DISSERTATION

submitted to the
Combined Faculty of Mathematics, Engineering and Natural Sciences
of Heidelberg University, Germany

for the degree of
DOCTOR OF NATURAL SCIENCES

Put forward by
LUKAS WITOLA
born in Bonn

Oral examination: July 25, 2024

Commissioning, Calibration & Early Performance of the LHCb Scintillating Fibre Tracker

Referees: Prof. Dr. Ulrich Uwer
Prof. Dr. Hans-Christian Schultz-Coulon

Abstract

The LHCb detector has undergone a major upgrade for Run 3 of the LHC enabling it to operate at a five times higher instantaneous luminosity and to read out the entire detector at a frequency of 40 MHz, corresponding to the bunch crossing frequency of the LHC. The main tracking stations of the spectrometer were replaced by the Scintillating Fibre (SciFi) Tracker, a large, high granular tracking detector based on 250 µm scintillating fibres readout by arrays of silicon photomultipliers (SiPMs). The SiPM signals are processed and digitised by a custom ASIC at 40 MHz. Further digital electronics perform clustering and data compression before sending the data via optical links to the data acquisition system.
This thesis describes the LHCb SciFi Tracker and the commissioning of its front-end electronic. It highlights the calibration of the frontend ASICs' detection thresholds, a crucial task to ensure the detector's optimal performance. It is shown that, following the calibration of the detection thresholds, the detector achieves a hit detection efficiency of 99 % and a single hit resolution of <100 µm while maintaining a sufficiently low noise rate.

Zusammenfassung

Der LHCb-Detektor wurde für Run 3 des LHC einem umfassenden Upgrade unterzogen, um den Betrieb bei einer fünfmal höheren Luminosität und das Auslesen des gesamten Detektors mit einer Frequenz von 40 MHz, entsprechend der Kollisionsrate am LHC, zu ermöglichen. Die Hauptspurkammern des Spektrometers wurden durch den Scintillating Fibre (SciFi) Tracker ersetzt, einen großen, hochgranularen Detektor basierend auf 250 µm szintillierenden Fasern, die von Arrays von Silizium-Photomultipliern (SiPMs) ausgelesen werden. Die Signale der SiPMs werden von einem eigens dafür entwickeltem ASIC mit einer Frequenz von 40 MHz verarbeitet und digitalisiert. Weitere digitale Elektronik führt ein Clustering und Datenkompression durch, bevor die Daten über optische Verbindungen an das Datenerfassungssystem gesendet werden.
Die vorliegende Arbeit präsentiert den LHCb SciFi Tracker und die Inbetriebnahme seiner Frontend-Elektronik. Der Schwerpunkt liegt hierbei auf der Kalibrierung der Detektionsschwellen des Frontend-ASICs, welche eine entscheidende Aufgabe darstellt, um die optimale Funktionsweise des Detektors zu gewährleisten. Es wird gezeigt, dass der Detektor durch die Kalibrierung der Detektionsschwellen eine Nachweiseffizienz von 99 % und eine Ortsauflösung von <100 µm erreicht, während gleichzeitig eine ausreichend niedrige Rauschrate beibehalten wird.

Contents

Preface

The work presented in this dissertation was carried out within the SciFi Tracker working group within the LHCb collaboration. The collaboration currently consists of more than 1000 authors from 120 institutes worldwide. No part of this work would have been possible without the numerous contributions of other people; constructing, commissioning, and operating a detector on the scale of the LHCb SciFi Tracker is a collaborative effort of dozens of physicists, technicians, and engineers.

During the assembly of the SciFi Tracker at CERN, the author was, together with another student and under the supervision of a postdoctoral researcher from Heidelberg University, responsible for commissioning the frontend electronics. This includes the development of the measurement procedures and analysis methods, running and analysing the tests, and documenting the results in a dedicated database. These tests ensure the correct functionality of the complete signal chain from silicon photomultipliers to the data acquisition chain and also include the first calibration measurements. The author developed the software tool `ScurveFit` to analyse the calibration measurements and extract the calibration constants. Furthermore, the detector assembly was frequently assisted by installing detector modules or connecting and testing cables and optical fibres.

Following the installation of the SciFi Tracker in the LHCb experimental cavern, the author was instrumental in the commissioning of the entire detector, working alongside several members of the SciFi Tracker project and other LHCb collaboration members. The author's primary responsibility was the calibration of the detector, and they contributed to all steps of the calibration procedure. This included data collection using the LHCb experimental control and data acquisition system, analysis of the raw data using algorithms they implemented within the LHCb software framework, fitting of the calibration curves using their developed tool, and the storage of calibration constants in a dedicated database. Furthermore, the author supported the global data taking of the LHCb experiment, taking several week-long on-call shifts as the first point of contact for the SciFi Tracker in case of any issues.

The author, under the guidance of a postdoctoral fellow at CERN, conducted the measurements of the hit detection efficiency, single hit resolution, and noise rates. While the author was solely responsible for recording the various data sets and analysing them, it is important to note that this work was not possible without the contributions of several current and former members of the collaboration, as they utilised or were based on previously existing algorithms in the LHCb software framework.

1 | Introduction

The field of particle physics has made significant progress since its emergence in the early 20th century, thanks to a close interplay between theoretical ideas and precise experimental data. Our current state of knowledge is summarised in a well-tested theory called the Standard Model of Particle Physics [1–6], which describes all known fundamental particles and their interactions. The rapid progress the field has seen has often been made possible by employing novel ideas, concepts, and technologies in the design of particle detectors.

The first particle detectors employed photographic plates or scintillating screens [7] and used the human eye as a light detection device to uncover the nature of radioactivity in the early 20th century. Photographic emulsions [8], cloud chambers [9], and later bubble chambers [10] made particle trajectories visible, leading to the discoveries of the positron [11], the pion [12], and many new particles in the following decades. The development of the first electronic detectors started with the invention of the Geiger-Müller tube [13], based on the ionisation of a gas by a particle. This concept was further developed into photomultiplier tubes [14], significantly increasing the sensitivity of scintillation counters. With the development of the multi-wire proportional chamber in 1968 [15], electronic detectors were equipped with an increasing number of readout channels, finer granularity, and higher rate capability. This trend was only amplified by the advancements in semiconductor technologies and the development of the first silicon tracking detectors in the early 1980s [16]. The higher rate capabilities of these new technologies allowed for studying increasingly rare processes which, amongst other particles, led to discovery of the gluon [17–19], the carrier of the strong force, and the Z and $W^{\pm}$ bosons [20–22] which mediate the weak force.

Today's particle physics experiments employ various detector concepts and technologies to measure the properties of fundamental particles. Arguably, the most complex detector systems have been constructed and are operated at the Large Hadron Collider (LHC), where the four experiments ALICE, ATLAS, CMS, and LHCb investigate the properties of fundamental particles in collisions of highly energetic protons with unprecedented precision. In searching for exceedingly rare processes and higher precision to probe the Standard Model, the LHC experiments push the detector technologies to even finer granularities, higher rate capability, increased radiation tolerance, and larger data rates. At the forefront of this development is the LHCb experiment, which has recently completed an ambitious upgrade of its detector and readout system.

The LHCb experiment is designed to precisely reconstruct decays of heavy-flavour hadrons[1] and to shed light onto physics beyond the Standard Model by studying CP-violation [23], discovering new particles directly [24] or indirectly through their contribution to loop processes [25]. From 2010 to 2018, the experiment collected the world's largest sample of events containing b and c hadrons. Despite the unprecedented

---

[1] Hadrons are composite particles made of two or more quarks bound by the strong force.

size of the data set, many of the key flavour physics observables remain statistically limited [26] and require a significantly larger data set to reach the precision of today's theory calculations and sensitivity to possible new physics effects.
To collect this larger data set, the upgraded LHCb detector, described in Chapter 2, is operated at a five times higher instantaneous luminosity[2] and collects events at the LHC collision frequency of 40 MHz. The new running conditions required an upgrade of all detectors, readout electronics, and the entire data acquisition and trigger system to cope with the higher multiplicities, irradiation levels, and data rates. The main tracking stations have been replaced by the Scintillating Fibre (SciFi) Tracker described in Chapter 3. With its 250 µm diameter scintillating fibres, it provides a sufficiently small granularity to cope with the increased number of particles in each collision. The scintillation light is detected by arrays of silicon photomultipliers (SiPMs), and their signals are processed and digitised by a complex chain of readout electronics at a frequency of 40 MHz, matching the LHC bunch crossing frequency.

In the scope of this dissertation, significant contributions have been made to the commissioning and calibration of the readout electronics while supporting various steps during the construction and installation and participating in the operation of the detector. A well-documented quality assurance during construction and meticulous commissioning is the foundation of a future stable operation of a large and complex detector, as it allows for spotting potential issues early on and to gain first experience in operating the detector on a large scale. The commissioning of the frontend electronics during the assembly of the SciFi Tracker is summarised in Chapter 4 and first experiences in operating the detector are given in Chapter 5. The main work presented in this dissertation focuses on the calibration of the detector, an integral part of the commissioning and operation of every detector. Precise and careful calibration is necessary to convert digital values to meaningful information and to ensure the detector's best performance. The calibration of the SciFi frontend electronics, based on the pulse height spectra of SiPMs, is discussed in Chapter 6. In order to achieve optimal tracking performance, the SciFi Tracker must maintain a high hit detection efficiency and single hit resolution while minimising the occurrence of noise hits. Chapter 7 contains measurements of these metrics based on data from LHCb Run 3, using different settings derived from the previously mentioned calibration. Finally an outlook for the remainder of the data taking in Run 3 is given.

---

[2] Luminosity $\mathcal{L} = \frac{1}{\sigma}\frac{\mathrm{d}N}{\mathrm{d}t}$ is a measure of the rate of events ($\frac{\mathrm{d}N}{\mathrm{d}t}$) relative to their cross section (σ).

2 | The LHCb Experiment

The LHCb experiment is one of the four large experiments at the LHC at CERN. Designed initially to primarily study decays of hadrons containing b and c quarks, it has shown remarkable capabilities in various other fields, including but not limited to electroweak, heavy ion, and fixed target physics [27]. This broader physics programme is reflected in the design of the upgraded detector, which will be discussed in this chapter. The original detector operated successfully from 2009 to 2018 and is described in detail in References [28, 29].

2.1 The Large Hadron Collider

The LHC, located at the border between France and Switzerland near Geneva, is the world's largest and most powerful particle accelerator. It consists of two superconducting storage rings with a circumference of 26.7 km and is designed to accelerate and collide protons up to a centre-of-mass energy of 14 TeV and an instantaneous luminosity in the order of $2 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ [30]. The two counter-rotating proton beams comprise up to 2800 bunches, each containing about 1.8×10^{11} protons [31]. The nominal bunch spacing is 25 ns, resulting in a collision rate of 40 MHz. As shown in Figure 2.1, the bunches are collided at four points along the accelerator where the large experiments ALICE, ATLAS, CMS, and LHCb are located.
The data-taking campaigns at the LHC are divided into Runs, the first two of which took place from 2009 to 2013 and 2015 to 2018. Run 3 has started in 2022 and is scheduled to continue until the end of 2025 [32]. During Run 3, the LHC collides protons with a centre-of-mass energy of $\sqrt{s} = 13.6\,\mathrm{TeV}$ and a peak instantaneous luminosity of $2 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. While ATLAS and CMS use the total but continuously decaying luminosity delivered by the LHC, LHCb uses luminosity levelling [33] to keep the luminosity lower but constant throughout a fill[1]. In Run 3, the LHCb detector is operated at a five times higher instantaneous luminosity compared to Run1 and Run2 of $\mathcal{L} = 2 \times 10^{33}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$, which results in about five proton-proton (pp) collisions on average per bunch crossing [34]. In special fills, the LHC also delivers collisions of heavy ions, such as lead and xenon, as well as proton-ion collisions.

2.2 The LHCb Detector

The LHCb detector, shown in Figure 2.2, is a single-arm forward spectrometer covering angles relative to the beamline of 10 mrad to 300 mrad in the x-z plane and up

---

[1] A fill refers to the period where protons are injected into the LHC, accelerated to the desired energy, and then collided until the beams are dumped, which usually occurs after several hours.

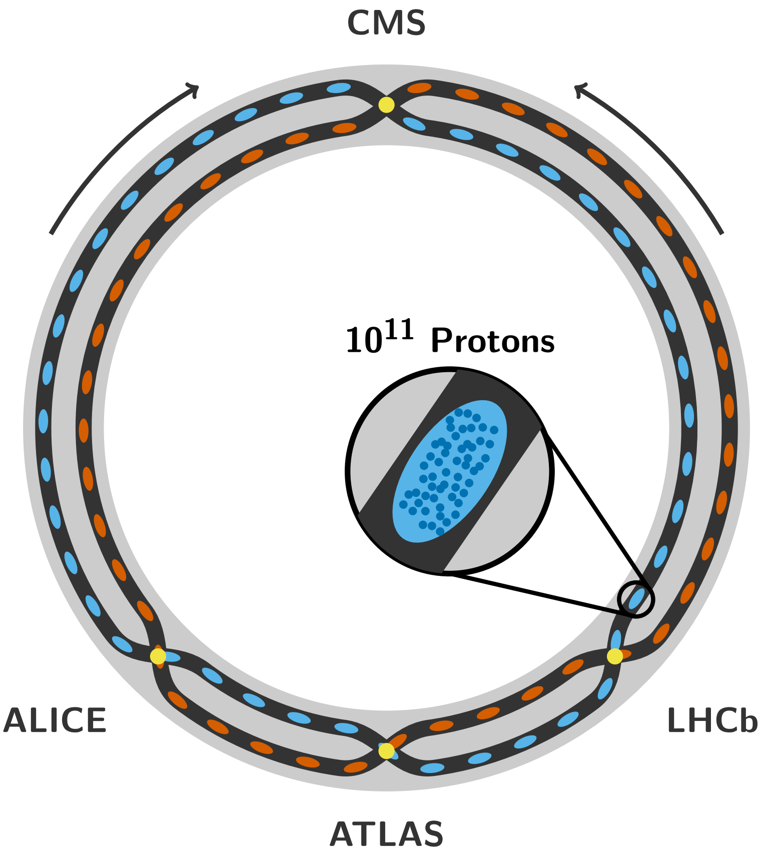

Figure 2.1 Schematic view of the LHC accelerator ring with its two beam pipes and the four crossing points with their experiments. The two beams are made up of up to 3546 individual bunches, each containing ≈10^{11} protons. Adapted from Reference [35].

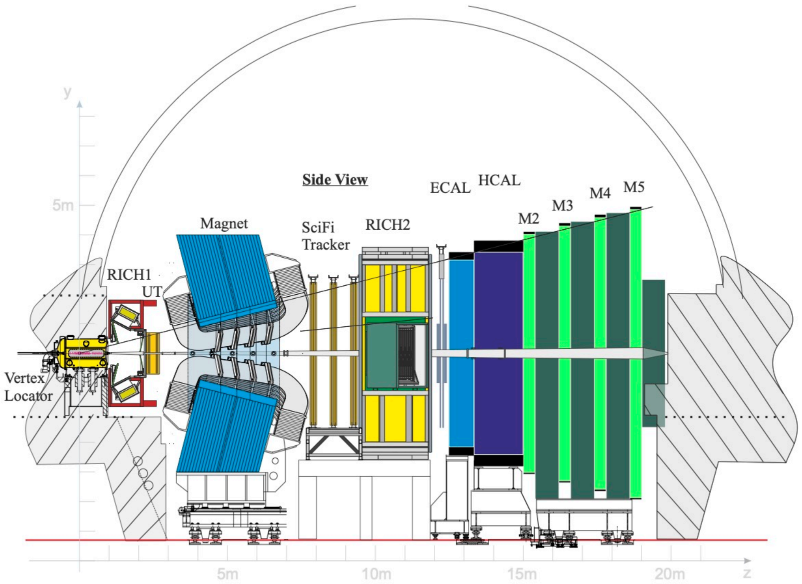

Figure 2.2 Layout of the LHCb detector. Particles collide inside o the Vertex Locator and their collission products are detected in the forward region by the other sub-detectors [34].

to 250 mrad in the y-z plane, corresponding to a pseudorapidity[2] range of $2 < \eta < 5$ [28]. The right-handed coordinate system has its origin at the nominal pp interaction point inside the Vertex Locator (VELO), the z-axis pointing along the beamline in the direction of the muon stations, and the y-axis pointing vertically upwards. The geometry of the LHCb detector is unique among the LHC experiments and was chosen because, at high energies, pairs of $b\bar{b}$ quarks are predominantly produced with small opening angles relative to the beam axis [36]. Due to the dominant production mechanism and the low mass of the b quark compared to the centre-of-mass energy, the $b\bar{b}$ pair is highly boosted with respect to the laboratory frame. The resulting hadrons, for example B mesons, will therefore typically fly for about a centimetre from their point of production before they decay, leading to a distinctive displaced vertex. Separating this secondary vertex from the primary vertex of the pp collision is a key feature of LHCb and the purpose of the first sub-detector, the VELO. The VELO, together with the dipole magnet, the Upstream Tracker (UT), and the SciFi Tracker, forms the tracking and spectrometer system of LHCb responsible for precisely reconstructing charged particle trajectories and measuring their momenta, it is described in detail in Section 2.3. Particle identification is provided by a set of ring-imaging Cherenkov detectors (RICH1, RICH2), the calorimeter system (ECAL, HCAL), and the muon stations (M2-M5), which are described in Section 2.4. The data acquisition system (DAQ) and the software-based trigger are both described in Section 2.5

---

[2] Pseudorapidity describes the angle of a particle relative to the beam axis. It is defined as $\eta = -\ln\left[\tan\left(\frac{\theta}{2}\right)\right]$.

2.3 Tracking System

The tracking system is at the heart of LHCb and consists of the VELO around the interaction region, the UT just before the magnet, the dipole magnet, and the main tracking detector, the SciFi Tracker, after the magnet. The magnet provides a vertical magnetic field along the y-axis and has a bending power of 4 T m [37]. By measuring the trajectories of charged particles and their deflection in the magnetic field, their momentum can be estimated. Each sub-detector has been specifically designed for the conditions during Run 3 of the LHC.

2.3.1 Vertex Locator

The Vertex Locator, or VELO [38], is a silicon pixel detector that surrounds the interaction region at LHCb. It is designed to provide an excellent vertex resolution to separate the primary pp vertex from displaced vertices in the collision. The detector consists of 52 modules arranged in two moveable halves as shown in Figure 2.3 that can be closed to a radius of 5.1 mm around the beam. In order to sit so close to the interaction point, the VELO is closely integrated into the LHC vacuum system and only separated from the primary vacuum by a 150 μm thin corrugated aluminium foil.

The L-shaped modules are composed of four sensor tiles, each of them read out by three VeloPix ASICs [39], as shown in Figure 2.3b. Each ASIC has an active matrix of 256 × 256 pixels with a pixel size of 55 μm × 55 μm leading to a total of over 41 million pixels for the entire VELO. The sensors and ASICs are designed to withstand high particle fluxes of up to 900×10^6 particles per second in the busiest region and tolerate a high integrated ionising radiation dose of 4 MGy, as well as a total non-ionising[3] dose of $8 \times 10^{15}\,\mathrm{n_{eq}\,cm^{-2}}$. Pixel hits are grouped into clusters by the VeloPix and encoded into 30 bit giving a peak data rate of 15.1 Gbit/s for the most central ASIC and up to 2.85 Tbit/s for the whole VELO [34].

