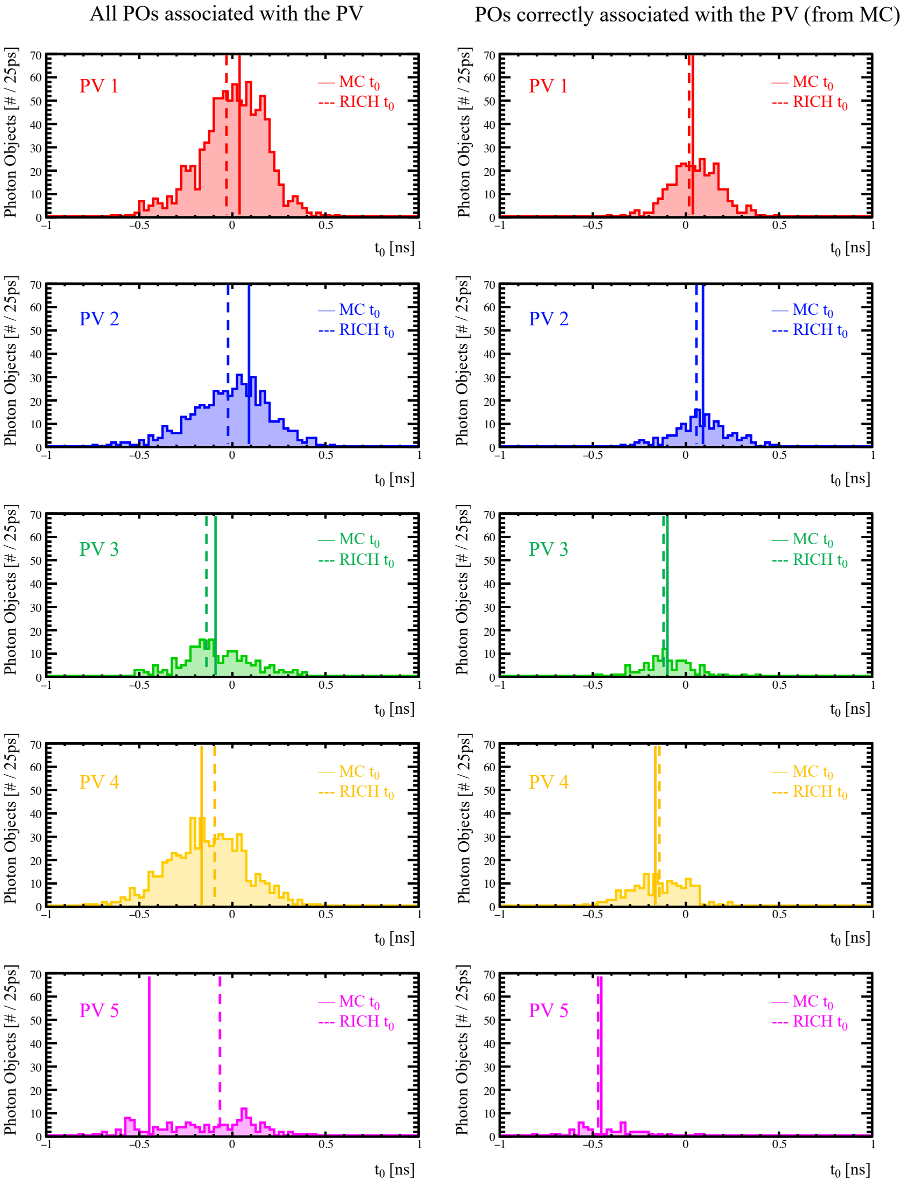

Figure 3.10: t_0 distribution of all the POs associated with each PV (left column) and only the correctly associated POs (right column) for an event with five PVs in RICH 1. On each plot, the vertical dotted line represent the reconstructed RICH PV t_0 (hence the average of the distribution), while the vertical plain line is the Monte Carlo PV t_0.

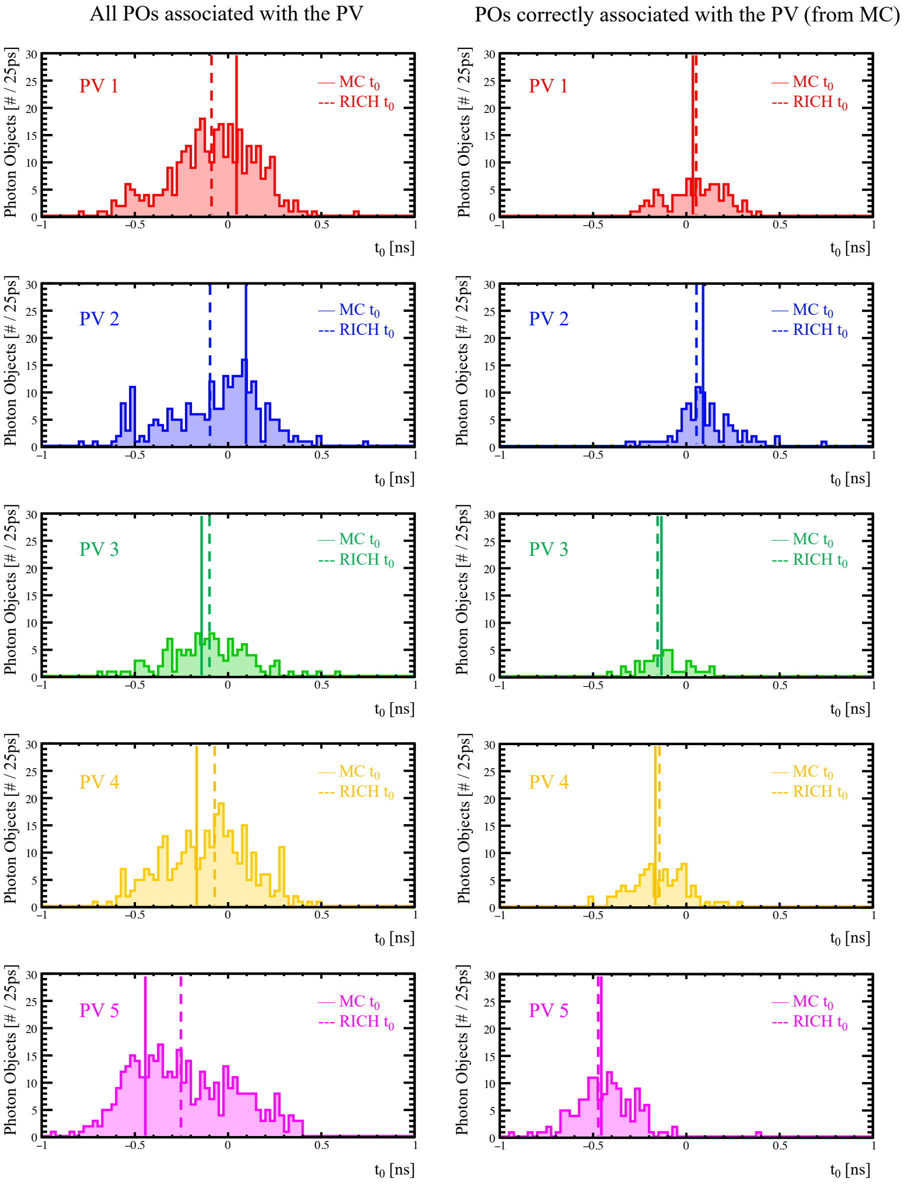

Figure 3.11: t_0 distribution of all the POs associated with each PV (left column) and only the correct POs associated with the PV (right column) for an event with five PVs in RICH 2. On each plot, the vertical dotted line represent the reconstructed RICH PV t_0 (hence the average of the distribution), while the vertical plain line is the Monte Carlo PV t_0.

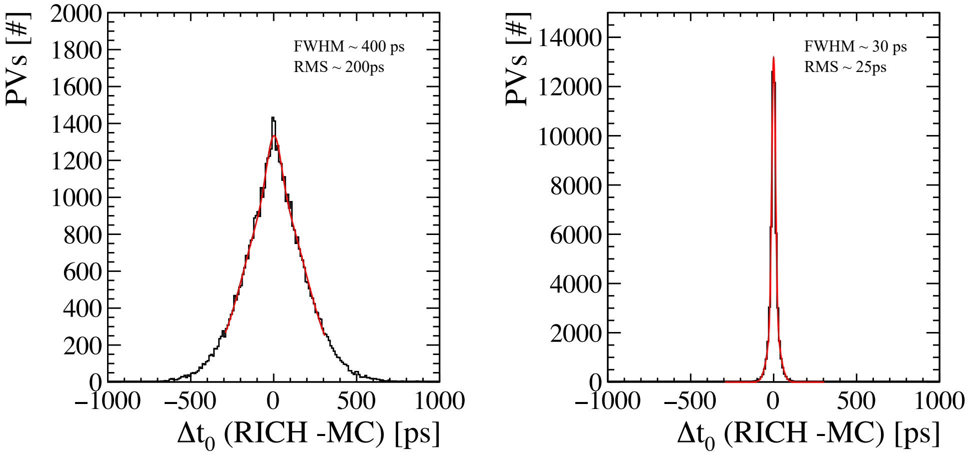

Figure 3.12: Difference between the RICH PV t_0 and MC PV t_0 for all the PVs in a sample of ten thousand bunch crossing events at LHC Run 4 beam conditions. The RICH PV t_0 is estimated using all the POs (left) and selecting only the correct POs using MC information (right).

PO selection is therefore essential to improve purity without reducing the PO sample too strongly. During LHC Run 4, PV purity will not be available, as it contains MC information, but it serves as a useful metric to assess and compare different selection criteria in simulation. Instead, an uncertainty will be assigned to each estimate of the PV t_0 based on the event topology and applied selections. This uncertainty will be useful for the subsequent application of the software time gate on the POs, described in Section 3.1.2.

Two approaches have been evaluated to improve PV purity and consequently the PV t_0 resolution. These are described in Section 3.3.

3.3 Photon-object selection criteria and PV t_0 resolution

High-quality POs selection is crucial for improving the PV purity and consequently the RICH PV t_0 resolution. This section evaluates two approaches for the PO selection. The first approach is described in Section 3.3.1 and consists of selecting tracks, and their associated POs, based on global track properties such as momentum and pseudorapidity. This method has limitations, since each track can contain POs of different quality, and this level of details is lost selecting all POs associated to a track. For this reason, Section 3.3.2 introduces selection criteria based on PO properties and achieves significant improvements in PV purity and PV t_0 resolution. Section 3.3.3 shows how the PV t_0 resolution is affected by the event topology and the number of PVs in the event.

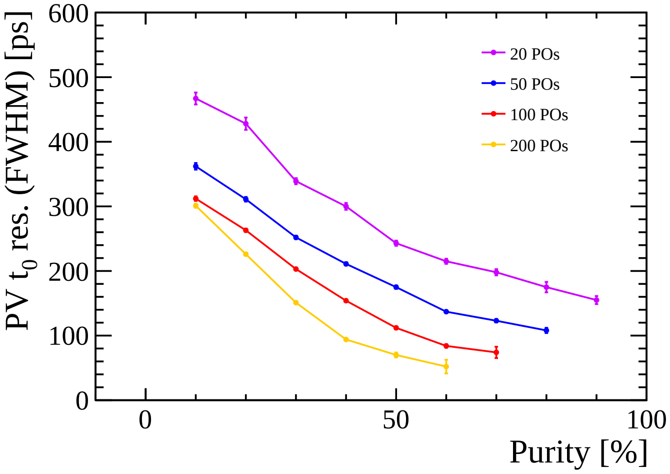

Figure 3.13: The RICH PV t_0 resolution as a function of the PV purity and of the number of POs per PV.

3.3.1 Selection criteria based on track properties

This section investigates tracks properties for the selection of tracks that are likely to yield high-purity POs. Analogous to PV purity, track purity is defined as the number of correct POs divided by the total number of POs associated with the track:

$$\text{Track purity} = \frac{N\text{correct POs associated with the track}}{N\text{all POs associated with the track}} \tag{3.12}$$

As with PV purity, track purity contains Monte Carlo information and can only be used as a figure of merit in simulation to evaluate the effectiveness of selection criteria. Since track purity contributes directly to PV purity, selecting POs from high-purity tracks only can improve PV purity and, consequently, enhance the PV t_0 resolution. Two track properties have been studied: momentum and pseudorapidity. Figure 3.14 (left) shows the average track purity as a function of momentum. The momentum distribution of all tracks (right) for the ten thousand bunch-crossings sample is shown in Figure 3.14 (right). No significant dependence on momentum is observed for track purity in either RICH 1 or RICH 2. The average purity remains relatively constant across the full momentum range. Only for low-momentum tracks in RICH 2 the average purity reaches 35 %, which is still insufficient to significantly improve the PV t_0 resolution. This indicates that track momentum is not a discriminator for selecting high-purity POs.

The second track property explored in this section is the track pseudorapidity (η), defined in Equation 1.6. Pseudorapidity is a spatial coordinate that describes the displacement of a particle relative to the beam axis. High η values correspond to tracks travelling at small angles relative

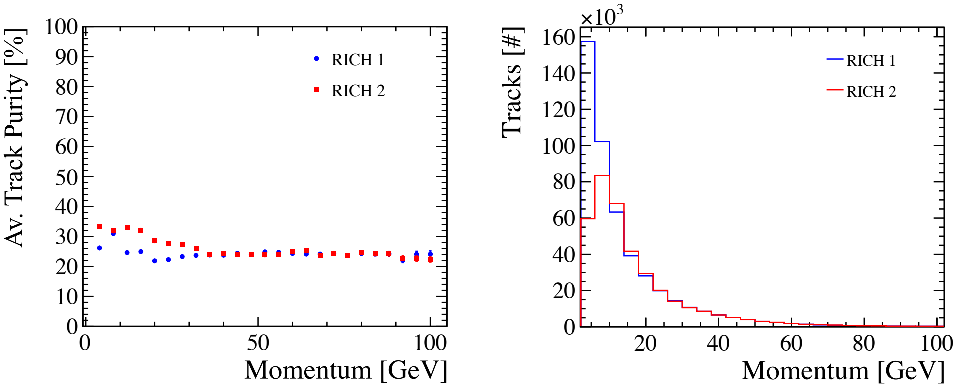

Figure 3.14: (Left) Average track purity as a function of track momentum. (Right) Distribution of track momentum for all the tracks in ten thousand bunch-crossing events.

to the beam direction, whereas low values correspond to tracks that diverge more from the beam axis. Figure 3.14 (left) shows the average track purity as a function of η for both RICH 1 and RICH 2. The average PV purity exceeds 50 % for tracks with η lower than 2.25 in RICH 1. These tracks, being more transverse, produce Cherenkov rings in the outer region of the detector, where the occupancy is lower. In this region, isolated rings and reduced overlap with other tracks are more likely, resulting in higher track purity. As shown in Figure 3.14 (left), the fraction of tracks with η lower than 2.25 is approximately 5 % of the total, from which the PV t_0 estimation for ~25 % of the PVs in the sample can be performed. The objective remains to develop selection criteria that enhance the PV t_0 resolution for as many PVs as possible. Therefore, this selection alone is therefore not sufficient. However, this result points to the presence of isolated tracks in low-occupancy regions of the detector. Isolated rings from isolated tracks could be geometrically identified in the event hitmap and directly used for the estimation of their PV t_0. This would provide a quick method for the ~25 % of PVs mentioned above, reducing the load on the t_0 estimate of the remaining PVs. A similar effect is not observed in RICH 2, possibly because its low-occupancy region is instrumented with 2-inch MAPMTs, whose larger pixel size increases the likelihood of rings overlap.

Track-based cuts cannot achieve high PO purity across a large fraction of PVs in the sample, as the quality of POs within a track depends on several factors beyond main global track parameters. This motivates a different strategy, based on per-PO selection rather than per-track.

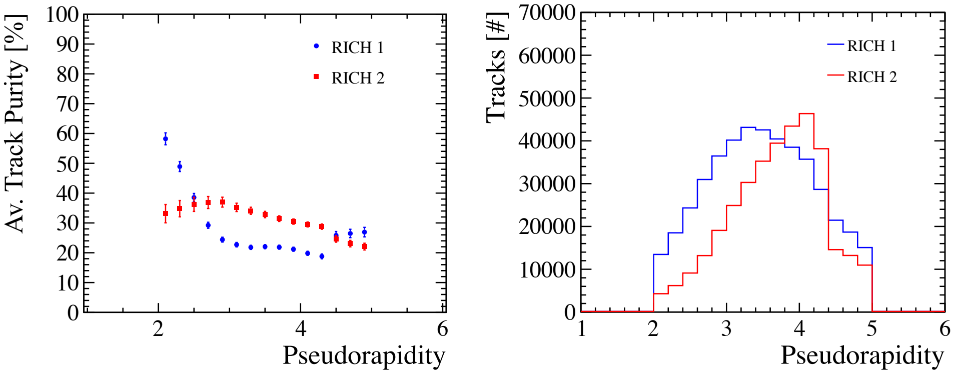

Figure 3.15: (Left) Average track purity as a function of track pseudorapidity. (Right) Distribution of the track pseudorapidity for all the tracks in ten thousand bunch-crossing events.

3.3.2 Selection criteria based on PO properties

As discussed in Section 3.2.2, PO selection plays a key role in the PV t_0 resolution. Track-based selection criteria have shown limited effectiveness in isolating correct POs and are therefore not sufficient. This section presents two selection criteria based on PO properties. Both demonstrate clear improvements in PV purity and, as a result, in the PV t_0 resolution.

Photon-objects uniquely associated with a pixel hit

The first criterion selects POs that are uniquely associated with a pixel hit. An example is shown by the green hits shown in Figure 3.9 (left). These hits come from portions of Cherenkov rings that do not overlap with others. They are associated with only one track, that is most likely the one from which they originated. The resulting PO has a higher probability to be a correct association and to produce a correct estimate of the PV t_0.

Figure 3.16 (top-row) shows the hitmap of pixel hits in RICH 1 and RICH 2 for an event with five PVs. Overlapped Cherenkov rings cause a single pixel hit to be associated with more than ten tracks in the high-occupancy regions of the detectors. This can be seen in the hitmap of POs in the middle row of Figure 3.16, where more than ten POs can be generated from the same pixel hit in both detectors. The bottom row of Figure 3.16 shows the same hitmaps after the selection criterion is applied, with the correct POs highlighted in red. As expected, the effect of the selection is more pronounced in the high-occupancy regions, where track multiplicity is high, while outer regions remain largely unaffected. Most of the remaining incorrect POs originate from secondary particles

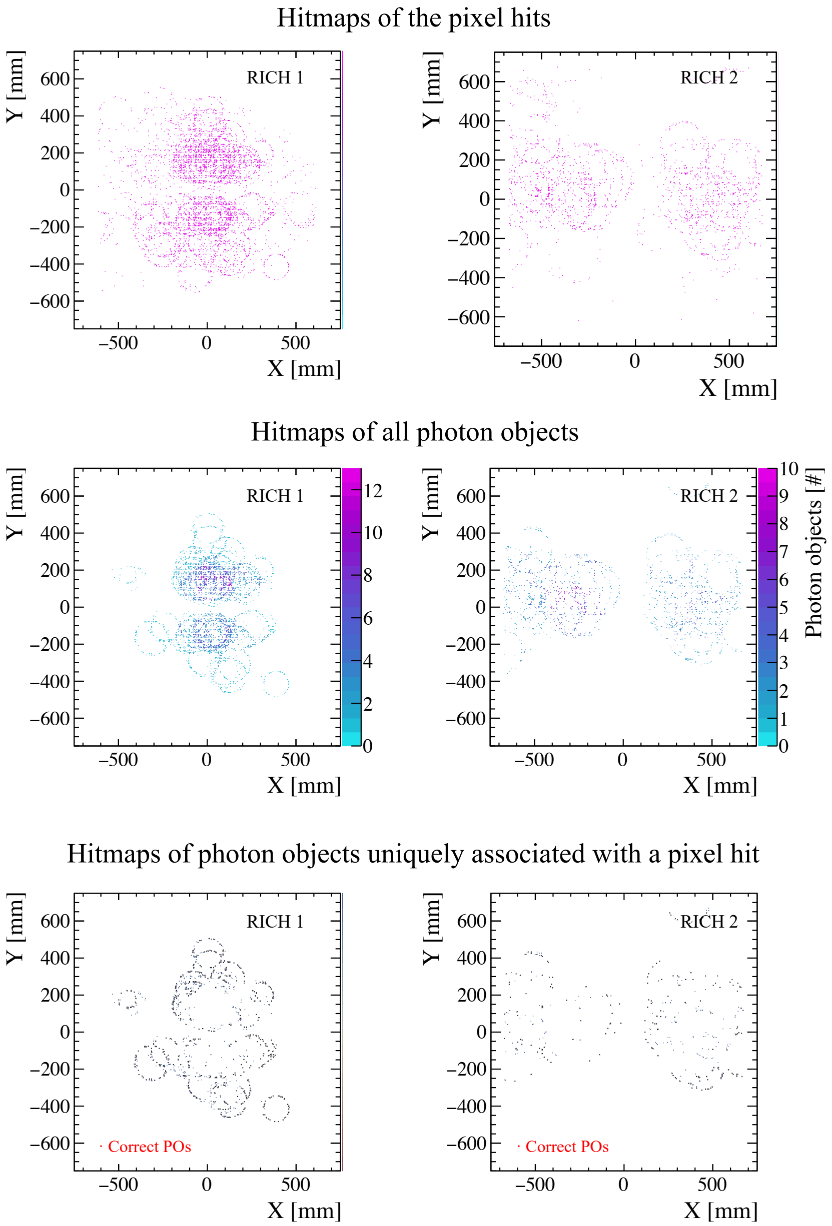

Figure 3.16: Hitmap distribution of the pixel hits (top), of the POs (middle), and of the POs uniquely associated with a pixel hit (bottom), in RICH 1 (left) and RICH 2 (right) for an event with 5 PVs.

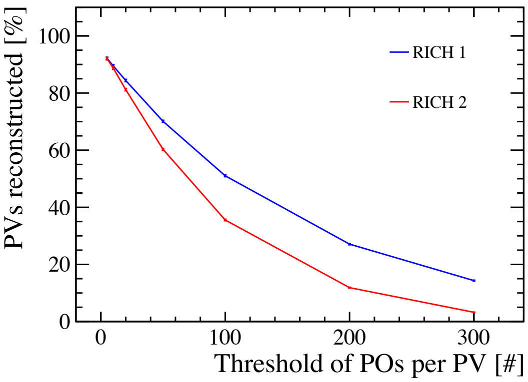

Figure 3.17: Fraction of PV reconstructed as a function of the minimum number of POs required per PV, using RICH 1 (blue) and RICH 2 (red).

produced in interactions of primary particles with the LHCb tracking stations or with the RICH detector entrance window. These secondary particles are mainly electrons with a momentum from a few MeV to hundreds of MeV, enough to produce Cherenkov photons in the RICH detectors. Given the low mass compared to hadrons, the electron Cherenkov threshold corresponds to a momentum of 20 and 40 MeV/c in the gas radiators of RICH 1 and RICH 2, respectively. The resulting photon hits from these secondary electrons are spread across the entire photodetector plane, with a higher concentration in the high-occupancy region. When one of these hits falls within the geometrical constraints of a reconstructed track, it may be incorrectly associated with that track, resulting in a fake PO.

The selection has been applied to the simulation sample of ten thousand bunch-crossing events at LHC Run 4 beam conditions. As a result of the selection, the average PV purity increases to 53 % in RICH 1 and 55 % in RICH 2. The selection has a strong impact on POs statistics. The total number of POs in the sample is reduced by approximately 80 %, retaining on average 160 POs per PV in RICH 1 and 100 in RICH 2. Figure 3.17 shows the percentage of the PVs reconstructed as a function of the minimum number of POs per PV required for the PV t_0. Despite the reduction in the number of POs per PV, 84 % of the PVs in the sample retain more than 20 POs in RICH 1 and 81 % in RICH 2. Combining the POs from the two detectors, the fraction of PVs with more than 20 POs increases to 95 %.

Figure 3.18 shows the distributions of the difference between the estimated RICH PV t_0 and MC PV t_0 after the PO selection is applied. The distributions are obtained using the POs from RICH 1 (left) and RICH 2 (right), including PVs with at least 20 POs per PV. As per Figure 3.12,

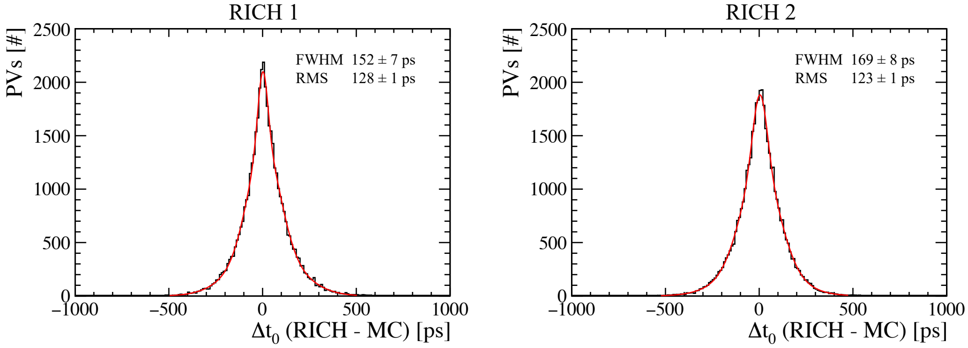

Figure 3.18: Difference between the RICH PV t_0 and Monte-Carlo PV t_0 using the POs uniquely associated with a pixel hit in RICH 1 (left) and RICH 2 (right). In red the double-Gaussian fit applied to the distribution to extract the FWHM as the PV t_0 resolution.

the PV t_0 resolution is extracted as the FWHM of a double-Gaussian fit, and the RMS is also used as an indicator of the spread, accounting for contributions from the tails. The PV t_0 resolution is 151 ± 3 ps FWHM (128 ± 1 ps RMS) for RICH 1 and 169 ± 3 ps FWHM (123 ± 1 ps RMS) for RICH 2. This demonstrates already a significant improvement compared to the distribution without PO selection (~400 ps FWHM), shown in Figure 3.12 (left). In addition, the values indicate that the RICH 1 distribution has a narrower core than RICH 2, although the overall spread is comparable between the two detectors. The narrower core of RICH 1 can be attributed to the higher average number of POs per PV compared to RICH 2. In some cases, the lower PO multiplicity in RICH 2 may limit the PV t_0 resolution, even for high PV purity.

Figure 3.19 shows the PV t_0 resolution, expressed as FWHM (left) and RMS (right), as a function of the minimum number of POs required per PV. As the threshold on the number of POs increases, the PV t_0 resolution improves for both detectors. With minimum of 300 POs per PV, the PV t_0 resolution reaches approximately 100 ps, although for only 20 % of the PVs in RICH 1 and 5 % in RICH 2. This shows that the selection criterion is effective, but the achievable PV t_0 resolution is limited by the available PO statistics. A compromise is therefore needed between the fraction of PVs that can be reconstructed and the resulting PV t_0 resolution.

Signal amplitude cut

The second selection criterion is based on the PO signal amplitude, which is computed as part of the RICH reconstruction. Defined in Equation 3.1, the signal amplitude is calculated for each PO

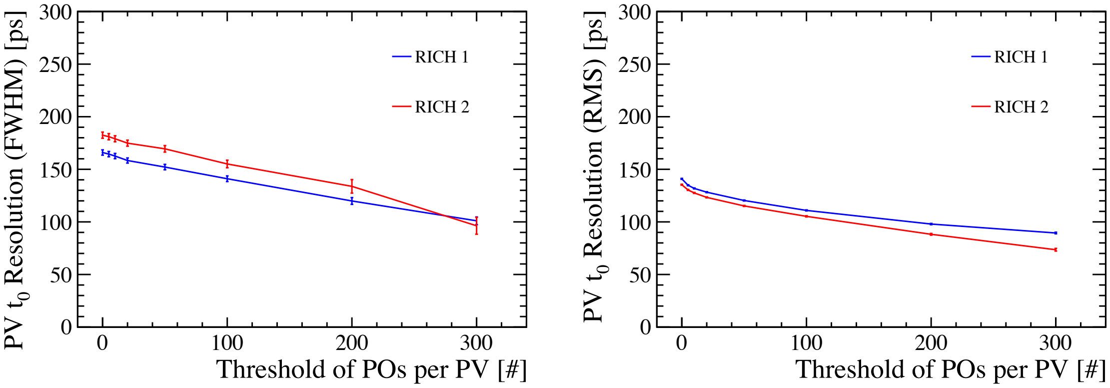

Figure 3.19: PV t_0 resolution as a function of the minimum number of POs required per PV. The resolution is shown as FWHM (left) and RMS (right) for RICH 1 (blue) and RICH 2 (red).

and for every mass hypothesis. It represents the probability that the pixel hit associated with the PO originates from the associated track with the given mass hypothesis. Since this selection is applied before the likelihood maximisation stage of the reconstruction, the mass hypothesis is set to the pion mass, the most abundant particle type in the LHCb experiment. Therefore, the signal amplitude selection consists of discarding all the POs with pion signal amplitude below a chosen value.