2.3.2 Upstream Tracker

The Upstream Tracker (UT) [40] is located just in front of the dipole magnet. Being so close to the magnet, it experiences a significant magnetic field, which is exploited to get a first momentum estimate from matching VELO tracks to hits in UT. This information is used to speed up the matching of VELO tracks to SciFi Tracker hits. The UT also plays a vital role in reconstructing tracks from particles decaying after the VELO, *e.g.* long-lived kaons and lambda particles.

The detector consists of four layers of silicon strip sensors, as illustrated in Figure 2.4, arranged in a so-called *x-u-v-x* geometry, where the two innermost layers are rotated by +5° and −5°, respectively, to provide a measurement of the y coordinate. Each layer is composed of several staves that house the sensors and readout electronics. Depending on the expected occupancy, sensors with different granularities are mounted throughout the detector to keep the average occupancy below 1 %. Most of the detector is

---

[3] The non-ionising radiation dose is typically expressed as *neutron equivalent* dose ($\mathrm{n_{eq}\,cm^{-2}}$), which is defined as the dose of 1 MeV neutrons producing the same radiation damage as induced by an arbitrary particle dose with a specific energy distribution.

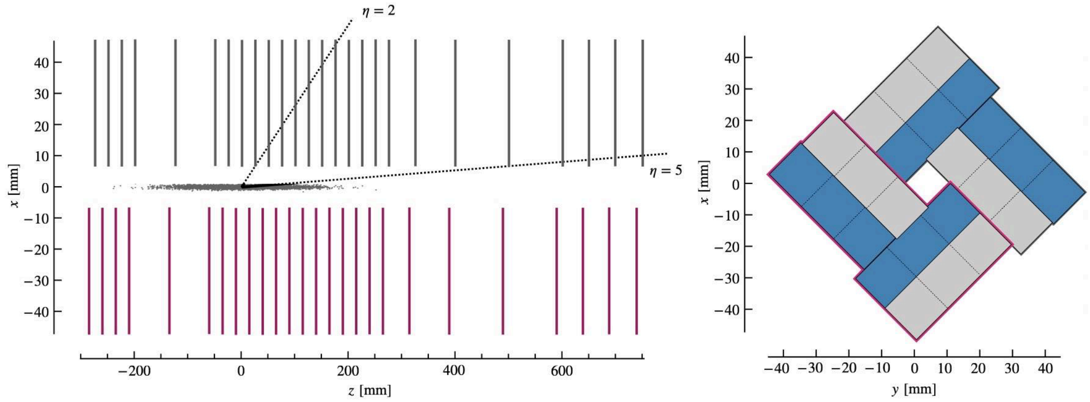

a

b

Figure 2.3 Left: top view of the VELO in the *z*-*x* plane showing the 26 modules on each side. In addition the luminous region and the nominal pseudorapidity acceptance of LHCb are shown. Right: sketch of the arrangement of ASICs around the *z* axis for the closed VELO. Half of the ASICs are located on the front side (grey) and half on the back side (blue) [34].

equipped with sensors with a strip pitch of 187.5 µm and strip length of 99.5 mm, while the central region uses two different sensor types with a pitch of 93.5 µm but different strip length of 99.5 mm and 51.45 mm for the innermost sensors [34]. The sensors are read out by up to eight custom frontend ASICs called SALT [41] that digitise the signals at 40 MHz. Both the sensors and SALT ASICs are designed to withstand an integrated non-ionising dose of up to $4 \times 10^{14}\,\mathrm{n_{eq}\,cm^{-2}}$ [34].

2.3.3 Scintillating Fibre Tracker

The last detector in LHCb's tracking system is the Scintillating Fibre (SciFi) Tracker [40], located behind the dipole magnet. It is responsible for charged particle tracking and provides the best momentum estimate by extending tracks from the VELO or UT and measuring their bending in the magnetic field. The SciFi Tracker, shown in Figure 2.5, consists of 12 layers arranged in three stations (T1, T2, T3) with the four layers of each station following the same x-u-v-x geometry as the UT where the stereo layers (u, v) are rotated by $\pm 5°$. Each layer covers an area of about $30\,\mathrm{m^2}$ and is built from 250 µm diameter scintillating fibres arranged in multilayer fibre mats. The scintillation light is detected by silicon photomultipliers (SiPMs) located at the top and bottom of the detector, outside of the acceptance of LHCb. The SiPM signals are processed and digitised by custom frontend electronics at the LHC bunch crossing frequency of 40 MHz leading to a total data rate of about 20 Tbit/s. Over the lifetime of the detector, the fibres are expected to see an integrated ionising dose of 35 kGy and the SiPMs a collected non-ionising fluence of $6 \times 10^{10}\,\mathrm{n_{eq}\,cm^{-2}}$ while still maintaining the required hit efficiency of 99 % and a hit resolution better than 100 µm [34]. As the SciFi Tracker, and especially its readout electronics, are the subject of this dissertation, a comprehensive overview is given in Chapter 3.

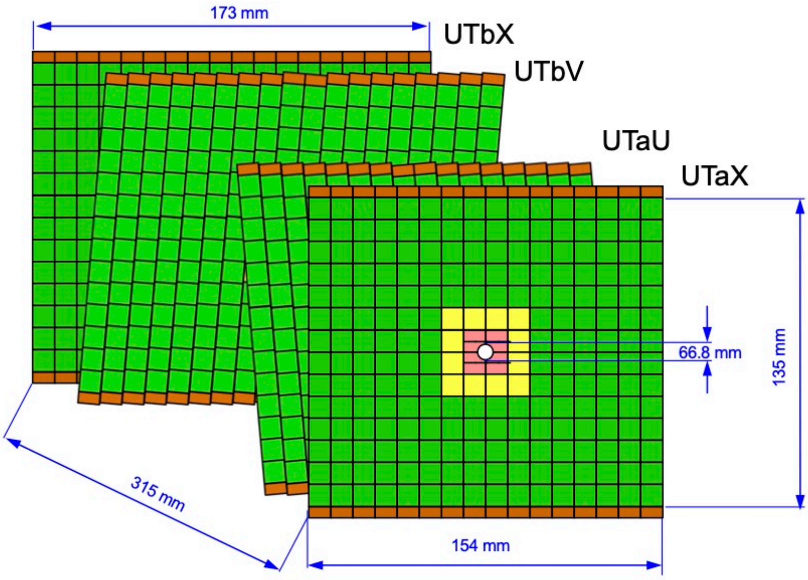

Figure 2.4 Schematic of the four UT planes showing the rotation of the inner two layers. The colours indicate the different sensor types utilising different strip pitches and lengths [34].

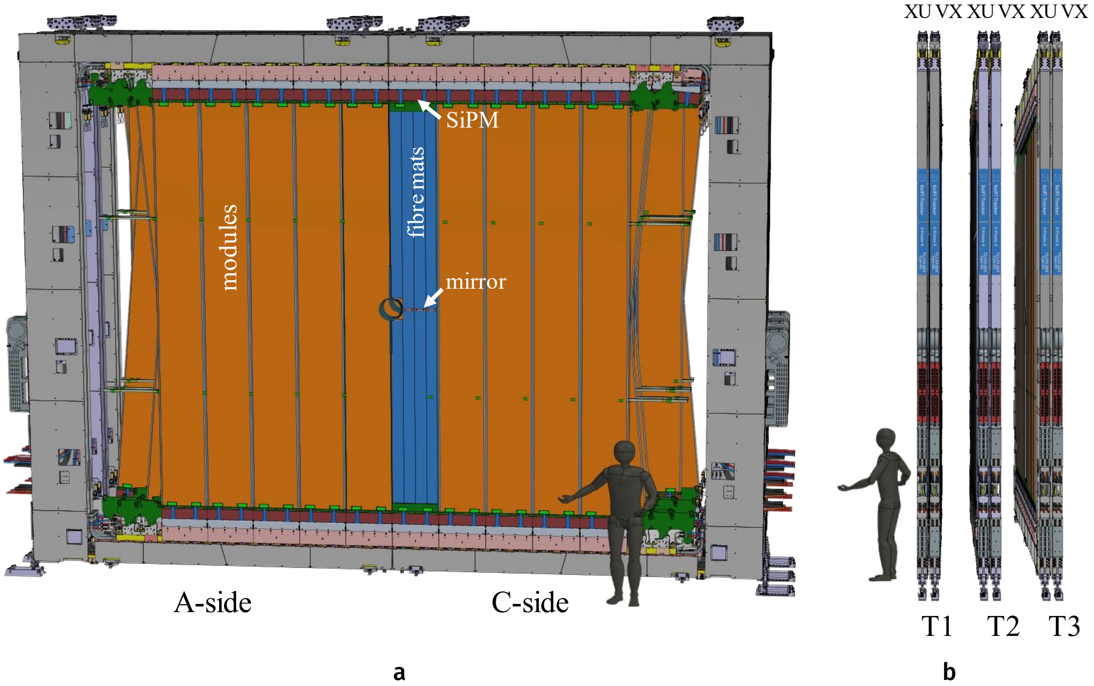

Figure 2.5 Illustration of the SciFi Tracker [34]. Left: front view of the SciFi Tracker. The module are shown in orange and the fibre mats of one module are highlighted in blue. Right: side view of the three detector stations.

2.4 Particle Identification

The particle identification (PID) system provides crucial information to suppress combinatoric backgrounds and control misidentification rates in many LHCb analyses. The two ring-imaging Cherenkov detectors (RICH1, RICH2) allow for distinguishing between pions, kaons, and protons, while photons and electrons are identified in the calorimeter system (ECAL, HCAL); lastly, muons are detected in the muon stations (M2-M5) [34, 42].

2.4.1 Ring-imaging Cherenkov Detectors

The two RICH detectors provide charged hadron discrimination between pions, kaons, and protons over a momentum range of 2.6 to 100 GeV/c. Both detectors use fluorocarbon gas radiators but with different refractive indices. Charged particles traversing the radiator emit Cherenkov photons under a characteristic angle depending on their mass, momentum, and the refractive index of the radiator. The Cherenkov light is reflected by a set of flat and spherical mirrors, focusing the ring images onto planes of photon detectors. By reconstructing the opening angle from the ring images and combining it with the momentum from a track, particles with different masses can be distinguished by their Cherenkov angle.
The first RICH detector, RICH1, is located before the dipole magnet and uses a C_4F_{10} gas radiator allowing for PID in the momentum range of 2.6 to 60 GeV c^{-1}, while RICH2, located after the magnet behind the SciFi Tracker, uses CF_4 covering the higher momentum range of 15 to 100 GeV c^{-1}. The Cherenkov photons are detected using multi-anode photomultiplier tubes (MaPMTs) of different granularities depending on the occupancy requirements. RICH1 and the inner region of RICH2 are equipped with MaPMTs with a pixel size of 2.88 mm × 2.88 mm while the outer areas of RICH2 use devices with a pixel size of 6 mm × 6 mm [34]. A custom ASIC called CLARO [43] has been designed that provides single photon counting at rates of up to 10^7 hits per second and can tolerate an integrated non-ionising dose of $10^{13}\,n_{eq}\,cm^{-2}$.

2.4.2 Calorimeter System

The electromagnetic calorimeter (ECAL) is located just behind RICH2 and is used to identify electrons and photons. Electrons are identified by matching clusters in the ECAL to reconstructed tracks, which also take into account the emission of bremsstrahlung photons. Consequently, photons are identified by clusters in the ECAL without an associated track. The hadronic calorimeter (HCAL) sits directly behind the ECAL, providing additional PID information on hadrons.
Both the detectors employ the same basic principle of alternating layers of plastic scintillator and absorber tiles with wave-length shifting fibres running through to collect and transmit the scintillation to photomultiplier tubes (PMTs). In the case of the ECAL lead with a total of 25 radiation lengths is used as an absorber, while the HCAL uses iron with a total of 5.6 interaction lengths. The detectors also share the same readout electronics that process and digitise the PMT signals at a rate of 40 MHz [34, 42].

2.4.3 Muon Stations

The four muon stations M2 to M5 are the last detectors in LHCb and provide identification of muons, which are essential in many LHCb analyses. At the LHC's energies, muons can easily traverse the entire detector as they do not feel the strong interaction and do not develop electromagnetic showers as bremsstrahlung is suppressed compared to electrons. Consequently, muons are identified by filtering out all other particles using thick layers of absorbers and matching the muon hits to reconstructed tracks.

The muon stations are composed of multi-wire proportional chambers (MWPCs) and interleaved with 80 cm of iron absorber with a total interaction length of about 20. The stations are divided into four regions, each one equipped with MWPCs of a different granularity to cope with the large variation in particle rates from the centre to the edge of the detector [34].

2.5 Data Acquisition and Trigger System

The LHCb data acquisition system (DAQ), shown in Figure 2.6, uses a synchronous readout where all sub-detectors send their data every bunch crossing at a frequency of 40 MHz, making the system essentially trigger-less [34]. Data is sent from the detector frontends via optical links from the experimental cavern to the data centre at the surface. The data centre hosts the event builder (EB) servers that contain the data acquisition boards (TELL40) and the graphics processing units (GPU) running the high-level trigger application (HLT1). Each TELL40 is connected only to a single sub-detector and is therefore configured with a specific firmware that decodes the data, orders it by bunch crossing, and creates a sub-detector-specific event fragment, before sending it to the network connecting all EBs. The EBs assemble all sub-detector data fragments and, once completed, push the events to the HLT1 GPUs. Events that are accepted are stored on a buffer before being picked up by the second high-level trigger (HLT2) running on general-purpose CPU servers called event-filter farm. Data processed by HLT2 is subsequently sent to permanent storage.

At the nominal luminosity in Run 3 data rates can reach up to 4 TB/s, whereas only about 10 GB/s can be recorded to permanent storage [34]. In addition, the rates of potentially interesting events involving decays of hadrons containing b and c quarks are in the order of 300 kHz and 1 MHz, respectively. To reduce the data rate by a factor of 400 and maximise the number of recorded signal events, a full offline-quality reconstruction of the events needs to be performed. Subsequently, only a subset of the event information will be saved while discarding the rest. This *real-time analysis* [44] was already pioneered during the Run 2 data taking of LHCb [45] and was developed into a two-stage software trigger system for Run 3, consisting of HLT1 and HLT2. HLT1 performs an inclusive selection based on charged particle reconstruction, reducing the data rate by a factor of 20. The data is temporarily stored on a large disk buffer while the real-time alignment and calibration are performed [46]. The data is then picked up by HLT2, where the full offline-quality reconstruction and final selection of signals is performed.

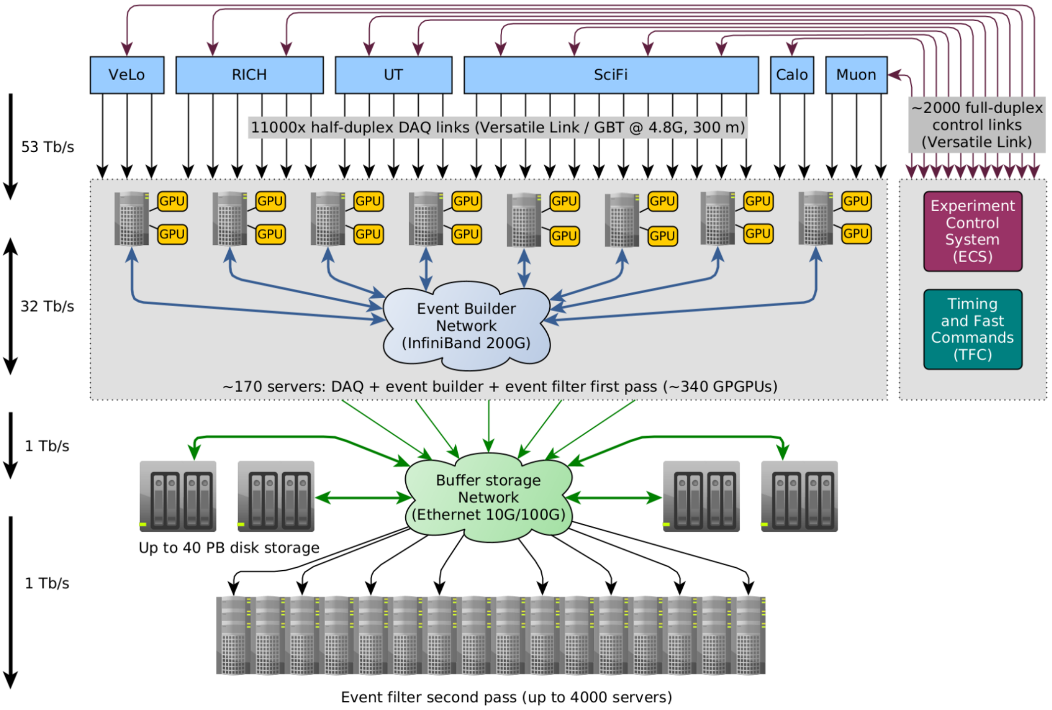

Figure 2.6 Schematic of the LHCb DAQ system. The detector frontends are located on the top. The event builder nodes and the event filter farm are separated by a large buffer. Adapted from Reference [34].

2.5.1 Readout & Control Boards

The backend electronics of the DAQ system are based on a common custom board called PCIe40 [47]. The board is based on an FPGA[4] and features 48 bidirectional optical links with a bandwidth of up to 10 Gbit/s connected to MiniPod[5] transceivers. The board's functionality is configured by flashing the device with one of three different firmware types: SODIN, TELL40, and SOL40. The readout supervisor SODIN is responsible for controlling and synchronising the entire DAQ system by distributing the 40 MHz LHC clock as well as timing and trigger information. The PCIe40 cards configured as TELL40 receive the data from the frontend electronics. They are responsible for decoding it and merging events from the same bunch crossing before shipping the data to the DAQ network. The TELL40 firmware, therefore, also contains logic that is specific to each sub-detector. The SOL40 boards act as interface boards between the various components of the DAQ system. They forward the clock signal, timing and fast control (TFC) commands from the SODIN to all frontend electronics and readout boards via optical links, establishing a synchronous readout of the whole detector. They also provide communication between the experiment control system and the frontend electronics by relaying the slow control information.

2.5.2 Experiment Control System

The experiment control system (ECS) [48] is the central interface between the operator of the detector and the various subsystems. It allows to configure, monitor, control, and operate all aspects of the experiment. This includes the high and low voltages, cooling systems, the DAQ with the frontend electronics and backend TELL40s boards, the TFC system consisting of SODIN and SOL40 boards, the high-level triggers HLT1 and HLT2, as well as different safety systems. The system builds on the Joint Controls Project (JCOP) [49] that defines a common architecture and framework for the four LHC experiments. Within the architecture, the different subsystems are organised in a tree-like structure where the nodes can either be logical systems or actual equipment, allowing for the concurrent use of different systems during commissioning tests or calibration phases.

---

[4] Intel Arria10. Intel Corporation, Santa Clara, United States.

[5] Broadcom MiniPod. Broadcom Inc., Palo Alto, United States.

3 | The LHCb Scintillating Fibre Tracker

Scintillating fibres have been used in tracking detectors for several decades as they allow the instrumentation of large areas with intrinsically fast, low-mass detectors in various geometries[50]. However, only in the past decade, with the advent of silicon photomultipliers (SiPMs) in combination with small-diameter scintillating fibres, new interest has been sparked in the technology and led to the development of the LHCb SciFi Tracker.