Figure 3.20 shows the average PV purity (left) and the number of POs per PV (right) as a function of the signal amplitude cut. As the cut becomes more stringent, the PV purity increases while the number of POs per PV decreases. Notably, even a cut of 0.005 leads to a significant improvement in purity for both detectors. Due to differences in detector geometry, hit occupancy and background level, RICH 2 achieves higher PV purity than RICH 1. The PV purity can reach up to 83 % in the RICH 2 detector, while it saturated at approximately 50 % for RICH 1 However, the gain in the PV purity results in a stronger reduction in POs per PV for RICH 2 compared to RICH 1. A signal amplitude cut of 0.1 for RICH 1 and 0.05 for RICH 2 increases the PV purity to 51 % and 61 %, respectively. These cuts represent a reasonable trade-off between improving PV purity and preserving PO statistics, as they retain an average of 450 POs per PV in RICH 1 and 170 POs in RICH 2.

Figure 3.21 (top) shows the percentage of PV reconstructed as a function of the signal amplitude cut for RICH 1 (left) and RICH 2 (right). The figure includes curves for different thresholds on the number of required POs per PV, ranging from 0 to 100. As the signal amplitude cut increases, fewer POs are retained per event, and the fraction of reconstructed PVs gradually reduces. This reduction is more pronounced for RICH 2 due to its lower average number of POs per PV. The

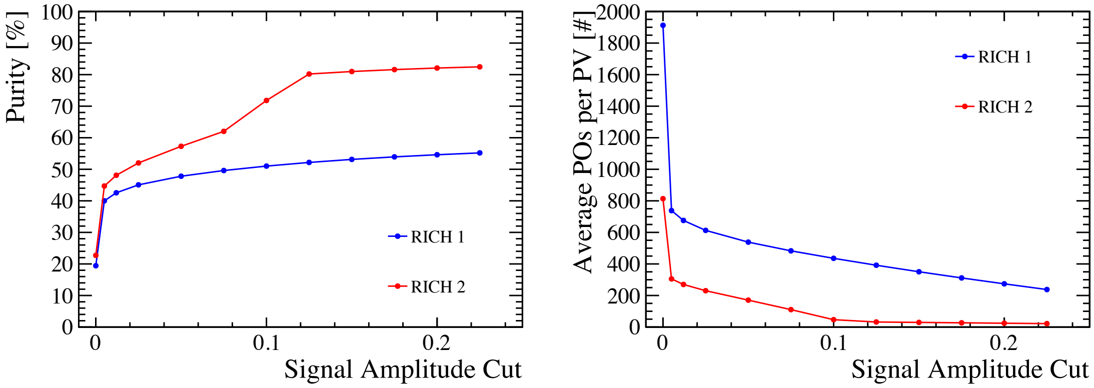

Figure 3.20: Average PV purity (left) and average number of PO per PV (right) as a function of a PO signal amplitude cut.

reduction also becomes more significant when the minimum number of POs per PV is increased. For example, requiring 100 POs per PV, the fraction of reconstructed PVs already shows a sharp drop even at a signal amplitude cut as low as 0.005.

The middle row of Figure 3.21 show the PV t_0 resolution, expressed in FWHM, as a function of the signal amplitude cut for RICH 1 (left) and RICH 2 (right). The PV t_0 resolution is consistent with the behaviour observed for the PV purity in Figure 3.20 (left). As the signal amplitude cut becomes more stringent, the resolution improves for both detectors due to the increased of PV purity, and even a cut of 0.005 results in a significant improvement of the PV t_0 resolution. RICH 2 achieves better PV t_0 resolutions than RICH 1. A plateau of approximately 170 ps is observed for RICH 1, where the FWHM stabilises beyond a signal amplitude cut of 0.1. This is consistent with the saturation of the PV purity observed in Figure 3.20 (left). In RICH 2, a resolution of 80 ps can be achieved with a signal amplitude cut of 0.125 and a threshold of 50 POs per PV, although for only 15 % of the PVs. For most cut values in both detectors, the number of POs per PV has no statistical effect on the PV t_0 resolution. The curves corresponding to different PO thresholds overlap closely, except for signal amplitude cuts above 0.1 in RICH 2. This behaviour is explained by the average PO statistics. In RICH 1, the number of POs per PV is above 100 for all cuts, while in RICH 2 this is valid only for cuts below 0.1. In addition, the PV t_0 resolution trends observed with the FWHM are confirmed by the RMS of the $\Delta t_0(RICH - MC)$ distributions, as shown in the bottom row of Figure 3.21.

Figure 3.21 confirms that a signal amplitude cut of 0.1 for RICH 1 and 0.05 for RICH 2, combined with a minimum of 5 POs per PV, provide the best balance between PV t_0 resolution and fraction of PVs reconstructed. With this selection, a PV t_0 resolution of 173 ± 4 ps FWHM (132 ± 1 ps

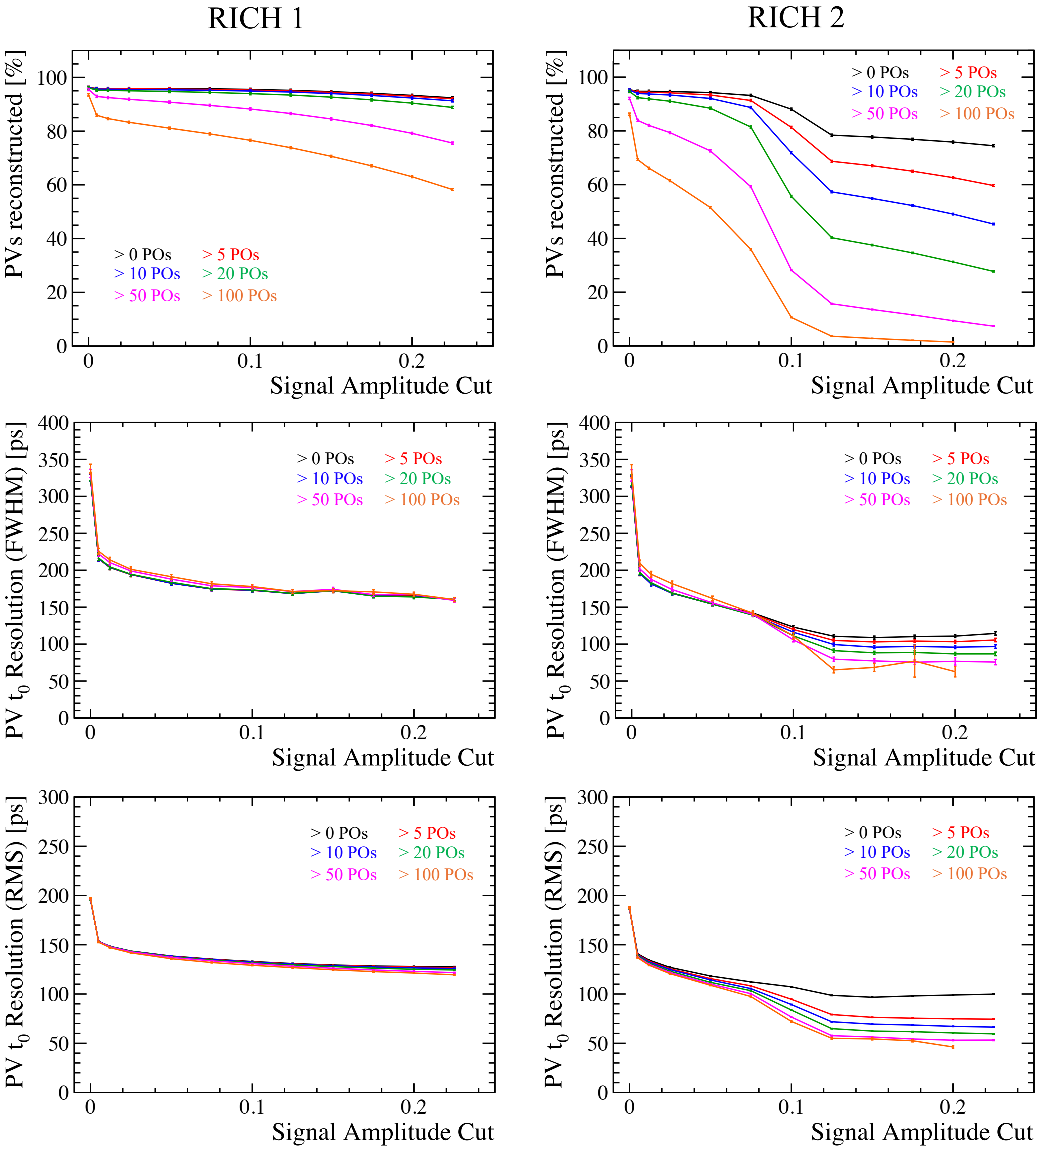

Figure 3.21: Fraction of PVs reconstructed (top), PV t_0 resolution expressed as FWHM (middle) and PV t_0 resolution expressed as RMS (bottom) as a function of the signal amplitude cut for different minimum number of POs per PV. The plots are obtained using POs from RICH 1 (left) and RICH 2 (right).

RMS) is obtained in RICH 1 for 94 % of the PVs, and 154 ± 3 ps FWHM (115 ± 1 ps RMS) in RICH 2 for 88 % of the PVs.

These results are similar to those obtained from the pixel-based PO selection. Both approaches alone do not achieve a PV t_0 resolution below 100 ps for a large fraction of PVs. In the next paragraph, the combination of the two selection criteria is evaluated to further improve the PV t_0 resolution.

Combination of the two PO-based selection criteria

In order to combine the two selection criteria described above, it is important to show first that the two selections do not target the same subset of POs. Figure 3.22 shows the PO hitmaps in RICH 1 for the ten thousand events after the two selections: on the left only POs uniquely associated with a pixel hit, on the the right POs with a signal amplitude greater than 0.10. The two selections affect the spatial distribution of POs in different ways. As seen in Figure 3.16, selecting POs uniquely associated with a pixel hit has a significant impact in the high-occupancy region, where track multiplicity is large. By contrast, the effect signal amplitude cut is less pronounced in this region. In areas where many rings overlap, an incorrect PO may still have a high signal amplitude under the pion hypothesis if its associated track is indeed a pion with adequate momentum. The signal amplitude cut, however, effectively removes POs generated by photon hits from secondary electrons. These POs typically have a high signal amplitude under the electron mass hypothesis, but a much lower value under the pion hypothesis. The two selection criteria therefore act on different subsets of POs and are complementary. Applying both selections simultaneously — POs uniquely associated with a pixel hit and a signal amplitude cut of 0.1 for RICH 1 and 0.05 for RICH 2 — the average PV purity increases to 83 % in RICH 1 and 81 % in RICH 2, without excessively reducing the number of usable POs.

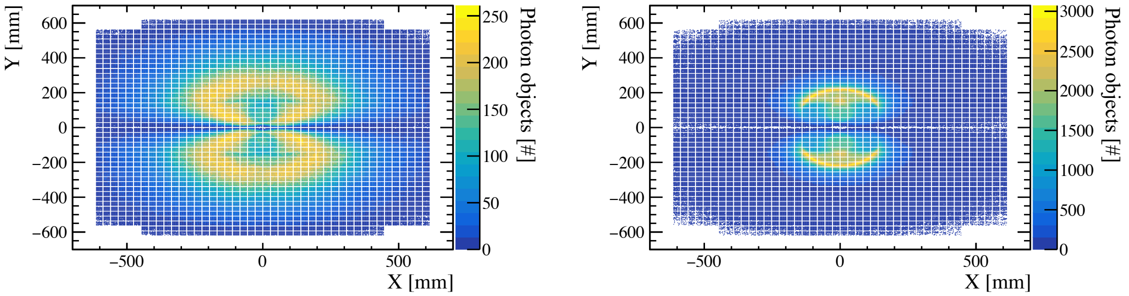

Figure 3.22: PO hitmap in RICH 1 for ten thousand events: (left) selecting POs uniquely associated with a pixel hit, (right) applying a signal amplitude cut of 0.10.

Figure 3.23 shows the same set of plots as Figure 3.21, with the additional selection of POs uniquely associated with a pixel hit. With this combined selection, the PV t_0 resolution improves significantly. A small signal amplitude cut already provides a large improvement in PV t_0 resolution, after which the performance gradually reaches a plateau beyond 0.05, with no further gain for higher signal amplitude cut values. With a minimum of 5 POs per PV, the PV t_0 resolution reaches 88±3 ps FWHM (70 ± 2 ps RMS) for 85 % of PVs in RICH 1, and 93 ± 3 ps FWHM (69 ± 2 ps RMS) for 81 % of PVs in RICH 2.

Combining the two selection criteria has a significant impact on the available POs statistics, limiting the achievable PV t_0 resolution. This limitation is shown by the resolution curves for different thresholds on the minimum number of POs per PV in Figure 3.23. Increasing the required POs per PV further improves the PV t_0 resolution but significantly reduces the fraction of PVs reconstructed. For example, requiring at least 100 POs per PV yields a PV t_0 resolution of about 50 ps, but only for 20 % of the PVs in RICH1 and 5 % in RICH2.

A possible solution to improve the PO statistics is to combine the POs associated with the same PV from both detectors. Given the similar behaviour of the two detectors in terms of PV t_0 resolution, this solution could increase the fraction of PVs reconstructed and possibly further improve the PV t_0 resolution. Figure 3.24 shows the fraction of PVs reconstructed combining the POs from both RICH detectors. The curves are obtained with the two criteria applied and for different thresholds on the minimum number of POs per PV. This combination clearly improves the fraction of PVs reconstructed compared to treating the detectors separately, as in Figure 3.23. For a signal amplitude cut of 0.05, the fraction of PVs reconstructed with at least 5 POs increases to 94 %. In addition, the PV t_0 resolution improves to 85±3 ps FWHM (66±2 ps RMS), as shown in Figure 3.25. This configuration represents the optimal balance between PV t_0 resolution and number of PVs reconstructed. Although combining information from both detectors increases the number of PVs reconstructed, the available POs statistics remains a limitation for the PV t_0 resolution. With a signal amplitude cut of 0.10 and a minimum of 100 POs per PV, the PV t_0 resolution reaches 54 ± 3 ps FWHM (34 ± 2 ps RMS) for 25 % of the PVs in the sample. This result is significant, as it approaches the PV t_0 resolution of 25 ps (RMS) obtained using only correctly associated POs from MC information (Figure 3.12, right). Figure 3.26 shows the $\Delta t_0(RICH\text{–}MC)$ distributions obtained with the optimal selection (left) and the stricter cut (right). Comparison with Figure 3.12 highlights the improvement with respect to the case without any PO selection (Figure 3.12, left), while approaching the ideal PV t_0 resolution obtained using only correct POs from MC information (Figure 3.12, right).

These results show that the PO selection strategy based on PO properties is effective in isolating correctly-associated POs. They demonstrate that the PV t_0 can be estimated using only RICH information under LHC Run 4 conditions, achieving a sub-100 ps PV t_0 resolution for almost all PVs in the sample. These results provide the basis for the application of the software time gate

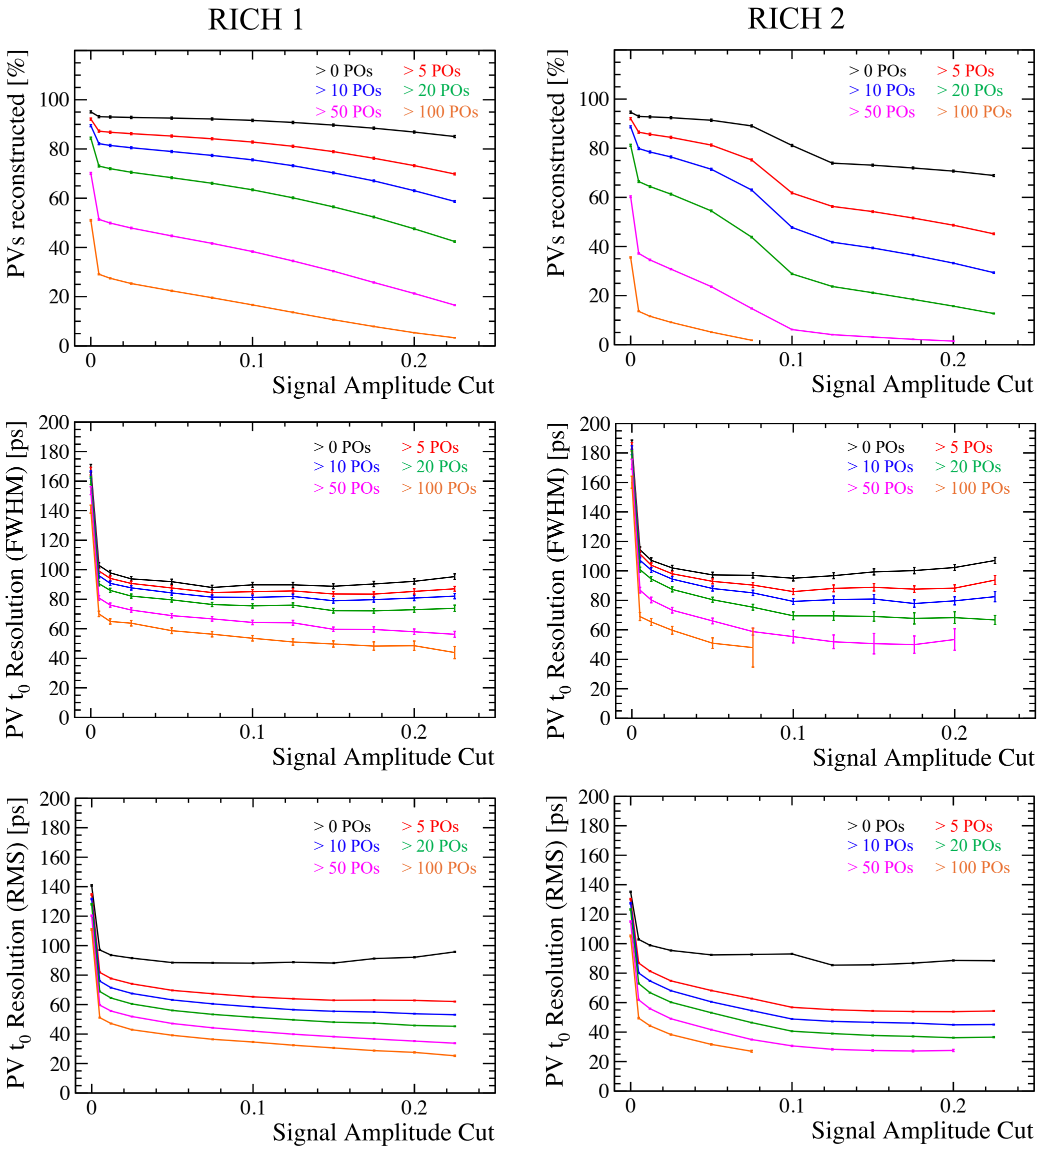

Figure 3.23: Fraction of PVs reconstructed (top), PV t_0 resolution expressed as FWHM (middle) and PV t_0 resolution expressed as RMS (bottom) as a function of the signal amplitude cut for different minimum number of POs per PV, with the selection of POs uniquely associated with a pixel hit. The plots are obtained using POs from RICH 1 (left) and RICH 2 (right).

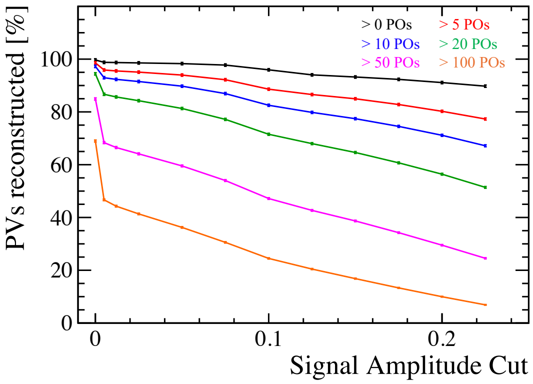

Figure 3.24: Fraction of PVs reconstructed as a function of the signal amplitude cut for different minimum numbers of POs per PV. The selection of POs uniquely associated with a pixel hit is also applied. For each PV, POs from both RICH detectors combined.

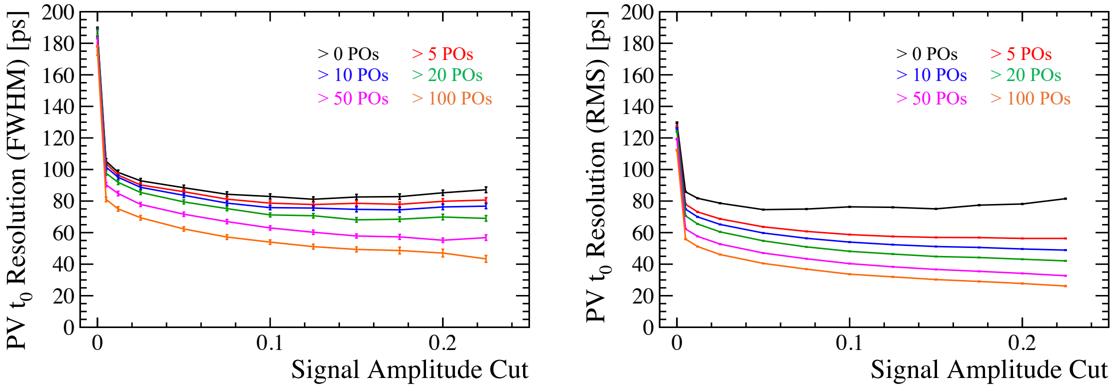

Figure 3.25: PV t_0 resolution expressed as FWHM (left) and PV t_0 resolution expressed as RMS (right) as a function of the signal amplitude cut for different minimum number of POs per PV, with the selection of POs uniquely associated with a pixel hit. For each PV, the POs are taken from both RICH detectors.

described in Section 3.1.2. Since MC information on the PV t_0 will not be available during LHC Run 4, an uncertainty margin may be assigned to each PV t_0 estimate. This margin would then propagate to the software time gate width. Possible strategies for the definition and implementation of this margin are discussed in Section 3.4. These results were obtained using the full sample of ten thousand bunch-crossing events with minimum bias. However, the event topology can have a significant impact on the PV t_0 resolution. For example, the number of PVs per event is an

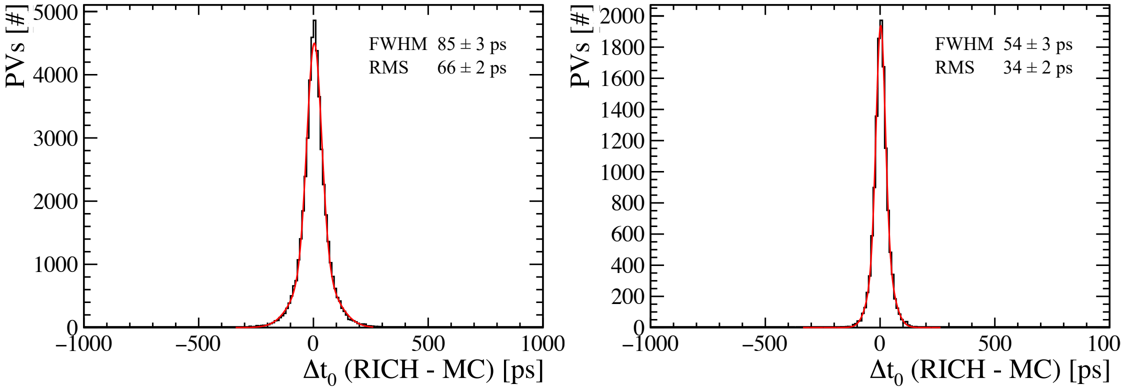

Figure 3.26: Distributions of Δt_0 (RICH–MC) for the optimal PO selection (left) and for a stricter cut (right). In the first distribution each PV t_0 is reconstructed with at least 5 POs with a signal amplitude greater than 0.05. In the second, the required POs are 100 with a signal amplitude greater than 0.10. In both cases, the POs are uniquely associated with a pixel-hit. With the optimal selection the PVs reconstructed are 94 % of the total, while with the stronger selection 25 %.

important factor. Events with many PVs have higher track and photon occupancy, which could affect the PO selection and, as a consequence, the PV t_0 resolution.

3.3.3 PV t_0 resolution and event topology

The PV t_0 resolution strongly depends on event topology, in particular on the number of PVs in the event. More PVs result in higher number of tracks and photons, which raise PO statistics but also the probability of incorrect hit–track associations. In addition, the selections described in this section rely on spatial information, and their efficiency may decreases as occupancy grows.

Figure 3.27 (left) shows the distribution of events as a function of the number of PVs per event. The average number of PVs per event at LHC Run 4 luminosity is 5.3. The fraction of PVs reconstructed as a function of the number of PVs in the event is shown in Figure 3.27 (right) for two PO selections. The first selection (red) corresponds to the optimal solution discussed in Section 3.3.2. Each PV t_0 is reconstructed with at least 5 POs, uniquely associated with a pixel hit and with signal amplitude greater than 0.05. With this selection, the fraction of PVs reconstructed decreases only slightly with PV multiplicity, and remains near 85 % even for events with 10 PVs. The second, stricter selection (green) requires at least 100 POs per PV, uniquely associated with a pixel hit and with signal amplitude greater than 0.10. This selection rejects more incorrect POs, improving the PV t_0 resolution but also reducing the number of PVs reconstructed. In this case, the fraction of reconstructed PVs shows a strong dependence on PV multiplicity. On average only

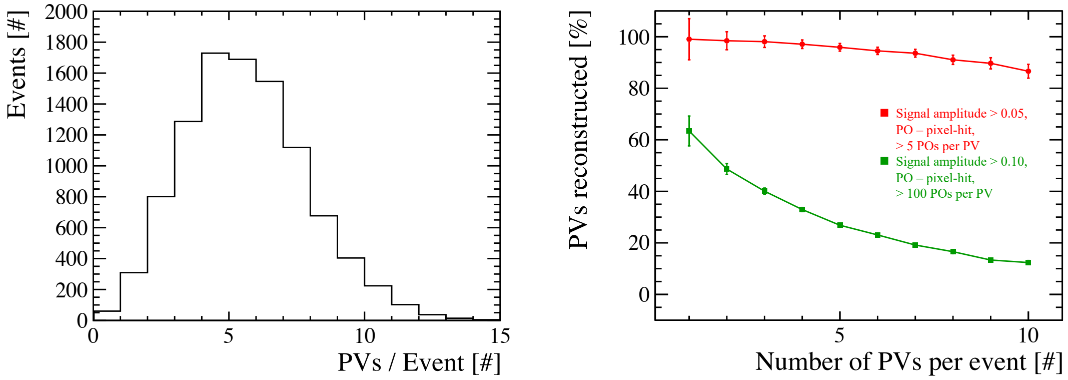

Figure 3.27: (Left) Distribution of events as a function of the number of PVs per event with the LHC Run 3 luminosity. (Right) Fraction of PVs reconstructed with the PV t_0 algorithm as a function of the number of PVs per event for two different PO selections: the optimal selection (red) and a stricter one (green).

25 % of PVs are reconstructed, with the fraction reaching about 50 % for events with 2 PVs and falling to 12 % for events with 10 PVs.

Figure 3.28 shows the PV t_0 resolution as a function of the number of PVs per event for the two PO selection solutions and for the case without any selection. The resolution is shown as FWHM (left) of a double-Gaussian fit and the RMS (right) of the $\Delta t(RICH - MC)$ distributions, as described in Section 3.2.2. For events with a single PV, the PV t_0 resolution is approximately 25 ps RMS in all three configurations, consistent with the resolution obtained using the correct POs from Monte Carlo. This result is expected given that tracks from the same PV arrive almost simultaneously to the RICH detector. Without PO selection, the resolution degrades rapidly with PV multiplicity. For events with more than two PVs, the PV t_0 resolution exceeds 200 ps FWHM (and 150 ps RMS) and rapidly saturates at the level of the intrinsic PV t_0 spread in the sample. This confirms that, without PO selection, the PV t_0 estimate is highly sensitive to the detector occupancy and PV purity, even in relatively simple events. With the optimal selection (red), the dependence of the PV t_0 resolution on the PV multiplicity is significantly reduced. The resolution stays below 120 ps FWHM (85 ps RMS) for events with up to 10 PVs. This demonstrates that the selection is not merely identifying PVs from simple events, but it is also effective in high-multiplicity events. This confirms the selection as optimal compromise between PV t_0 resolution and fraction of PVs reconstructed. With the stricter selection (green) the resolution improves further and remains below 90 ps FWHM (50 ps RMS) at 10 PVs. However, this improvement is obtained at the cost of a substantial reduction in PVs reconstructed, as shown in Figure 3.27.

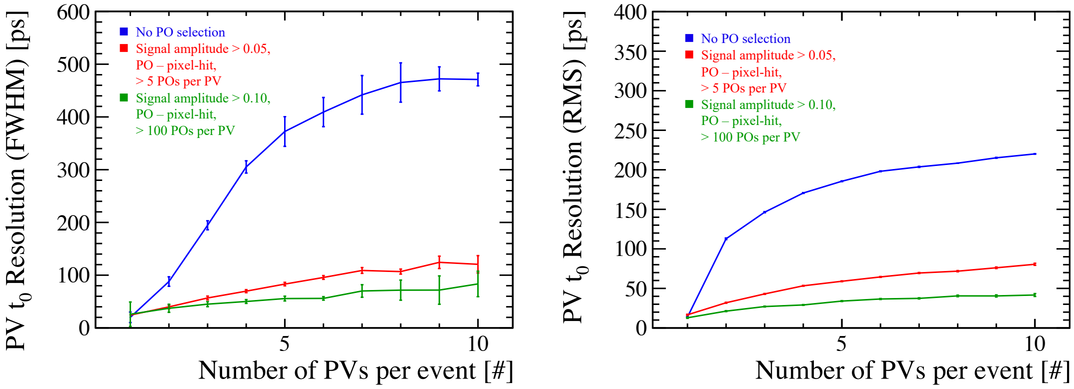

Figure 3.28: PV t_0 resolution as a function of the number of PVs per event, expressed as FWHM (left) and RMS (right). Results are shown without PO selection (blue), for the optimal selection (red), and for a stricter one (green).

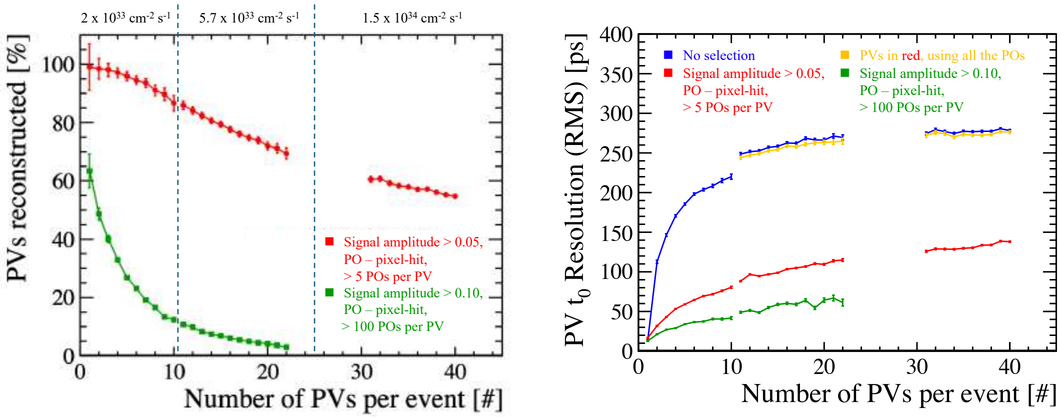

Figure 3.29: Fraction of PVs reconstructed (left) and PV t_0 resolution (right) as a function of the number of PVs per event extended to pile-up 40. In order to make such extension, two additional simulations samples were produced at the luminosities of 5.7×10^{33} and 1.5×10^{34} $\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. Results are shown without PO selection (blue), with the optimal selection (red), with the stricter selection (green), and for the subset of PVs reconstructable with the optimal selection but re-evaluated using all POs (orange).

Figure 3.29 extends the study to higher event multiplicities, up to 40 PVs per event, corresponding to the average pile-up expected for LHC Run 5. Two additional samples were produced with luminosities of 5.7×10^{33} and $1.5 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$, using the same detector geometry described in Section3.2.1. The latter corresponds to the baseline luminosity for LHC Run 5. As expected, the fraction of reconstructed PVs decreases with PV multiplicity, while the PV t_0 resolution worsens. With the stricter selection the reconstruction efficiency drops below 10 % already at 15 PVs and vanishes by 25 PVs. With the optimal selection, instead, more than 50 % of PVs can still be reconstructed at LHC Run 5 luminosity with a PV t_0 resolution better than 150 ps RMS. The orange curve in Figure 3.29 (left) shows the PV t_0 resolution for the subset of PVs reconstructable with the optimal selection, but re-evaluated using all POs instead. In this case the resolution degrades to nearly the no-selection level. This demonstrates that the optimal selection enhances PO purity for these PVs, rather than simply selecting the intrinsically clean ones. In addition, these results at high luminosity are obtained with the Run3 RICH detector geometry. The upgraded geometry, finer segmentation, and enhanced photodetector timing foreseen for the RICH Upgrade II detector are expected to further improve the performance of the PV t_0 estimation. However, the reduced fraction of PVs reconstructed represents a limitation for these PO selections in high multiplicity events. During LHC Run 5, the PV time will be provided by the tracking system, while the RICH estimate can serve as a complementary input.

It is important to remark that the present study is based on minimum-bias events, without targeting specific physics interactions. The PV t_0 resolution reported here therefore represents a general performance across all event types. As discussed in the introduction, however, LHCb is optimised for the study of beauty hadron decays. A natural continuation of this work will be to investigate the PV t_0 resolution in physics-specific samples, such as selected B-decay channels, in order to assess whether the event topology influences the performance of the algorithm.

3.4 Implications for real experiment application

The PV t_0 estimation algorithm will be integrated into the RICH reconstruction sequence. The algorithm is applied on a selected subset of photon objects (POs). The PV t_0 estimate is then used for a software time gate based on the predicted time of arrival of the photons and applied to the original set of POs. Finally, a likelihood maximisation is performed on the final set of POs to compute the PID.

As demonstrated, the PO selection plays a key role in the PV t_0 estimate resolution. The two PO selection criteria identified in Section 3.3 serve as fast indicators of PO quality for the PV t_0 estimation, and improve significantly the PV t_0 resolution. A possible strategy to further improve the selection is using the information from an initial likelihood maximisation. Performing a first likelihood maximisation allows the identification of the most probable mass hypothesis for each

PO. This enables a more accurate signal amplitude cut, tailored to the most probable particle mass rather than assuming all tracks as pions. Moreover, knowledge of the mass hypothesis allows a more precise prediction of the photon hit time, increasing the effectiveness of the software time gate in suppressing background POs. To implement this strategy, the likelihood maximisation must be run twice: once to extract the optimal set of POs, and a second time, after the software time gate, for final PID. However, given the limited CPU resources available in the LHCb HLT, both approaches requires careful evaluation. A pre-likelihood PO selection based on PO properties increases moderately the CPU load, while running the likelihood maximisation twice has a much larger impact on the HLT resources. Part of this overhead can be recovered through the software time gate, which reduces the number of POs entering the final likelihood step and therefore lowers the CPU load. Quantitative studies will be carried out to assess the computational cost of the two approaches and to identify the optimal balance between PV t_0 resolution and CPU resources.

Another crucial aspect of this study is that the uncertainty associated to the PV t_0 estimate propagates to the predicted time of arrival of the photons. An addition width must be included to the software time gate to account for this uncertainty. This width can be defined in two ways. The first option is to use a global value, equal for every PV and derived from the overall resolution of the PV t_0 algorithm. An alternative strategy is to assign an uncertainty value to each PV t_0 estimate, based on the event topology, the applied PO selection criteria, and the number of POs used in the PV t_0 estimate. This approach requires storing additional information for every PV, and increases the CPU resources needed for the PV t_0 algorithm. However, a per-PV uncertainty on the PV t_0 allows for a tailored time gate, which improves its efficiency and reduces the CPU required for the subsequent likelihood maximisation. Both strategies will be investigated within the RICH reconstruction framework, with particular attention to their impact on the detector PID performance and CPU resources.

An additional source of uncertainty may arise from the association between tracks and primary vertices. In the present study, this association is performed using Monte Carlo truth information, which provides an unambiguous matching. In the real experiment, however, the association relies on the tracking reconstruction, which may introduce mis-associations depending on the event topology and tracking performance. These mis-associations can degrade the accuracy of the PV t_0 estimate. They are not included in the present study, as the focus is on assessing the intrinsic capability of the RICH detector to estimate the PV t_0. The impact of tracking performance will be evaluated and, if necessary, accounted for in the software time gate.

During LHC Run 5, the expected pile-up in LHCb will reach about 40 PV per event, leading to a substantial increase in detector occupancy. As shown in Section3.3.3, under these conditions, spatial information alone cannot provide PV t_0 estimates for all PVs with sufficient resolution. In the LHCb Upgrade II baseline, the PV t_0 will be provided by the tracking system. Nevertheless, the experience gained by integrating timing information into the RICH reconstruction already in LHC

Run 4 will be of great value. It will allow the optimisation of the software time gate, the validation of PO-based selection strategies, and the development of algorithms that can exploit RICH timing as an additional source of information for the event reconstruction during LHC Run 5.

3.5 Chapter conclusions

The method presented in this chapter for estimating the PV t_0 during LHC Run 4 represents a first-time measurement for the LHCb experiment. The method relies solely on the RICH reconstruction, exploiting the photon hit time information that will be available in the LHC Run 4 RICH detectors. The PV t_0 is estimated using a selected subset of reconstructed photon objects (POs) associated with each PV. The selection is based on PO properties, among which the signal amplitude and the unique association to a pixel hit proved most effective. Using these criteria, the PV t_0 can be reconstructed with a precision of $85 \pm 3\,\mathrm{ps}$ (FWHM) and $66 \pm 2\,\mathrm{ps}$ (RMS) for 94% of the PVs in a sample of ten thousand bunch-crossing events. The number of selected POs was identified as the main limiting factor to the final PV t_0 resolution. For a smaller subset of PVs (about 25%), the PV t_0 uncertainty reaches 34 ps (RMS), consistent with the intrinsic resolution obtained from Monte Carlo truth information. Studies performed at different instantaneous luminosities demonstrated an approximately linear dependence of both the PV t_0 resolution and reconstruction efficiency on the event pile-up.

The presented method is a crucial step towards the use of timing information in the RICH reconstruction during LHC Run 4. The PV t_0 is a key ingredient for applying a software time gate on the photon hit times, which suppresses combinatorial background and improves the detector PID performance. In addition, the PV t_0 information may also benefit other sub-detectors, such as the LHCb tracking system.

Chapter 4

Flex-PCB prototype for operating SiPM arrays at cryogenic temperatures

Silicon Photomultipliers (SiPMs) are a promising photodetector candidate for the LHCb RICH Upgrade II detector, as discussed in Section 1.4.2. The increased luminosity foreseen for LHC Run 5 implies higher event rate and particle multiplicity, resulting in greater hit occupancy for the RICH detectors. To cope with these challenges, Upgrade II RICH detectors will require time information, improved rate capabilities and increased granularity. SiPMs arrays meet these requirements due to their excellent single-photon detection efficiency, time resolution and fine granularity. However, their high dark-count rate (DCR), which increases after irradiation, raises concerns about their operation in the detector environment. Operation at cryogenic temperatures (~100 K) can significantly mitigate the SiPM DCR, as discussed in Section 1.4.2. Building a cryogenic system for large photodetector areas presents significant challenges. One of these consists of transmitting fast analogue signal from the SiPMs at cryogenic temperatures to the readout electronics at room-temperature, maintaining signal integrity and limiting heat transmission. A possible solution is to use a fully flexible Printed Circuit Board (flex-PCB) [108, 109] with high-density, centimetre-long traces from the SiPMs to the readout electronics. This approach raises concerns about the impact of such traces on signal integrity and timing performance.

To address this, a flex-PCB prototype has been produced for the readout of a 64-channel SiPM array. This chapter presents its design, characterisation and testing. Section 4.1 illustrates the design of the flex-PCB and the rationale behind the main choices. Section 4.2 presents the signal integrity measurements, while Section 4.3 describes the full characterisation of the flex-PCB SiPM array under single-photon illumination, with particular focus on time resolution. Section 4.4 reports the performance of the flex-PCB SiPM array under beam conditions, obtained within the LHCb RICH beam test campaign at the CERN SPS North Area facility. Finally, Section 4.6 discusses prospects for future iterations of the flex-PCB.

Before diving into the chapter it is important to mention that, while the design of the flex-PCB targets operation in a cryogenic system, all measurements reported in this chapter were performed at room temperature. These measurements were essential to demonstrate the correct operation of the flex-PCB prototype, assessing the SiPM array signal integrity and timing performance, and providing a strong reference for future measurements at cryogenic temperatures. These measurements are planned for 2026, integrating the flex-PCB into a dedicated cryostat demonstrator currently under development at CERN, as described in Section 4.1.3.

4.1 Flex-PCB prototype design

This section describes the design of a flex-PCB prototype developed to transmit fast analogue signals from an SiPM array at cryogenic temperature to the readout electronics at room temperature through high-density, centimetre-long traces. The flex-PCB solution was adopted to meet several requirements: preserving the integrity of the analogue signals, limiting heat transfer from the readout electronics to the SiPM array, and ensuring efficient thermal coupling between the SiPMs and the cryogenic cooling system. The flexibility of flex-PCBs allows operation in compact or confined environments, while the use of centimetre-long traces minimises thermal exchange between the two temperature regions and avoids mechanical stress in some directions as a result of CTE mismatches. The SiPM array is directly soldered onto the flex-PCB, removing the need for intermediate connectors that could degrade signal integrity. All aspects of the design — from trace geometry to layer stack-up and mechanical integration — were optimised to balance electrical, thermal, and mechanical performance. These features are introduced and discussed in the following sections.

Figure 4.1 shows a top view of the flex-PCB prototype, with the main features and components annotated. A 64-channels SiPM array is placed on the left side of the flex-PCB. The array is negatively biased via an on-board biasing circuit, which includes per-channel decoupling RC filters. The SiPM analogue signals are routed through traces ranging from 14 cm to 16 cm, depending on the channel position, to the connectors for the readout electronics. Additional features, such as test pulse injection circuits, a temperature monitoring sensor, and mechanical mounting holes, are included in the flex-PCB to facilitate its testing and integration into different setups.

The flex-PCB stack-up and layout are described in Section 4.1.1. Section 4.1.2 describes the SiPM array and its biasing scheme. Finally, Section 4.1.3 describes the integration of the flex-PCB into a cryostat demonstrator, currently under development at CERN.

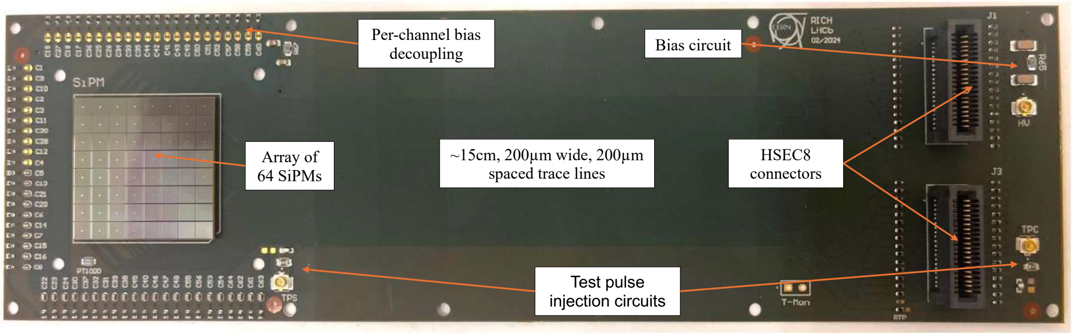

Figure 4.1: Top-view of the flex-PCB prototype with the main features annotated.

4.1.1 Layer stack-up and layout

This section presents the stack-up and layout of the flex-PCB prototype. The flex-PCB prototype has three copper layers. The top layer is mainly used for the SiPM array analogue signals traces, while the bottom layer contains the copper planes for the SiPM bias distribution. An inner layer, consisting of a ground plane, separates the two active layers and provides isolation between the signal traces and the biasing circuit. The copper thickness of the top and bottom layers is 35 μm, while for the inner layer it is 17 μm. These values correspond to standard copper weights, commonly used in flex-PCB fabrication. They were chosen to ensure the compatibility with standard fabrication processes and to simplify the production. The total copper thickness was however kept sufficiently low to preserve the flexibility of the flex-PCB prototype. Layers of Epoxy glue [110] and Kapton polyimide [111] are used as dielectric between the copper layers. This combination of materials provides the mechanical flexibility typical of flex-PCBs. The thickness of the Epoxy glue is 50 μm, while for the Kapton is 100 μm. Logically, lower thickness of the dielectric corresponds to higher flexibility. However, the dielectric thickness plays also a crucial role in determining the impedance of the SiPM analogue signal traces. Therefore, these thickness values were the result of a trade-off between the board flexibility and the signal traces impedance. Given that the trace impedance depends also on the trace width, this will be discussed later in the section. The stack-up of the flex-PCB is summarised in Table 4.1. The total thickness of the flex-PCB is 450 μm, resulting in a bending radius of 20 mm.

The layout of the three copper layers is shown in Figure 4.2: in red the top layer, in gold the inner layer and in blue the bottom layer. All components are placed on the top layer of the flex-PCB. Restricting the component placement to a single side of the board simplifies the assembly process and reduces the likelihood of errors. The SiPM array is placed on the left side of the top layer.

Table 4.1: Layer stack-up of the flex-PCB.

Layer	**Thickness**	**Material**	**Target**
Soldermask top	30 μm	–	–
Top layer	35 μm	Copper	Signal mainly
Dielectric	100 μm	Polyimide	–
Glue	50 μm	Epoxy glue	–
Inner layer	17 μm	Copper	GND plane
Glue	50 μm	Epoxy glue	–
Dielectric	100 μm	Polyimide	–
Bottom layer	35 μm	Copper	Bias mainly
Soldermask bot.	30 μm	–	–

Two Samtec HSEC8 connectors [112] are located on the right side, interfacing with the readout electronics developed for the LHCb RICH beam tests, described in Section 1.4.3.

Analogue signal traces from the SiPM array to the HSEC8 connectors are routed primarily on the top layer. Half of them transition to the bottom layer through vias for a few centimetres to reach the farthest connector pins. Each signal trace includes a 0 Ω series resistor, allowing individual channels to be disconnected from the readout electronics in the event of malfunction or excessive noise. Considering the minimum bending radius of 20 mm, the trace length of 15 cm allows the board to be safely bent by 180° without mechanical stress on the copper layers. A distance of 200 μm between signal traces and a trace width of 200 μm were initially chosen to reproduces the constraints of a high-density detector environment. As discussed in Section 4.5, future versions of the flex-PCB are expected to feature even higher channel density, potentially requiring narrower spacing. This prototype therefore serves as a first step toward that goal, allowing the investigation of signal integrity in long analogue traces under realistic detector-layout conditions. The signal trace width was selected together with the dielectric thickness of the Kapton–epoxy stack to achieve a characteristic impedance close to 50 Ω, which matches standard input impedance of readout electronics for photon detectors. Impedance matching between the traces and the readout electronics minimises signal reflections and preserves the shape and timing of the analogue pulses. The dielectric thickness was therefore defined as 50 μm for the epoxy adhesive and 100 μm for the Kapton layer. The trace impedance (Z_0) can be calculated using the Hammerstad–Jensen microstrip model [113]:

$$Z_0 = \frac{87}{\sqrt{\varepsilon_{eff} + 1.41}} \ln\left(\frac{5.98h}{0.8w + t}\right), \tag{4.1}$$

where $\varepsilon_{eff} \sim 2.8$ is the effective dielectric constant of Kapton, Epoxy glue and solder-mask, $h = 150\,\mu$m is dielectric thickness between the trace and the ground plane, $w = 200\,\mu$m is trace width and $t = 35\,\mu$m the trace thickness. The resulting $\sim 58\,\Omega$ is sufficiently close to targeted 50 Ω,

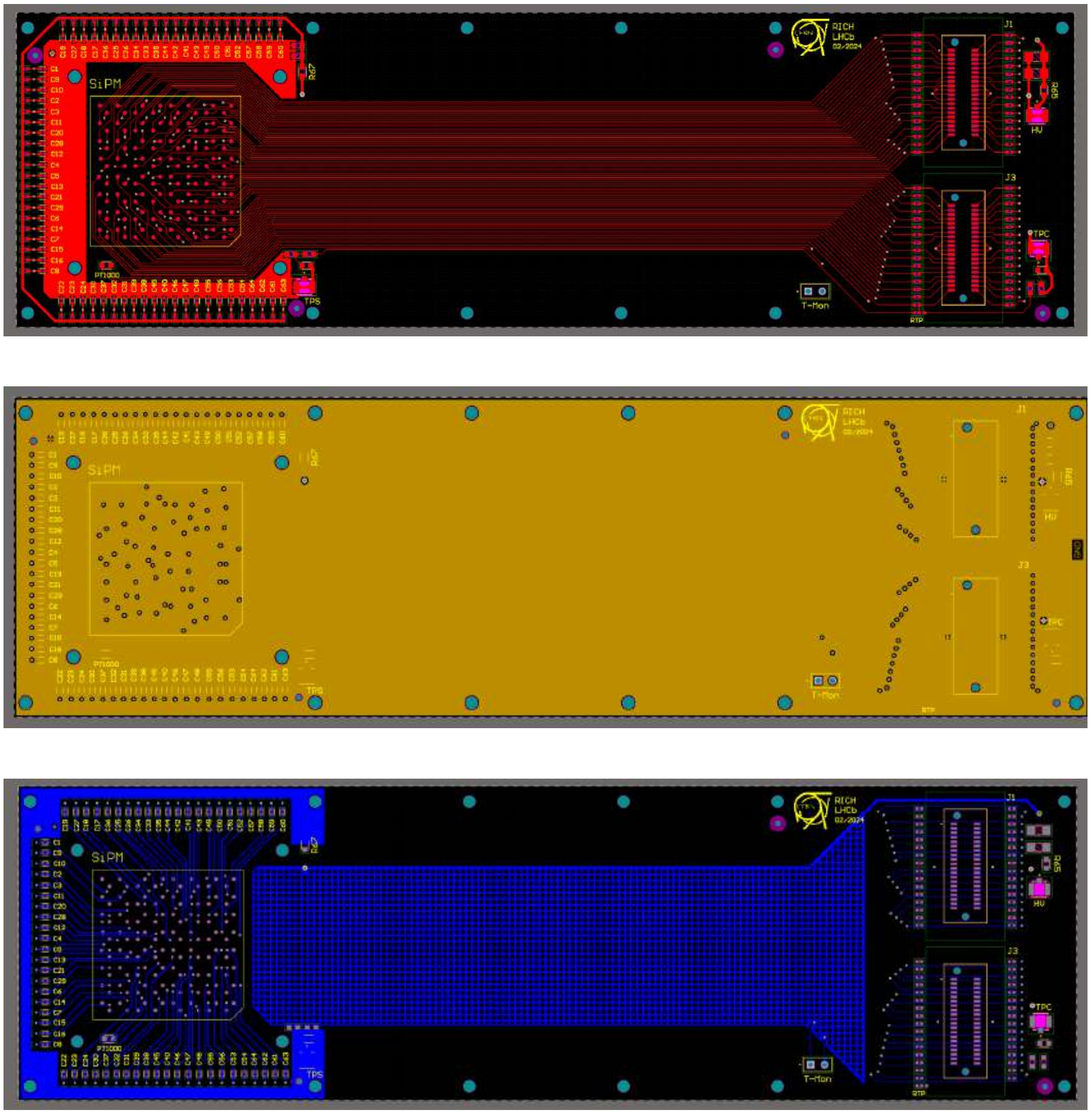

Figure 4.2: Layout of the Flex-PCB with top layer, ground (middle) layer and bottom layer, in order from top to bottom.

confirming that the chosen geometry provides adequate impedance matching for the readout electronics. Additionally, this configuration lies within typical manufacturing tolerances for flex-PCBs and could be reliably produced at the CERN PCB manufacturing workshop without specialised processes.

Additional features and components are placed around the SiPM array for biasing, monitoring, and mechanical purposes. Each of the 64 SiPMs is connected to an individual RC filter, as part of the biasing circuit. These filters provide per-channel decoupling of the bias voltage. Their design and function are discussed in more detail in Section 4.1.2. A ground plane was introduced around the SiPM array to locally increase the mechanical stiffness of the flex-PCB and to improve the structural stability of the array during assembly and operation. The top layer also carries the transmission lines for two test-pulse injection circuits and the routing for a PT1000 temperature sensor[1]. The test-pulse circuits are used to measure the signal integrity of the analogue traces and are described in Section 4.2.

The inner copper layer consists of a continuous ground plane, as shown in Figure 4.2 (middle). This layer serves multiple purposes: it provides electromagnetic shielding between the analogue signal traces on the top layer and the biasing lines on the bottom layer and it ensures controlled impedance and a low-inductance return path for the analogue signals, essential to minimise high-frequency noise and preserve signal integrity. In addition, the continuous ground plane improves the mechanical stability of the flex-PCB, while its impact on flexibility remains marginal.

The bottom layer is mainly used for the SiPM bias distribution. It consists of a copper grid covering almost the entire layer, as shown in Figure 4.2 (bottom). This layout balances the copper between the top and bottom layers, preventing the warping of the flex-PCB during assembly and thermal cycles. For the same reason, a solid copper plane has been added on the SiPM-side in the bottom layer to compensate for the RC filter network on the top layer. Additionally, this plane increases the stiffness of the flex-PCB on the SiPM-side, reducing the risk of mechanical damage to the SiPM connections during handling or operation.

4.1.2 SiPM array and biasing scheme

This section describes the SiPM array mounted on the flex-PCB prototype and its biasing circuit. The SiPM array mounted on the flex-PCB prototype is a 64-channel Hamamatsu S13361-3050NE-08 SiPM array with $3 \times 3\,\mathrm{mm}^2$ channel size [114]. A close up of the top side of the flex-PCB is shown in Figure 4.3. Each SiPM consists of 3584 microcells (SPADs) of $50 \times 50\,\mu\mathrm{m}^2$ with a fill factor of 74 %. The breakdown voltage is approximately 51 V, with recommended operation at 3 V overvoltage. At this overvoltage, the typical gain is 1.7×10^6 Me, with a peak photon detection efficiency (PDE) of 40 % at 450 nm, in the blue–green spectral region. The nominal DCR at this biasing condition

---

[1] Platinum resistance temperature sensor (1000 Ω at 0 °C), supplied by Reichelt: `https://www.reichelt.com/`

is 0.5 MHz per channel, but it can vary significantly with temperature. Given the SiPM size of $3 \times 3\,\text{mm}^2$, the equivalent capacitance for this SiPM is approximately 320 pF. It is demonstrated that the timing performance of SiPMs depends on the equivalent capacitance, which increases with channel area [115]. Higher equivalent capacitance broadens the output pulse and degrades the time resolution. However, the choice of $3 \times 3\,\text{mm}^2$ channel was driven by two factors. First, the area matches that of the MAPMTs used in the current RICH detector, ensuring compatibility with the existing mechanical design and channel density, and enabling a direct comparison of detection efficiency and timing performance. Second, the same SiPM model had already been used in previous LHCb RICH beam test campaigns, mounted on a rigid PCB with short analogue traces [91]. This would allow the comparison of the same SiPM array performance between short and long analogue traces to the readout electronics. As described in Section 4.5, future iterations of the flex-PCB will target higher granularity with reduced SiPM dimensions, in line with the RICH Upgrade II baseline specifications (see Section 1.4). This will also result in improved intrinsic timing performance for the next flex-PCB version.

A negative-polarity biasing scheme was used for the SiPM array, as illustrated in Figure 4.4. The bias input connector is located on the HSEC8 connector side of the flex-PCB. A common RC filter at the input suppresses high-frequency noise from the voltage power supply. The bias voltage is distributed through the copper grid in the bottom layer. An additional per-channel RC network is placed around the SiPM array, as shown in Figure 4.3. These filters were added to decouple the bias line of each SiPM and improve signal quality. To assess their effectiveness, only

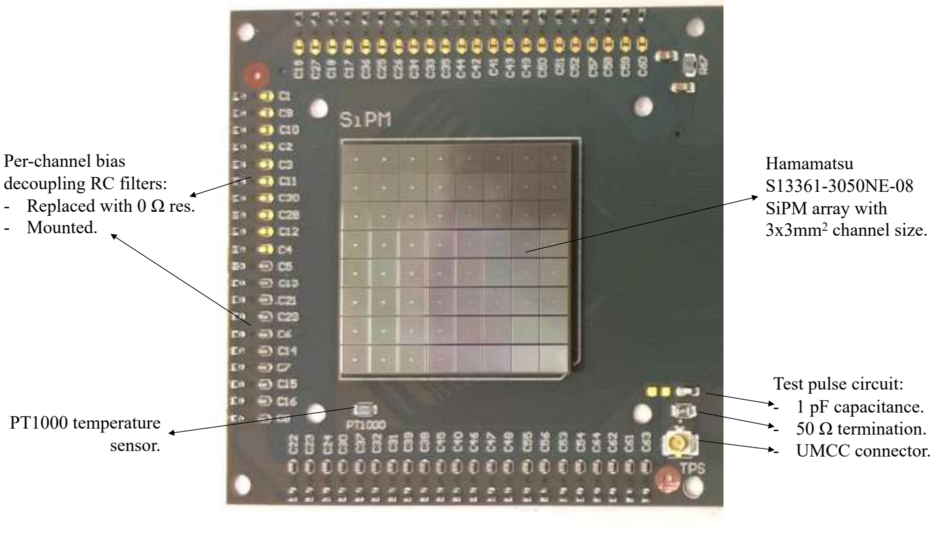

Figure 4.3: Zoom-in on the SiPM side of the flex-PCB prototype.