The Scintillating Fibre (SciFi) Tracker is the new main tracking detector of LHCb for Run 3 and beyond. It replaces the former hybrid detector consisting of the Inner Tracker (IT) [51] and the Outer Tracker (OT) [52]. The IT was a cross-shaped silicon microstrip detector that covered the innermost area of the tracking stations and provided a spatial resolution of 50 µm and a hit efficiency of greater than 99 %. The OT covered most of the 30 m^2 detector area of each layer and was constructed from 2.4 m long gas straw tubes with a diameter of 4.9 mm able to measure hits with a resolution better than 200 µm and hit detection efficiency above 99 % [53]. Initially designed for a peak instantaneous luminosity of $2 \times 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ with occupancies of 10 %, it was later successfully operated at a luminosity of $5 \times 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ leading to an occupancy of 25 %. However, with the further increase in luminosity for Run 3, it was shown that the OT performance would degrade, and some modules would need to be replaced with detectors of higher granularity [26]. In addition, the switch to a trigger-less 40 MHz readout would have required the replacement of all frontend electronics of the IT and OT. Several possible scenarios and technologies were studied, including increasing the area of the IT along with shorter OT modules. However, ultimately, the decision was made for a scintillating fibre tracker as a cost-effective, lightweight, and uniform detector. This chapter gives an overview of the design and working principle of the LHCb SciFi Tracker.

3.1 Performance Requirements

Like any detector, the SciFi Tracker must meet various performance requirements while considering geometric and budgetary constraints. To ensure an excellent tracking performance, the detector must have a hit detection efficiency of over 99 % and a spatial resolution of less than 100 µm in the bending plane of the dipole magnet. The detection layers must have an area of approximately 6 m × 5 m to cover the full acceptance and provide mechanical rigidity to ensure that the detector elements are stable within 50 µm or 300 µm along the x- and z-axis, respectively. The detector elements must also be straight along their length within a precision of 50 µm [40].

The detector granularity must be small enough to keep the occupancy low and avoid compromising hit efficiency at an instantaneous luminosity of $2 \times 10^{33}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. To limit multiple scattering and secondary particle production, the material budget in the

acceptance should not exceed 1 % per layer [40].
The readout electronics must comply with the specifications of the new DAQ system and operate at a frequency of 40 MHz without dead time [40].
Finally, the required performance must be maintained for the lifetime of the detector, during which an integrated luminosity of 50 fb^{-1} is expected to be collected [40].

3.2 Detector Design

The SciFi Tracker consists of three stations (T1, T2, T3) located between the dipole magnet and RICH2 as shown in Figure 2.2. Each of the three stations is composed of four detection layers instrumented with a single detector technology based on 250 µm diameter plastic scintillating fibres arranged in multilayer fibre mats. The detector covers an acceptance ranging from 20 mm from the beam pipe up to distances of ± 3186 mm in the horizontal and ± 2425 mm vertical directions. The layers are arranged in a x-u-v-x geometry with the u and v layers rotated by $+5°$ and $-5°$, respectively, as shown in Figure 2.5. The x layers have their fibres oriented vertically to provide measurements of the deflection of particles in the bending plane of the magnet, while the inner two stereo layers, u and v, are used to reconstruct the vertical position of a track hit.
Each station comprises four individually moveable structures called C-Frames, with two C-Frames on either side of the beam pipe. The C-Frames are equipped with approximately 52 cm wide identical modules spanning the full height of the detector, except for the innermost modules, which feature a cut-out for the beam pipe. The C-Frames in station T3 are equipped with six modules, whereas in T1 and T2, only five modules are needed to cover a similar angular acceptance.
The detector modules are constructed using a sandwich of honeycomb and carbon fibre and contain eight 2424 mm long and 130.4 mm wide fibre mats made from six layers of staggered fibres. A mirror is glued to the fibre ends at one end of the fibre mat to reflect light on the readout side. The mats are oriented such that the mirror is located at $y = 0$ and the readout side points upwards or downwards. A more detailed description of the fibre mats and modules is given in Section 3.4.
The light emitted by the scintillating fibres is read out by arrays of SiPMs with 128 channels with a pitch of 250 µm, and their signals are processed by a chain of custom frontend electronics at a rate of 40 MHz. The analogue signals of each SiPM channel are first processed and digitised by a custom ASIC called PACIFIC and later grouped into clusters on a dedicated FPGA before being sent out to the DAQ system. The SiPMs are discussed in Section 3.5, while the frontend electronics are described in detail in Section 3.6.
The signal generation in the SciFi Tracker is shown in Figure 3.1. A charged particle traversing a fibre mat produces scintillation photons (black dots) in each fibre along its trajectory, which are detected by the SiPM arrays (yellow squares) after propagation through the fibre. As the fibres and SiPM channels are not aligned, multiple channels typically detect the light. By clustering these channels, the location at which the particle crossed the fibre mat can be reconstructed from the barycentre of the cluster.

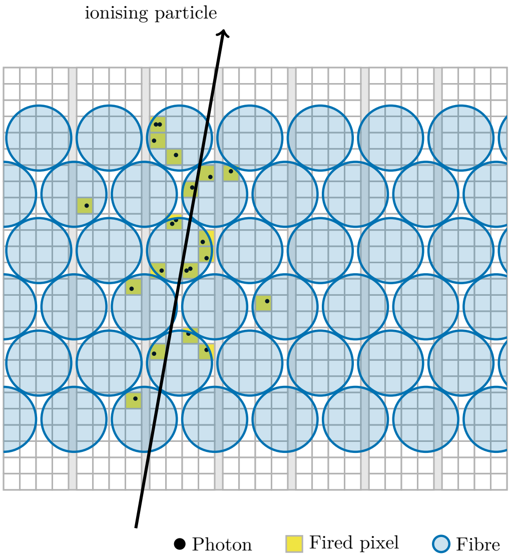

Figure 3.1 Illustration of the detection principle of the SciFi Tracker. The six layers of scintillating fibres are shown in front and the SiPM channels and pixels in the back. The scintillating photons created by the traversing particle are shown as black dots and the subsequently fired SiPM pixels as yellow squares. Adapted from Reference [40].

3.3 Scintillating Fibres

The SciFi Tracker uses double-clad blue-green emitting plastic scintillating fibres[1] with a circular cross-section and a diameter of 250 µm. The fibres were chosen due to their fast decay time of 2.4 ns and long attenuation length of 3.5 m [54].

When a charged particle traverses the fibre, it loses energy due to ionisation; typically, an energy deposit of a few eV is needed to produce an excitation in the polystyrene base polymer. However, polystyrene has a poor quantum efficiency and long relaxation time, which is why a first dopant, an organic fluorescent dye (p-Terphenyl [40]) with matched excitation energy levels and high quantum efficiency of >95 %, is added. Energy is transferred from the polystyrene base to the primary dye predominantly via radiation-less dipole-dipole interactions (Förster transitions), where the excited energy state of the dye will quickly relax (less than a few ns) by emission of a photon [55]. In order to increase the transparency of the fibre, a secondary dye (tetraphenylbutadiene, TPB [40]) acting as a wavelength shifter is added. It absorbs the excitation of the primary dye, either radiative or non-radiative, and emits photons at a longer wavelength such that re-absorption in the fibre is less likely [55]. The resulting photon emission spectrum peaks in the range of 450 to 500 nm.

A schematic of the fibre is shown in Figure 3.2. The fibre's core is made of doped polystyrene surrounded by two claddings of decreasing indices of refraction. The inner cladding is made of Polymethylmethacrylate (PMMA, $n = 1.49$) and the outer cladding of a fluorinated polymer (FP, $n = 1.59$). The scintillation light is transported through the fibre via total internal reflection at the interface between the fibre core

---

[1] Kuraray SCSF-78MJ. Kuraray Co., Ltd., Tokyo, Japan.

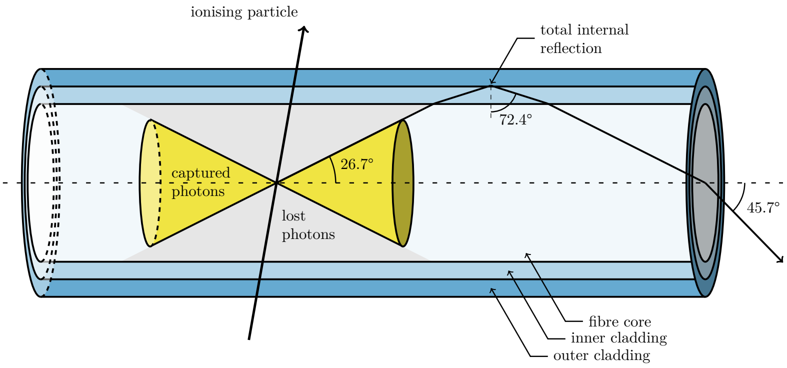

Figure 3.2 Sketch of a scintillating fibre with two claddings. The emission of light due to the ionising particle is shown as yellow cone. In addition the critical angles for the propagation of light through fibre due to total internal reflection are shown. Adapted from Reference [55].

and the cladding system. Total internal reflection at the boundary between two media occurs if the angle of incidence exceeds a critical angle given by $\theta_{crit} = \arcsin\left(\frac{n_2}{n_1}\right)$. For the interface between the inner and outer claddings, this gives $\theta_{crit} = 72.4°$, resulting in an opening angle of 26.7° at the centre of the fibre and trapping efficiency of 5.3 %.

The scintillating fibres have an intrinsic light yield of about 8000 photons per MeV of deposited ionisation energy. This results in about 18 to 20 detected photons at the SiPMs for a minimum ionising particle[2] traversing one detection layer of the SciFi Tracker close to the beam pipe [34]. Irradiation of the scintillating fibres over the lifetime of the detector will reduce the light yield. Simulations have shown that the fibres in the innermost region of the detector will be exposed to up to 35 kGy of ionising dose. The dose profile, however, falls off quickly to about 50 Gy at the readout end of the fibres. Several irradiation campaigns have been performed and have shown that a signal loss of about 40 % is expected for hits in the most irradiated regions of the detector due to a loss of transparency in the fibres [34]. In addition, oxygen diffusion into the fibres further modifies the transparency, resulting in a loss of attenuation length of about 1 to 2 % per year [56].

3.4 Fibre Mats & Modules

The scintillating fibres are wound into fibre mats on threaded winding wheels with a diameter of about 82 cm as described in [57]. The winding ensures that fibres are placed in a regular pattern with a pitch of 275 µm forming a matrix of six staggered layers, as shown in Figure 3.3a. Fibres are glued together using epoxy loaded with 20 % (by weight) of titanium dioxide to shield neighbouring fibres from crosstalk. A thin black polyimide foil is glued to both sides of the fibre mat to provide mechanical stability

---

[2] The mean energy loss rate of particles in matter due to ionisation is decreasing with momentum until it reaches a minimum from where it starts increasing again. A particle with a momentum close to the minimum mean energy loss rate is called a minimum ionising particle. The minimum is typically found at $\beta\gamma = \frac{p}{Mc} \approx 3$.

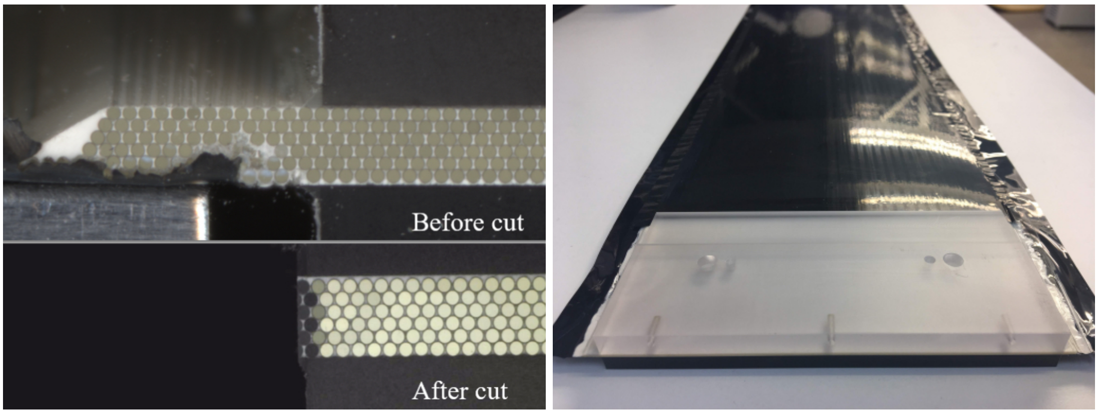

a b

Figure 3.3 Pictures of a fibre mat [57]. Left: microscope picture of the fibre mat before and after the mat is cut to the correct width. Right: picture of the fibre mat with the polycarbonate end-piece.

and shielding from light. The fibre mats are cut to the desired length (2424 mm) and width (130.6 mm), and a transparent plastic end-piece is glued to the readout side as shown in Figure 3.3b, while a thin foil mirror is applied to the other side.

Eight fibre mats, arranged in two rows with four lying side by side in each row, are assembled to form one module, as shown in Figure 3.4. The mats are sandwiched between two half-panels made from about 20 mm of polyaramid honeycomb laminated with a thin (0.2 mm) carbon-fibre reinforced polymer layer that provides the necessary mechanical rigidity while keeping a low material budget.

3.5 Silicon Photomultipliers

Silicon photomultipliers (SiPMs) are a relatively new technology for the detection of photons [58, 59], and consist of arrays of small (15 to 70 μm) pixels of avalanche photodiodes (APDs) connected in parallel. They combine single photon sensitivity, high detection efficiency and high granularity, making them ideal devices for photon detection in a high-resolution scintillating fibre tracker.

In its simplest form, a photodiode is a reversed-biased p-n junction. A p-type semiconductor contains a large concentration of positive charge carriers, called holes, whereas an n-type semiconductor has a surplus of electrons. At a p-n junction, the majority charge carriers of either side diffuse to the other side and recombine, creating a zone void of free charge carriers, called depletion zone. The diffusion process continues until an equilibrium with the electric potential created by the separation of charges is reached. When applying an external voltage to the junction in the same direction as the intrinsic electric field, the depth of the depletion zone increases, and in this case, the junction is called reverse-biased.

Photons with an energy above the semiconductor's bandgap can be absorbed in the depletion region and create electron-hole pairs. These pairs are separated in the electric field and drift towards their respective electrodes, inducing a small current. The

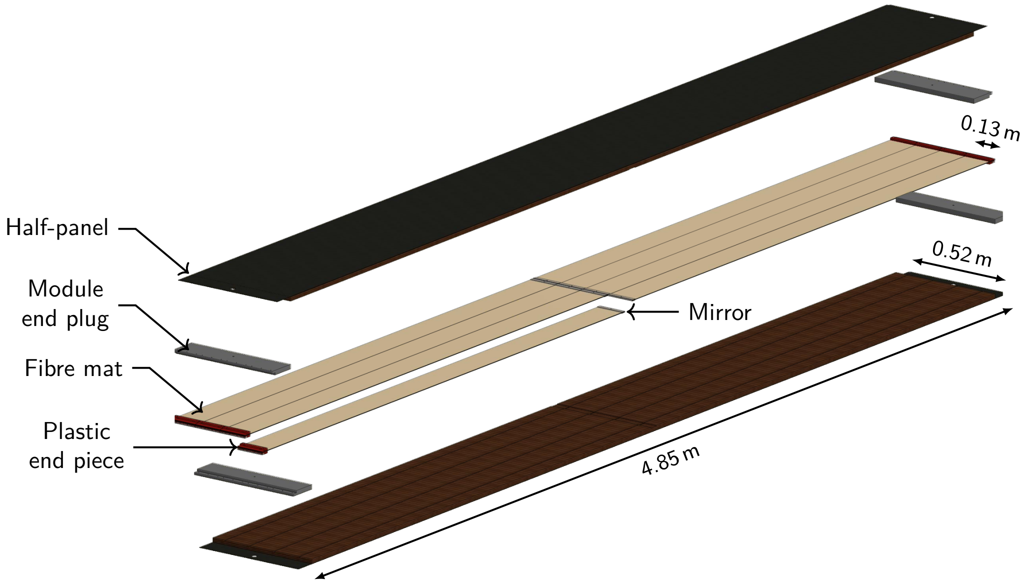

Figure 3.4 Exploded view of one fibre module of the SciFi Tracker showing the eight fibres mats, two half-panels, and aluminium end plugs. Adapted from Reference [57].

primary electron is typically referred to as *photoelectron* (pe).

The APD extends the principle of the p-n diode by adding a multiplication of the primary charge carriers. The doping profile, as shown in Figure 3.5 is modified such that a high electric field is obtained within a small region between the highly doped n^+ and p-doped layers. In this region, very high electric fields in the order of $10^5\,\mathrm{V\,cm^{-1}}$ are formed such that a single free charge carrier, either produced by the absorption of a photon or thermally, can trigger an avalanche of newly created electron-hole pairs. If the reverse bias voltage exceeds the so-called breakdown voltage (V_{bd}), the avalanche becomes self-sustaining, and the APD is operated in Geiger mode. In this mode the APD delivers a large discharge current that is independent of the magnitude of the primary ionisation and the APD becomes a counting device, sensitive to single photons. The avalanche is stopped when the voltage falls below breakdown due to an external resistor called quench resistor.

In an SiPM, multiple APD pixels operated above breakdown are connected in parallel and form a single channel. The signal is, therefore, the analogue sum of all pixels and is proportional to the number of detected photons. This is illustrated in Figure 3.6, where the pulse height spectrum of an SiPM under pulsed illumination is shown. The first peak, called *pedestal*, corresponds to the noise level of the readout electronics, whereas all following peaks are due to one or more pixels firing. The distance between the peaks is a measure of the signal amplitude per photoelectron and corresponds to the multiplication factor of the electrons in an avalanche called *gain*. The width of the peaks is given by the width of the noise and slight variations in the capacitance of the individual pixels that lead to variations in the signal amplitude.

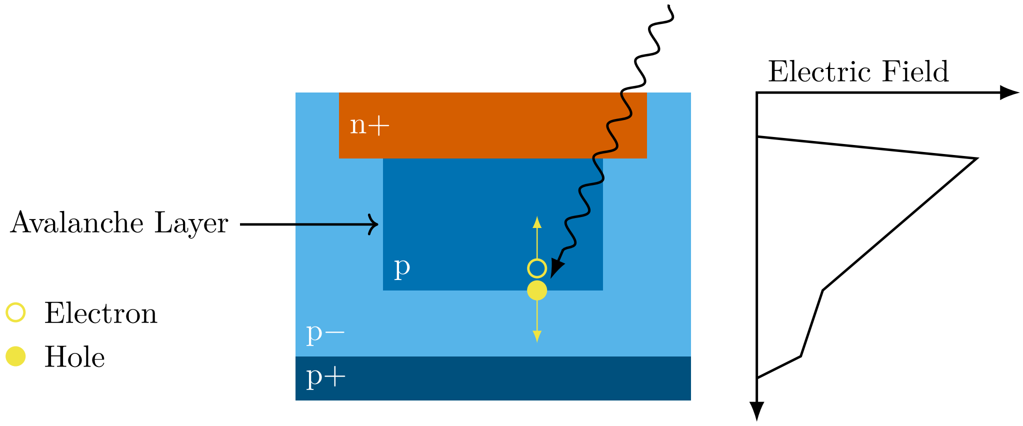

Figure 3.5 Schematic cross section of an avalanche photo diode with a n-on-p doping structure. The approximate electric field strength across the diode is shown on the right. Adapted from Reference [60].

The dynamic range of an SiPM starts from the detection of single photons and is limited by the total number of pixels in the SiPM. At high light intensities, the linearity is lost as each pixel is more likely to be hit by multiple photons.