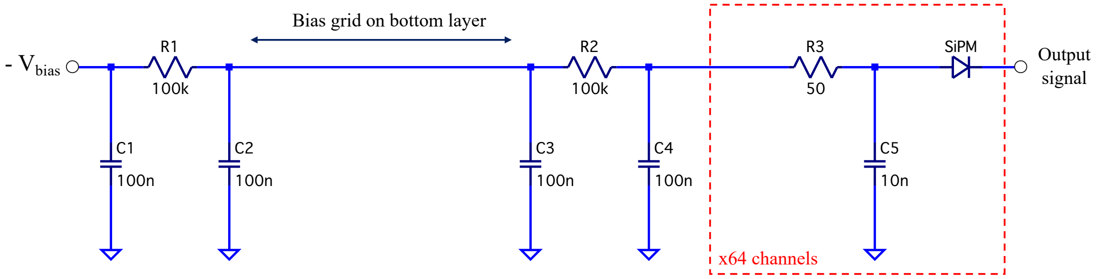

Figure 4.4: Biasing scheme for the 64-channels SiPM array, including per-channel RC filters for bias decoupling.

half were populated in the prototype, while the others were replaced with 0 Ω resistors. Because the filters occupy a significant area on the layout, their impact must be evaluated before adopting a high-density design in future flex-PCB versions. The performance comparison between the two configurations is discussed in Section 4.3.

4.1.3 Integration into a cryostat demonstrator

This section describes the integration of the flex-PCB prototype into a cryostat demonstrator developed by the CERN RICH group to operate SiPMs at cryogenic temperatures. Cooling SiPMs to cryogenic temperatures effectively mitigates their high DCR, which increases significantly after irradiation. However, implementing a cryogenic system for large photodetector areas introduces several challenges. These include efficient thermal coupling between the SiPMs and a cold volume, maintaining vacuum insulation, operating multi-channel readout electronics in vacuum, and transmitting analogue signals over centimetre-scale distances while preserving signal integrity. The modular cryostat demonstrator developed at CERN is designed to study some of these challenges. This demonstrator provides a compact and controlled test environment to gain valuable insights into the scalability of the system to large photodetector areas.

A cross-sectional view of the cryostat is shown in Figure 4.5. The system consists of a vacuum enclosure that houses a copper cold volume, mechanical supports for the readout electronics and several service ports. The cold volume is thermally anchored to a liquid-nitrogen cooling circuit and reaches an operational temperature of ~100 K. The SiPM-side of the flex-PCB is mounted onto the cold volume using the mechanical holes along the board edge, as shown in Figure 4.6. An intermediate copper mount is placed between the flex-PCB and the main cold volume. This mount includes the mechanical mounting holes for the flex-PCB and can accommodate future versions of the board. As described in Section 4.1.1, the overall thickness of the flex-PCB was minimised to

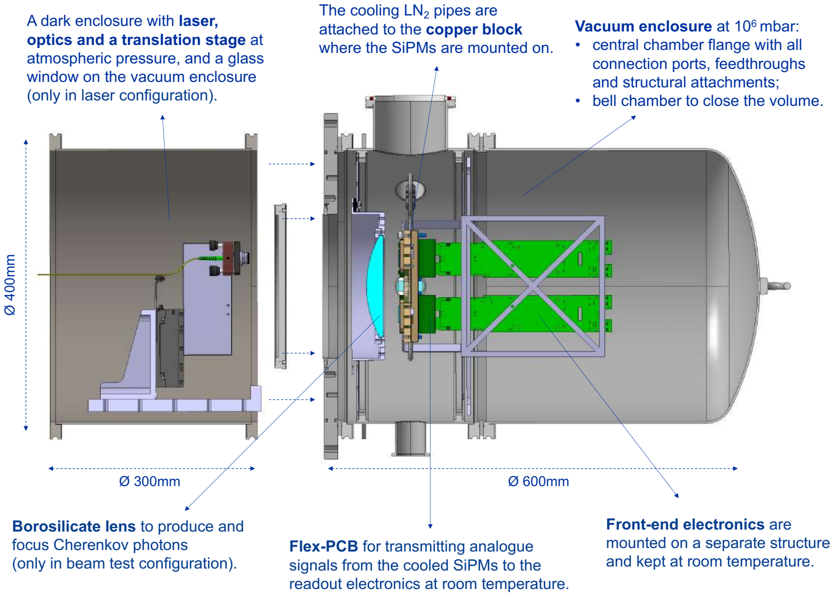

Figure 4.5: Cross-sectional view of the cryostat demonstrator developed by the CERN RICH group for operating SiPMs at cryogenic temperatures.

improve thermal contact. In addition, four holes near the SiPM array accommodate a clamping frame that increases the contact pressure between the board and the cold volume, enhancing heat exchange and mechanical stability.

On the other side of the cold volume, a mechanical structure supports the readout electronics. An independent water-cooling system maintains the readout electronics at room temperature. Vacuum inside the enclosure acts as a thermal insulator, ensuring decoupling between the two temperature zones. The 15 cm-long analogue traces in the flex-PCB provide sufficient distance to minimise heat transfer from the readout electronics to the SiPM array. At the same time, the board flexibility allows a 180° bend. This improves the cryostat compactness by reducing the space required for signal routing.

The enclosure also includes ports for the liquid-nitrogen pipes, vacuum pumping, electronics power lines, optical fibre feedthroughs and sensor control. The system allows two test configurations: one for beam test and one for pulsed-laser illumination. In the beam-test configuration, the front panel is made of aluminium and a borosilicate lens is placed in front of the SiPM arrays plane. As in the current LHCb RICH beam test setup [91], the lens will act as Cherenkov radiator and will focus the emitted photons onto the photodetector plane. In the laser-test configuration, a glass window replaces the front panel, and a dark enclosure is mounted externally, as shown in

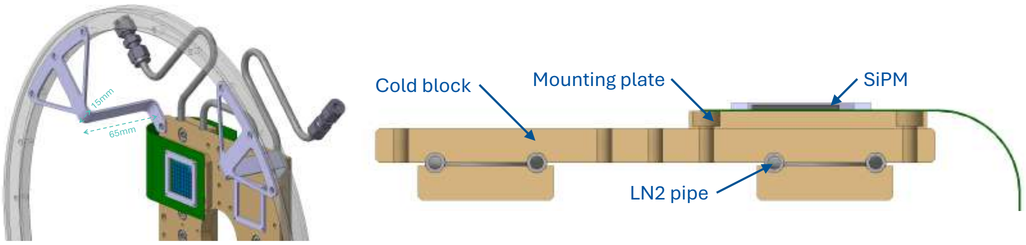

Figure 4.6: Close-up of the flex-PCB mounted on the cold volume of the cryostat demonstrator.

Figure 4.5. A picosecond-pulsed laser delivers light via an single-mode optical fibre mounted on translation stages, allowing well-controlled, position-resolved illumination of the SiPM array. At the time of writing, the assembly and commissioning of the cryostat demonstrator are being completed by other members of the CERN RICH group. First operation of the cryostat demonstrator is expected in 2026.

4.2 Signal integrity measurements

The flex-PCB includes two identical test-pulse injection circuits. One circuit is located near the SiPM array (see Figure 4.3) and routed to the HSEC8 connector through a 15 cm-long trace. The second circuit is placed on the readout side of the flex-PCB, with a shorter 2.5 cm trace to the connector. Both traces have the same 200 μm width as the analogue traces used for the SiPM readout. Each circuit consists of a Hirose UMCC connector [116] followed by a 1 pF capacitor. The two circuits were used to measure the signal integrity of the flex-PCB by injecting identical pulses and comparing the outputs from the long and short traces. Figure 4.7 shows the setup used for these measurements. A voltage step of 160 mV with a 70 ps (20–80 %) rise time is produced with a pulse generator [2] and injected into the circuit through a 10 cm SMA-to-UMCC coaxial cable with 9 GHz bandwidth. The voltage step is terminated with a 50 Ω resistor in the flex-PCB. Passing through the capacitor, the voltage step reproduces the main characteristics of a single-photon signal: sub-nanosecond rise time, amplitude of a few tens of millivolts, and integrated charge of about 1 Me. A copy of the voltage step is sent to the oscilloscope as time reference. The output signals are measured across a 56 Ω termination, matching the 58 Ω impedance of the flex-PCB

---

[2] Active Technology Pulse RIDER PG-1072 pulse generator (`https://www.activetechnologies.it/products/signal-generators/pulse-generators/pg-1000/`).

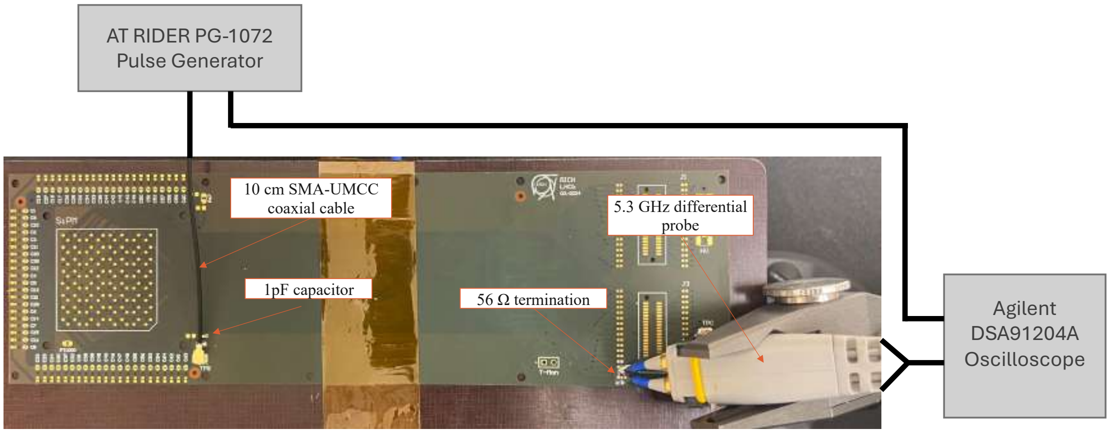

Figure 4.7: Setup for the signal integrity measurements using the test-pulse circuit on the flex-PCB. The configuration shown here is for a test-pulse injected into the 15 cm signal trace.

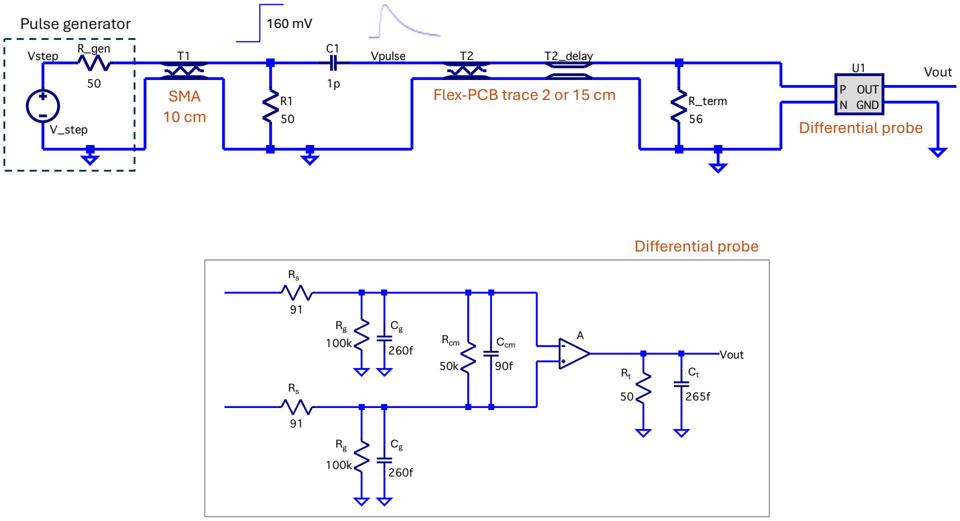

Figure 4.8: Equivalent electronics circuit of the test-pulse setup implemented in the LTSpice simulator.

analogue traces. The waveforms are acquired with an oscilloscope[3] using a differential probe[4]. The differential probe has a bandwidth (BW) of 5.3 GHz, compatible with an intrinsic rise time of ~66 ps (10-90 %) estimated using $t_r \approx 0.35/\mathrm{BW}$ [117]. The setup is therefore sensitive to small changes in the signal rise time induced by the transmission line.

The equivalent electronics circuit of the setup is shown in Figure 4.8. The circuit was implemented in the LTSpice[5] software to simulate the output pulse including first-order non-ideal effects. The coaxial cable and signal traces are modelled as lossy transmission lines including resistive losses and parasitic capacitance and inductance. Standard transmission-line models and manufacturer specifications were used to parametrise the lossy line of the 50 Ω coaxial cable [117, 118]: 100 Ω/km for the DC resistance, 100 pF/m for the capacitance and 200 nH/m as inductance. The signal traces on the flex-PCB were parametrised using per-unit-length resistance, inductance, conductance, and capacitance parameters (RLGC model) derived from a microstrip line approximation [117]. With the trace impedance $Z_0 = 58\,\Omega$ and the effective dielectric constant $\varepsilon_{\text{eff}} \approx 2.8$, the signal velocity in the trace is $v = c/\sqrt{\varepsilon_{\text{eff}}} \approx 1.79 \times 10^8$ m/s. From the relations

$$L = \frac{Z_0}{v}, \qquad C = \frac{1}{Z_0 v}, \tag{4.2}$$

the inductance and capacitance per-unit-length are $L \approx 325$ nH/m and $C \approx 96$ pF/m. The resistive losses were estimated from the copper surface resistance. For a good conductor at frequency f, the surface resistance is given by

$$R_s = \sqrt{\frac{\pi f \mu_0}{\sigma}}, \tag{4.3}$$

where $\mu_0 = 4\pi \times 10^{-7}$ H/m is the vacuum permeability and $\sigma = 5.8 \times 10^7$ S/m is the electrical conductivity of copper. At $f = 5$ GHz this gives $R_s \approx 18\,\mathrm{m}\Omega$. For a microstrip trace width of 200 μm, the resistance per-unit-length can be approximated as

$$R \approx \frac{R_s}{w_{\text{eff}}}, \tag{4.4}$$

where w_{eff} is the effective width of the trace. Assuming that the current is confined to the bottom surface of the trace (the closest to the ground plane) and that the contributions from edges and sidewalls are negligible, $w_{\text{eff}} \approx w = 200\,\mu$m. This yields a resistance per-unit-length of $R \approx 90\,\Omega$/m, which was therefore used as the per-unit-length series resistance in the RLGC model of the flex-PCB trace. The dielectric loss was considered negligible for the centimetre-scale traces and was omitted

---

[3] Keysight Agilent DSA91204A oscilloscope (`https://www.keysight.com/us/en/product/DSA91204A/infiniium-high-performance-oscilloscope-12ghz.html`).

[4] Agilent 1169A probe with an E2675A differential browser (`https://www.keysight.com/us/en/assets/9018-05040/quick-start-guides/9018-05040.pdf.`).

[5] LTSpice: `https://www.analog.com/en/resources/design-tools-and-calculators/ltspice-simulator.html`

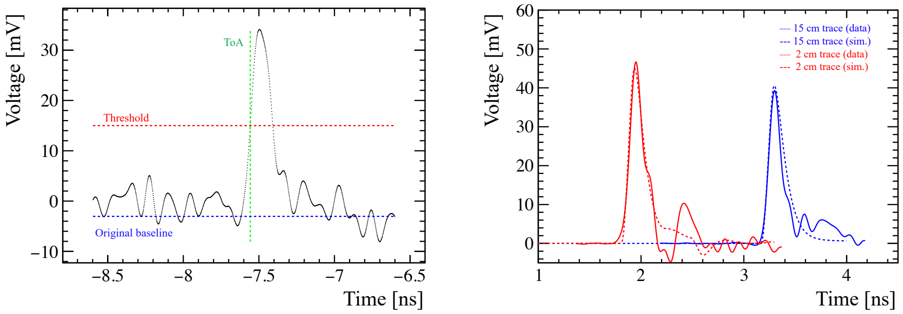

Figure 4.9: (Left) Example of a waveform recorded using the test-pulse circuit with the 15 cm signal trace. The waveform is corrected for the baseline and a threshold is applied at 50 % of the peak to extract the timestamp. (Right) Average waveforms from data and simulated output pulses for both trace lengths.

from the simulation model. An equivalent circuit for the differential probe was also implemented, as shown in Figure 4.8, using parameters from the probe datasheet [119].

A sample of 1000 waveforms was recorded for each trace length. Each waveform was baseline-corrected and a threshold was applied at a fixed fraction of the peak to extract the time-of-arrival (ToA), as shown in Figure 4.9 (left). Figure 4.9 (right) compares the average output waveforms from data with the simulated output pulses for both trace lengths, showing excellent agreement. The rise time is accurately reproduced by the simulation, while small ripples in the waveform tails are due to higher-order effects not included in the model. The measured rise time is about 70 ps (20–80 %) in both cases, as shown in Table 4.2. This value matches the injected signal rise time, indicating no bandwidth limitation or signal degradation along the traces. The only noticeable difference between the two trace lengths is the higher attenuation observed for the 15 cm trace, consistent with the resistive losses predicted by the simulation. According to the simulation model, longer trace lengths or smaller trace widths lead to stronger attenuation. This behaviour must be carefully considered in future flex-PCB iterations, especially if thinner traces are adopted to achieve higher channel density. In that case, the increased resistive losses could lead to additional signal attenuation, which should be compensated for in the design.

Figure 4.10 shows the ToA distributions for the two trace lengths. A Gaussian fit is applied to each distribution and the ToA jitter is extracted as the standard deviation of the fit. The resulting jitters are 10.7 ± 0.5 ps for the 15 cm trace and 10.8 ± 0.5 ps for the 2 cm trace. The two values agree within uncertainties, demonstrating that the longer trace does not introduce additional timing jitter relative to the shorter one. Since the oscilloscope was triggered on the reference signal from the pulse generator, these values include the intrinsic jitter of the readout system. The latter was measured independently by probing the injected pulse directly at the injection point, yielding 8.8 ± 0.5 ps.

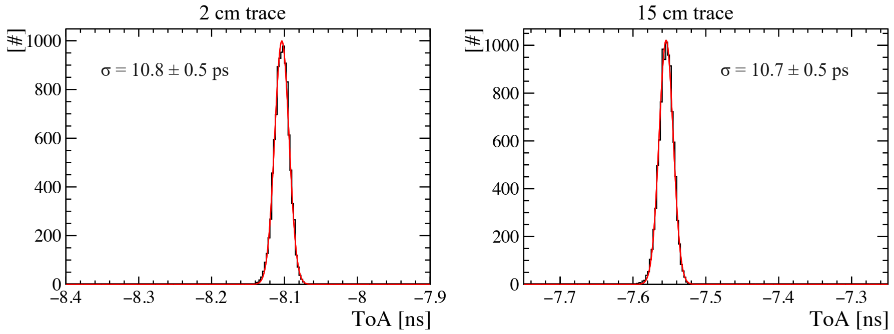

Figure 4.10: ToA distributions of the waveforms recorded with the two test-pulse circuits. The jitter is extracted as the σ of a Gaussian fit applied to the distributions.

Table 4.2: Rise time and jitter for the two trace lengths.

	Rise time	Jitter
2 cm trace	71 ± 10 ps	10.8 ± 0.5 ps
15 cm trace	69 ± 10 ps	10.7 ± 0.5 ps

Subtracting this contribution in quadrature gives an intrinsic trace jitter of about 6 ps, negligible compared to the >50 ps single-photon time resolution (SPTR) of SiPMs. These results confirm that the 15 cm-long traces preserve signal integrity and can reliably transmit SiPM analogue signals to the readout electronics.

4.3 Single-photon characterisation of the SiPM array

This section describes the characterisation of the SiPM array on the flex-PCB under controlled single-photon illumination. Section 4.3.1 presents the experimental setup. Section 4.3.2 describes the procedure used to determine the optimal single-photon working point. Section 4.3.3 outlines the analysis method used to extract the SPTR. Section 4.3.4 summarises the obtained SPTR results.

4.3.1 Experimental setup for single-photon pulsed-laser illumination

The setup described in this section was used to characterise the single-photon response of the flex-PCB SiPM array. The setup is shown in Figure 4.11, while Figure 4.12 presents the schematic of all connections. The SiPM array is illuminated with a picosecond-pulsed laser[6]. The laser delivers

---

[6] PicoQuant TAIKO pulsed laser (`https://www.picoquant.com/products/category/picosecond-pulsed-driver/taiko-pdl-m1#specification.`).

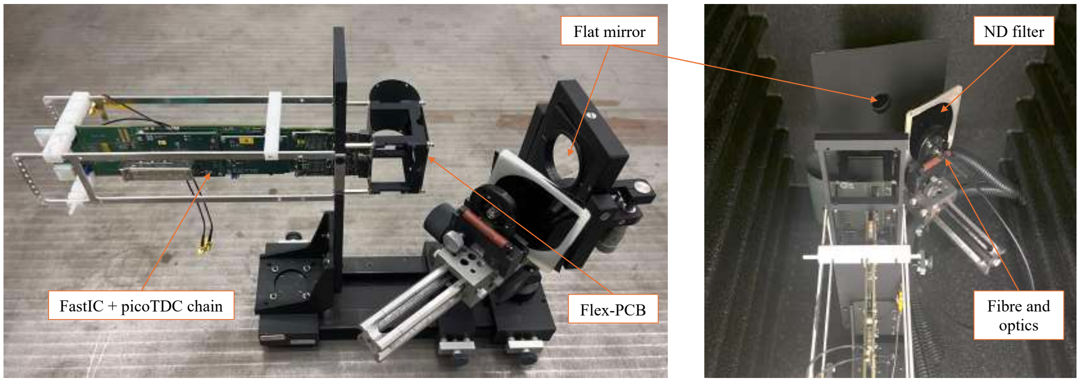

Figure 4.11: Experimental setup for picosecond-pulsed laser illumination of the flex-PCB SiPM array. (Left) The setup lies on a bench to better display the components. (Right) the setup is placed inside the light-tight box where the measurements were conducted.

50 ps (FWHM) pulses, corresponding to a Gaussian $\sigma \sim 21$ ps, at a wavelength of 405 nm through a single-mode optical fibre. Neutral-density filters are placed between the fibre and the SiPM array to attenuate the light and ensure single-photon illumination. The light is directed onto the SiPM array using a flat mirror. The mirror is covered by a mask with a central aperture to suppress reflections from the mirror edges that could otherwise affect the timing measurement. The flex-PCB is bent by 180° and connected to the readout electronics. A mechanical structure holds the mirror, filters, fibre head, flex-PCB, and readout electronics in fixed positions.

The flex-PCB is coupled to the FastIC + PicoTDC readout electronics developed for the RICH beam tests and described in Section 1.4.3. Each SiPM in the array is connected to a FastIC channel, hereafter referred to as a channel for simplicity. The data acquisition is handled by the MuDAQ system [96], which also provides the 40 MHz clock distribution and trigger synchronisation. The trigger signal is internally generated by the MuDAQ at 20 kHz and distributed to the readout board via the optical link (VTRX+ and lpGBT chipset). In addition, the trigger signal is routed to a fan-out board to generate three identical copies. One copy is converted to TTL format and used to trigger laser pulses. The other two copies are sent to separate picoTDC input channels and used as time reference. The photon signals produced by the SiPM array are shaped and discriminated by the FastIC. For each signal, the FastIC generates a differential digital output defined by a leading and a trailing edge. This output is sent to the picoTDC, operated in *Full TDC* mode with a 12 ps bin size. The picoTDC assigns a timestamp to both leading and trailing edges and stores them within a configurable readout window. Upon receiving a readout trigger, the content of the

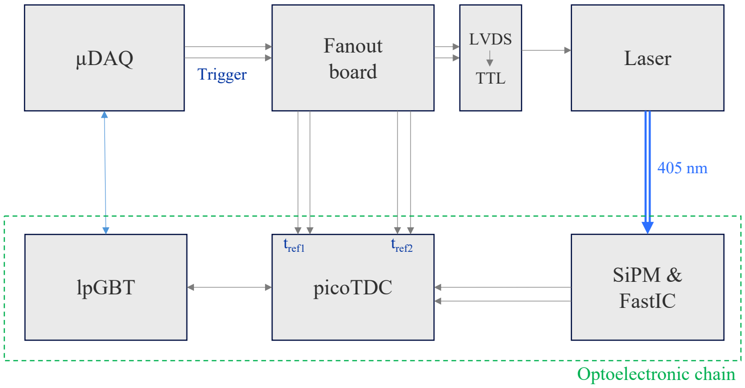

Figure 4.12: Schematic of the connections for the picosecond-pulsed laser setup.

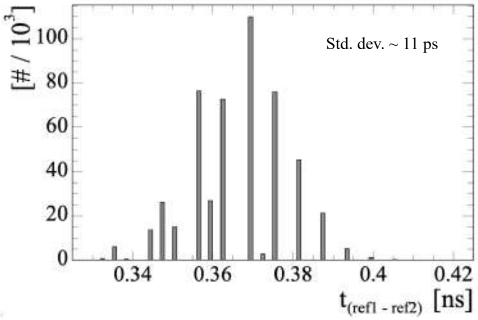

Figure 4.13: Time-difference distribution of the two reference signals measured by the picoTDC. The distribution spread of 11 ps consistent with the picoTDC bin size.

picoTDC readout window is transmitted to the MuDAQ via the lpGBT interface. For each photon hit, the leading-edge time corresponds to the photon ToA ($\mathrm{ToA_{ch}}$), while the signal time-over-threshold (ToT) is obtained from the interval between leading and trailing edges. The average ToA of the two time-reference channels, $\mathrm{ToA_{ref}} = (\mathrm{ToA_{ref1}} + \mathrm{ToA_{ref2}})/2$, is subtracted from $\mathrm{ToA_{ch}}$. The resulting time difference ($\mathrm{ToA_{ch}} - \mathrm{ToA_{ref}}$) is analysed to determine the SPTR of the SiPM channels, as described in Section 4.3.3.

The time difference between the two time reference signals is shown in Figure 4.13. The measured spread of 11 ps, consistent with the picoTDC bin size, confirms that the trigger distribution introduces no additional timing jitter. Other intrinsic contributions to the SiPM SPTR, such as jitter from the FastIC comparator and the laser, are corrected a posteriori. To obtain the best single-photon timing performance, the optimal working point has to be determined for all the SiPMs in the array. This requires tuning the SiPM bias voltage, the FastIC comparator threshold and the laser illumination. The optimisation procedure for single-photon acquisition is described in Section 4.3.2.

4.3.2 Single-photon detection working point

Single-photon timing performance of SiPMs depends on the selected working point. The bias voltage and the FastIC comparator threshold affect key SiPM parameters such as gain, DCR and time resolution. Laser intensity is an additional crucial parameter to ensure single-photon illumination of the SiPM array. A dedicated set of measurements was performed to characterise these dependencies and to define the optimal operating conditions for the best single-photon timing performance of the SiPM array.

As a first step, the SiPM array current was measured as a function of the applied bias voltage to verify the correct operation of the device and estimate the breakdown voltage. The bias voltage was supplied by a power unit[7], which also monitored the total current drawn by the SiPM array with a precision of 0.05 μA. Figure 4.14 shows the measured current as a function of the applied bias voltage, both in complete darkness and under diffuse illumination. In dark conditions, the current increases exponentially with the applied bias, as expected for SiPMs. No difference is observed between SiPMs with and without RC bias decoupling. Under diffuse light, the current rises more steeply with bias voltage, showing a sharp increase above breakdown. This behaviour indicates a breakdown voltage of V_{bd} = 51.3 V for the SiPM array. Since the measured current is the sum over all SiPMs, channels with lower breakdown voltages dominate the behaviour. Nevertheless, the estimated V_{bd} agrees with the average breakdown voltage reported by Hamamatsu and serves as a reliable reference. This measurement confirms the correct operation of the SiPM array on the flex-PCB.

As shown in Chapter 2, the threshold level affects the single-photon timing performance. Unlike MAPMTs, SiPMs produce a discrete response proportional to the number of detected photons. They are therefore not affected by time-walk, and constant-fraction discrimination (CFD) is not required. Nevertheless, noise-level variations between FastIC channels must be taken into account. A scan was performed to characterise all channels, varying the FastIC comparator threshold in steps of 0.49 μA. At each threshold step, about ten thousand triggers were sent, and the number of leading edges within an 800 ns trigger windows was recorded. With the SiPMs unbiased, leading

---

[7] CAEN DT8031 power unit (`https://www.caen.it/products/dt8031/`).

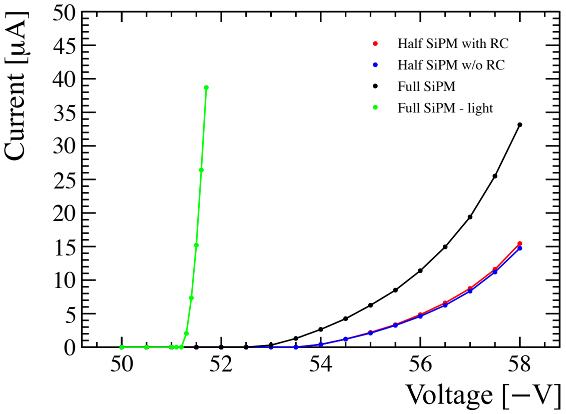

Figure 4.14: SiPM array current as a function of the bias voltage applied. The curves are obtained in two configurations: complete darkness and with diffuse light in the enclosure. In the dark configuration, separate measurements for the channels with and without RC bias decoupling are shown.

edges are produced only by electronic noise when the FastIC threshold approaches the noise pedestal. Figure 4.15 (left) shows the threshold scan output for one channel. The peak corresponds to the noise pedestal. The relatively large SiPM capacitance broadens this peak over a few threshold steps. Figure 4.15 (right) shows the pedestal position for all channels of the eight FastICs on the readout board. The spread of the pedestal values is about 10 μA, comparable to the expected SiPM single-photon signal amplitude at −54 V. Determining the noise pedestal is therefore essential to define the working point of each channel.

Repeating the threshold scan with the SiPMs biased yields the staircase response shown in Figure 4.16. In this case, leading edges originate from both electronics noise and dark-count events. The discrete nature of the SiPM response can be appreciated, with plateaux corresponding to one, two or more detected photoelectrons (pe). Because dark counts are indistinguishable from single-photon signals, the width of the first plateau corresponds to the single-photon gain. The comparator threshold can therefore be set to about 20 % of the single-photon gain, which provides the best timing performance, as observed for the MAPMTs in Chapter 2. In addition, the number of leading edges can be used to estimate the DCR using the formula:

$$\mathrm{DCR} = \frac{N_{\mathrm{edges}}}{N_{\mathrm{trig}}\, t_{\mathrm{win}}\, A_{\mathrm{SiPM}}}, \tag{4.5}$$

where N_{edges} is the total number of detected leading edges, N_{trig} is the number of triggers of one threshold-scan step, t_{win} is the trigger window width and $A_{\mathrm{SiPM}} = 9\,\mathrm{mm}^2$ is the area of an SiPM. Figure 4.17 shows the single-photon gain and DCR for all channels at a bias of −54 V. The channel-

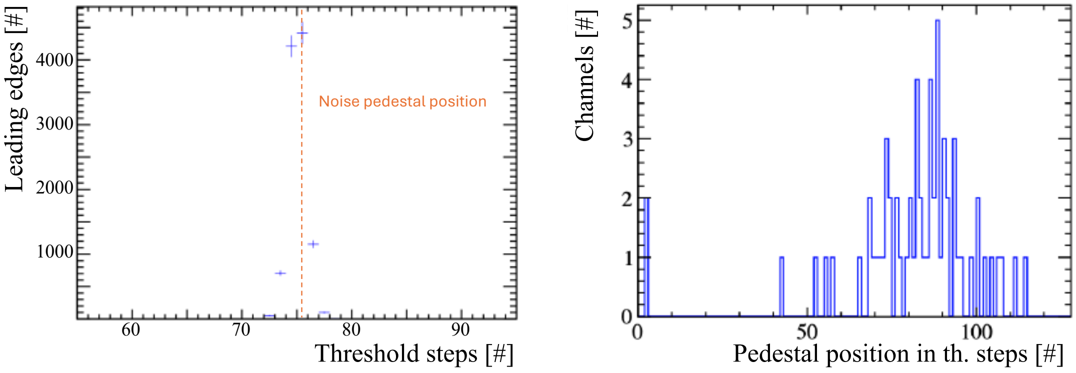

Figure 4.15: (Left) Number of detected leading edges as a function of the FastIC comparator threshold (1 step = 0.49 μA) for a single channel. The peak represent the noise pedestal level for the FastIC channel. (Right) Distribution of the noise pedestal values for all the channels of 8 FastICs.

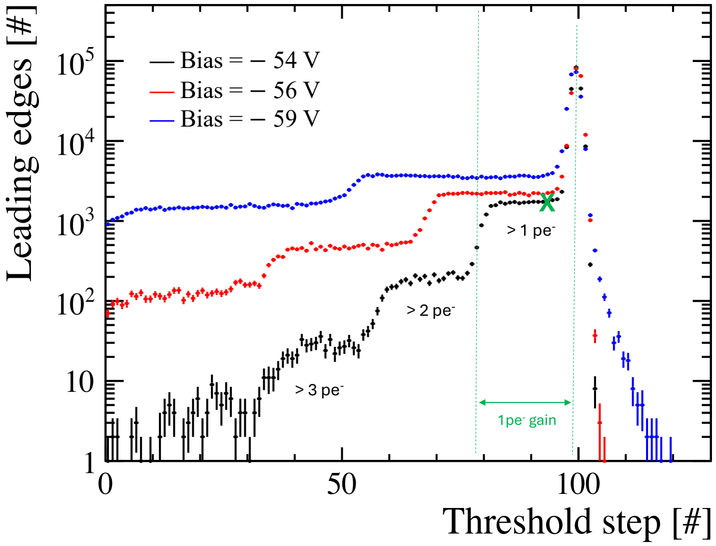

Figure 4.16: Number of detected leading edges as a function of the FastIC comparator threshold (1 step = 0.49 μA) for a single channel with the sensor biased at different voltages. The discrete response of the SiPM can be appreciated, with plateaux corresponding to one, two, or more detected photoelectrons. These distributions can be used to measure single photon gain and DCR of each channel of the SiPM array individually, and to estimate the optimal working point for each bias voltage. The green "X" marks the optimal working point at −54 V.

to-channel spread remains below 10 %, indicating uniform gain and DCR values across the entire SiPM array. The same analysis was repeated for other bias voltages, and the average single-photon gain and DCR values are reported in Figure 4.18. Since both parameters depend on temperature, all measurements were performed at 21 °C, monitored with the PT1000 sensor on the flex-PCB. As expected, both the single-photon gain and DCR increase linearly with the applied bias voltage. Between −54 V and −59 V, the gain doubles, while the DCR rises by a factor of approximately 1.5. Although a higher gain generally improves photon detection efficiency and time resolution, the associated rise in DCR can degrade the performance if not properly controlled.

Once the optimal FastIC comparator thresholds were identified, the SiPM array was illuminated with the picosecond-pulsed laser. Initial measurements were performed to align the mirror, centre the laser spot on the SiPM array and ensure operation in the single-photon illumination regime. A practical figure of merit to verify operation in the single-photon regime is the photon occupancy, defined as

$$\text{Occupancy} = \frac{N_{\text{photons}}}{N_{\text{triggers}}}, \tag{4.6}$$

where N_{photon} is the number of detected photons and N_{triggers} is the number of triggers. The number of detected photons per trigger in a given channel follows a Poisson distribution:

$$P(n) = \frac{\lambda^n}{n!} e^{-\lambda}, \tag{4.7}$$

where λ is the expected number of detected photons per trigger. The probability of detecting two

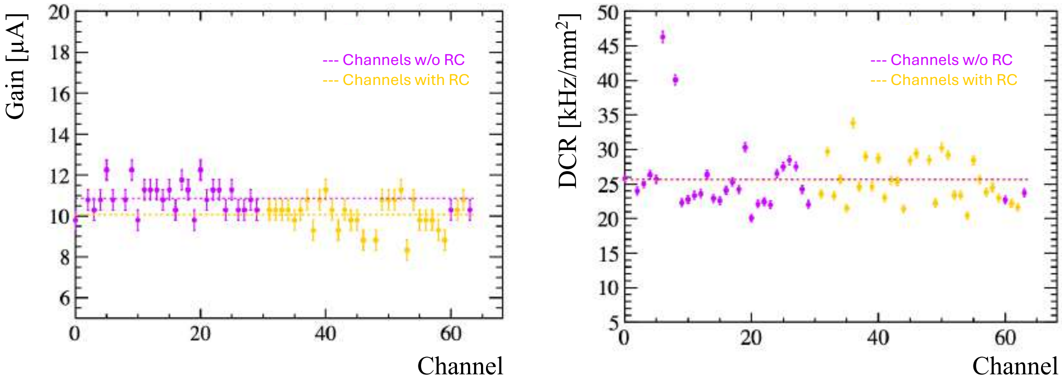

Figure 4.17: Gain (left) and DCR (right) for all the channels with a bias of −54 V. The measurement was taken at 21 °C.

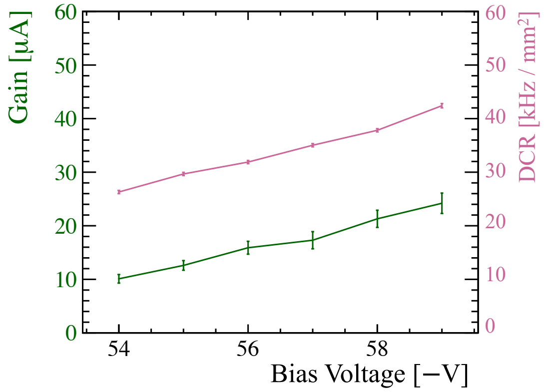

Figure 4.18: Average gain (left) and DCR (right) as a function of the applied bias voltage. The measurements were taken at 21 °C.

or more photons can therefore be calculated as

$$P(n \geq 2) = 1 - e^{-\lambda}(1 + \lambda). \tag{4.8}$$

For an average occupancy of 10 % ($P(0) = 0.9$), λ is obtained from

$$\lambda = -\ln(P(0)) \approx 0.105, \tag{4.9}$$

yielding $P(n \geq 2) \approx 0.5\,\%$. This confirms that photon a occupancy of 10 % ensure negligible multi-photon contamination. The laser intensity was therefore adjusted to achieve this occupancy across the SiPM array.

To estimate the photon occupancy from the SiPM measurements, the dark-count contribution must be subtracted. Figure 4.19 shows the distribution of leading-edge times (ToA_{ch}) for one channel. The in-time photon peak is clearly separated from the uniform dark-count background. The picoTDC readout trigger latency was adjusted to centre the photon peak within an 800 ns acquisition window. Leading edges in the first 300 ns of the window were used to estimate the DCR, following the approach described in Equation 4.5. The expected number of dark counts within the signal peak was then subtracted to obtain the in-time photons only:

$$N_{\mathrm{photon}} = N_{\mathrm{leading}}^{T_{\mathrm{signal}}} - N_{\mathrm{leading}}^{T_{\mathrm{DCR}}} \cdot \frac{T_{\mathrm{signal}}}{T_{\mathrm{DCR}}}. \tag{4.10}$$

Here, $N_{\mathrm{leading}}^{T_{\mathrm{signal}}}$ and $N_{\mathrm{leading}}^{T_{\mathrm{DCR}}}$ are the numbers of leading edges detected within the signal and dark-count time windows (T_{signal} and T_{DCR}), respectively. The corrected occupancy is then calculated

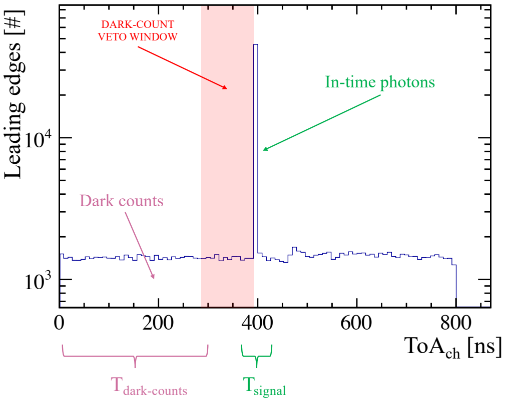

Figure 4.19: Distribution of leading edges recorded in one channel over the 800 ns acquisition window. The in-time photon peak around 400 ns is clearly visible above the uniform dark-count background.

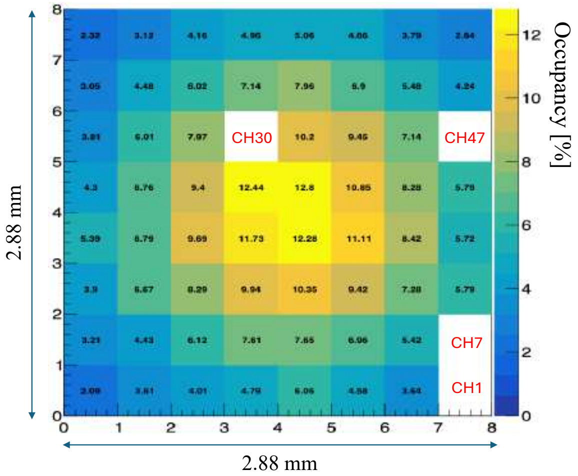

Figure 4.20: Occupancy map of the SiPM array obtained with the pulsed laser. The laser spot is visible with a maximum occupancy of 12 % at the centre and an average occupancy of 6 % across the array. Four channels are absent: CH1 and CH7 (used for time reference), CH47 (disconnected for test-pulse measurements), and CH30 (non-operational).

using Equation 4.6.

The resulting occupancy map of the SiPM array is shown in Figure 4.20. The distribution reproduces the laser spot profile, with a maximum occupancy of 12 % at the centre and an average occupancy of 6 % across the array. This confirms operation in the single-photon regime. Four channels are absent from the map. CH1 and CH7 correspond to the picoTDC channels used for the time-reference signals. Their associated FastIC channels were therefore not connected to the picoTDC. CH47 was disconnected by removing its 0 Ω resistor in the flex-PCB, as its trace was initially used for the test-pulse measurements. CH30 was not operational, likely due to a faulty contact between the SiPM and the flex-PCB or damage to the channel. Since the same channel worked correctly in other flex-PCBs, the problem was attributed to the assembly.

Then, several datasets were then collected under the same illumination conditions. Each dataset corresponds to a specific combination of comparator threshold and SiPM bias voltage and contains 500 thousand triggers. These data were used to validate the working point and to study the SPTR of the SiPM array.

4.3.3 Single-photon time resolution analysis method

The SPTR is measured as the spread of the single-photon hits ToA for each channel. The ToA of each photon hit ($\mathrm{ToA}_{\mathrm{ch}}$) is measured relative to the average ToA of the two time-reference signals ($\mathrm{ToA}_{\mathrm{ref}}$) to compensate for trigger jitter. Figure 4.21 (left) shows the relative ToA distribution of all photon hits in a single channel. The distribution is obtained for a central channel of the SiPM array biased at −54 V. Two distinct populations can be observed. The main population corresponds to single-photon signals. The second population appears at earlier ToA and with larger ToT pulses, as shown in Figure 4.21 (right). This behaviour is consistent with double-photon events, which produce pulses of higher amplitude and longer duration, as shown in Figure 4.23. The fraction of these events is less than 1 % of the total triggers, in agreement with the expected fraction of double-photon hits given by Equation 4.8 for a photon occupancy of 12 %. A ToT selection was therefore applied to remove double-photon events and retain only single-photon hits (Figure 4.22).

In addition to the ToT selection, a dark-count veto is applied. This veto discards the trigger event if a dark count occurs within 100 ns before the signal window (T_{signal} in Figure 4.19). This prevents baseline distortions caused by dark counts close to the expected signal time, that would affect the photon-signal ToA.

After these photon-hit selections, the $\mathrm{ToA}_{\mathrm{ch}} - \mathrm{ToA}_{\mathrm{ref}}$ distribution is fitted with a double-Gaussian function. The main Gaussian describes the core of the distribution, while the second accounts for the tails. An iterative process selects the best fit based on the reduced χ^2 value:

$$\chi^2_{\mathrm{red}} = \frac{1}{N_{\mathrm{dof}}} \sum_i \frac{(y_i - f_i)^2}{\sigma_i^2}. \tag{4.11}$$

Here, y_i is the measured number of counts in bin i, f_i is the corresponding value from the fit

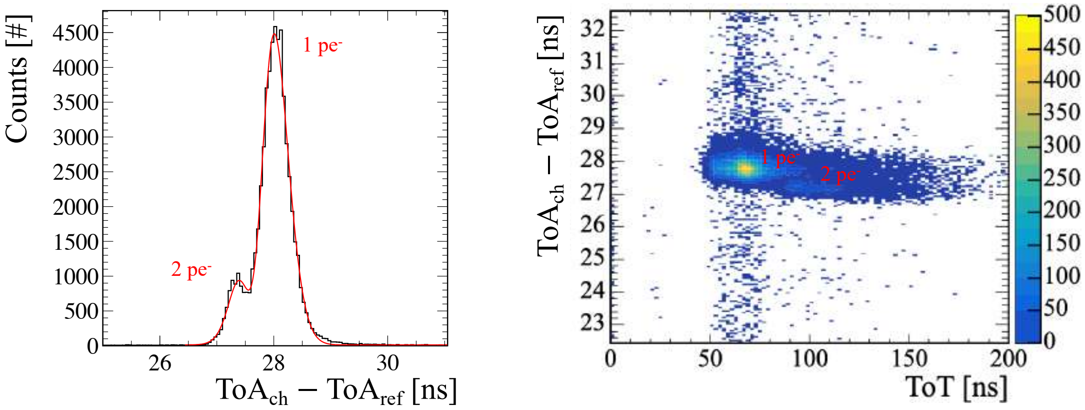

Figure 4.21: (Left) Distribution of the difference between the channel ToA and the average time-reference ToA for the photon hit registered in 500 thousand triggers. (Right) Difference between the channel ToA and the average time-reference ToA as a function of the pulse ToT. The main population arises from single pe pulses, the smaller one from double pe pulses.

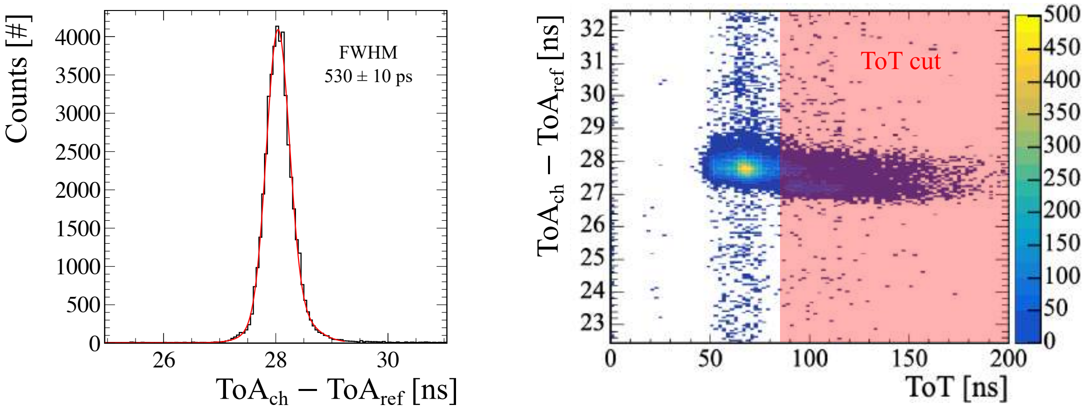

Figure 4.22: Same as Figure 4.21 with a ToT cut applied to select single pe pulses only. The relative ToA distribution is fitted with a double-Gaussian function and the SPTR is extracted as the FWHM of the fit.

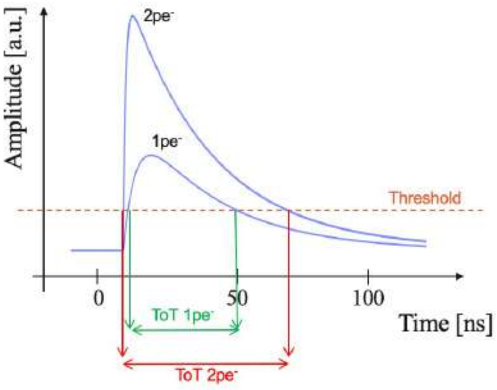

Figure 4.23: Example of 1 pe and 2 pe pulses showing the difference in the expected ToA and ToT.

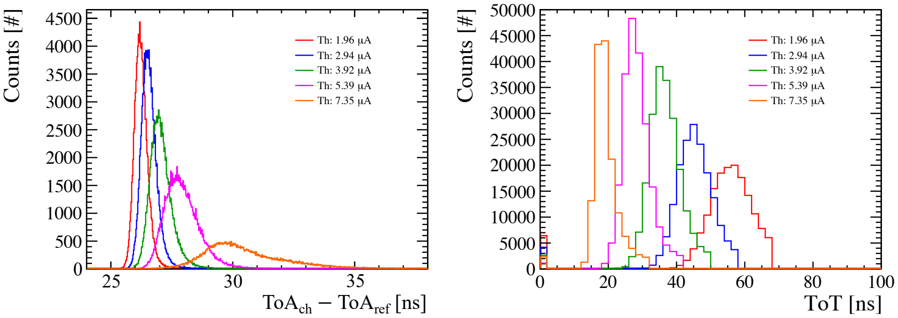

Figure 4.24: Distributions of single-photon $\mathrm{ToA}_{\mathrm{ch}} - \mathrm{ToA}_{\mathrm{ref}}$ (left) and ToT (right) for different FastIC discriminator thresholds at −54 V. The shift of the mean ToA towards later times and the reduction of ToT with increasing threshold are consistent with the expected comparator behaviour. The ToT peak position and spread for each threshold was used to define the single-photon selection window.

function, σ_i is the statistical uncertainty of bin i, and N_{dof} is the number of degrees of freedom. The fit with the lowest reduced χ^2, closest to unity, is used. The SPTR is defined as the FWHM of the best fit. The uncertainty associated to the SPTR is calculated as:

$$\sigma_{\mathrm{FWHM}}^2 = \sigma_{\sigma_1}^2 + \sigma_{\sigma_2}^2 + \sigma_{\mu_1}^2 + \sigma_{\mu_2}^2 + \frac{1}{N}\sum_i (\mathrm{FWHM}_i - \bar{\mu})^2. \tag{4.12}$$

where $\sigma_{1,2}$ and $\mu_{1,2}$ are the standard deviation and mean values of the two Gaussian components, N is the number of iterations performed during the fitting procedure, FWHM_i is the FWHM of the i^{th} fit and $\bar{\mu}$ is the average FWHM of the fits. The last term of the sum accounts for the systematic uncertainty introduced by the fit procedure. The best fit for the channel presented so far is shown

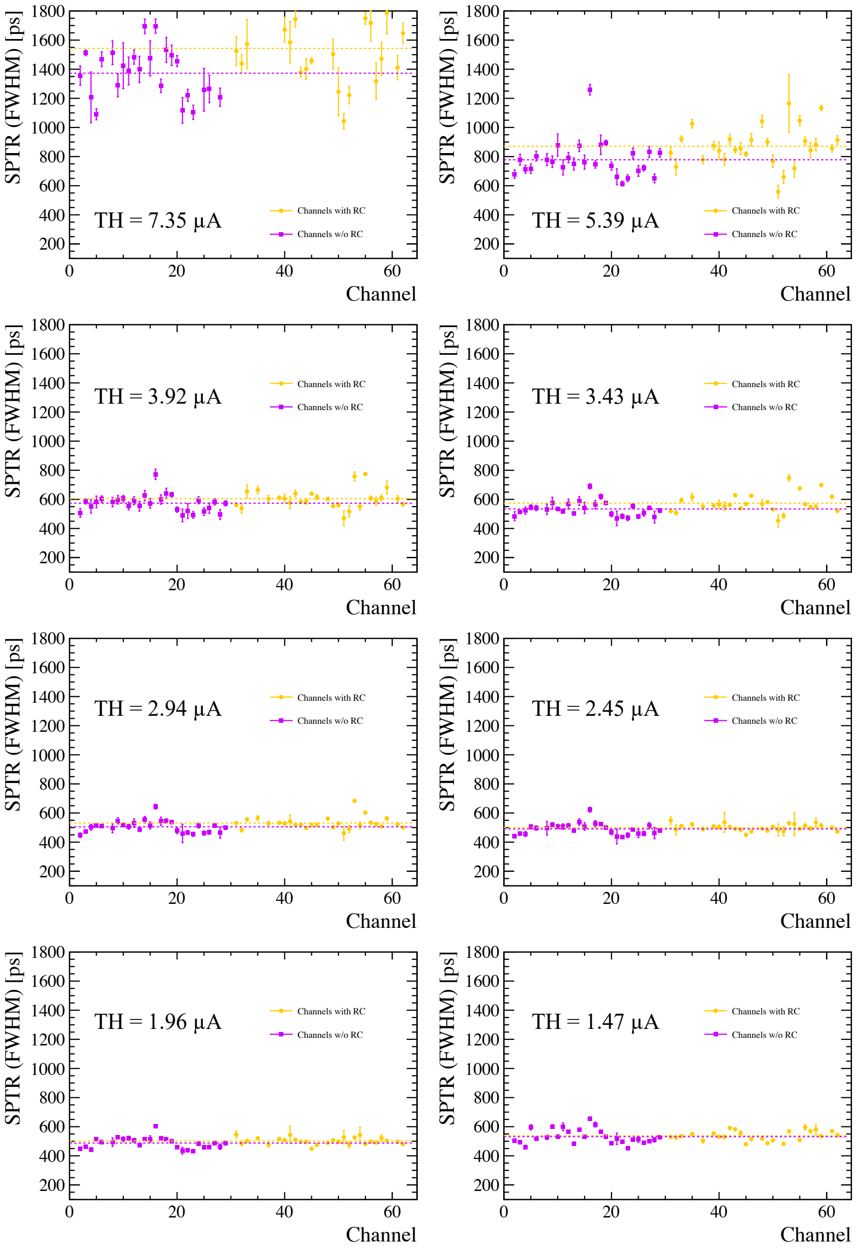

Figure 4.25: SPTR of all channels at −54 V bias for different discriminator thresholds. Channels with and without the bias RC decoupling network are shown separately. The horizontal dotted lines define the average SPTR.

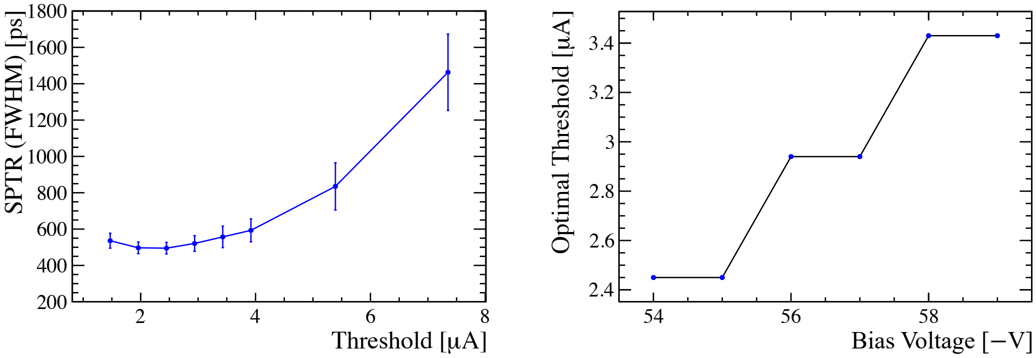

Figure 4.26: (Left) Average SPTR value of the 60 SiPMs as a function of the FastIC comparator threshold and at a bias voltage of −54 V. The optimal threshold is found to be 2.45 μA. (Right) Optimal FastIC threshold as a function of the SiPM bias voltage.

in Figure 4.22 (left), yielding an SPTR of 530±10 ps.

To refine the optimal working point, the SPTR was measured as a function of the FastIC comparator threshold. The corresponding ToA and ToT distributions for different thresholds are shown in Figure 4.24. Increasing the threshold shifts the average ToA to later times and reduces the ToT. This observation confirms the expected correlation between the SiPM response and the FastIC discriminator, validating their correct operation. The position and width of the ToT distribution were used to define the ToT cut value for each threshold. Figure 4.25 shows the SPTR measured for all channels at −54 V for different discriminator thresholds. Both the SPTR value and its uniformity depend on the threshold. Channels with and without the RC bias decoupling network are shown separately, with their respective averages indicated by the horizontal dotted lines. Since no significant difference is observed between the two groups, all channels are included in the final SPTR average, summarised in Figure 4.26 (left). The SPTR remains relatively stable for thresholds up to about 4 μA, while for higher thresholds the resolution degrades rapidly. The best SPTR for a bias voltage of −54 V is found to be 492±28 ps with an optimal threshold of 2.45 μA.

The same procedure was repeated for different bias voltages up to −59 V. The resulting optimal thresholds are shown in Figure 4.26 (right). The discrete steps are the result of the minimal threshold step size, as shown in Figure 4.16 (right). The values are consistent with the expectations from the threshold scans described in Section 4.3.2. They correspond to 10–20 % of the single-photon gain for the respective bias voltages, in agreement with the CFD values found to provide the best timing performance in the MAPMT TTS analogue measurements presented in Section 2.1.