3.5.1 Pulse Shape

The key characteristics of the pulse shape of an SiPM can be understood by looking at a simplified electrical model following [61]. In the circuit, shown in Figure 3.7a, the diode capacitance C_d is initially fully charged and biased at V_{bias}, and the conceptual switch S is open. No current flows through the series resistor R_d or the quench resistor R_q. When an electron-hole pair is created due to photo absorption or thermal generation, the switch closes, and the capacitance C_d begins to discharge through the series resistance R_d, causing a voltage drop over R_q. As R_d is typically small, a rapid surge in current, given by

$$I_d(t) \approx 1 - \exp\left(-\frac{t}{R_d C_d}\right) = 1 - \exp\left(-\frac{t}{\tau_d}\right), \qquad [3.1]$$

occurs with a time constant τ_d, that reaches its maximum

$$I_{max} = \frac{V_{bias} - V_{bd}}{R_q + R_d} \qquad [3.2]$$

after about 1 ns. At this time, t_{max}, the voltage over the diode drops to V_{bd}, which is insufficient to sustain the discharge, and quenching occurs, closing the conceptual switch S. The diode capacitance C_d begins to recharge causing the voltage over R_q to decrease and leading to an exponentially decreasing current

$$I_q \approx \exp\left(-\frac{t}{R_q C_d}\right) = \exp\left(-\frac{t}{\tau_q}\right) \qquad [3.3]$$

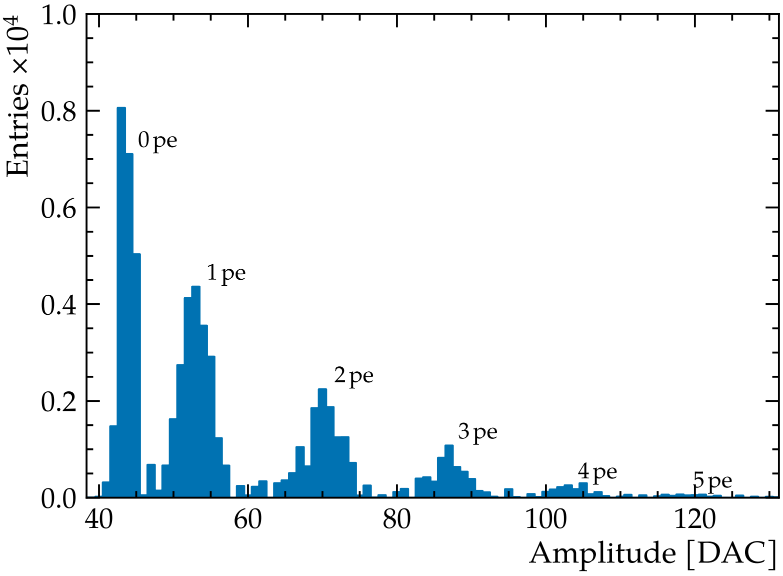

Figure 3.6 Typical pulse height spectrum of an SiPM. Every peak corresponds to a certain number of fired pixels.

with a time constant τ_q, also called recovery time, in the order of several 10 ns. The resulting pulse shape is shown in Figure 3.7b. The total amount of charge released by the discharge of the APD is approximately given by

$$Q = I_{max} \cdot \tau_q = \frac{V_{bias} - V_{bd}}{R_q + R_d} \cdot R_q C_d = \Delta V \cdot C_d. \qquad [3.4]$$

where $\Delta V = V_{bias} - V_{bd}$ is called over voltage. The multiplication factor or gain (G) is obtained by dividing the total amount of charge by the elementary charge e

$$G = \frac{Q}{e} = \frac{\Delta V \cdot C_d}{e}. \qquad [3.5]$$

As C_d is fixed by the architecture of the SiPM, the gain factor is thus only controlled by a single parameter ΔV, called overvoltage. Typical gain values that can be achieved are in the range of 10^5 to 10^7. In addition, there is an implicit dependence on the temperature as V_{bd} does depend on it. SiPMs are typically operated at stable temperatures to eliminate this dependency, so ΔV is kept constant.

3.5.2 Dark Counts

Dark counts are avalanches triggered without incoming and detected photons but caused by the diode's thermal excitations or tunnelling effects. The signals are, therefore, the same as those created by a photon-induced discharge and cannot be distinguished. The rate at which dark counts appear is called the dark count rate (DCR). Due to its thermal origin, the DCR is strongly temperature-dependent but also increases with the applied overvoltage. Additionally, silicon lattice defects created through irradiation drastically increase the DCR.

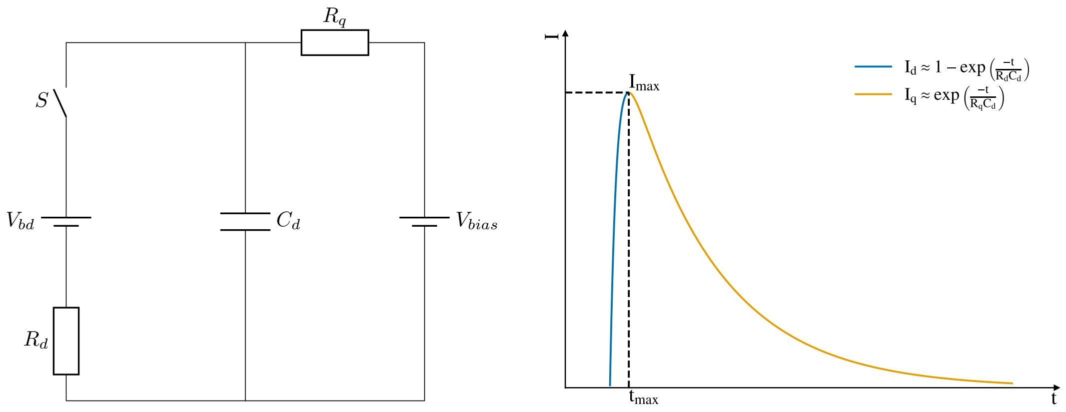

a **b**

Figure 3.7 Left: simple equivalent circuit for an externally biased SiPM pixel. Right: graph of the current flow in the circuit after closing of the switch *S*. Adapted from Reference [61].

The DCR can be suppressed by exploiting the fact that a true signal, e.g. a particle crossing a scintillator generating multiple photons, will fire multiple SiPM pixels in coincidence, thus creating a large signal. In contrast, dark counts only produce single-pixel (1 pe) signals. Applying a threshold to the signal above the one photon signal, the DCR can effectively be suppressed. As dark counts are of thermal origin, the DCR can be further reduced by operating the SiPM at lower temperatures.

3.5.3 Optical Crosstalk

In the discharge of a pixel, infrared photons can be generated and travel to neighbouring pixels. These photons can then trigger an avalanche themselves. This effect is called *direct optical crosstalk* and is a form of correlated noise, as secondary pulses are produced simultaneously with the primary pulse. Crosstalk can also be observed with significant delay with respect to the primary pulse. This phenomenon is then called *delayed optical crosstalk* and is explained by infrared photons being absorbed in the silicon substrate of the neighbouring pixel. The subsequently released electron first slowly diffuses out of the substrate before drifting to the amplification zone where the secondary pulse is induced with a delay in the order of ≈20 ns [62]. Each crosstalk-induced avalanche again has the chance to produce further crosstalk photons. Optical crosstalk can be reduced by introducing an optically opaque material around the pixels, so-called trenches; a method that has been applied for the SiPM arrays of the SciFi Tracker.

Due to crosstalk, dark counts can reach amplitudes above 1 pe, requiring larger discriminator thresholds for suppression and possibly impacting the detection efficiency for real signals.

3.5.4 After Pulsing

After pulsing is another form of correlated noise found in SiPMs. The charge carriers in an avalanche can become trapped at impurities or defects in the crystal lattice and

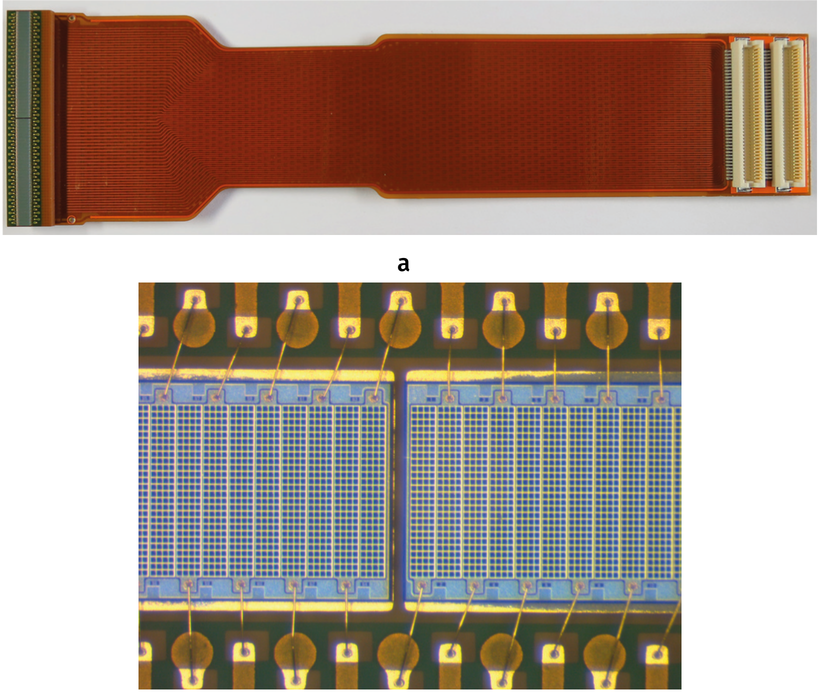

Figure 3.8 Top: picture of an SiPM array assembly for the SciFi Tracker with flex cable and connector. Bottom: microscope picture of the active area between the two dies of one SiPM array showing the individual channels and pixels. Adapted from Reference [63].

be released with some delay, triggering a secondary avalanche. Typically, after-pulsing occurs with delays of a few 100 ns.

3.5.5 SiPM Arrays

The SciFi Tracker uses a total of 524 288 SiPM channels to read out the light from the scintillating fibres. The channels are grouped into 128-channel arrays produced by Hamamatsu[3] and feature 4 × 26 APD pixels with a size of 57.5 µm × 62.5 µm and a pitch between channels of 250 µm [34]. The array comprises two dies of 64 channels with a 220 µm gap between them. The geometry was chosen such that four arrays cover the entire width of the fibre mat and that the height of the channels extends by about 20 % over the height of the mat.

The full assembly of the array bonded to a custom flex PCB is shown in Figure 3.8. A 100 µm thin epoxy window protects the front of the silicon while a stiffener is glued to the back to provide mechanical support. The PCB also hosts a Pt1000 temperature sensor to monitor the temperature.

The average breakdown voltage of the SiPM channels is measured in Reference [64] to be 51.75 V with variations in the range of ±300 mV. The temperature dependence of the breakdown voltage is $(60 \pm 2)\,\mathrm{mV\,K^{-1}}$, requiring a cooling system with a temperature stability of 1 K. In addition, the arrays were grouped according to their average breakdown voltage. Since four arrays are always biased by a common power

---

[3] Hamamatsu Photonics KK, Hamamatsu, Japan.

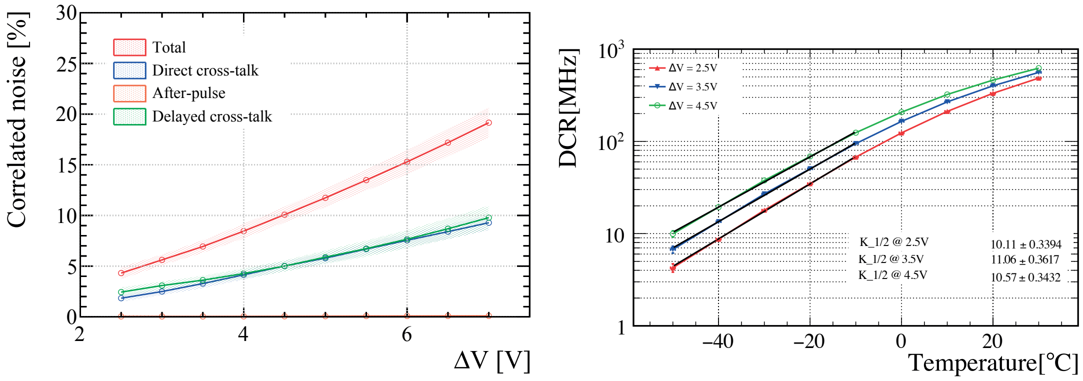

a **b**

Figure 3.9 Left: measurement of the correlated noise as a function of the overvoltage [62]. Right: measurements of the dark count rate per channel as a function of the temperature for three different overvoltages. The SiPMs for the SciFi Tracker are nominally operated at 3.5 V overvoltage. A linear fit is used to extract the temperature coefficient. [63].

supply channel, this ensures a homogenous over-voltage and, thus, performance. At the nominal overvoltage for the SciFi Tracker of 3.5 V, the photon detection efficiency of the SiPM was measured to be $(43.5 \pm 3.5)\ \%$. [64]

The probability of secondary avalanches due to crosstalk has been significantly reduced in the development of the SiPMs by the addition of optical trenches in between the individual pixels that absorb the infrared photons emitted in the primary avalanche. The measurement of the correlated noise is shown in Figure 3.9a. At the nominal over voltage and a temperature of −40 °C direct crosstalk was measured to be ≈3.3 % and delayed crosstalk to be ≈3.5 % while after-pulsing is negligible [34, 62]. The recovery time τ_q is (84.0 ± 0.2) ns [64] .

The SiPMs are expected to accumulate only a relatively low ionising dose of about 50 Gy as they are located at the edge of the acceptance. However, of greater concern is the damage due to non-ionising energy loss, where the integrated fluence is expected to be up to $6 \times 10^{11}\ \mathrm{n_{eq}\,cm^{-2}}$ and will significantly increase the DCR from about 40 kHz to 550 MHz per channel at room temperature. Coupled together with crosstalk, the single-pixel dark counts then have a significant probability of creating signals with amplitudes similar to those coming from real tracks. As shown in Figure 3.9b, the DCR decreases by about half for every 10 K decrease in temperature. By operating the SiPMs at −40 °C the increase in DCR due to irradiation over the lifetime of the detector can be mitigated and kept to an acceptable level of about 14 MHz per channel at the nominal overvoltage of 3.5 V as shown in Figure 3.9b. In addition to cooling the SiPMs, the rate of noise clusters due to the DCR will be further reduced by appropriate detection thresholds and a cluster algorithm that will be described in the following section.

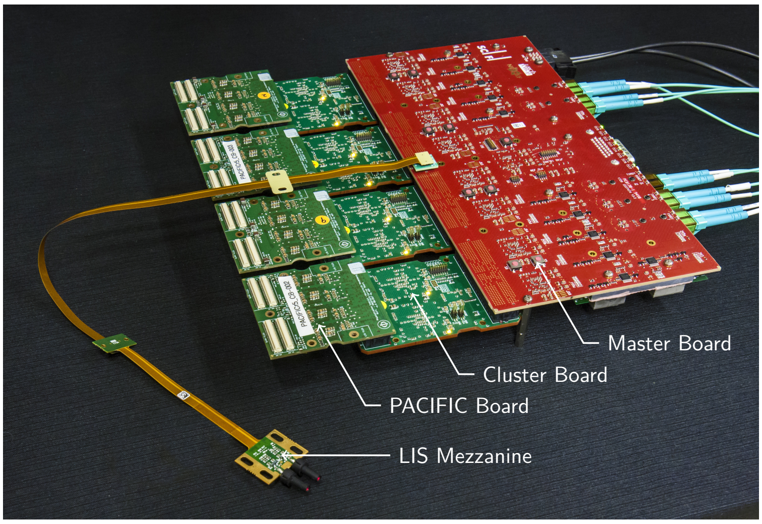

Figure 3.10 Photograph of a complete set of frontend electronics of one SciFi Readout Box including the light injection system mezzanine board. Adapted from Reference [67].

3.6 Frontend Electronics

The frontend electronics of the SciFi Tracker are an integral part of the detector as they digitise and process the SiPM signals at the LHC bunch crossing frequency of 40 MHz. The electronics are located at the top and bottom of each module in so-called Readout Boxes (ROBs) split into two equal halves called HalfROBs. Each HalfROB comprises three modular elements, as shown in Figure 3.10, each providing a specific functionality. The PACIFIC Board houses the PACIFIC frontend ASICs that process and digitise the SiPM signals with three comparators whose combined output is encoded in a 2-bit word. The signals of one SiPM array are always read out by two PACIFIC ASICs, resulting in a data rate of 10.24 Gbit/s. The data is then shipped to the Cluster Board where two FPGAs perform a cluster search over the 128 channels of one array and calculate the hit position. Finally, the Master Board serialises the data using eight GBTx ASICs [65], each with a bandwidth of 4.8 Gbit/s, and transmits it via VTTx [66] optical links to the DAQ system.

3.6.1 PACIFIC ASIC

The PACIFIC[4] has been specifically designed for the requirements of the SciFi Tracker to process and digitise the analogue SiPM signals at 40 MHz. The ASIC allows to read out 64 SiPM channels and was developed using TSMC's[5] 130 nm CMOS[6] process. Each channel contains an analogue processing, digitisation, slow control, and serialisation synchronised with the 40 MHz bunch clock of the LHC. A simplified schematic of

---

[4] low-**p**ower **A**SIC for the **Sci**ntillating **Fi**bre Tra**c**ker.

[5] Taiwan Semiconductor Manufacturing Company Limited, Hsinchu, Taiwan.

[6] Complementary metal–oxide–semiconductor (CMOS) is a type of metal–oxide–semiconductor field-effect transistor (MOSFET) fabrication process that uses complementary and symmetrical pairs of p-type and n-type MOSFETs for logic functions.

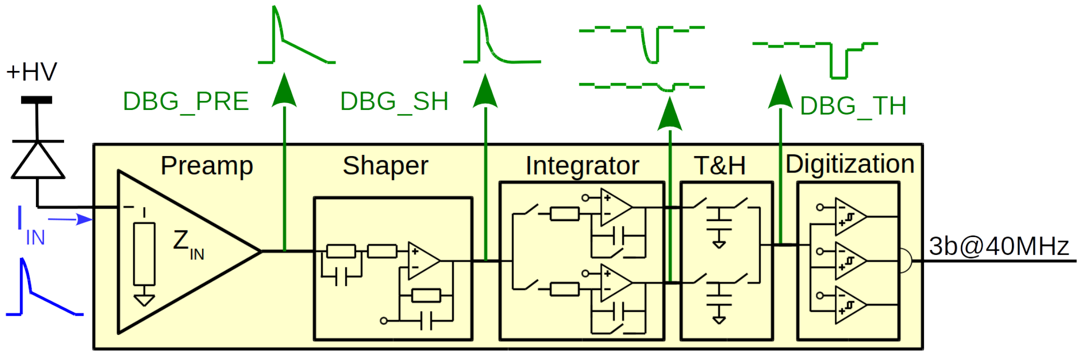

Figure 3.11 Simplified block diagram of one PACIFIC channel consisting of five stages where the analogue SiPM signals processed and digitised [68].

the processing blocks in each channel is shown in Figure 3.11. The five processing stages consisting of the pre-amplifier, shaper, integrator, track and hold (T&H), and the digitisation are described in the following sections.