4.3.4 Single-photon time resolution results

The SPTR measured for all channels at different bias voltages is shown in Figure 4.27. For each voltage, the FastIC comparator threshold was set its optimal value from Figure 4.26 (right). The SPTR improves steadily with increasing bias, as the higher signal amplitude enhances the signal-to-noise ratio and reduces timing jitter. The spread among channels also decreases with bias, indicating more uniform timing performance across the array. No significant difference is observed between channels with and without the RC bias decoupling, indicating that there is no need to individually decouple each SiPM in the array for the signal quality or timing performance

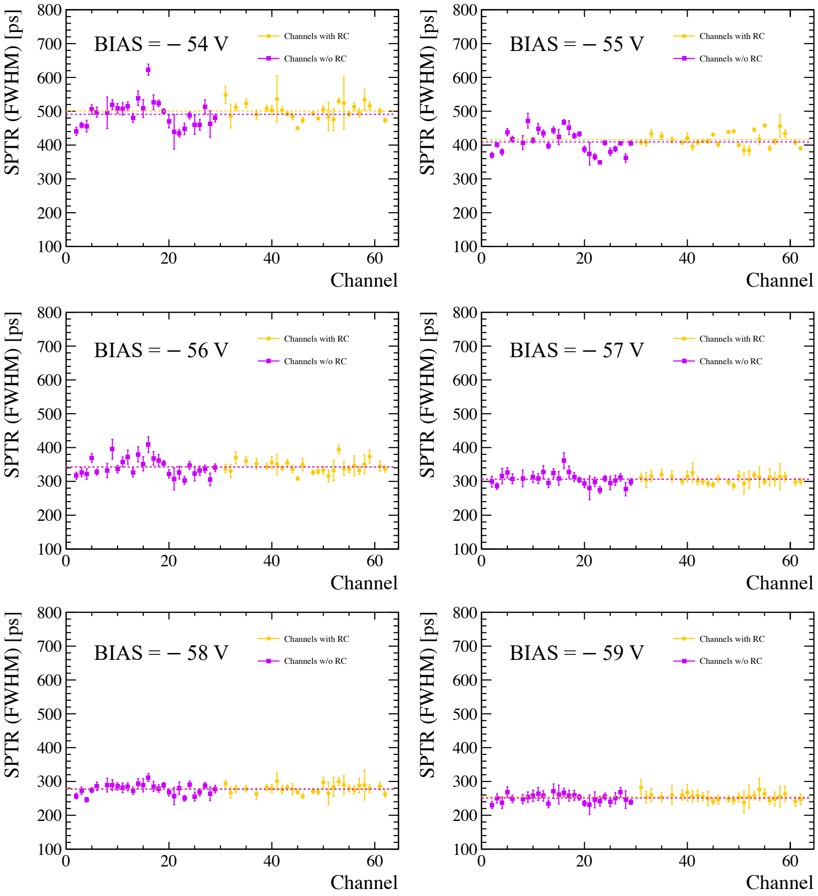

Figure 4.27: SPTR of all channels at different bias voltages for the respective optimal working point. Channels with and without the bias RC decoupling network are shown separately. The horizontal dotted lines define the average SPTR.

The average SPTR as a function of bias voltage is summarised in Figure 4.28 (left). The resolution improves from 492 ± 28 ps (FWHM) at −54 V to 255 ± 10 ps (FWHM) at −59 V. For comparison with other studies, the FWHM values are converted into Gaussian standard deviations using $\sigma_{meas} = \mathrm{FWHM}_{\mathrm{meas}}/2.355$. The resulting σ_{meas} is corrected for the jitter contributions of the laser ($\sigma_{\mathrm{laser}} \approx 21\,\mathrm{ps}$) and the FastIC comparator ($\sigma_{\mathrm{FastIC}} \approx 25\,\mathrm{ps}$) according to:

$$\sigma_{\mathrm{res}} = \sqrt{\sigma^2_{\mathrm{meas}} - \sigma^2_{\mathrm{laser}} - \sigma^2_{\mathrm{ref}} - \sigma^2_{\mathrm{FastIC}}}. \tag{4.13}$$

After correction, the average SPTR at −59 V is $\sigma_{\mathrm{res}} = 103 \pm 5\,\mathrm{ps}$. The best-performing channel reaches an SPTR of 220 ± 10 ps (FWHM), corresponding to a Gaussian sigma of 88 ± 4 ps, as shown in Figure 4.28 (right).

These results agree with literature values for the same type of SiPM under similar testing conditions. They represent the first demonstration that the LHCb RICH time-resolved readout electronics prototype can achieve 100 ps time resolution across multiple channels. These measurements, performed under controlled laser illumination, serve as reference for the flex-PCB SiPM array performance, allowing direct comparison with charged-particle beam measurements (see Section 4.4). As described in Section 4.1.2, the main limitation to further improve the SPTR is the intrinsic capacitance of the SiPM, which scales with its active area and the number of microcells connected in parallel. For this reason, subsequent flex-PCB versions will employ smaller SiPMs,

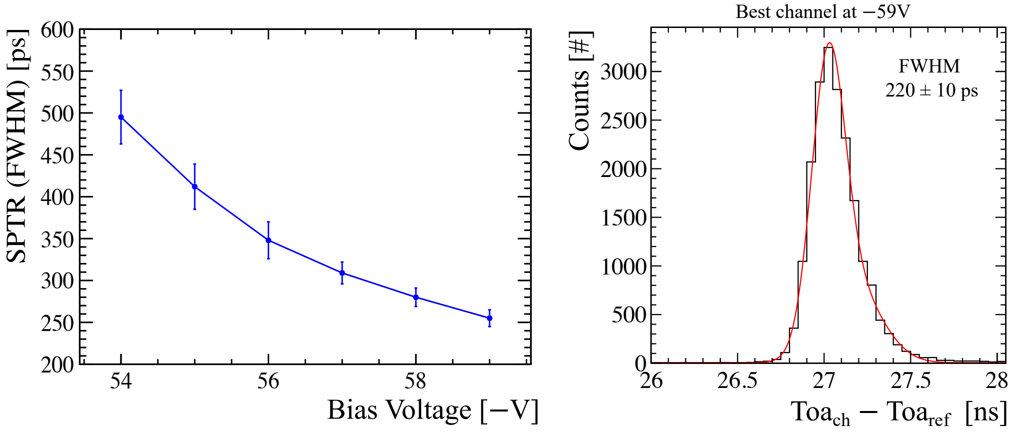

Figure 4.28: (Left) Average SPTR value of the 60 SiPMs as a function of the bias voltage. Each dataset was recorded with the optimal FastIC threshold for the corresponding bias voltage. (Right) Relative ToA distribution of the most performing SiPM at −59 V. The double-Gaussian fit on the distribution yields an SPTR of 220±10 ps (FWHM), corresponding to a Gaussian standard deviation of 88 ± 4 ps, after corrections.

with active areas of 2 mm or less, as described in Section 4.5. This choice is also consistent with the increased granularity required for the RICH photon detector foreseen for Upgrade II.

4.4 The flex-PCB performance at SPS charged particle beam

This section presents the measurements performed with the flex-PCB prototype at the CERN SPS North Area facility, as part of the 2025 LHCb RICH beam test campaign. The flex-PCB was operated under detector-like conditions, using an optical setup for Cherenkov photon illumination and a dedicated track-based time reference system. The goal was to evaluate the SiPM array performance under realistic operating conditions, focusing on the SPTR.

Section 4.4.1 describes the experimental setup used in the beam test. Section 4.4.2 presents the analysis method and the optimisation of the working point. Section 4.4.3 reports the SPTR results obtained with the flex-PCB prototype.

4.4.1 Experimental setup

The experimental setup used to measure the flex-PCB performance in the charged-particle beam test is described in this section. The measurements were performed at the CERN SPS North Area. The facility provides a high-energy charged-hadron beam, primarily composed of pions and protons with a momentum of 180 GeV/c. The beam profile was measured for each particle spill using multi-wire proportional chambers positioned a few metres downstream of the test setup. The Gaussian beam profile has typical spreads of $rms_x = 8.1\,\mathrm{mm}$ and $rms_y = 8.8\,\mathrm{mm}$. The beam intensity varied between 5×10^5 and 10^6 particles per spill.

Two flex-PCB prototypes were integrated in the beam test setup, referred to as *Flex-PCB Top* and *Flex-PCB Bottom*. In addition, a 2-inch MAPMT and four Hamamatsu S13360-3050 SiPM arrays were included in the setup, as shown in Figure 4.29 (right). The SiPM arrays are identical to those mounted on the flex-PCB but installed on a RICH baseboard-like PCB [91], featuring short traces between the sensors and the readout electronics. Since these sensors were previously used in other beam test campaigns, their performance shows some degradation. A detailed comparison with the flex-PCB results is foreseen for future studies. The 2-inch MAPMT was included in the setup as a reference sensor. Extensive datasets were collected with these MAPMT in previous RICH beam test campaigns, where its integration with the readout electronics was already well established. Consistent performance from the MAPMT therefore provided a benchmark for validating the measurement conditions.

The optical system follows the same geometry as in the 2022 LHCb RICH beam test [91], as shown in Figure 4.30. Cherenkov photons are produced by the beam particles traversing a

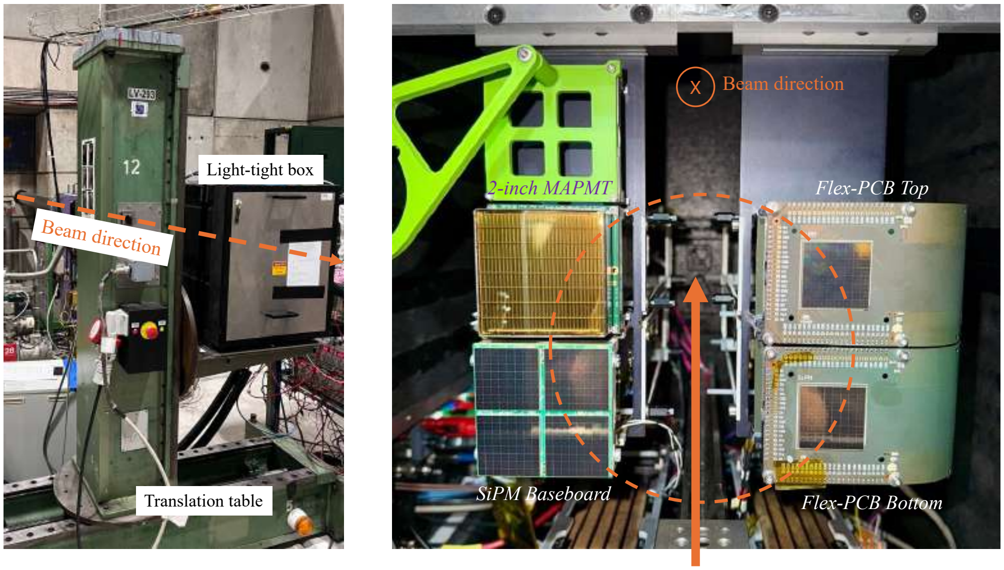

Figure 4.29: (Left) Image of the light-tight box mounted on the remotely-operated translation table for alignment with the charged particle beam. (Right) Photodetectors viewed from the beam direction. The orange-dotted line represents the expected position of the Cherenkov photon ring.

borosilicate lens[8], which focuses the emitted light onto the photodetector plane. The expected position of the Cherenkov ring on the photodetector plane is indicated by the orange-dotted line in Figure 4.29 (right). The sensors, readout electronics, and optical system are installed inside a light-tight box mounted on a translation table. The table allows precise movement in the plane perpendicular to the beam direction, as illustrated in Figure 4.29 (left). The temperature inside the box is maintained at 16 °C by a water-cooling system coupled with air fans. To align the setup with the beam axis, two pairs of crossed scintillator detectors with a $5 \times 5\,\text{mm}^2$ active area perpendicular to the beam are installed inside the box. One pair is placed upstream of the lens, while the second pair downstream at a distance of 50 cm. The beam position is determined by maximising the coincidence rate between the upstream and downstream scintillators. The relative displacement of the two pairs at their respective maxima is used to calculate the beam angle with respect to the setup, which was found to be below 1 mrad. The coincidence signals from the scintillators are also used as the trigger to the data acquisition and to monitor the beam alignment during operation.

The data acquisition is based on the MuDAQ system introduced in Section 1.4.3 and and further detailed for the laser measurement setup. The photodetectors are read out by the FastIC + PicoTDC electronics chain with the same timestamp processing described in Section 4.3.1. Differently from

---

[8] Edmund Optics PCX condenser lens no.27509 with 150 mm diameter and 300 mm focal length.

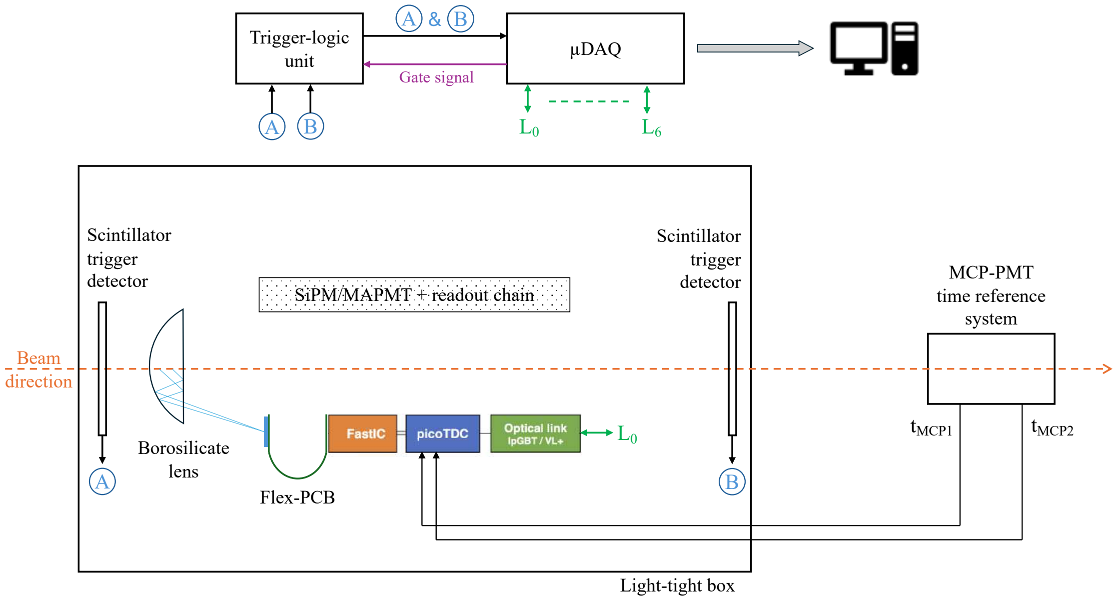

Figure 4.30: Schematic of the beam test setup.

the laser setup, the trigger is provided externally by a dedicated FPGA-based Trigger Logic Unit (TLU). The TLU sends a trigger signal to the MuDAQ upon receiving a coincidence from the two pairs of crossed scintillator detectors. This ensures that a valid readout trigger is generated only by particles traversing the full optical system and producing a Cherenkov photon ring on the photodetectors. To prevent data loss, the MuDAQ issues a busy-veto signal to the TLU whenever its internal buffer approaches full capacity, temporarily disabling new triggers. The MuDAQ distributes the 40 MHz system clock both to the trigger board and to the readout electronics.

Since the SPS beam is asynchronous with respect to the DAQ clock, a dedicated time-reference system is positioned along the beam line downstream of the light-tight box, as shown in Figure 4.31. The system consists of two Photonis PP2365Y MCP-PMTs [120], each optically coupled to a quartz bar. Particles crossing the quartz produce Cherenkov light detected by the MCP-PMTs. A custom amplification circuit amplifies the MCP-PMT signals, which are then read out by a CFD module. The CFD output of each MCP-PMT is fed to separate picoTDC channels of the *Flex-PCB Bottom* readout chain and stored in the data as a time reference for each trigger event. In addition, this system provides an particle detection system. The number of signals detected by each MCP-PMT is compared with the number of gated coincidences to estimate the efficiency of the trigger system. The fraction of MCP-PMT signals per trigger remained consistently above 95 % throughout the beam test duration. The time jitter of the reference system is measured from the time-difference

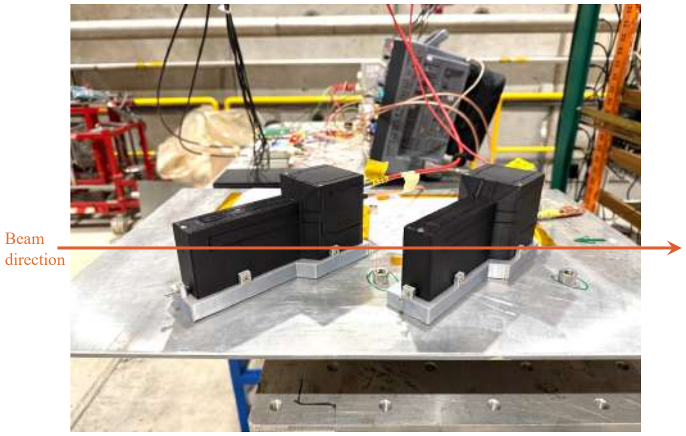

Figure 4.31: Time-reference system consisting of two Photonis PP2365Y MCP-PMTs coupled to quartz bars. Each MCP–quartz pair is mounted in a dedicated 3D-printed housing.

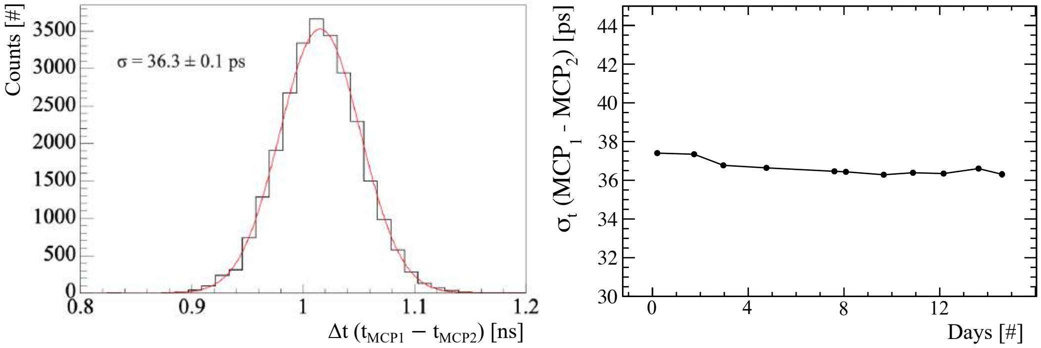

Figure 4.32: (Left) Distribution of the time difference between the two MCP-PMTs. The spread of the distribution extracted with a Gaussian fit is used for measuring the jitter of the time-reference system. (Right) Trend of the time-reference jitter during the beam test period.

distribution between the two MCP-PMTs. Figure 4.32 (left) shows the distribution obtained from a run at the beginning of the beam test. The Gaussian fit yields a standard deviation of ~37 ps, which corresponds to the combined jitter of the two MCP-PMTs. The intrinsic jitter introduced by the time-reference system is therefore $36/\sqrt{2} \approx 25$ ps. This value remained stable throughout the beam test, as shown in Figure 4.32 (right).

Several datasets were collected with the flex-PCBs at different FastIC comparator thresholds and SiPM array bias voltages. This allowed the chosen working points to be validated and the results to be compared with those obtained from the laser measurements. In addition, each configuration was recorded twice: once during the first days of the beam test and once at the end. This enabled the study of variations in the SiPM performance caused by beam exposure. The analysis method followed the same approach used for the laser data, with the main differences highlighted in Section 4.4.2. The results presented in this chapter concern *Flex-PCB Bottom*, which was extensively characterised in the laboratory and used for the laser measurements described in Section 4.3. *Flex-PCB Top* was installed shortly before the beam test and underwent only limited laboratory characterisation. The data collected with this board showed signs of a picoTDC misconfiguration that affected the timing performance. Despite this limitation, the data from *Flex-PCB Top* were still used to verify the position of the Cherenkov ring and to monitor the DCR.

4.4.2 Working point and analysis method

The working point definition and SPTR analysis build on the extensive characterisation of the flex-PCB performed with the laser setup, described in Sections 4.3.2 and 4.3.3. Figure 4.33 shows the distribution of the leading-edge ToA for a single channel. In contrast to the laser measurements, the in-time photon peak is spread over approximately 25 ns. This results from the asynchronous

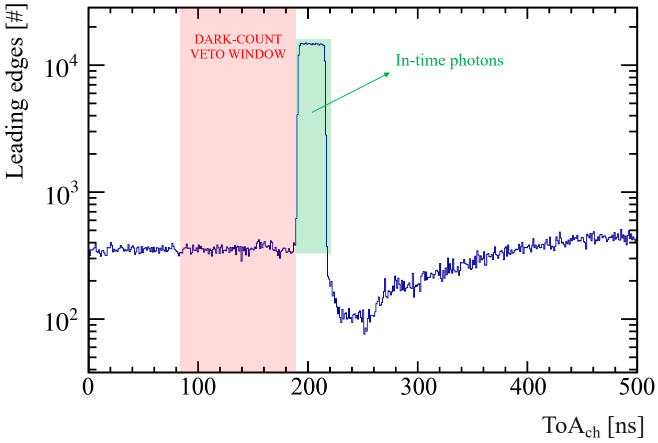

Figure 4.33: Distribution of leading-edge times (ToA) for a representative channel during beam operation. The green region indicates the in-time Cherenkov photon signals, clearly distinguishable over the dark counts background. The 25 ns spread of the Cherenkov photon peak arises from the asynchronous beam arrival with respect to the 40 MHz acquisition clock. The shaded red region marks the dark-count veto window: the trigger event is discarded if a dark count is detected in the 100 ns before the expected signal peak.

nature of the SPS beam with respect to the 40 MHz acquisition clock, which causes the photon signals to be distributed over one clock period. To correct for this effect, the ToA of each channel (ToA_{ch}) is measured relative to the time reference, defined as the average ToA of the two MCP-PMT signals ($\mathrm{ToA}_{\mathrm{ref}} = (\mathrm{ToA}_{\mathrm{MCP1}} + \mathrm{ToA}_{\mathrm{MCP2}})/2$). The SPTR for each channel is then extracted from the $\mathrm{ToA}_{\mathrm{ch}} - \mathrm{ToA}_{\mathrm{ref}}$ distribution. The ToA_{ch} distribution is also used to measure the photon occupancy. As in the laser data, the number of in-time photons is corrected for dark counts and normalised to the number of triggers, following Equations 4.6 and 4.10. For the beam test configuration, the acquisition window was reduced to 500 ns, with photon signals time-aligned around 200 ns. This optimisation reduces the number of dark-count hits per trigger, improving acquisition efficiency and limiting the data volume. Therefore, the DCR is estimated from the first 150 ns of the acquisition window. A dark-count veto is also applied to the 100 ns preceding the signal peak in this analysis. A drop in the number of leading edges is visible immediately after the photon peak. This behaviour originates from the relatively long recovery time of the SiPMs, which translates into large ToT values. After the detection of a hit, the signal remains above threshold for several tens of nanoseconds, temporarily disabling the binary readout channel. The effect is more pronounced in the beam test data than in the laser measurements, indicating a higher photon occupancy in the beam test than in the laser data.

Figure 4.34 (left) shows the hitmaps for *Flex-PCB Top* and *Flex-PCB Bottom*, recorded at a bias voltage of -54 V and a threshold of $2.94\,\mu$A from 500 thousand triggers. The Cherenkov ring is clearly visible and located in the expected position, as shown in Figure 4.29. The SPTR analysis focuses on the channels located on the ring, defined as those with an occupancy above 10 %. The average occupancy of the ring channels is approximately 40 %, while the maximum occupancy reaches up to 65 %. Figure 4.34 (right) show the time-difference distribution $\mathrm{ToA}_{\mathrm{ch}} - \mathrm{ToA}_{\mathrm{ref}}$ as a function of the ToT for two representative channels: one located on the ring with 65 % occupancy, and one on the ring edge with 15 % occupancy. For both channels, multiple populations are observed, corresponding to one, two, or more than three photon hits. These populations are more pronounced for the channel on the ring, because of its higher photon occupancy. Table 4.3 compares the expected and measured probabilities of detecting one, two, or more than three photons. The expected probabilities are calculated using the Poisson distribution described in Equations 4.7 and 4.9, while the measured ones result from the integration of the two-dimensional distributions around each population peak. The table shows excellent agreement between the expected and measured values, confirming the correct operation of the SiPM array and readout chain.

To extract the SPTR of each channel, an iterative double-Gaussian fitting procedure is applied to the $\mathrm{ToA}_{\mathrm{ch}} - \mathrm{ToA}_{\mathrm{ref}}$ distribution after a ToT selection, as shown in Figure 4.35. The SPTR is defined as the FWHM of the fitted distribution, yielding 572 ± 18 ps for the representative channel on the ring at -54 V. The same procedure was applied to all channels on the ring for each combination of bias voltage and discriminator threshold. Figure 4.36 (left) shows the SPTR distribution across the channels at a bias voltage of −54 V and a threshold of 2.96 μA. The mean and standard deviation of this distribution define the average SPTR and its uncertainty for this configuration. The average SPTR values obtained at different thresholds were compared to identify the optimal working point for each bias voltage, as shown for −54 V in Figure 4.36 (right). The optimal thresholds defined in this way are consistent with those obtained from the laser data for each bias voltage, within

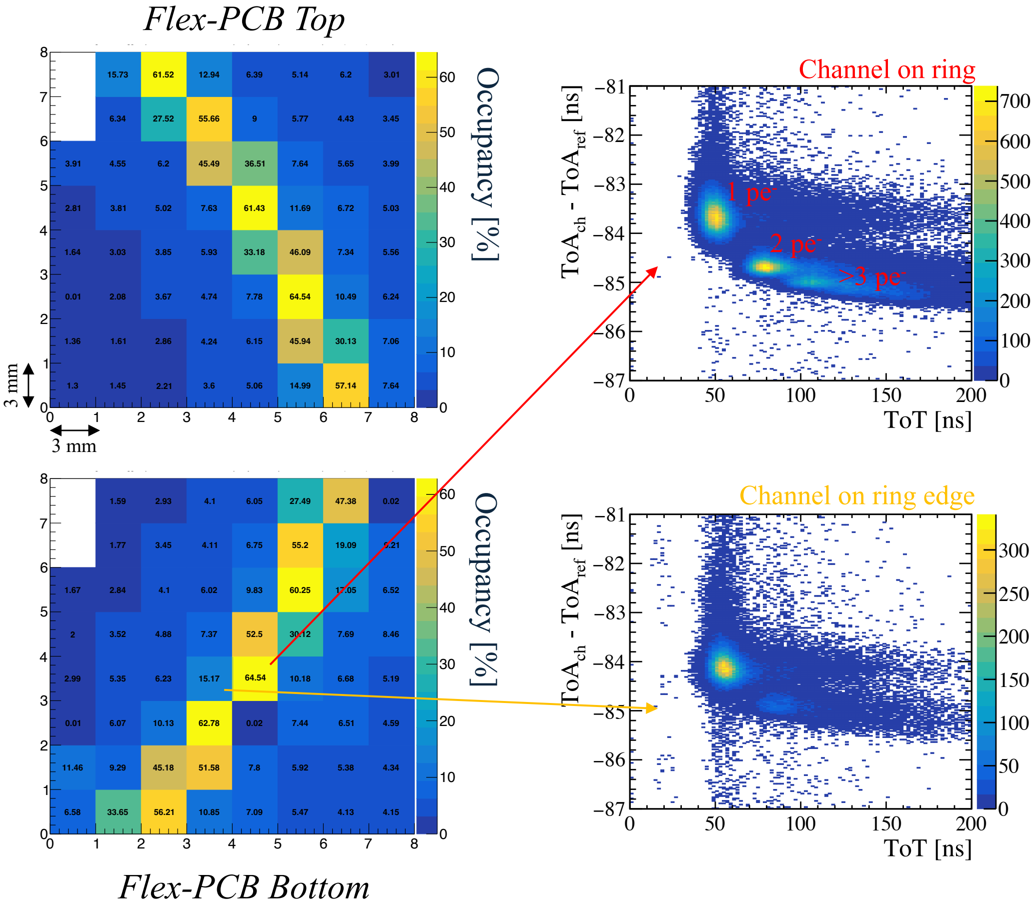

Figure 4.34: (Left) Hitmaps of the *Flex-PCB Top* and *Flex-PCB Bottom* recorded at a bias voltage of −54 V and a threshold of 2.94 μA from 500 thousand trigger events. The Cherenkov ring is clearly visible and located in the expected position. (Right) Distribution of the time difference between the channel ToA and MCP-PMT time reference (ToA$_{ref}$) as a function of the ToT for a channel on the ring (top, 60 % occupancy) and a channel on the edge of the ring (bottom, 15 % occupancy). Distinct populations corresponding to one, two, or more than three photon hits are visible for the channel on the ring, consistent with its higher occupancy.

Table 4.3: Expected and measured photon multiplicity probabilities for channels at different occupancies.