Amplification

The first stage in the PACIFIC signal processing is the pre-amplifier, or input stage, that directly connects to the SiPM anodes. The input stage uses a novel double feedback current conveyor [69] with a low input impedance of 50 Ω and a broad bandwidth of ≈250 MHz and allows for the selection of four output gains. In addition it is possible to adjust the SiPM bias voltage in the range from 100 to 700 mV in steps of 40 mV. The current conveyor is followed by a trans-impedance amplifier to convert the current signal to a voltage. Figure 3.12a shows the output of the input stage for several SiPM signals recorded with an oscilloscope. The tails of the signals can been extending over several x-axis divisions, each measuring 20 ns.

Shaping

As discussed in Section 3.5.5, the SiPM signals have a recovery time of ≈84 ns leading to a long tail that extends over several 25 ns LHC clock cycles. The long signal tails seen at the pre-amplifier output shown in Figure 3.12a would lead to unwanted spillover from one bunch crossing to the following. In addition, the light from particles that traverse the scintillating fibres at different locations along their length will arrive with delays of up to 15 ns [71]. In order to minimise the spillover and the fluctuation of the integrated signal depending on the signal arrival time, a fast shaper is used.

The implementation of the shaper is based on a pole-zero filter circuit as shown in Figure 3.13. To analyse the frequency dependence of the circuit it is instructive to employ the mathematical treatment by means of the transfer function $H(s)$ in Laplace

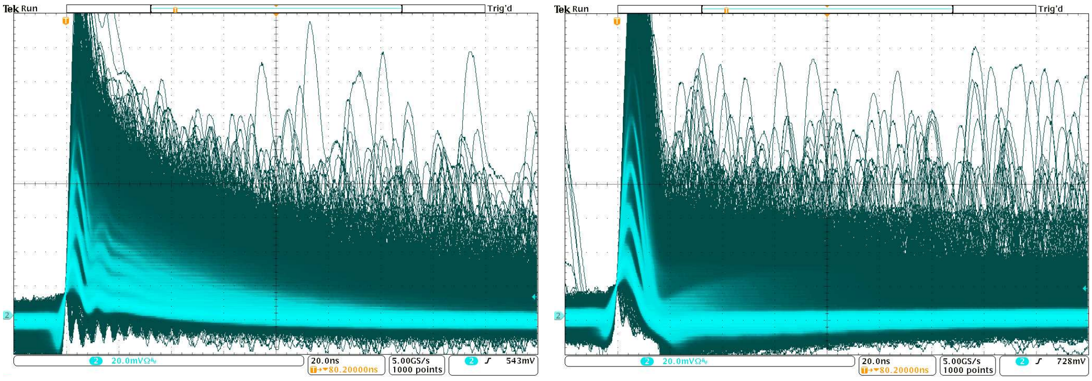

a b

Figure 3.12 Oscilloscope measurements of the PACIFIC pre-amplifier and shaper output from light pulses where the 1 to 4 pe signals are visible as distinct bands [70]. One division on the horizontal axis corresponds to 20 ns in both images. Left: output of the pre-amplifier showing the long tail of the SiPM signals. Right: measurement of the shaper output. The long tail of the signals has been reduced.

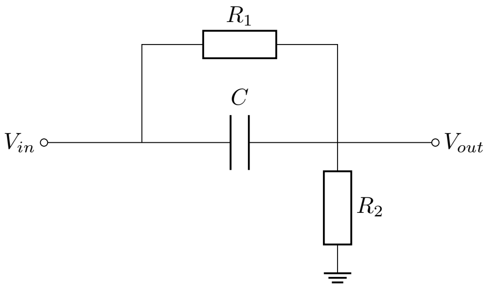

Figure 3.13 Pole-zero circuit consisting of a capacitor and two resistors as used within the PACIFIC shaper to cancel the slow component of the SiPM signals. Adapted from Reference [72].

space[7]. The frequency dependent output of a circuit is then given by

$$V_{out}(s) = H(s)V_{in}(s) \tag{3.6}$$

with the complex variable s related to the frequency ω by $s = \sigma + i\omega$ where σ is a constant ensuring convergence. For the pole-zero circuit in Figure 3.13 the transfer function is

$$H_{pz}(s) = \frac{s + \frac{1}{\tau_1}}{s + \frac{1}{\tau_2}} = \frac{\tau_2}{\tau_1}\frac{1 + s\tau_1}{1 + s\tau_2}, \tag{3.7}$$

with $\tau_1 = R_1 C > \tau_2 = \frac{R_1 R_2}{R_1 + R_2} C$ [55, 73]. An exponential input signal

$$V_{in}(s) = \mathcal{L}\left[\exp\left(-\frac{t}{\tau_1}\right)\right] = \frac{1}{s + \frac{1}{\tau_1}} \tag{3.8}$$

with a time constant equal to τ_1 will result in an exponential output signal

$$V_{out}(s) = H_{pz}(s) \cdot V_{in}(s) = \frac{1}{s + \frac{1}{\tau_2}} \tag{3.9}$$

with a time constant τ_2. Since $\tau_1 > \tau_2$ the pole-zero filter shortens the exponential tail of the input signal as can be seen in Figure 3.12b. Both time constants can be tuned in the PACIFIC via special registers that control the resistances R_1, R_2, and the capacitance C.

Integration

After the amplification and shaping stage, the signals are integrated in each 40 MHz clock cycle. A design based on two interleaved gated integrators, as shown in Figure 3.14, has been chosen to avoid any dead time in the data acquisition. This is achieved by switching between the integrators at a frequency of 20 MHz, which means that while one integrator is working, the other is being reset.

Each integrator comprises an amplifier with capacitive feedback, a resistor, and a switch. An input voltage V_{in} will lead to a current through the resistor R_1 that charges the capacitor C. The current is given by

$$I = \frac{V_{in}(t)}{R} = -C\frac{\mathrm{d}V_{out}(t)}{\mathrm{d}t}, \tag{3.10}$$

which, after integration with respect to time, yields

$$V_{out}(t) = -\frac{1}{RC}\int_0^t V_{in}(t)\,\mathrm{d}t + const. \tag{3.11}$$

Thus, the output voltage is proportional to the integral of the input voltage.

---

[7] The Laplace transform of a time dependent function $f(t)$ with $f(t = 0) = 0$ is defined by $F(s) = \mathcal{L}[f(t)] = \int_0^\infty f(t)\exp(-st)\mathrm{d}t$.

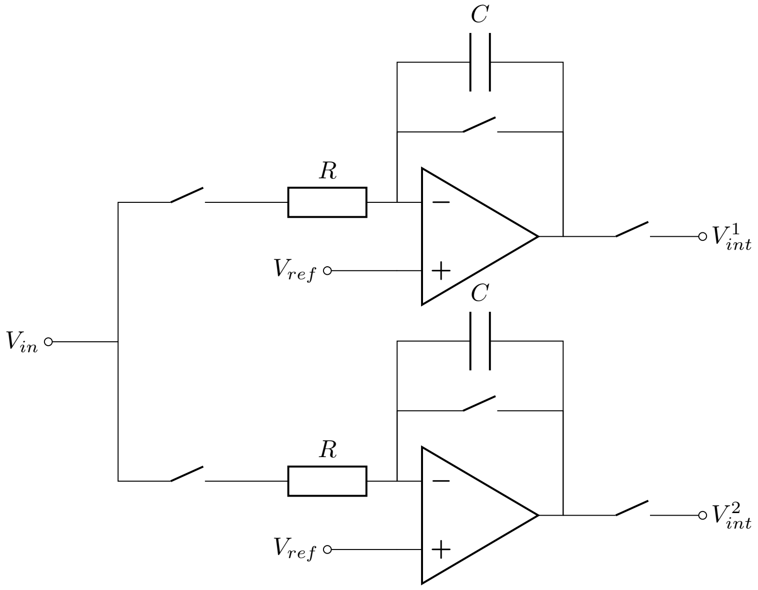

Figure 3.14 Schematic of the dual interleaved integrator used in the PACIFIC. Each integrator consists of an operational amplifier with capacitive feedback. Switching between the two integrators allows for an operation without dead time. Adapted from Reference [72].

Track and Hold

The integrators are followed by a passive track-and-hold stage shown in Figure 3.15a consisting of two capacitors C_1 and C_2 and a series of switches. It is responsible for stabilising the signals from the two integrators and merging them before digitisation. The track-and-hold stage samples the outputs of the integrators in an interleaved manner with a frequency of 20 MHz. While integrator 1 is working, its output V_{int}^1 is connected to the capacitor C_1 but disconnected from the shared output V_{out}; at the same time, integrator 2 is disconnected from the capacitor C_2, but C_2 is connected to V_{out}. Additional switches have been implemented that disconnect the capacitors from ground for a brief period (≈1 ns) during the switching process between the two integrator inputs. This arrangement significantly reduces the amount of signal spillover from one clock cycle to the next [68].

The output of the track-and-hold stage is shown in Figure 3.15b, where the signals from 0 to 4 pe are visible as separate bands. The integration and track and hold stages are very sensitive to the relative time between the signals and the system clock. The delay needs to be carefully tuned so that the signals are maximally integrated. This will be further discussed in Section 6.2.

Digitisation

In the last processing stage, the analogue signals are digitised by three comparators. The comparator thresholds can be adjusted in the range of 0 to 750 mV with three independent 8 bit DACs per channel. A hysteresis of 10 mV has been added to prevent repetitive switching if the signal is close to the threshold voltage [72]. The digital

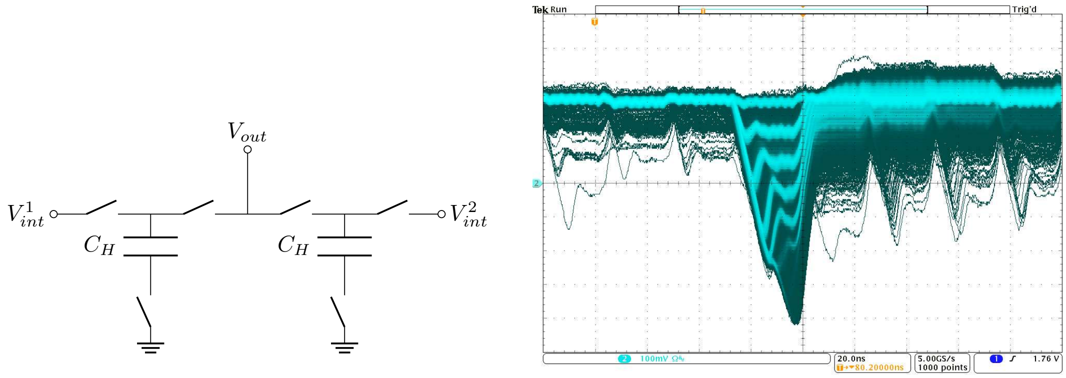

a

b

Figure 3.15 Left: track and hold circuit. Adapted from Reference [68]. The input is switched alternating between the output of the two integrators $\mathrm{V}_{\mathrm{int}}^1$ and $\mathrm{V}_{\mathrm{int}}^2$. Additional switches between the capacitors and ground help to reduce spillover by momentarily disconnecting the capacitors [68]. Right: oscilloscope measurement of the track and hold output from light pulses [70]. The 0 to 4 pe signals are visible as well separated bands.

outputs of all three comparators are added, resulting in a 2 bit output. The output of four channels is serialised together at 320 MHz leading to a data rate of 5.12 Gbit/s per PACIFIC ASIC.

Digitisation based on three comparators has been chosen as it drastically reduces the required bandwidth with minimal to no loss in resolution when compared to a 6 bit ADC [74]. The three comparators effectively cut away the 1 pe noise of the SiPMs while still providing enough pulse-height information for a precise position reconstruction. This requires the thresholds to be calibrated with respect to the pulse height spectrum of the connected SiPM channel, which allows for the conversion of signal amplitudes from photoelectrons to DAC values. In practice, this is done by performing threshold scans under pulsed illumination, as will be discussed in Chapter 6.

Serialisation

After the digitisation of the SiPMs signals by the three comparators the digital output is serialised. The output of the three comparators of each channel is sampled with a frequency of 40 MHz. To reduce the required bandwidth the three bits are arithmetically added together and encoded in a 2 bit value. The outputs of four channels are then written to a register and from there single bits are pushed to the output with the help of a shift register sampled at 320 MHz.

To ensure a synchronous data taking across all 8192 PACIFIC ASICs, the serialiser receives a special *SYNC* TFC signal. This signal is evaluated at the rising edge of the external 320 MHz clock and resets the clock generation in the ASIC. It thereby ensures that the additionally generated 20 MHz and 40 MHz clocks are generated at the same predetermined time.

3.6.2 PACIFIC Board

The PACIFIC Board serves as the carrier board for four PACIFIC ASICs and provides the connection for two SiPM arrays (256 channels in total) via four high-density connectors. The connectors feature 80 pins, of which 64 are used for the 64 channels of one SiPM die, and the remaining pins are connected to the bias voltage and ground. Additionally, one line is reserved for the Pt1000 temperature sensor housed on the SiPM assembly's backside. The analogue signals of one SiPM (128 channels) are routed to two PACIFICs whose outputs are then forwarded via a high-speed interconnect to the Cluster Board. In addition, the board also features two Pt1000 temperature sensors and a monitoring circuit for the SiPM bias voltage, which are all read out by ADCs located on the Cluster Board.

3.6.3 Cluster Board

The Cluster Board houses two radiation tolerant FPGAs[8] that each process the output of two PACIFICs (one SiPM array) [75]. The FPGAs run an online cluster-finding algorithm that groups neighbouring channels from the same SiPM array into clusters and encodes their position in 8 bit [34]. This significantly reduces the data rate from $10.24\,\mathrm{Gbit\,s^{-1}}$ to $4.8\,\mathrm{Gbit\,s^{-1}}$ and, in combination with the PACIFIC comparator thresholds, suppresses the rate of dark clusters while maintaining a high hit detection efficiency.

Besides the two FPGAs, the Cluster Board also contains a GBT-SCA [76] Slow Control Adapter ASIC that enables communication to the FPGAs and the connected PACIFIC Board via the JTAG and I^{2}C interfaces. In addition, the SCA features 31 analogue inputs that are multiplexed to a common ADC. Each input has a switchable current source to read out temperature sensors such as Pt1000s. They are connected to the temperature sensors on the Cluster Board (3), PACIFIC Board (2), and SiPMs (2). Four further lines are used to monitor the bias voltage of the individual SiPM dies.

Cluster Finding

The cluster-finding algorithm is responsible for combining signals in multiple channels belonging to the same particle and rejecting signals due to noise. The principle of the algorithm is illustrated in Figure 3.16. The figure shows the analogue signal in each SiPM channel as well as the three PACIFIC thresholds Vth_1, Vth_2, and Vth_3. If calibrated the thresholds correspond to the number of detected photoelectrons. Typical threshold settings are 1.5 pe, 2.5 pe, 4.5 pe. Each channel is assigned a weight w_i corresponding to the threshold value in photoelectrons that has been passed. A cluster is formed when the sum of the weights of two or more neighbouring channels is larger than the sum of the first and second threshold, as is the case for clusters (a) and (c) in Figure 3.16. Additionally, a single channel can also form a cluster if the amplitude exceeds Vth_3, such as cluster (b). The channels around (e) do not fulfil these conditions and, therefore, do not produce a cluster. The position of the cluster is computed from the weighted average of all included channels, rounded to

---

[8] Microchip Technology Igloo2. Microchip Technology Inc., Chandler, United States.

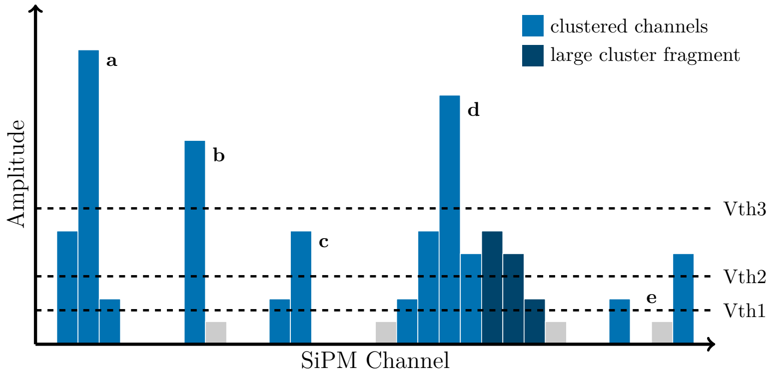

Figure 3.16 Illustration of the cluster-finding algorithm. The height of each rectangle represents the amount of charge measured in one channel. The PACIFIC comparator thresholds are indicated by dashed lines. Six clusters are identified by the algorithm as indicated by the blue colours. Signals below the lowest threshold are not considered and are shown in grey. Adapted from Reference [34].

half a channel position and encoded in an 8-bit word. This half-channel precision is enough to provide the required hit resolution of less than 100 µm [34]. If the number of channels in a cluster exceeds four, the cluster is flagged as *large* and combined with the subsequent cluster fragment, such as cluster (d) [77]. The position of these clusters is simply the median position.

Data Format

The maximum number of clusters that can be sent by the cluster FPGA is limited by the bandwidth of the GBTx ASICs that are used to serialise the data of a bunch crossing. The GBT transport protocol uses 112 bit wide frames to pack the data. Within these 112 bits, 20 bits are used for the header and 92 bits for the clusters. The header contains the 12-bit bunch crossing number (ranging from 0 to 3653) that is required by the TELL40s to associate the clusters from different parts of the SciFi Tracker to the same event. A further 5 bits are used to encode the number of formed clusters, while the remaining 3 bits contain additional meta information. The clusters are encoded in 9-bit words, resulting in a maximum of ten clusters that can be send out per bunch crossing. The position of the cluster is encoded in 8 bits. As each Cluster FPGA processes the data of one SiPM (128 channels) this allows for encoding the cluster position within the SiPM with half-channel precision. For large clusters consisting of more than four channels only the first and last 9-bit cluster fragments are encoded, with the last cluster having the ninth bit set to 1 to indicate the end of the large cluster. Since both the header and data sizes are fixed this data format is called *fixed-header fixed-data* (FF). If the clustering algorithm finds more than ten clusters, only the first ten are sent. In the high occupancy regions of the SciFi Tracker, a *fixed-header variable-data* (FV) data format is therefore used to frame the clusters which allows for sending up to 16 clusters per link by extending the frame into the next bunch crossing when needed.

In addition special variants of the FF and FV formats, called *non-zero-suppressed* (NZS), exist that not only send the found cluster information but also transmit the raw 2-bit output of each PACIFIC channel. As the data volume exceeds the capacity of a single frame, multiple frames are needed to send the data. This also means that the data taking can not take place at the nominal rate of 40 MHz. The NZS data formats are used when calibrating the detector, as will be discussed in Section 6.1, and are also invaluable for debugging purposes.

3.6.4 Master Board

The data from the clustering FPGAs is serialised in the Master Board by the GBTx ASIC [65] and sent to the DAQ system via optical VTTx [66] links. The frontend electronics data stream is organised so that the clustered data from one SiPM is serialised by one GBTx ASIC and transmitted over a single optical fibre. Each Master Board, therefore, houses eight GBTx ASICs and four VTTx transmitters with two fibres per transmitter. A separate GBTx configured for bidirectional communication, called Master GBTx, together with a bidirectional VTRx [66] transceiver, provides the communication to the experimental control (ECS) and timing and fast control (TFC) systems. The GBTx ASIC features a set of tunable clock outputs used to drive the different parts of the frontend electronics. The tuning of these clocks during the assembly of the detector is discussed in Section 4.8.

The Master Board also contains an additional FPGA, the so-called Housekeeping (HK) FPGA. It is used to operate the light injection system (LIS) that is discussed in Section 3.6.5 and is responsible for monitoring various status registers of the Master GBTx. The Master GBTx is connected to a GBT-SCA, referred to as the Master SCA, which is used to control the eight data GBTx, the HK FPGA, and the LIS. It also reads out the eight temperature sensors (NTC thermistors) located on the Master Board and monitors the input voltage.

Power is supplied to the Master Board by an external 8 V power supply (see Section 3.7) and transformed by 13 FEASTMP [78] DC/DC converters to the various voltages required by the PACIFIC and Cluster Boards. An additional connection supplies the bias voltage to the SiPMs.

3.6.5 Light Injection System

The light injection system (LIS) allows for injecting light into the SiPM arrays for calibration and commissioning purposes. It is therefore central to the work presented in this thesis (see Chapters 4 and 6), and considerable time was spend to commission it and tune its parameters. The system is shown in Figure 3.17. It consists of a VCSEL laser diode driven by the same radiation tolerant GigaBit Laser Driver (GBLD) [79] ASIC that is used in the VTTx and VTRx modules. The laser diode is connected to a plastic optical fibre, where the last 13 cm of the fibre, corresponding to the width of one fibre mat, are embedded in an aluminium bar and scratched to allow the light to escape. The aluminium bar with the embedded scratched fibre is called *light injection bar* and sits directly on top of the clear plastic fibre mat end piece, as shown in Figure 3.17, through which the light reaches the SiPMs. Four light injection bars and two LIS mezzanine boards, each hosting two GBLDs and VCSELs, are embedded within the aluminium end plugs of each module.

The GBLD settings such as the bias and modulation currents are configured via the Master SCA. The input to the GBLD, from here on called *LIS pulse*, is generated in the HK FPGA and determines the delay and duration of the light pulse. The LIS pulse is generated based on a dedicated TFC calibration command from the control system called *Calib C* that is sent in one of the 25 ns bunch crossing periods. The light pulses need to be in phase with the PACIFIC integration period to be fully integrated. For this purpose, two clocks are generated on the Master Board that allow to set the delay and width of the light pulse relative to the Calib C pulse. The interplay of the signals involved in the generation of the LIS pulse is illustrated in Figure 3.18. Upon arrival of the Calib C pulse, the start of the LIS pulse is given by the first rising edge of the *Start Clock*, and similarly the first rising edge of the *Stop Clock* determines the end of the LIS pulse as indicated by the solid black lines in Figure 3.18. The delay of the LIS pulse can be adjusted in 512 steps with a granularity of 48.8 ps by shifting the phase of both clocks simultaneously. The pulse width is given by the phase difference between the start and stop clocks [80]. Typically, a pulse width of ≈15 ns is used. When the Start Clock sampling of the Calib C pulse approaches the clock edge (<1 ns), the order of the signals is not guaranteed due to jitter. As a result, the generated LIS pulse becomes metastable and jumps by 25 ns in either direction. To avoid this scenario, the Calib C pulse can be sampled on the falling instead of the rising edge of the TFC clock, effectively shifting it by 12.5 ns.

As the clocks are all operated at a frequency of 40 MHz it is impossible to set delays larger than 25 ns. To shift the pulse by a full 25 ns period, the Calib C pulse needs to be issued in the previous or next bunch crossing period. The latency of the Calib C pulses as well as the LIS delay need to be adjusted to inject the light at the correct time into the SiPMs, so that the signals are optimally integrated by the PACIFIC. Both are crucial steps before the calibration of the PACIFIC thresholds and will be discussed in Section 6.2.

3.7 Infrastructure & Services

The operation of the SciFi Tracker is enabled by a set of services that include the power supplies for the electronics and the bias voltage for the SiPMs, as well as the cooling systems. All of the systems work together to ensure a stable and efficient operation of the detector throughout the data-taking periods.

3.7.1 Low Voltage

The frontend electronics of the SciFi Tracker are powered by radiation tolerant W-IE-NE-R[9] MARATON power supplies that have been specifically developed for the operation at the LHC and have been used successfully in the previous iteration of the LHCb detector. Each unit features 12 individual low voltage channels that can provide up to 300 W of power [82]. One channel supplies the voltage of 8 V for up to two ROBs.

The power supply consists of three main components: the power box, a primary rectifier, and a control module. The rectifier converts the 230 V AC mains voltage to a 385 V

---

[9] W-IE-NE-R Power Electronics GmbH, Burscheid, Germany.

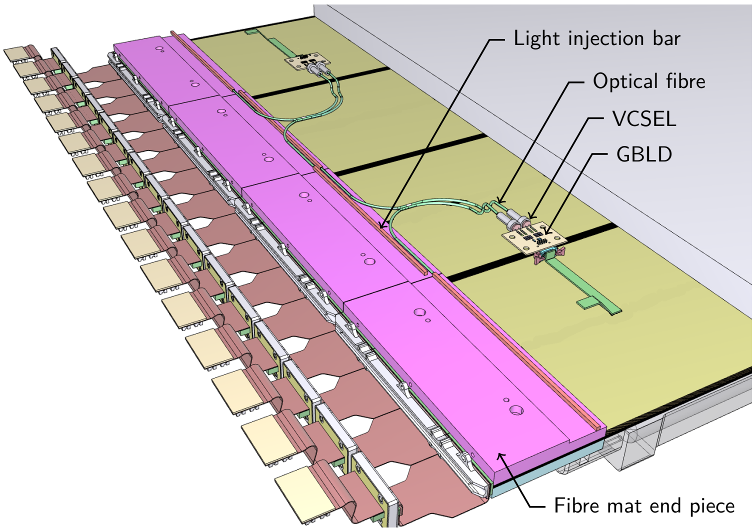

Figure 3.17 Schematic of a fully assembled detector module showing the interface between the light injection system and the fibre mat end pieces. The top half of the module end plug is removed to show the routing of the LIS optical fibres and light injection bars. The light is send from the mezzanine board to the fibre mats via optical fibres. Inside the light injection bar the fibre is scratched to allow the light to escape. Adapted from Reference [81].

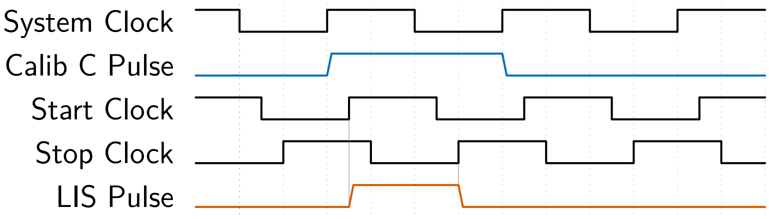

Figure 3.18 Timing diagram of the LIS pulse generation in the HK FPGA. Adapted from Reference [80]. The start and stop clocks are provided by one of the Data GBTxs on the Master Board. Upon reception of the Calib C pulse (blue), the LIS pulse (orange) is generated from the first rising edge of the start and stop clocks within the Calib C pulse.

DC voltage that is then fed to the power boxes. The control module is connected to the LHCb experiment control system and allows for the remote control of the power supplies and monitoring of their status. The rectifiers and control modules are located far away from the detector in a dedicated counting room, whereas the power boxes producing the low voltage are located close to the detector to limit power losses due to the high currents ($\approx$20 A) combined with the long distance between the experiment and the counting room.

3.7.2 SiPM Bias Voltage

The SiPM arrays are biased by CAEN[10] power supplies that deliver a voltage in the range of 0 to 100 V with up to 20 mA, and an accuracy of $\pm$0.05 % ($\pm$50 mV) [83]. Each CAEN channel supplies four SiPM arrays. In case of defects, single SiPM dies (64 channels) can be disconnected through an intermediate switch panel. The power supply modules are housed in a crate that contains the primary power supply, allowing remote control and monitoring of the system's status.

In addition the bias voltage can be adjusted for each SiPM channel individually via the PACIFIC input stage within a small range 600 mV.

3.7.3 Electronics Cooling

Each Readout Box produces a heat load of close to 100 W during operation that needs to be cooled away in order to limit the temperature of the electronics to a maximum of 50 °C [84]. The ROBs are, therefore, tightly mounted to aluminium cooling blocks with water pipes running through them that dissipate the heat.

3.7.4 SiPM Cooling

In order to keep the dark count rate to an acceptable level after irradiation of the detector, the SiPMs need to be cooled to −40 °C, as mentioned in Section 3.5.2. The SiPM arrays are glued to 3D-printed titanium cooling bars and housed in an insulating shell called Cold Box. The single phase coolant[11] is circulated through vacuum insulated supply and return lines and connected via flexible bellows to the Cold Boxes. The cooling liquid temperature can be adjusted in the range of +30 to −50 °C[12] [34].

To prevent condensation and ice formation on the inside of the Cold Box, the whole assembly is constantly flushed with dry air at a dew-point of −70 °C. In addition, the Cold Boxes and bellows are wrapped in thin heating wires to maintain a stable outer temperature above the cavern dew-point [34].

---

[10] CAEN S.p.A., Viareggio, Italy.

[11] Either an environmental friendly Fluoroketone or a Perfluorohexane (C_6F_{14}) [34].

[12] To mitigate radiation damage in the SiPMs a coolant temperature of 30 °C is used to anneal the SiPMs during periods without beam at LHC.

4 Frontend Electronics Commissioning during Detector Assembly

The SciFi Tracker was assembled at the LHCb experimental site at LHC Point 8 between 2019 and 2022. While all individual components passed quality assurance (QA) tests during their production, integrating the various components into the C-Frames introduces new complexities and potential points of failure. The commissioning during the assembly of the detector before installing it in the LHCb experimental cavern, therefore, serves as a crucial step in identifying and addressing any issues that may arise. This includes verifying the functionality of the whole readout chain, as well as mechanical compatibility and the interface to the LHCb data acquisition system. In addition, the commissioning enables the validation of the calibration procedure and the first optimisation of detector settings.

After the complete assembly of each of the 12 C-Frames, a series of tests is carried out to ensure the correct functioning of the frontend electronics and to validate the whole signal chain from the SiPMs to the DAQ system. The test sequence is split into four stages, as shown in Figure 4.1. In the first stage, the basic functionality of the electronics is tested. This includes verifying the communication to the Master GBTx through the optical control links and configuring the various components within each ROB. In the second stage, the readings of the different temperature and voltage sensors on the electronics are checked. The last two stages involve testing the full data-taking chain. The third stage includes a careful tuning of the internal clocks of all components within a ROB to ensure error-free data transmission and subsequent bit error rate measurements. The final stage tests the entire data-taking chain by recording signals injected by the LIS. The data recorded in this final stage also allows for the initial calibration of the PACIFIC comparator thresholds.

While the initial commissioning procedure was already established, significant contributions have been made in the scope of this thesis to implement the last stage of tests involving the whole readout chain using the LIS. In addition, the responsibility to carry out the entire test procedure was gradually taken over. Therefore, a summary of the frontend commissioning of all 12 C-Frames is presented here. A comprehensive discussion of each step in the commissioning procedure and a detailed description of the setup and control system can be found in Reference [85].

In total 236 ROBs were tested during the assembly of the detector, including 472 control links, 3776 data links, and 483 328 channels. The remaining 20 ROBs were only installed after installation of the C-Frames in the experimental cavern due to time constraints.

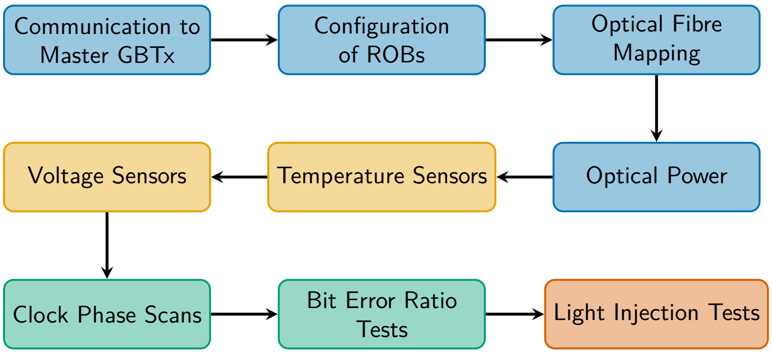

Figure 4.1 Test sequence of the frontend electronics commissioning during the assembly of C-Frames. The sequence is divided into four stages, progressing from basic functional tests to tests of the entire readout chain.

4.1 Commissioning Setup during Detector Assembly

The commissioning setup represents a slice of the final system that is installed in the LHCb cavern, see Section 3.7, and allows for the commissioning of a single C-Frame. It consists of a MARATON power supply with 12 channels that deliver the low voltage for the frontend electronics. Three CAEN power supplies with a total of 96 channels are used to bias the SiPM arrays, meaning each channel supplies four neighbouring SiPM arrays corresponding to one fibre mat. Cooling for the electronics is provided by a portable water chiller, while a small cooling plant cools the SiPMs to $-40\,^{\circ}\mathrm{C}$. A dedicated server is equipped with two CPUs and nine PCIe40 cards, of which eight are configured as TELL40s while the remaining one is configured as combined SODIN/SOL40. The TELL40s receive the data from the frontend electronics and merge the data from events belonging to the same bunch crossing. Each card supports up to 48 physical links; however, due to firmware limitations, the TELL40s could only process data from 24 links. As a result, only half of the, up to 384 links of one C-Frame can be read out simultaneously, and the commissioning steps that require recording data are performed twice. To avoid re-patching optical fibres for the two halves, two TELL40 firmware versions have been compiled, each with 24 active data links. The active links are split into the lower and upper half of the C-Frame. No such limitation exists for the SODIN/SOL40 firmware, and all ROBs can be controlled by one card.

4.2 Communication to Master GBTx

After powering up the frontend electronics with the MARATON power supply, the first test is to check the communication between the ROBs and backend electronics via the bi-directional control links. The Master GBTx on each Master Board recovers the clock from the incoming data stream on the control link and replies by sending its own GBT frames. If the header of the transmitted frames is recognised 15 times by the SOL40, the frames are locked, and communication between the frontend and backend has been established. If the header is missed more than four times without detecting

eight valid headers in-between, the communication is lost again [86]. The state of the links is monitored throughout the commissioning of the ROBs by dedicated counters on the SOL40 that count the number of clock cycles without frame-lock.
During the commissioning of all C-Frames, corresponding to 472 HalfROBs, only once the communication to the Master GBTx could not be established, and the corresponding HalfROB was replaced.

4.3 Configuration of ROBs

After the connection of the SOL40 to the ROBs has been successfully established, the internal communication protocols and control registers are tested by configuring the various components of the frontend electronics. The complete sequence to prepare the ROBs for data taking consists of a reset followed by a configure command. The reset puts the electronics in a well-defined state. First, the MasterGBTx is configured, and the reference clocks for all Data GBTx are set; after that, power to the Cluster and PACIFIC Boards is provided by enabling the DC/DC converters. Next, the Data GBTx are configured to provide the external clocks to the Cluster FPGAs and PACIFICs. Finally, a reset of the Cluster FPGAs and PACIFICs is issued, and the initial configuration of the PACIFICs is loaded.
During the commissioning of the first C-Frame, stability issues during the reset procedure were discovered that led to the loss of communication with the HalfROBs. It was found out that the low voltage would fall below the required minimum operational voltage of the DC/DC converters of 6.5 V due to excessive ringing. In order to stabilise the input voltage, specifically designed capacitor boards were added to each of the low-voltage supply lines on the C-Frames.
While the reset programs all ROBs with the same minimal configuration, the configure command prepares the ROBs for data taking by loading the individual settings for each of the Master Boards, Cluster Boards, and PACIFICs. These values are determined from the quality assurance tests during production and stored in a dedicated database. This includes the LIS bias and modulation currents. The configuration also contains the serial numbers of all boards, which are checked before applying the configuration. This is important to verify that the optical control links are connected to the intended HalfROBs. In addition, a 10-bit location identifier is written to a register of the HK FPGA that is used to verify the correct mapping of the data links that will be discussed in Section 4.4.
In total, the configuration failed for two HalfROBs due to the inability to write or read the LIS registers on the HK FPGA. The two affected HalfROBs were exchanged for spare ones.

4.4 Optical Fibre Mapping

The correct mapping of optical fibres to the control and data links on the HalfROBs is facilitated by the location identifier that is programmed on every HK FPGA during configuration. The ID of each data link is constructed in the Cluster FPGA from the

10-bit of the HalfROB ID[1] followed by three additional bits. These bits are set by dedicated pins on the connector between the Cluster Board and Master Board, which are either set to 1.5 V or ground depending on the location on the Master Board. The same principle is used to identify the two FPGAs on each Cluster Board.

The location ID is transmitted when initiating the data taking in special GBT frames called SYNC frames. These frames are used to synchronise the bunch crossing IDs in the TELL40s. They consist of the 12-bit bunch crossing ID followed by a fixed SYNC pattern of which the 13 least significant bits are used to transmit the data link ID. The TELL40s capture the SYNC frames, and the data link ID is stored in a dedicated register, where it can be read out and compared to the expected ID.

Like this, the mapping of the optical fibres for all C-Frames was verified, and only one issue was found. For one HalfROB, two neighbouring data links were sending incorrect IDs due to a faulty connection between the Cluster Board and Master Board, which caused one of the bits to be stuck at 0. The whole ROB was replaced.

4.5 Optical Power

Sufficient optical power on the control and data links is crucial for the detector's operation as it guarantees the signal integrity over the long optical fibres that connect the LHCb cavern and the DAQ system. Great care was taken during the production and QA of ROBs to keep the VTRxs and VTTxs clean. In addition, every optical fibre on the C-Frame was cleaned and inspected before being connected.

Based on a previous study [87] that evaluated the power budget for the LHCb experiment assuming 400 m long fibres from the detector to the data centre and three breakpoints, a limit of 224 μW is taken for the commissioning. In practice, a stronger limit of 300 μW was applied, and fibres that fell below it were inspected and cleaned. The optical power is measured by the optical transmitter and receiver modules of the PCIe40 cards with a precision of 10 % [85]. In total, only 8 out of the 472 tested control links and 11 out of 3776 data links measured an optical power of below 300 μW, corresponding to less than 0.5 % of all tested links. However, as can be seen from the histogram in Figure 4.2, all of the links are above the critical limits of 224 μW and therefore deemed operational.

4.6 Temperature Sensors

The frontend electronics are equipped with various temperature sensors, as mentioned in Section 3.6, totalling a number of 82 sensors for one ROB. The readings of all 19 352 sensors are checked to ensure that all of them are functional and provide accurate values. The estimated precision is in the range of ±2 °C for the Pt1000 sensors readout by the SCA on the Cluster Board; this includes the sensors found on the SiPMs, PACIFIC Board and Cluster Board itself [85]. For the NTC on the Master Board, the precision is ±3 °C. The uncertainty is mainly governed by the integral non-linearity of the ADC, the knowledge of the current source, and the tolerance of the sensors.

---

[1] The HalfROB ID is constructed from the station (2 bit), layer (2 bit), quarter (2 bit), module (3 bit), and HalfROB (1 bit) numbers.