Occupancy	P(n)	Expected P(n)	Measured P(n)
65%	P(1)	0.367	~ 0.325
	P(2)	0.192	~ 0.164
	P(3+)	0.091	~ 0.113
15%	P(1)	0.138	~ 0.137
	P(2)	0.012	~ 0.010
	P(3+)	<0.001	—

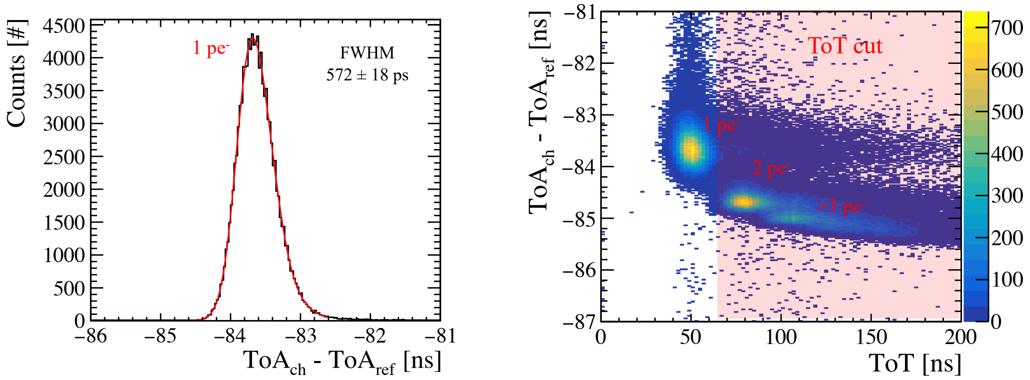

Figure 4.35: (Left) Distribution of the difference between the channel ToA and the average ToA from the MCP-PMT time-reference after the ToT cut to select single-pe hits. The relative ToA distribution is fitted with a double-Gaussian function and the SPTR is extracted as the FWHM of the fit. (Right) Distribution of the time difference between the channel ToA and the average reference MCP-PMT (ToA_{ref}) as a function of the signal ToT. The ToT cut isolates single-pe hits.

approximately one threshold step (~ $0.49\,\mu$A).

4.4.3 Single-photon time resolution results

The average SPTR as a function of the bias voltage is shown in Figure 4.37. The trend follows the trend observed with the laser data: the SPTR improves with increasing bias voltage. The best SPTR, 327 ± 30 ps (FWHM), is obtained at −59 V, consistent with the laser results. The measured FWHM is converted to the equivalent Gaussian sigma ($\sigma_{meas} = \text{FWHM}_{\text{meas}}/2.355$) and corrected for the jitter contributions from the setup:

$$\sigma_{\text{res}} = \sqrt{\sigma^2_{\text{meas}} - \sigma^2_{\text{FastIC}} - \sigma^2_{\text{ref}}}\,. \tag{4.14}$$

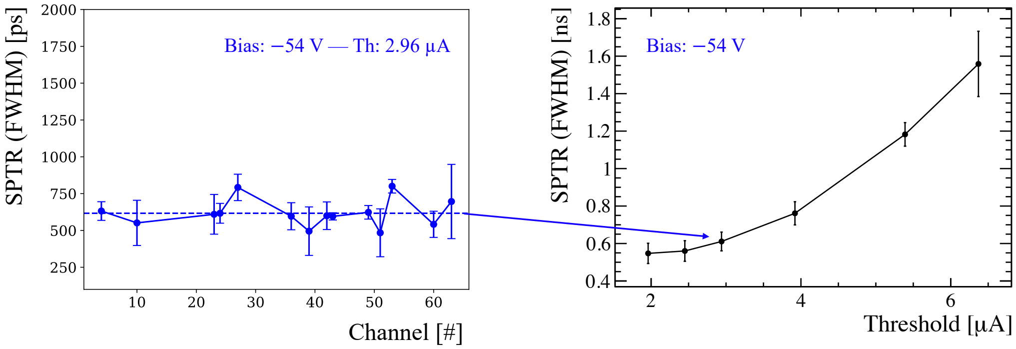

Figure 4.36: (Left) SPTR values for all channels on the SiPM ring at a bias voltage of −54 V and threshold of 2.96 μA. (Right) Average SPTR as a function of discriminator threshold for the same bias voltage.

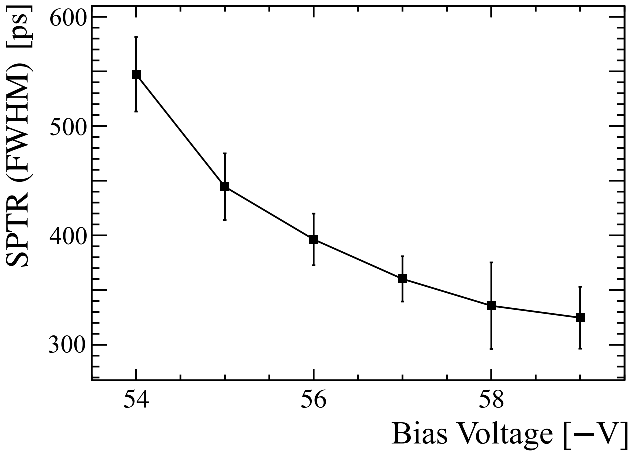

Figure 4.37: Average SPTR of the channels on the ring as a function of the bias voltage applied to the SiPM array.

Here, $\sigma_{\text{FastIC}} \sim 25\,\text{ps}$ is the FastIC comparator jitter and $\sigma_{\text{ref}} \sim 25/\sqrt{2}\,\text{ps} \sim 18\,\text{ps}$ is the time-reference jitter contribution. The corrected SPTR at −59 V is therefore $\sigma_{\text{res}} = 134 \pm 15\,\text{ps}$. This value is slightly higher than that obtained under controlled laser illumination, likely due to the higher photon occupancy observed in the beam test. The increased activity in the readout electronics may introduce additional noise, degrading the timing performance. To support this hypothesis, the expected number of simultaneously active channels within a 25 ns trigger window is estimated. The

Table 4.4: Expected number of active channels within a 25 ns trigger window ($N_{\text{ch}}/25\,\text{ns}$), including dark count and photon contributions. For beam test data the number of channels are weighted for the on-ring and off-ring occupancy.

Data set	DCR [kHz/ch]	Occupancy [%]	$N_{\text{ch}}/25\,\text{ns}$
Laser data	390	6	4.4
Beam test (day 2)	560	40 (ring) / 4 (rest)	8.4
Beam test (day 14)	2510	40 (ring) / 4 (rest)	11.1

calculation includes both dark counts and photon-induced hits, assuming Poisson statistics for each process (Equation 4.7). The probability that a channel records a dark count is

$$P_{\text{dark}}(\geq 1) = 1 - e^{-R\,\Delta t}, \tag{4.15}$$

where R is the DCR in Hz and Δt is the 25 ns window. The DCR at −59 V is measured to be 390 kHz/channel for the laser data and 560 kHz/channel in the beam test (Day 2). This is a first indication an increased DCR due to beam exposure. The probability for a channel to be hit by at least one photon is

$$P_{\text{photon}}(\geq 1) = 1 - e^{-\lambda}, \qquad \lambda = -\ln(1 - \text{occupancy}). \tag{4.16}$$

For beam test data, an average photon occupancy of 40 % is considered for the channels on the ring and 4 % for the remaining ones. For laser data, an average photon occupancy of 6% is considered for all channels. The total probability to record a pulse in a channel is therefore

$$P(\geq 1\,\text{hit}) = 1 - (1 - P_{\text{dark}}(\geq 1))(1 - P_{\text{photon}}(\geq 1)). \tag{4.17}$$

Multiplying this probability by the 64 channels in the readout chain gives the expected number of active channels within a 25 ns trigger window. The estimated values are summarised in Table 4.4. The expected number of active channels in the laser data is 4.4, while for the beam test data it is 8.4. These results support the hypothesis that higher occupancy increases the electronics activity and noise, degrading the timing performance.

To further illustrate the impact of photon occupancy on the optoelectronic timing performance, Figure 4.38 compares the SPTR measured for the same channel in the beam test and in the laser setup. In the beam test, the channel is located on the edge of the Cherenkov ring with an occupancy of about 18 %, comparable to the 12 % in laser data for the same channel. The two distributions show comparable SPTR values, both consistent with the 255 ± 10 ps average SPTR (FWHM) obtained with the laser measurements. This comparison confirms that the slight degradation of the timing performance observed in the beam test is likely due to the higher overall occupancy.

The same SPTR analysis was repeated on data collected at the end of the two-week beam test

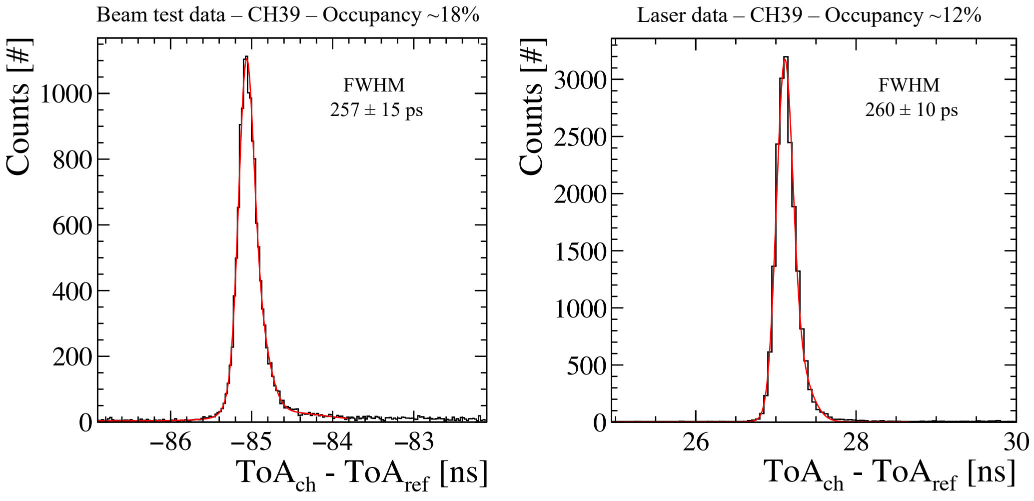

Figure 4.38: (Left) SPTR distribution for a channel at the edge of the ring in the beam test data, with an occupancy of ~18 %. (Right) Corresponding distribution for the same channel in the laser data, with an occupancy of ~12 %. Both measurements yield comparable FWHM values, consistent with the average SPTR obtained from the laser.

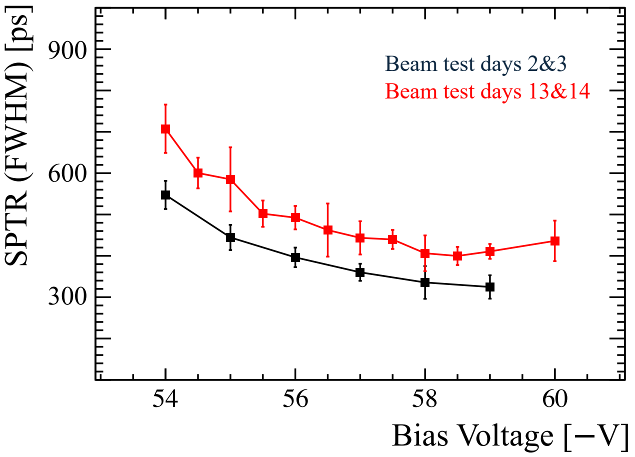

Figure 4.39: Average SPTR as a function of bias voltage measured at the beginning (days 2–3, black) and at the end (days 13–14, red) of the beam test. A degradation of approximately 70 ps (FWHM) is observed over the two-week period.

to allow a comparison with the initial measurements. Figure 4.39 shows the SPTR as a function of bias voltage for the two sets of data. A degradation of approximately 70 ps (FWHM) is observed

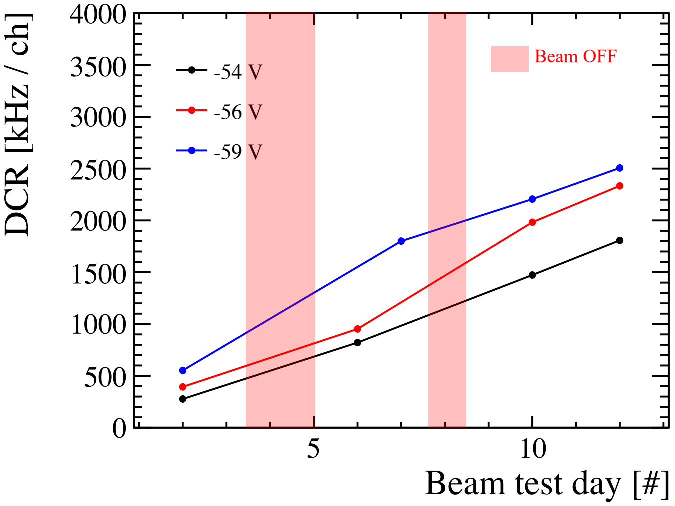

Figure 4.40: Average DCR per channel as a function of beam test day for three bias voltages. The shaded areas correspond to periods when the beam was off. A linear increase in DCR is observed, reaching about 2.5 MHz/channel at −59 V by the end of the test.

across all bias voltages between the beginning and the end of the beam test. Since the reference timing system remained stable (Figure 4.32) and no significant variations were observed in the beam size or intensity, the most likely explanation is an increase in the SiPM DCR during the beam test. Figure 4.40 shows the average DCR per channel as a function of the beam test day for three bias voltages. All DCR values were measured at a stable temperature of 16 °C maintained by the cooling system installed in the light-tight box. The DCR exhibits a linear increase over time for all bias voltages, reaching a final value of approximately 2.5 MHz/channel at −59 V. Applying this DCR value to Equation 4.17, the expected number of simultaneously active channels per 25 ns trigger window increases to 11.1. This result supports the hypothesis that the timing performance degradation is consistent with the higher activity in the readout electronics caused by the increased DCR.

The rise in DCR is attributed to radiation-induced damage to the SiPM array. Although the array was not directly exposed to the beam, it was placed only a few centimetres away and was likely affected by scattered or secondary particles in the 180 GeV/c beam. As discussed in Section 1.4.2, SiPMs are known to be highly sensitive to radiation damage, which explains the observed behaviour. This observation reaffirms the original motivation behind the flex-PCB concept: the operation of

SiPMs at cryogenic temperatures to strongly mitigate DCR. The measurements presented in this chapter will therefore be repeated with the SiPM array at cryogenic temperatures once the flex-PCB is integrated into the cryostat demonstrator. As an intermediate step, the flex-PCB will be installed in the high-precision single-photon characterisation setup described in Section 2.2.5. This will allow laser measurements at higher occupancy and repetition rate for a direct validation of these beam-test results. In addition, the setup will enable a sub-millimetre characterisation of the SiPM array SPTR both at room temperature and at −20°C.

4.5 Towards a second flex-PCB prototype

This section discusses the design ideas for the next flex-PCB prototype. The goal is to maximise the photosensitive area and minimise dead area, making the design suitable for large photodetector assemblies. Figure 4.41 shows the proposed layout, where the SiPM-side of the flex is populated with four 8×8 arrays of $3{\times}3\,\mathrm{mm}^2$ SiPMs to extend the active area. This configuration maintains the compatibility with the present elementary cell geometry of the LHCb RICH detector. However, as the detector granularity will increase for Upgrade II, prototype versions could be equipped with arrays of $2{\times}2\,\mathrm{mm}^2$ or even $1.4{\times}1.4\,\mathrm{mm}^2$ SiPMs [33].

Part of the current flex-PCB area in Figure 4.41 is occupied by RC filters for the per-channel bias decoupling. The results from Section 4.3 show that these components are not necessary for stable SiPM operation and have no measurable effect on performance. Their removal frees space and improves the scalability of the design. A limited number of RC filters may still be retained for small SiPM groups (8, 16, or 32 channels) to preserve operability in case of damage. If an SiPM fails, the filter of the corresponding group can be disconnected, disabling only that section of SiPMs instead of the entire board. The same logic applies to the bias distribution: a single bias line simplifies the layout, while multiple lines allow selective powering and remote isolation of damaged regions.

Increasing the number of channels within the same area introduces new challenges for routing the signals to the readout connectors. In the current version of the flex-PCB, all channels are routed on the top layer. A possible improvement is to also use the bottom layer. In this configuration, the bias distribution must be redesigned and analogue traces on the bottom layer must be properly shielded against interference from the bias lanes. One option to simplify the design is to route signals from two pairs SiPM arrays in opposite directions, creating two symmetric routing blocks, as illustrated in Figure 4.42. An alternative (or complementary) solution is to use $100\,\mu$m traces instead of the current $200\,\mu$m ones. This technology is already employed in the LHCb SciFi detector [46], which uses flex-PCB technology to transmit analogue SiPM signals for the readout scintillating fibres. However, reducing the trace width increases both its characteristic impedance and the signal attenuation. Dedicated tests are therefore required to verify the suitability of such narrow

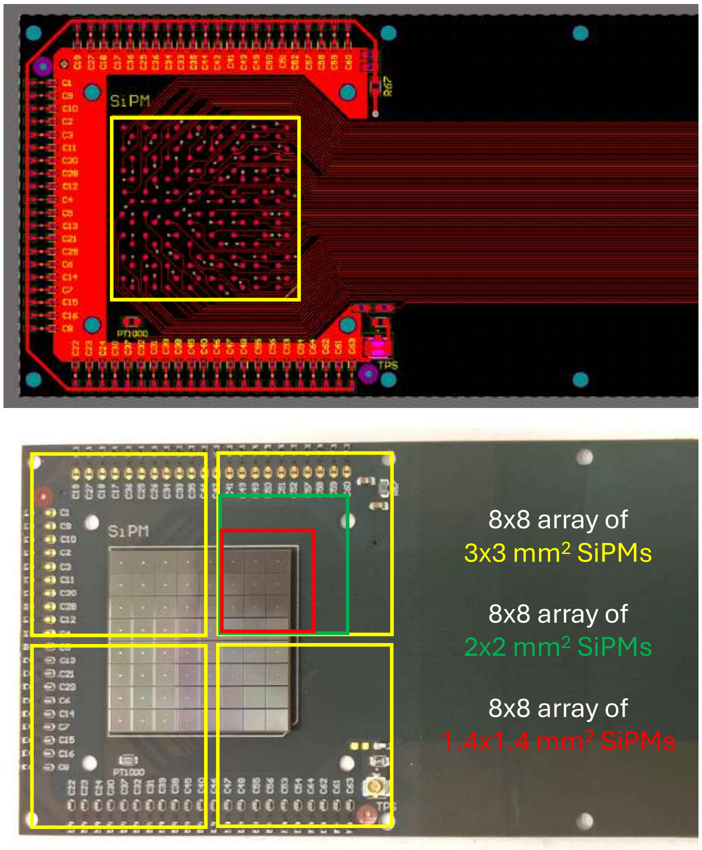

Figure 4.41: (Top) Layout of the new flex-PCB version showing the routing optimisation for increased active area. (Bottom) Prototype concept with four 8×8 arrays of SiPMs: 3×3 mm^2 (yellow), 2×2 mm^2 (green), and 1.4×1.4 mm^2 (red) pixels. This configuration maximises the photosensitive coverage while maintaining compatibility with the current elementary cell geometry.

traces for single-photon timing in the RICH detector. This solution is particularly relevant for high-granularity versions of the next flex-PCB, where the space for signal routing is further reduced.

Another important aspect to consider in close-packed SiPM arrays is the clamping of the flex-PCB to the cold volume. In the current version, dedicated mounting holes accommodate a mechanical frame around the SiPM array, ensuring good thermal contact between the flex-PCB and the cold surface. In the new flex-PCB, a similar frame might still be used but with a reduced footprint to minimise the dead area. Part of the loss can be recovered by using micro-lenses, which focus the incident photons towards the SiPM centre, as shown in Figure 4.43. Micro-lenses have been shown to enhance both photon detection efficiency and timing performance [82], also mitigating the impact of DCR in SiPMs. Although still at an early design stage, this approach represents a

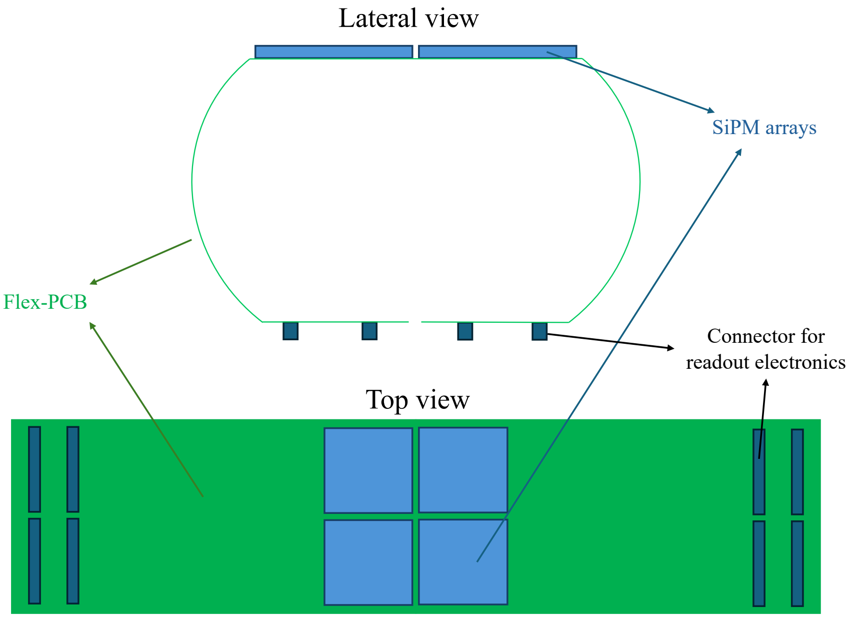

Figure 4.42: Schematic concept for the new flex-PCB prototype routing strategy.

practical compromise between mechanical stability and optical efficiency, and can serve as a starting point for future flex-PCB designs.

The next flex-PCB version also aims to improve mechanical flexibility. To scale the design and make it compatible for a photodetector module, the flex-PCB should achieve a smaller bending radius, allowing a nearly 90° bend immediately after the SiPM arrays. The Kapton and Epoxy glue thicknesses used in the current design are standard values, chosen to simplify and speed up the fabrication of the first prototype. Discussions with the CERN manufacturing service have indicated that both the Kapton and adhesive layers can be reduced to $25\,\mu$m, lowering the total thickness of the flex to approximately $250\,\mu$m and significantly improving its flexibility. The thickness of the dielectric layers and the width of the traces must be carefully chosen to ensure $50\,\Omega$ impedance matching. Non-optimised values could degrade the signal integrity and affect the SiPM timing performance. For this reason, test-pulse circuits should be included in the next version to repeat the characterisation measurements described in Section 4.2.

The temperature monitoring system should also be revised. The PT1000 sensor must be relocated and its traces should be better shielded from the analogue ones. Moreover, since these sensors are not radiation-hard, an alternative should be identified for future RICH detector applications.

Finally, the new flex-PCB must be designed to be fully compatible with the cryostat demonstrator currently under development. This will ensure straightforward integration into the existing

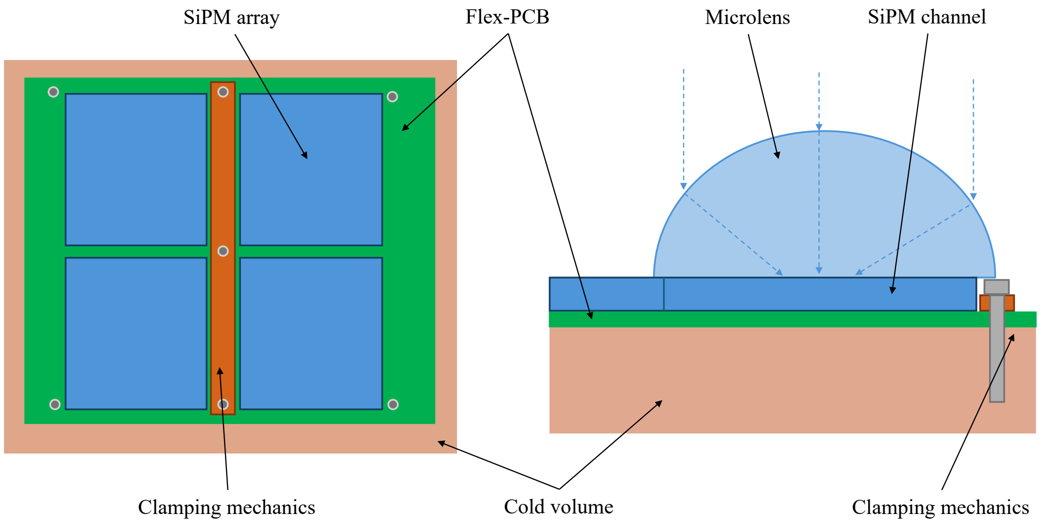

Figure 4.43: Concept for the clamping mechanism of the new flex-PCB version. (Left) Top view showing the SiPM array and possible clamping structures between the SiPMs. (Right) Lateral view illustrating the contact between the flex-PCB and the cold volume, with the potential use of micro-lenses to recover photosensitive area and enhance optical performance

system for SiPM array measurements at cryogenic temperatures.

4.6 Chapter conclusion

This chapter presented the design, characterisation, and testing of a flex-PCB for the readout of a SiPM array at cryogenic temperatures. The prototype was designed to transmit fast analogue signals over centimetre-long traces while preserving signal integrity and timing performance. Laboratory measurements demonstrated that analogue traces of ~ 15 cm preserve signal integrity and timing performance. A detailed single-photon characterisation of the flex-PCB SiPM array was performed under picosecond-pulsed laser illumination, using the FastIC + picoTDC-based readout electronics. The SiPM array achieved an average SPTR of $\sigma_{\text{res}} = 103$ ps at a bias voltage of −59 V, with the best channel reaching 88 ps. These results represent the best performance so far achieved by coupling the LHCb RICH time-resolved readout electronics with a single-photon sensor. The flex-PCB was subsequently integrated into the LHCb RICH setup for charged-particle beam tests at the CERN SPS North Area. An average SPTR of 134 ps was achieved under realistic detector-like conditions. The small degradation compared to the laser measurements was attributed to the different photon occupancies of the two configurations. An increase in DCR was also observed over the beam time, which may have contributed to the performance degradation during the beam test.

These results confirm that the flex-PCB technology is well suited to the transmission of analogue signals over centimetre-scale distances. They also provide a strong reference for the integration of the flex-PCB into the cryostat demonstrator currently under development at CERN, which will enable SiPM operation at cryogenic temperatures in 2026.

Conclusions

In this thesis, I presented studies supporting the upgrade of the LHCb RICH detectors for the LS3 Enhancements and the Upgrade II programme. I performed analogue and digital measurements of the MAPMT single-photon response under picosecond laser illumination. These results demonstrated the intrinsic timing performance of the MAPMTs, provided a reference for subsequent beam tests results and supported the development of the time-resolved optoelectronic readout chain for the LHC Run 4 RICH detector. Based on the expected time performance of the LHCb Run 4 RICH detector, I developed an algorithm for the estimation of the primary vertex time (PV t_0) using only RICH reconstruction, a first-time measurement for the LHCb experiment. The PV t_0 is a key element for applying a software time gate on the photon hit times, which suppresses combinatorial background and improves the PID performance of the RICH detector during LHCb Run 4. Finally, I designed a flex-PCB for the readout of an SiPM array. The board features ~15 cm analogue traces, allowing operation of the SiPMs at cryogenic temperature while keeping the electronics at room temperature. Laboratory studies demonstrated stable signal integrity and timing performance of the centimetre-long traces. Single-photon characterisation of the SiPM array with pulsed-laser illumination resulted in an average SPTR of 103 ps, and an SPTR of 88 ps for the best channel. These results mark the best SPTR obtained so far with an LHCb RICH time-resolved optoelectronic chain. The board was also tested on a particle beam at the CERN SPS facility, integrated into the LHCb RICH beam-test setup. The SPTR values measured under particle-beam conditions were consistent with those obtained with laser illumination.

This thesis forms part of the ongoing upgrade programme of the LHCb RICH detectors and supports the future developments, in preparation for the integration of sub-nanosecond timing information into the detector readout and reconstruction systems.

Bibliography

[1] Lyndon Evans and Philip Bryant. “LHC Machine”. In: *Journal of Instrumentation* 3.08 (Aug. 2008), S08001. doi: `10.1088/1748-0221/3/08/S08001`. url: `https://dx.doi.org/10.1088/1748-0221/3/08/S08001`.

[2] T. Massam et al. “Experimental observation of antideuteron production”. In: *Nuov Cim A* 63 (1965), p. 10. doi: `10.1007/BF02898804`.

[3] F. J. Hasert et al. “Observation of neutrino-like interactions without muon or electron in the gargamelle neutrino experiment”. In: *Physics Letters B* 46.1 (1973), pp. 138–140. doi: `10.1016/0370-2693(73)90499-1`.