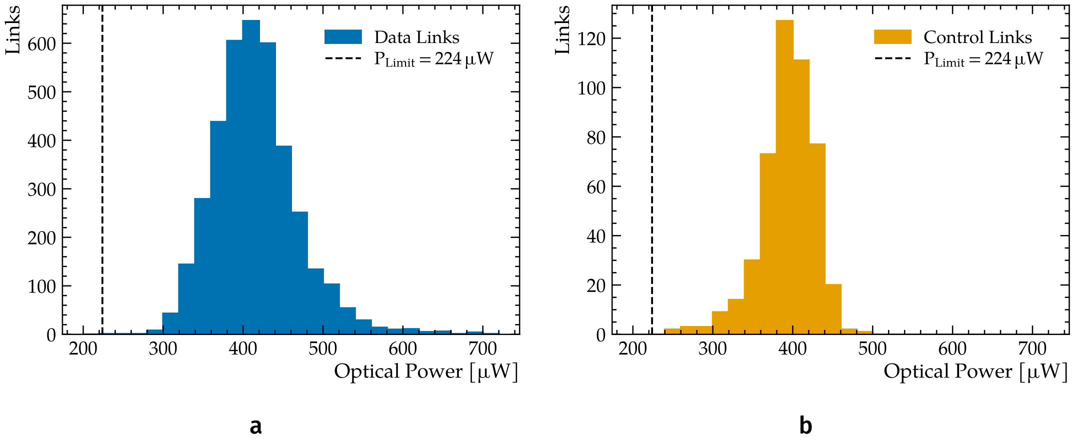

Figure 4.2 Results of the optical power measurements. Left: histogram of the optical power measured for each of the 3776 tested data links. Right: histogram of the optical power for all 472 tested control links. The dashed lines indicate the lower limit below which the optical transmission is expected to degrade.

The measured values of all sensors split up into Master Board, Cluster Board, PACIFIC Board, and SiPM temperatures, are shown in Figure 4.3. The Master Boards and Cluster Boards register similar average temperatures of around 27 °C, whereas the PACIFIC Boards measure higher temperatures of around 31 °C. This is due to the PACIFIC Boards being furthest away from the cooling block, which is located just behind the Master Board, and the PACIFIC drawing a significant amount of power. Importantly, none of the temperature sensors on the ROBs provide unreasonable values below the cooling water temperature of 20 °C, and the temperatures also do not exceed the maximum of 35 °C, indicating good thermal contact between the components and the cooling blocks. The temperature readings, therefore, also proved useful in detecting an incorrect mounting of the ROBs to the C-Frame.

The distribution of the SiPM temperatures shown in Figure 4.3d exhibits multiple peaks because no coolant was circulated during the commissioning of the electronics. The SiPMs were, therefore, subject to the temperature in the assembly building, with large variations throughout the year. Only one sensor was found to be defective, and upon inspection of the Pt1000 sensor mounted on the SiPM array, it was found that its resistance deviated by about 20 % from the nominal value leading to an unreasonable reading of 80 °C. It was decided to leave the affected SiPM on the C-Frame and mask the sensor in the control system.

4.7 Voltage Sensors

As introduced in Section 3.6, each HalfROB features one monitoring circuit on the Master Board for the supply voltage and 16 circuits to monitor the bias voltage of each SiPM die.

The input voltage of the HalfROBs is monitored through a voltage divider by the ADC of the Master SCA with a precision of about ±2 % [85]. A minimum voltage of 6.5 V is required to power the FEASTMP DC/DC converters and ensure a stable operation of

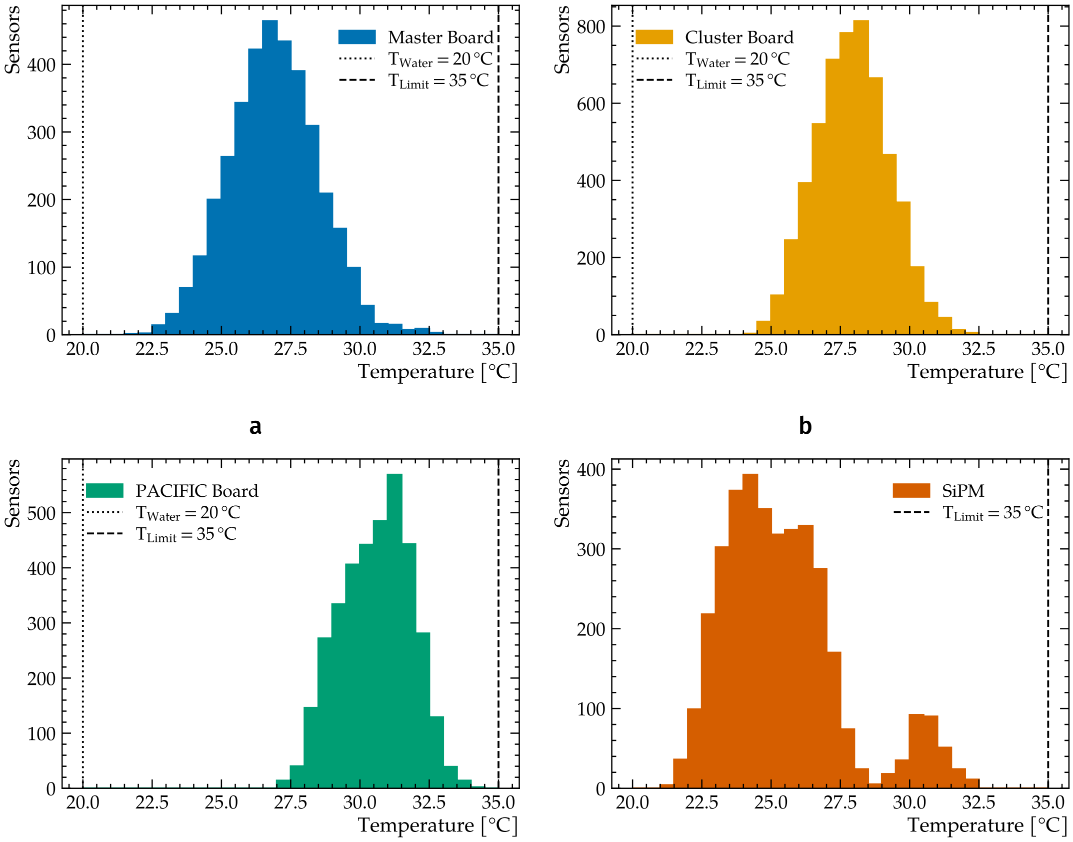

a

b

c

d

Figure 4.3 Histograms of the temperature readings for all temperature sensors on the tested Master Boards, Cluster Boards, PACIFIC Boards, and SiPMs. The dotted lines show the set point of the water cooling plant, while the dashed lines show the upper temperature limit.

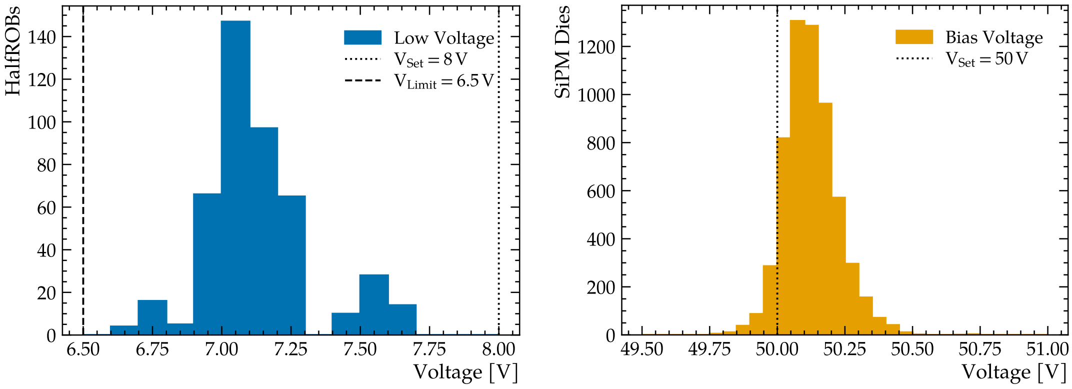

a

b

Figure 4.4 Left: histogram of the monitored low voltage supplied to the frontend electronics. The voltage set at the power supply is indicated as dotted line. The lower limit below which the frontend electronics are not able to operate is shown as dashed line. Right: histogram of the measured SiPM bias voltages with the set point indicated by the dotted line.

the frontend electronics. A histogram of the measured voltages is shown in Figure 4.4a. The voltage is always below the set point of 8 V due to the voltage drop in the cables. The vast majority of monitored voltages are around 7.1 V; in addition, two distinct peaks are observed. During detector assembly, one MARATON channel powers two ROBs, drawing a combined current of about 20 A. However, stations T1 and T2 are equipped with fewer modules, so the outermost ROBs do not share a low-voltage channel with their neighbours. As only half the current is drawn over the channel, only half the voltage drop occurs over the cables, and the monitored voltage is slightly larger. The small peak to the left in Figure 4.4a is due to a MARATON channel delivering a voltage of less than 8 V. As a result, the monitored voltage for the connected HalfROBs is slightly smaller but still above the limit and stable throughout the commissioning.
The SiPM bias voltage is monitored through voltage divider circuits located on the PACIFIC Boards and readout by the Master SCA ADC. Each SiPM die (64 channels) features its own bias line and monitoring circuit. In order to verify the correctness of the monitored voltage the bias voltage at the CAEN power supplies is set to 50 V. Figure 4.4b shows the distribution of the monitored voltages. As can be seen the distribution is not centred around the nominal voltage of 50 V, indicated by the dashed line. On average the monitored voltage is about 0.15 V higher. At this stage, however, the monitored voltages are considered okay if they do not differ by more than ±1 V. The calibration of the conversion factor of the Master SCA ADC values to the monitored voltage was postponed to a later stage; it is described in detail in Reference [85]. Only once a monitored voltage outside of the required range was measured. The monitoring circuit of one SiPM die reported a value of 0 V. The SiPM die was indeed not biased due to a broken bias line within the corresponding ROB; it was therefore replaced.

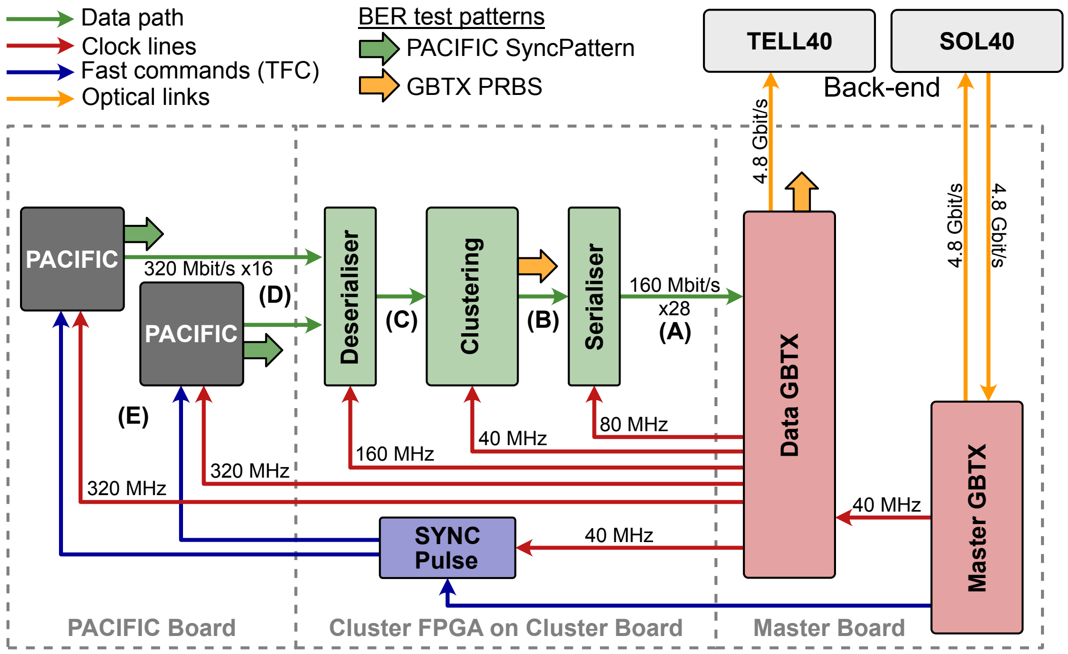

Figure 4.5 Illustration of the clocking scheme and data flow in one data link of the SciFi frontend electronics [85]. At each of the clock domain crossings labelled A to E the clock phases are adjusted to ensure an error-free data transmission.

4.8 Clock Phase Scans

The different parts of the frontend electronics are driven by a set of clocks of various frequencies that all need to be carefully tuned to achieve an error-free data transmission between the different digital processing stages. Figure 4.5 illustrates the distribution and interplay of clocks for one data link. Going from left to right, the SiPMs signals are digitised by the two PACIFICs and transferred at a rate of 320 Mbit/s to the Cluster FPGA. Inside the Cluster FPGA, the data is first deserialised before the clustering algorithm is performed. In the last stage, the clusters are serialised at 160 Mbit/s and transferred to the Data GBTx from where they are shipped to the TELL40 at 4.8 Gbit/s. The data taking is synchronised by the SYNC pulse that resets the PACIFIC clocks and ensures that the 40 MHz and 20 MHz in the PACIFIC are generated synchronously.
The various clocks provided to each stage in the frontend electronics are derived on the Data GBTx from the same 40 MHz reference clock of the Master GBTx. The clocks derived from the Data GBTx range from 40 to 320 MHz and need to be tuned to achieve the highest possible stability in the transfer of data between the different processing stages. The phase of the Master GBTx clock does not have an impact during this time of the commissioning but will be relevant during the actual data taking at the LHC, as it allows the fine-tuning of the phase with respect to the bunch crossings.

The clocks are tuned by injecting a known pattern into the data stream, scanning through the phase and counting the number of errors in the pattern received by the TELL40s. Practically, the clocks are scanned from right to left in Figure 4.5, fixing the

determined optimal phases one by one. Two types of test patterns can be injected at three different points, as indicated by the thick arrows in Figure 4.5: at the Data GBTx, the Cluster FPGA, and the PACIFIC.
The GBTx PRBS pattern consists of a 7-bit pseudo-random binary sequence (PRBS) that is copied 16 times to fill the 112 bit of the GBT frame. It can be generated by the Data GBTx and the Cluster FPGA and allows for probing of the processing chain starting from the clustering stage in the Cluster FPGA. A corresponding decoder in the TELL40 firmware compares the received data to the expected pattern and counts the number of mismatches, *i.e.* bit errors, and stores them in a register. For each of the 16 PRBS words within a GBT frame a separate 8-bit counter exists.
The second pattern called the PACIFIC SyncPattern, not to be confused with the SYNC TFC signal, is generated by the PACIFIC. It consists of a 4-bit pattern that can be set by the Cluster FPGA via special pads on the PACIFIC and is duplicated to produce the 8-bit word at the PACIFIC output. A special block in the TELL40 firmware decodes and analyses the received data and counts the number of bit errors. For each PACIFIC channel, corresponding to two output bits, a separate 8-bit counter exists. The PACIFIC SyncPattern is passed through the Cluster FPGA without performing clustering. Due to bandwidth limitations in the FPGA only the outputs of 32 PACIFIC channels can be tested at the same time. Tests involving the PACIFIC SyncPattern, therefore, need to be performed four times to cover all 128 channels connected to one Cluster FPGA.

An example of a clock phase scan between the Cluster FPGA serialiser and Data GBTx is shown in the left panel of Figure 4.6 for one 7-bit PRBS word inside the GBT frame. Within the phase range covered by the filled area, no errors in the pattern have been observed, while on the outside, the number of bit-errors is at the maximum value of 255. The results for all 16 data links (256 PRBS words) of one exemplary ROB are shown in the right panel of Figure 4.6, where variations in the width and centre of the error-free interval are observed. However, a common interval without errors, indicated by the dashed lines, can be identified, allowing for the operation of all data links stably with the identical clock phase setting.

4.8.1 FPGA Serialiser Clock (A)

The Cluster FPGA serialiser clock is the first clock that is tuned; it drives the transmission of data from the Cluster FPGA serialiser to the Data GBTx. As shown in Figure 4.5 the GBTx PRBS pattern is used for this scan. Since this pattern is generated in the Cluster FPGA, the clustering and serialiser clock must be moved synchronously.
Figures 4.7a and 4.7b show a summary of the centres and widths of the error-free intervals for the PRBS corresponding to all tested data links. As shown in Figure 4.7a, the central values are all within the common interval indicated by the dashed lines. It should be noted that the absolute values of the central values have no further significance; rather, the relative differences are important. The interval widths shown in Figure 4.7b measure the signal quality and have a theoretical limit given by the inverse of the transmission frequency shown by the dashed line. The distribution has a peak around 5.24 ns and a shoulder around 4.6 ns from electrical lines with a slightly worse signal quality.

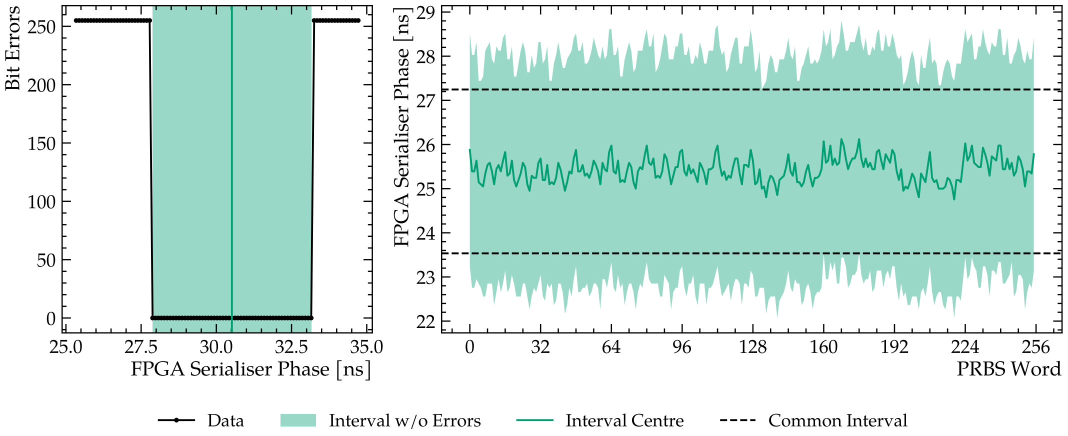

Figure 4.6 Left: bit errors for one PRBS word depending on the phase between the Cluster FPGA serialiser and the Data GBTx. Right: clock phase intervals without errors for all 256 PRBS words (16 words per data link) of one exemplary ROB. The common interval is indicated by dashed lines.

4.8.2 FPGA Clustering Clock (B)

After the FPGA serialiser clock has been tuned and set to the central value, the next clock that is scanned is the one between the Cluster FPGA clustering and serialiser blocks. The distributions shown in Figures 4.7c and 4.7d differ significantly from the previous ones. The different peaks can be attributed to the different words inside one GBT frame, and entirely depend on the routing of signals in the FPGA. Importantly, the central values lie well within the common overlap region with plenty of margins and thus allow for a stable and error-free transmission.

4.8.3 FPGA Deserialiser Clock (C)

Starting from the FPGA deserialiser clock, the PACIFIC SyncPattern is used for the clock scans, as the GBTx pattern is generated after the deserialiser block. The clustering clock is kept fixed to the previously determined optimal value, and only the deserialiser clock is varied. The PACIFIC and SYNC pulse clocks are varied synchronously to avoid any interference. The distributions of the interval centres and widths shown in Figures 4.8a and 4.8b indicate good quality of the data transmission, the central values are well within the common overlap, and the width is very close to the theoretical maximum of 6.25 ns. The variations in the centres and widths of the error-free interval are again due to the routing of the signals inside the FPGA and depend on the specific firmware.