[4] UA1 Collaboration. “Experimental observation of isolated large transverse energy electrons with associated missing energy at $\sqrt{s} = 540$ GeV”. In: *Physics Letters B* 122.1 (1983), pp. 103–116. doi: `10.1016/0370-2693(83)91177-2`.

[5] UA2 Collaboration. “Observation of single isolated electrons of high transverse momentum in events with missing transverse energy at the CERN pp collider”. In: *Physics Letters B* 122.5–6 (1983), pp. 476–485. doi: `10.1016/0370-2693(83)91605-2`.

[6] ALEPH Collaboration et al. “Precision electroweak measurements on the Z resonance”. In: *Physics Reports* 427.5–6 (2006), pp. 257–454. doi: `10.1016/j.physrep.2005.12.006`.

[7] ALEPH Collaboration. “Determination of the number of light neutrino species”. In: *Physics Letters B* 231.4 (1989), pp. 519–529. doi: `10.1016/0370-2693(89)90704-1`.

[8] G. Baur et al. “Production of antihydrogen”. In: *Physics Letters B* 368.3 (1996), pp. 251–258. doi: `10.1016/0370-2693(96)00005-6`.

[9] LHCb Collaboration. “Observation of $J/\psi p$ Resonances Consistent with Pentaquark States in $\Lambda_b^0 \rightarrow J/\psi\ K^- p$ Decays”. In: *Physical Review Letters* 115.7 (2015), p. 072001. doi: `10.1103/PhysRevLett.115.072001`.

[10] CMS Collaboration. “Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC”. In: *Physics Letters B* 716.1 (2012), pp. 30–61. doi: `10.1016/j.physletb.2012.08.021`.

[11] ATLAS Collaboration. “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”. In: *Physics Letters B* 716.1 (2012), pp. 1–29. DOI: `10.1016/j.physletb.2012.08.020`.

[12] CERN. *The World Wide Web Project.* First website created at CERN by Tim Berners-Lee. 1990. URL: `https://info.cern.ch` (visited on 07/30/2025).

[13] *An Accelerator-based Research Infrastructure for Cancer Therapy and Biomedical Sciences with Ion Beams.* CERN-2024-004. CERN Yellow Reports: Monographs, 2024. DOI: `10.23731/CYRM-2024-004`. URL: `https://doi.org/10.23731/CYRM-2024-004`.

[14] Alan Paul Jeavons et al. “A high-resolution proportional chamber positron camera and its applications”. In: *IEEE Trans. Nucl. Sci.* 25 (1978), pp. 164–173. DOI: `10.1109/TNS.1978.4329298`. URL: `https://cds.cern.ch/record/132740`.

[15] C. Reymond. “Magnetic Resonance Techniques”. In: *Measurement and Alignment of Accelerator and Detector Magnets.* Ed. by G. Arduini et al. CERN Accelerator School. Presented by METROLAB Instruments SA, Geneva, Switzerland. 1998, pp. 219–231.

[16] M. Benedikt et al. *Future Circular Collider – European Strategy Update Documents.* Tech. rep. CERN-ACC-2019-0005. CERN, 2019. URL: `https://fcc-cdr.web.cern.ch/`.

[17] M. Benedikt and other. *Future Circular Collider Feasibility Study Report Volume 1.* Tech. rep. Submitted to the European Physics Journal ST, a joint publication of EDP Sciences, Springer Science+Business Media, and the Società Italiana di Fisica. Geneva: CERN, 2025. DOI: `10.17181/CERN.9DKX.TDH9`. arXiv: `2505.00272`. URL: `https://cds.cern.ch/record/2928193`.

[18] ATLAS Collaboration. “The ATLAS Experiment at the CERN Large Hadron Collider”. In: *Journal of Instrumentation* 3 (2008), S08003. DOI: `10.1088/1748-0221/3/08/S08003`.

[19] CMS Collaboration. “The CMS experiment at the CERN LHC”. In: *Journal of Instrumentation* 3 (2008), S08004. DOI: `10.1088/1748-0221/3/08/S08004`.

[20] ALICE Collaboration. “The ALICE experiment at the CERN LHC”. In: *Journal of Instrumentation* 3 (2008), S08002. DOI: `10.1088/1748-0221/3/08/S08002`.

[21] LHCb Collaboration. “The LHCb Detector at the LHC”. In: *Journal of Instrumentation* 3 (2008), S08005. DOI: `10.1088/1748-0221/3/08/S08005`.

[22] Oliver Brüning et al. “70 years at the high-energy frontier with the CERN accelerator complex”. In: *Nature Reviews Physics* 6 (2024), pp. 628–637. DOI: `10.1038/s42254-024-00758-5`.

[23] O. Aberle et al. *High-Luminosity Large Hadron Collider (HL-LHC): Technical design report.* Geneva, 2020. DOI: `10.23731/CYRM-2020-0010`. URL: `https://cds.cern.ch/record/2749422`.

[24] Oliver Brüning and Lucio Rossi. *The High Luminosity Large Hadron Collider*. Ed. by Oliver Brüning. Advanced series on directions in high energy physics. World Scientific, 2024. DOI: `10.1142/13487`. URL: `https://cds.cern.ch/record/2892725`.

[25] G. Apollinari et al. *High-Luminosity Large Hadron Collider (HL-LHC) Technical Design Report V.0.1*. Tech. rep. CERN-2017-007-M. CERN, 2017. DOI: `10.23731/CYRM-2017-004`. URL: `https://doi.org/10.23731/CYRM-2017-004`.

[26] F. Keizer. "Sub-nanosecond Cherenkov photon detection for LHCb particle identification in high-occupancy conditions and semiconductor tracking for muon scattering tomography". In: *CERN-THESIS-2019-203* (2020). DOI: `10.17863/CAM.45822`.

[27] C. D. Anderson. "The Positive Electron". In: *Physical Review* 43.6 (1933), p. 491. DOI: `10.1103/PhysRev.43.491`.

[28] A. D. Sakharov. "Violation of CP invariance, C asymmetry, and baryon asymmetry of the universe". In: *Soviet Physics Uspekhi* 34.5 (1991). DOI: `10.1070/PU1991v034n05ABEH002497`.

[29] Makoto Kobayashi and Toshihide Maskawa. "CP-Violation in the Renormalizable Theory of Weak Interaction". In: *Progress of Theoretical Physics* 49.2 (1973), pp. 652–657. DOI: `10.1143/PTP.49.652`.

[30] L. Wolfenstein. "Parametrization of the Kobayashi–Maskawa Matrix". In: *Physical Review Letters* 51 (1983), pp. 1945–1947. DOI: `10.1103/PhysRevLett.51.1945`.

[31] A. Bevan et al. "The Physics of the B Factories". In: *Springer Nature* (2017). DOI: `10.1140/epjc/s10052-014-3026-9`.

[32] Belle II Collaboration. "The Belle II Experiment". In: *Nuclear and Particle Physics Proceedings* 260 (2015), pp. 233–237. DOI: `10.1016/j.nuclphysbps.2015.02.049`.

[33] LHCb LHCb collaboration. *LHCb Upgrade II Scoping Document*. Tech. rep. Geneva: CERN, 2024. DOI: `10.17181/CERN.2RXP.HDK0`. URL: `https://cds.cern.ch/record/2903094`.

[34] S. Bifani et al. "Review of lepton universality tests in B decays". In: *Journal of Physics G: Nuclear and Particle Physics* 46 (2019), p. 023001. DOI: `10.1088/1361-6471/aaf5de`.

[35] L. Canetti, M. Drewes, and M. Shaposhnikov. "Matter and antimatter in the universe". In: *New Journal of Physics* 14 (2012), p. 095012. DOI: `10.1088/1367-2630/14/9/095012`.

[36] LHCb Collaboration. *Physics case for an LHCb Upgrade II: Opportunities in flavour physics, and beyond, in the HL-LHC era*. Tech. rep. CERN-LHCC-2018-027. CERN, 2018. URL: `https://arxiv.org/abs/1808.08865`.

[37] LHCb Collaboration. "LHCb Detector Performance". In: *International Journal of Modern Physics A* 30.07 (2015), p. 1530022. DOI: `10.1142/S0217751X15300227`.

[38] R. Aaij et al. "The LHCb Upgrade I". In: *Journal of Instrumentation* 19.05 (May 2024), P05065. ISSN: 1748-0221. DOI: `10.1088/1748-0221/19/05/P05065`. URL: `http://dx.doi.org/10.1088/1748-0221/19/05/P05065`.

[39] Placido Fernandez Declara et al. "A Parallel-Computing Algorithm for High-Energy Physics Particle Tracking and Decoding Using GPU Architectures". In: *IEEE Access* 7 (2019). 15 pages, 13 figures, pp. 91612–91626. DOI: `10.1109/ACCESS.2019.2927261`. arXiv: `2002.11529`. URL: `https://cds.cern.ch/record/2689507`.

[40] LHCb Collaboration. *LHCb VELO Upgrade Technical Design Report*. Tech. rep. 2013. DOI: `10.17181/CERN.4DGI.MZN4`. URL: `https://cds.cern.ch/record/1624070`.

[41] LHCb Collaboration. "LHCb VELO upgrade". In: *Nucl. Instrum. Meth. A* 845 (2017), pp. 97–100. DOI: `10.1016/j.nima.2016.04.077`.

[42] R. Ballabriga et al. "ASIC developments for radiation imaging applications: The Medipix and Timepix family". In: *Nucl. Instrum. Meth. A* 878 (2018), pp. 10–23. DOI: `10.1016/j.nima.2017.07.029`.

[43] A Fernández Prieto. "The LHCb VELO Upgrade II: design and development of the readout electronics". In: *JINST* 19.05 (2024), p. C05011. DOI: `10.1088/1748-0221/19/05/C05011`. URL: `https://cds.cern.ch/record/2914039`.

[44] X. Llopart et al. "Timepix4, a large area pixel detector readout chip which can be tiled on 4 sides providing sub-200 ps timestamp binning". In: *Journal of Instrumentation* 17.01 (Jan. 2022), p. C01044. DOI: `10.1088/1748-0221/17/01/C01044`. URL: `https://dx.doi.org/10.1088/1748-0221/17/01/C01044`.

[45] Sandro Cadeddu et al. "Timespot1: a 28 nm CMOS Pixel Read-Out ASIC for 4D Tracking at High Rates". In: *Journal of Instrumentation* 18.03 (Mar. 2023), P03034. ISSN: 1748-0221. DOI: `10.1088/1748-0221/18/03/p03034`. URL: `http://dx.doi.org/10.1088/1748-0221/18/03/P03034`.

[46] Plamen Hopchev. *SciFi: A large Scintillating Fibre Tracker for LHCb*. 2017. arXiv: `1710.08325 [physics.ins-det]`. URL: `https://arxiv.org/abs/1710.08325`.

[47] "SciFi Performance Plots July 2024". In: (2024). URL: `https://cds.cern.ch/record/2904659`.

[48] LHCb RICH Collaboration. "Performance of the LHCb RICH detector at the LHC". In: *European Physical Journal C* 73 (2013), p. 2431. DOI: `10.1140/epjc/s10052-013-2431-9`.

[49] R. Manera et al. "The analog front end for FastRICH: an ASIC for the LHCb RICH detector upgrade". In: *Journal of Instrumentation* 19 (2024), p. C04030. DOI: `10.1088/1748-0221/19/04/C04030`.

[50] C. D'Ambrosio and othres. "The future of RICH detectors through the light of the LHCb RICH". In: *Nim A* 876 (2017), pp. 194–197. DOI: `10.1016/j.nima.2017.02.076`.

[51] N. Harnew et al. "Status of the TORCH time-of-flight project". In: *Nuclear Instruments and Methods in Physics Research Section A* 924 (2019), pp. 44–48. DOI: `10.1016/j.nima.2018.12.007`.

[52] T. Gys et al. "Latest results from the TORCH R&D Project". In: *Nuclear Instruments and Methods in Physics Research Section A* 912 (2018), pp. 53–56. DOI: `10.1016/j.nima.2017.10.034`.

[53] N. H. Brook et al. "Testbeam studies of a TORCH prototype detector". In: *Nuclear Instruments and Methods in Physics Research Section A* 908 (2018), pp. 256–268. DOI: `10.1016/j.nima.2018.07.023`.

[54] LHCb Collaboration. *Framework TDR for the LHCb Upgrade II*. Tech. rep. Geneva: CERN, 2021. DOI: `10.17181/CERN.NTVH.Q21W`. URL: `https://cds.cern.ch/record/2776420`.

[55] S Barsuk et al. *Design and construction of electromagnetic calorimeter for LHCb experiment*. Tech. rep. Geneva: CERN, 2000. URL: `https://cds.cern.ch/record/691508`.

[56] R Djeliadine, O Iouchtchenko, and V F Obraztsov. *LHCb hadron trigger and Hcal cell size and length optimization*. Tech. rep. Geneva: CERN, 1999. URL: `https://cds.cern.ch/record/691688`.

[57] S Amato et al. *LHCb calorimeters*. Technical design report. LHCb. Geneva: CERN, 2000. URL: `https://cds.cern.ch/record/494264`.

[58] Aleksandar Bordelius. "The LHCb PicoCal". In: *Nucl. Instrum. Methods Phys. Res., A* 1079 (2025), p. 170608. DOI: `10.1016/j.nima.2025.170608`. URL: `https://cds.cern.ch/record/2933345`.

[59] "Performance of a spaghetti calorimeter prototype with tungsten absorber and garnet crystal fibres". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 1045 (2023), p. 167629. ISSN: 0168-9002. DOI: `https://doi.org/10.1016/j.nima.2022.167629`. URL: `https://www.sciencedirect.com/science/article/pii/S0168900222009214`.

[60] P R Barbosa-Marinho and Other. *LHCb muon system*. Technical design report. LHCb. Geneva: CERN, 2001. URL: `https://cds.cern.ch/record/504326`.

[61] J Albrecht. "The LHCb Trigger System". In: *Nucl. Phys. B, Proc. Suppl.* 187 (2009), pp. 237–244. DOI: `10.1016/j.nuclphysbps.2009.01.033`. URL: `https://cds.cern.ch/record/1245164`.

[62] The LHCb collaboration. *LHCb Trigger and Online Upgrade Technical Design Report*. Tech. rep. 2014. DOI: `10.17181/CERN.5F5X.FDJM`. URL: `https://cds.cern.ch/record/1701361`.

[63] J Albrecht et al. "The upgrade of the LHCb trigger system". In: *Journal of Instrumentation* 9.10 (Oct. 2014), p. C10026. DOI: `10.1088/1748-0221/9/10/C10026`. URL: `https://dx.doi.org/10.1088/1748-0221/9/10/C10026`.

[64] "RTA and DPA dataflow diagrams for Run 1, Run 2, and the upgraded LHCb detector ". In: (2020). URL: `https://cds.cern.ch/record/2730181`.

[65] J.P. Cachemiche et al. "The PCIe-based readout system for the LHCb experiment". In: *Journal of Instrumentation* 11.02 (Feb. 2016), P02013. DOI: `10.1088/1748-0221/11/02/P02013`. URL: `https://dx.doi.org/10.1088/1748-0221/11/02/P02013`.

[66] R. Aaij et al. "A comprehensive real-time analysis model at the LHCb experiment". In: *Journal of Instrumentation* 14.04 (Apr. 2019), P04006. DOI: `10.1088/1748-0221/14/04/P04006`. URL: `https://dx.doi.org/10.1088/1748-0221/14/04/P04006`.

[67] The LHCb collaboration. *LHCb Particle Identification Enhancement 3 Technical Design Report*. Technical Design Report CERN-LHCC-2023-005. Submitted to the LHCC. CERN, 2024. URL: `https://cds.cern.ch/record/2860194`.

[68] C. Gotti. "An ASIC for fast single photon counting in the LHCb RICH upgrade". In: *Journal of Instrumentation* 12 (2017), p. C03016. DOI: `10.1088/1748-0221/12/03/C03016`.

[69] S. Gómez et al. "FastIC: a fast integrated circuit for the readout of high performance detectors". In: *Journal of Instrumentation* 17.05 (May 2022), p. C05027. DOI: `10.1088/1748-0221/17/05/C05027`. URL: `https://dx.doi.org/10.1088/1748-0221/17/05/C05027`.

[70] S. Gómez et al. "FastIC: A Highly Configurable ASIC for Fast Timing Applications". In: *2021 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*. Piscataway, NJ, U.S.A.: IEEE, 2021, pp. 1–4. DOI: `10.1109/NSS/MIC44867.2021.9875546`.

[71] Hamamatsu Photonics. *Multi-anode photomultiplier tube R13742 Tentative Data Sheet*. Private communication. 2015.

[72] M. Calvi et al. "Single photon time resolution of photodetectors at high rate: Hamamatsu R13742 MaPMT and R10754 MCP-PMT". In: *Journal of Instrumentation* 15.10 (Oct. 2020), P10031–P10031. ISSN: 1748-0221. DOI: `10.1088/1748-0221/15/10/p10031`. URL: `http://dx.doi.org/10.1088/1748-0221/15/10/P10031`.

[73] Stefan Gundacker and Arjan Heering. "The silicon photomultiplier: fundamentals and applications of a modern solid-state photon detector". In: *Physics in Medicine & Biology* 65.17 (2020), 17TR01. DOI: `10.1088/1361-6560/ab7a27`.

[74] S. Cova et al. "Avalanche photodiodes and quenching circuits for single photon-detection". In: *Applied Optics* 35.12 (1996), pp. 1956–1976. DOI: `10.1364/AO.35.001956`.

[75] D. Renker. “Geiger-mode avalanche photodiodes, history, properties and problems”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 567.1 (2006). Proceedings of the 4th International Conference on New Developments in Photodetection, pp. 48–56. ISSN: 0168-9002. DOI: https://doi.org/10.1016/j.nima.2006.05.060. URL: https://www.sciencedirect.com/science/article/pii/S0168900206008680.

[76] Fabio Acerbi and Stefan Gundacker. “Understanding and simulating SiPMs”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 926 (2019). Silicon Photomultipliers: Technology, Characterisation and Applications, pp. 16–35. ISSN: 0168-9002. DOI: 10.1016/j.nima.2018.11.118. URL: https://www.sciencedirect.com/science/article/pii/S0168900218317704.

[77] E. Garutti and Yu. Musienko. “Radiation damage of SiPMs”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 926 (2019). Silicon Photomultipliers: Technology, Characterisation and Applications, pp. 69–84. ISSN: 0168-9002. DOI: https://doi.org/10.1016/j.nima.2018.10.191. URL: https://www.sciencedirect.com/science/article/pii/S0168900218315055.

[78] M. Calvi et al. “Single photon detection with SiPMs irradiated up to $10^{14}cm^21-MeV-equivalentneutronfluence$”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 922 (Apr. 2019), pp. 243–249. ISSN: 0168-9002. DOI: 10.1016/j.nima.2019.01.013. URL: http://dx.doi.org/10.1016/j.nima.2019.01.013.

[79] M. Calvi et al. “Photon detectors and front-end electronics for RICH detectors in high particle density environments”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 952 (2020). 10th International Workshop on Ring Imaging Cherenkov Detectors (RICH 2018), p. 161788. ISSN: 0168-9002. DOI: https://doi.org/10.1016/j.nima.2019.01.015. URL: https://www.sciencedirect.com/science/article/pii/S0168900219300488.

[80] F. Acerbi et al. “Cryogenic Characterization of FBK HD Near-UV Sensitive SiPMs”. In: *IEEE Transactions on Electron Devices* 64.2 (2016), pp. 521–526. DOI: 10.1109/TED.2016.2641586. arXiv: 1610.01915.

[81] Dania Consuegra Rodríguez et al. “Characterization of neutron-irradiated SiPMs down to liquid nitrogen temperature”. In: *The European Physical Journal C* 84.970 (2024). Regular Article - Experimental Physics. DOI: 10.1140/epjc/s10052-024-13302-7. URL: https://doi.org/10.1140/epjc/s10052-024-13302-7.

[82] G. Haefeli et al. "Microlens-Enhanced SiPMs". In: *IEEE Transactions on Nuclear Science* 72.4 (Apr. 2025), pp. 1594–1601. ISSN: 1558-1578. DOI: `10.1109/tns.2025.3542597`. URL: `http://dx.doi.org/10.1109/TNS.2025.3542597`.

[83] T. Gys. "Micro-channel plates and vacuum detectors". In: *Nuclear Instruments and Methods in Physics Research Section A* 787 (2015), pp. 254–260. DOI: `10.1016/j.nima.2014.12.044`.

[84] T. Gys et al. "Performance and lifetime of micro-channel plate tubes for the TORCH detector". In: *Nuclear Instruments and Methods in Physics Research Section A* 766 (2014), pp. 171–172. DOI: `10.1016/j.nima.2014.04.020`.

[85] L. Castillo Garcia. "Testing micro-channel plate detectors for the particle identification upgrade of LHCb". In: *Nuclear Instruments and Methods in Physics Research Section A* 695 (2012), pp. 398–402. DOI: `10.1016/j.nima.2011.12.021`.

[86] N.H. Brook et al. "Testbeam studies of a TORCH prototype detector". In: *Nuclear Instruments and Methods in Physics Research Section A* 908 (2018), pp. 256–268. DOI: `10.1016/j.nima.2018.07.023`.

[87] M. Fiorini et al. "Single-photon imaging detector with $\mathcal{O}(10)$ ps timing and sub-10 μm position resolution". In: *Journal of Instrumentation* 13 (2018), p. C12005. DOI: `10.1088/1748-0221/13/12/C12005`.

[88] M.J. Minot et al. "Large Area Picosecond Photodetector (LAPPD) - Pilot production and development status". In: *Nuclear Instruments and Methods in Physics Research Section A* (2018). DOI: `10.1016/j.nima.2018.11.137`.

[89] A. V. Lyashenko et al. "HRPPD photosensors for RICH detectors with a high resolution timing capability". In: (2025). arXiv: `2505.10658 [physics.ins-det]`. URL: `https://arxiv.org/abs/2505.10658`.

[90] E.V. Antamanova et al. "Anode current saturation of ALD-coated Planacon MCP-PMTs". In: *Journal of Instrumentation* 13 (2018), T09001. DOI: `10.1088/1748-0221/13/09/T09001`.

[91] M. Bartolini et al. "LHCb RICH Fast-timing photon detection at the SPS charged particle beam". In: *Journal of Instrumentation* 20.03 (Mar. 2025), P03034. DOI: `10.1088/1748-0221/20/03/P03034`. URL: `https://dx.doi.org/10.1088/1748-0221/20/03/P03034`.

[92] L. Cojocariu et al. "A multi-channel TDC-in-FPGA with 150 ps bins for time-resolved readout of Cherenkov photons". In: *Nucl. Instrum. Meth. A* 1055 (2023), p. 168483.

[93] Samuele Altruda et al. "PicoTDC: a flexible 64 channel TDC with picosecond resolution". In: *Journal of Instrumentation* 18.07 (July 2023), P07012. DOI: `10.1088/1748-0221/18/07/P07012`. URL: `https://dx.doi.org/10.1088/1748-0221/18/07/P07012`.

[94] lpGBT Team. *lpGBT Manual.* 2019. URL: `https://cds.cern.ch/record/2809058/files/lpGBT_manual.pdf`.

[95] VTRx+ Project Collaboration. *Versatile Link Plus (Technical Documentations).* 2019. URL: `https://edms.cern.ch/project/CERN-0000149832.EDMS-0000149832`.

[96] Steve Wotton. *LHCb RICH electronics for R&D – RICH-Tb DAQ.* Presentation at DRD4 WG3 meeting on electronics: kick-off. Accessed: 2025-10-16. 2024. URL: `https://indico.cern.ch/event/1416002/contributions/5959505/subcontributions/484389/attachments/2860908/5005328/RICHTbDaq.pdf`.

[97] M. Baszczyk et al. *LHCb RICH upgrade - Production Readiness Review report for the Elementary Cell of the Photo-Detector.* Technical Report LHCb-2016-TBD. `http://lhcb-project-rich-upgrade.web.cern.ch/PRR_ofTheEC/Soi2ECka6KB1uORKd/EC-RR-v2r5.pdf`. LHCb Collaboration, 2016.

[98] S. Gambetta. "First results from quality assurance testing of MaPMTs for the LHCb RICH upgrade". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 876 (2017). The 9th international workshop on Ring Imaging Cherenkov Detectors (RICH2016), pp. 206–208. ISSN: 0168-9002. DOI: `https://doi.org/10.1016/j.nima.2017.02.079`. URL: `https://www.sciencedirect.com/science/article/pii/S0168900217302917`.

[99] Federica Borgato. "LHCb RICH and VELO upgrade studies and measurement of the ratio of branching fractions $BR(\Lambda_b^0 \to \Lambda_c(2625, 2595)^{\pm} D_s^{(*)\mp})/BR(\Lambda_b^0 \to \Lambda_c^{\pm} D_s^{(*)\mp})$". PhD thesis. U. Padua (main), 2024.

[100] Active Technologies. *PG-1000 Pulse Generator.* `https://www.activetechnologies.it/products/signal-generators/pulse-generators/pg-1000/`. Accessed: 2025-09-07. 2025.

[101] CERN Repository. *High-Precision Illumination Setup for Single-Photon Characterisation.* Preliminary report available in the CERN Open Data Repository, record ID: h08ws-13m73. 2025. URL: `https://repository.cern/records/h08ws-13m73`.

[102] R. Calabrese et al. "Performance of the LHCb RICH detectors during LHC Run 2". In: *JINST* 17 (2022), P07013. DOI: `https://doi.org/10.1088/1748-0221/17/07/P07013`. arXiv: `2205.13400 [physics.ins-det]`.

[103] Johannes Albrecht et al. *Upgrade trigger & reconstruction strategy: 2017 milestone.* Geneva, 2018. URL: `https://cds.cern.ch/record/2310579`.

[104] R. Forty. "RICH pattern recognition for LHCb". In: *NIMA* 433.1 (1999), pp. 257–261. ISSN: 0168-9002. DOI: `10.1016/S0168-9002(99)00310-1`. URL: `https://www.sciencedirect.com/science/article/pii/S0168900299003101`.

[105] *The GAUSS project.* `http://lhcbdoc.web.cern.ch/lhcbdoc/gauss/`. Accessed: 02-07-2025.

[106] *The LHCb Geant4 simulation toolkit.* `http://lhcbdoc.web.cern.ch/lhcbdoc/geant4/`. Accessed: 02-07-2025.

[107] *The BOOLE project.* `http://lhcbdoc.web.cern.ch/lhcbdoc/boole/`. Accessed: 02-07-2025.

[108] IPC Association Connecting Electronics Industries. *IPC-2223C: Sectional Design Standard for Flexible Printed Boards.* ISBN 978-1-61193-126-0. IPC. 2016.

[109] DuPont Electronics & Industrial. *Design Guide: Flexible and Rigid-Flex Printed Circuits.* Available at: `https://www.dupont.com/flexcircuits.html`. 2022.

[110] H. Lee and K. Neville. *Handbook of Epoxy Resins.* New York: McGraw–Hill Book Company, 1967.

[111] DuPont Electronics & Industrial. *DuPont™ Kapton® Polyimide Film – Technical Data Sheet.* Available at: `https://www.dupont.com/kapton.html`. 2023.

[112] Samtec Inc. *HSEC8 Series – Edge Rate® High-Speed Edge Card Socket.* Technical Data Sheet, available at: `https://www.samtec.com/products/hsec8`. 2023.

[113] E. Hammerstad and O. Jensen. “Accurate Models for Microstrip Computer-Aided Design”. In: (1980), pp. 407–409. DOI: `10.1109/MWSYM.1980.1124303`.

[114] Hamamatsu Photonics K.K. *S13361-3050NE-08: 64-channel Silicon Photomultiplier Array.* Datasheet, available at: `https://www.hamamatsu.com/resources/pdf/ssd/s13361_series_kapd1052e.pdf`. 2023.

[115] F. Acerbi et al. “Analysis of single-photon time resolution of FBK silicon photomultipliers”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 787 (2015), pp. 34–37. DOI: `10.1016/j.nima.2014.11.013`.

[116] Ltd. Hirose Electric Co. *UMCC Ultra-Miniature Coaxial Connector Series.* Datasheet. 2020. URL: `https://www.hirose.com/en/product/group/UMCC/`.

[117] David M. Pozar. *Microwave Engineering.* 4th. John Wiley & Sons, 2011. ISBN: 978-0-470-63155-3.

[118] Caledonian Cables. *Coaxial Cable Technical Datasheet.* `https://www.caledonian-cables.co.uk/DdFls/coaxial%20cable.pdf`. Accessed: 2025-09-07. 2025.

[119] Keysight Technologies. *Quick Start Guide: 1169A Probe with E2675A Differential Head.* `https://www.keysight.com/us/en/assets/9018-05040/quick-start-guides/9018-05040.pdf`. Accessed: 2025-09-07. 2025.

[120] ExoSens. *MCP-PMT.* Accessed October 17, 2025. 2025. URL: `https://www.exosens.com/products/mcp-pmt`.