4.8.4 PACIFIC Clock (D)

Fixing the value of the FPGA deserialiser clock, the 320 MHz PACIFIC clock can be scanned, and as before, the SYNC pulse clock needs to be moved accordingly. The theoretical upper limit for the interval widths is 3.125 ns. Due to the routing of the

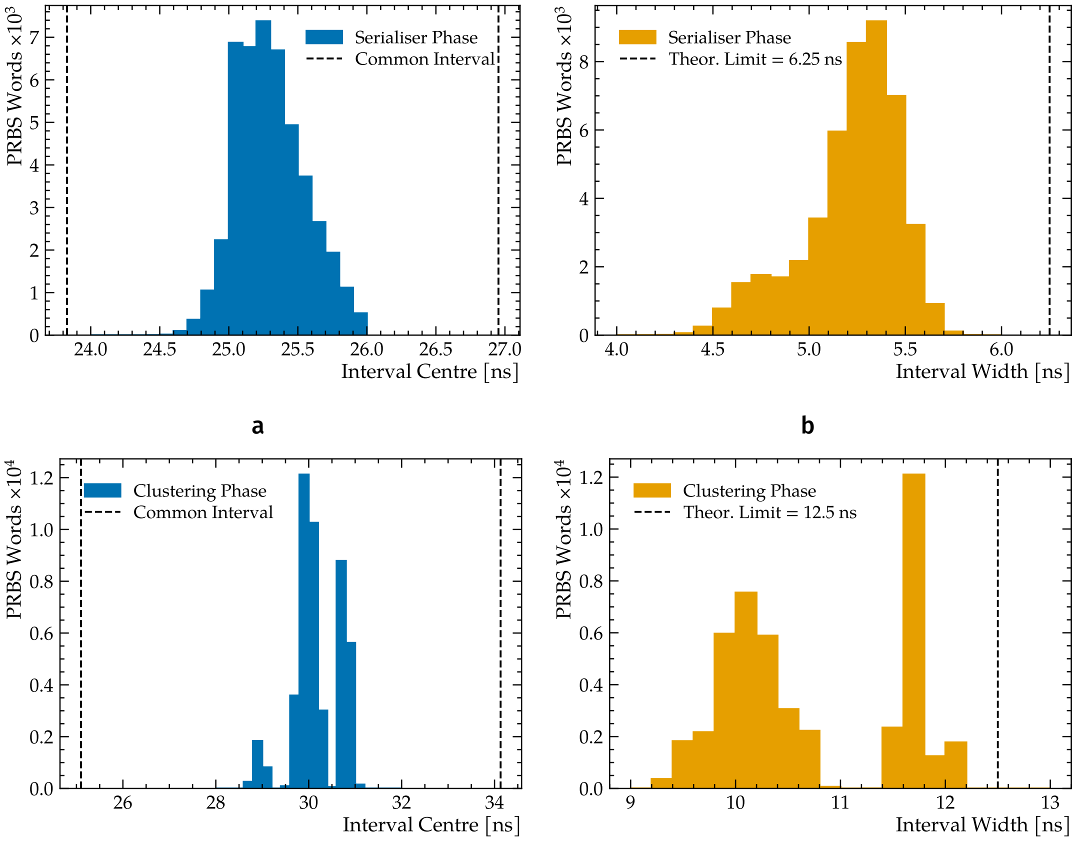

a b

c d

Figure 4.7 Results of the clock phase scans using the GBTx PRBS pattern for the 236 tested ROBs (3776 data links, 60 416 PRBS words). Left: histogram of the interval centres. The dashed indicate the common overlap of all intervals. Right: histogram of the interval widths with the dashed line indicating the theoretical limit derived from the respective transmission speeds.

signals across multiple PCBs with variation in the lengths of the signal lanes, the measured error-free intervals exhibit large variations and the common interval amounts to 1.2 ns as indicated by the dashed lines in Figure 4.8c. The distribution of the measured interval centres, shown in the same figure, extends over the common interval. However, this is not an issue, as the interval widths, shown in Figure 4.8d, are very close to the theoretical limit of 3.125 ns, leaving enough margin to operate all PACIFICs with a common phase setting.

4.8.5 SYNC Pulse Clock (E)

In the last step, the SYNC pulse clock is tuned. Although the SYNC pulse is not strictly part of the data transmission chain, it plays a critical role in synchronizing data taking across all PACIFIC ASICs. As mentioned in Section 3.6.1, the SYNC pulse is sampled by the PACIFIC at a frequency of 320 MHz. If the SYNC pulse clock is varied by more than 3.125 ns, the signal is evaluated on the previous or next rising edge. This causes the PACIFIC output and, therefore, the test pattern to shift by one clock cycle, leading to bit errors at the receiving end. Figure 4.8e shows that the measured centres of the error-free intervals have a large spread, due to variations in the trace length from the Cluster FPGA (where the SYNC pulse is generated) to the PACIFIC. This is also reflected in the small common interval of 0.9 ns. The interval widths, shown in Figure 4.8f, however, are close to the theoretical limit of 3.125 ns, so that a common phase setting can be used for all SYNC pulse clocks.

4.9 Bit Error Ratio Tests

After tuning the clocks of the frontend electronics, the data transmission quality is evaluated by measuring the bit error ratio (BER). This includes the transmission through the frontend electronics and the transmission via the optical links, whose output power has been validated in a previous step, presented in Section 4.5. The bit error ratio is defined as the number of bit errors divided by the total number of transmitted bits. Following the specifications of the Versatile Link in Reference [87] the BER should not exceed $<10^{-12}$.

The BER is measured by injecting the PACIFIC SyncPattern into the data stream and recording the number of total bits send and the number of bit errors detected at the TELL40. In order to acquire high enough statistics to probe the BER, the test is run for three hours. As the Cluster FPGA has limited bandwidth, four sets of measurements are recorded to cover all 128 channels of one data link. Additionally, the tests are carried out with two different versions of the PACIFIC SyncPattern: `0x5` (`0b10100101`) and `0xA` (`0b01011010`). Therefore, covering all channels using both SyncPatterns takes 24 hours. Figure 4.9 shows the distribution of the bit error ratios for all 3776 tested data links. Most data links measure a BER in the range of 10^{-15} to 10^{-14}, and only a few links are close to the limit of 10^{-12}. In total, about 8.005×10^{17} bits were probed, and 4054 errors were observed, giving an average BER over all data links of $\approx 5.1 \times 10^{-15}$. Importantly, none of the data links have a BER above the limit, indicating a high-quality data transmission in all 3776 tested data links.

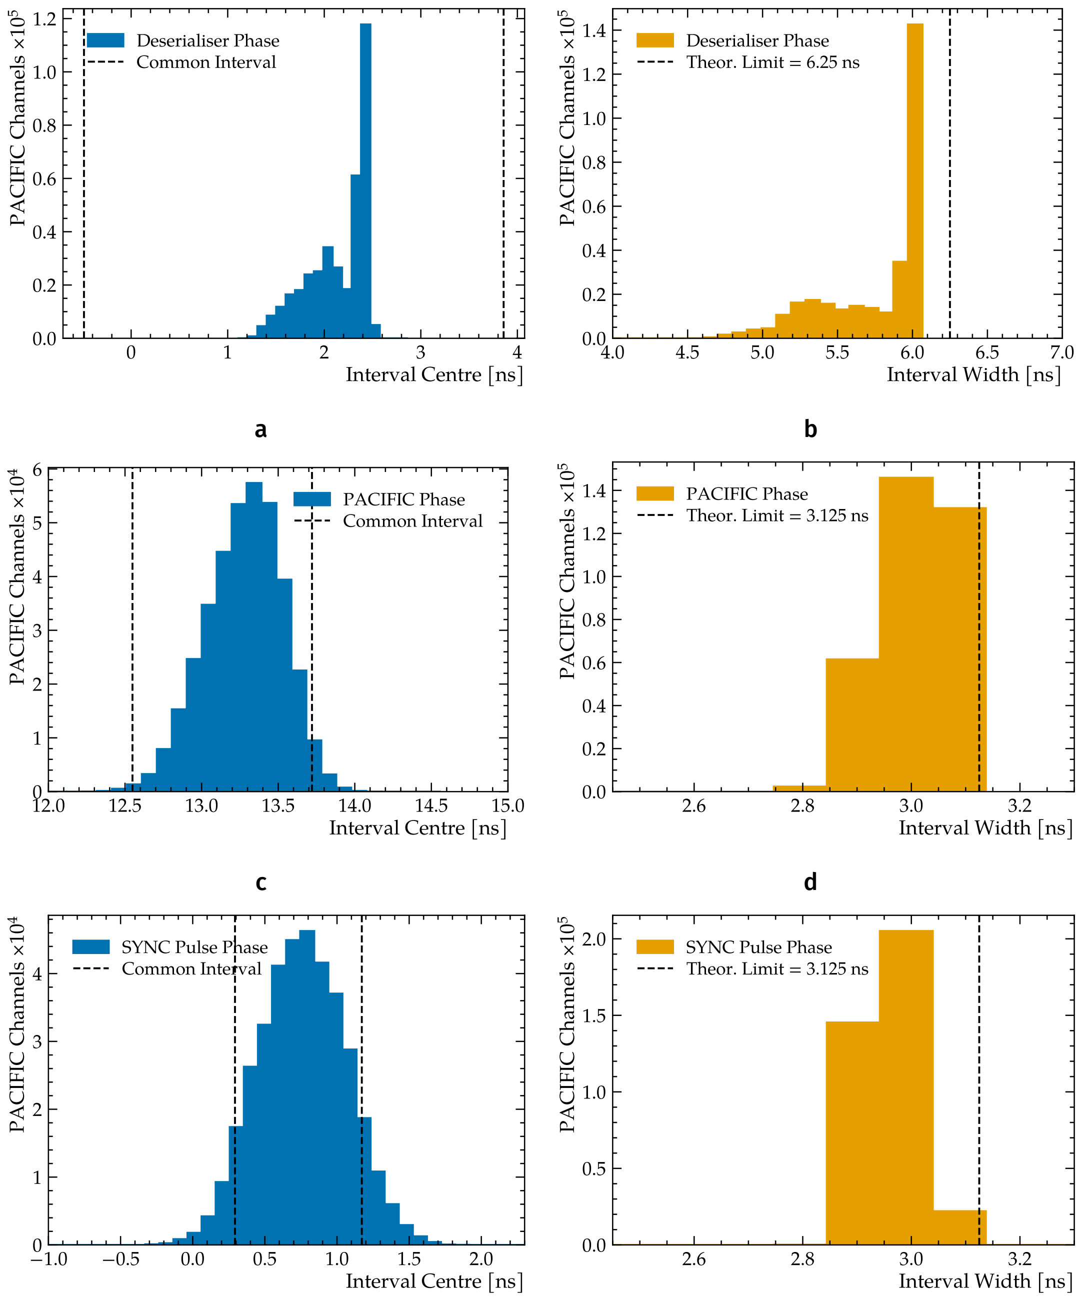

a

b

c

d

e

f

Figure 4.8 Results of the clock phase scans using the PACIFIC SyncPattern for the 236 tested ROBs. Left: histogram of the interval centres. Right: histogram of the interval widths.

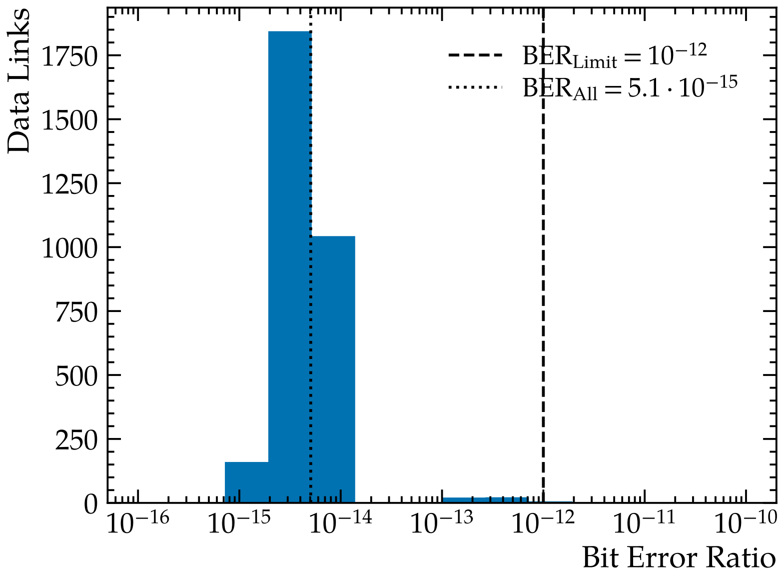

Figure 4.9 Distribution of the measured bit error ratios for all 3776 tested data links. The required operational limit of 10^{-12} is shown as dashed line while the combined bit error ratio of all links is indicated by the dotted line.

4.10 Light Injection System Tests

The last tests in the frontend electronics commissioning are designed to probe the entire signal chain. Whereas the digital processing has been tested and tuned during the clock phase scan discussed in Section 4.8, the light injection system tests also probe the analogue signal generation and processing. As described in Section 3.6.5, the light injection system (LIS) allows for injecting light pulses into the fibre mat endpieces from where the SiPM arrays detect them. The LIS tests, therefore, allow for identifying broken SiPM or PACIFIC channels. In addition, a preliminary calibration of the PACIFIC comparator thresholds can be performed.

4.10.1 LIS Mapping Test

In the first step, the correct functionality and mapping of the LIS drivers is checked. For this test, the SiPM arrays are biased at the nominal 3.5 V overvoltage, and their currents are monitored. The dark current of one SiPM bias channel, corresponding to four adjacent SiPMs, is typically around 100 µA and increases by about 200 µA when illuminated by the LIS with a frequency of 20 kHz.

In order to further tune the LIS, the correct mapping between the fibre mats and the LIS GBLDs needs to be confirmed to ensure that the settings determined during the QA of the LIS are written to the intended device. As shown in Figure 4.10, the LIS controlled by HalfROB *H0* is connected to the fibre mats 0 and 2 while *H1* is connected to mats 1 and 3. Unfortunately, it is not known which driver connects to which fibre mat due to an oversight during the production of the fibre modules. A first mapping check was done during the LIS QA, however, an additional test was performed here.

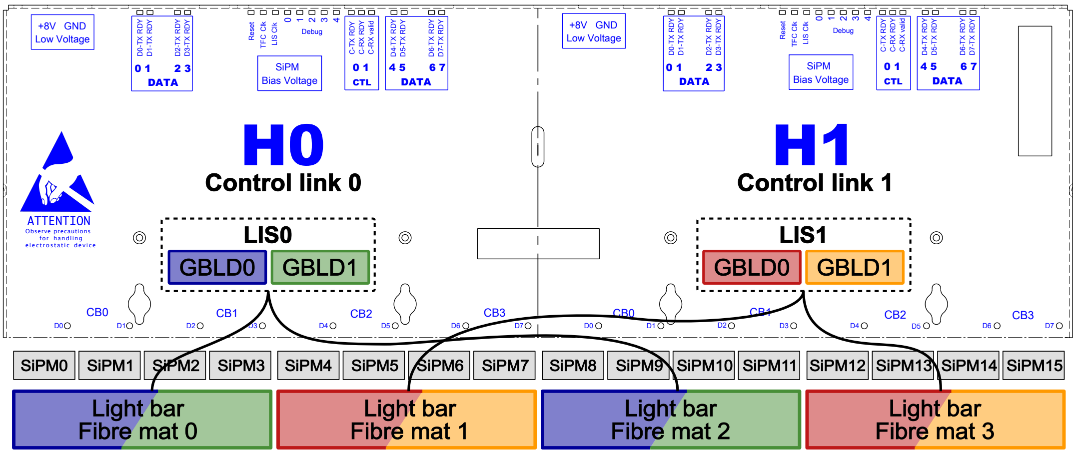

Figure 4.10 Schematic of the SciFi ROB cover plate with the routing of the LIS to the fibre mats [85].

To determine the mapping, only one of the two GBLDs on each LIS is enabled, and the currents are monitored again.

4.10.2 LIS Delay Scans

In order to verify that all channels detect incident light using the LIS, it is first necessary to ensure that the light pulse is triggered at the right time with respect to the PACIFIC clock so that the SiPM signals are fully integrated. This is done by varying the delay as described in Section 3.6.5 and measuring the signal amplitude in a LIS delay scan procedure.

The complete procedure consists of recording threshold scans under pulsed illumination from the LIS at different LIS delays and extracting the photoelectron gain for each of them. It is explained in detail in Chapter 6 and is crucial for properly calibrating the PACIFIC comparator thresholds. However, at the time of the detector assembly, a different approach to optimising the LIS delay was used as the complete procedure was too time-consuming due to limitations in the PCIe40 firmware and control software. In the simplified version of the LIS delay scan, the thresholds are fixed while the LIS delay is varied. As no proper calibration of the thresholds existed during the scan, any combination of thresholds is suitable as long as they are above the pedestal. For the delay scans, the three comparator thresholds were set to 15, 30, and 45 DAC above the pedestal. These threshold settings roughly correspond to 1.5 pe, 2.5 pe, and 3.5 pe, respectively, based on the typical gain of 15 DAC per photoelectron when operating the SiPM at an overvoltage of 3.5 V. The position of the pedestal is determined from a threshold scan without light injection. An example of such a scan is shown in Figure 4.11, where the ratio of events over threshold divided by the number of all events is shown as a function of the comparator threshold. The pedestal is given by the point where the ratio equals 0.5 as indicated by the black dashed line. In addition the three comparator thresholds calculated from the pedestal are shown as colourful dashed lines.

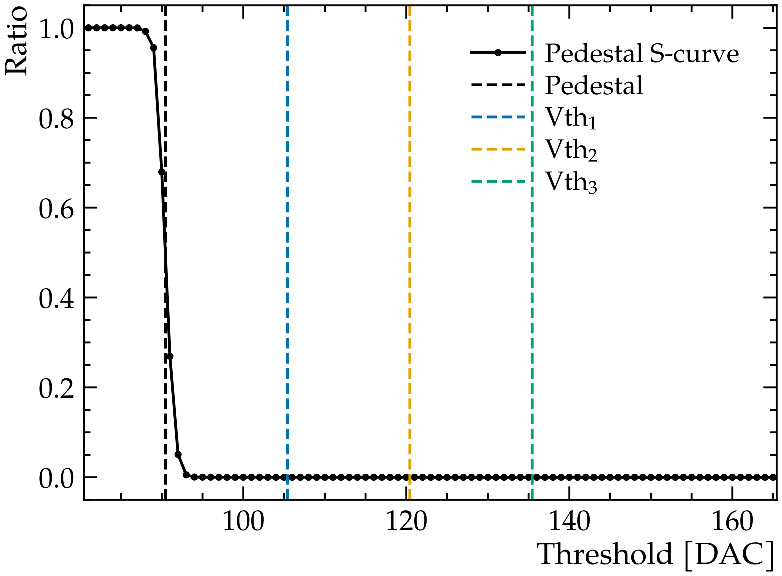

Figure 4.11 Example of a threshold scan without light injection. The pedestal position, indicated by the black vertical dashed line, is given by the point where the S-curve reaches a ratio of 0.5. The three threshold values are shown as colourful dashed lines.

Depending on the delay between the LIS pulse and the PACIFIC, the signals from the SiPMs will be more or less integrated. By keeping the thresholds constant, the signals have a higher or lower chance of passing the thresholds. Therefore, the ratio of events over the threshold divided by the total number of collected events can be used as a proxy for the signal amplitude. Figure 4.12a shows the ratio of events over the three thresholds Vth_1, Vth_2, and Vth_3 as a function of the LIS delay. Since one LIS driver injects the light over the width of one fibre mat, the ratio is averaged for the 512 channels corresponding to one fibre mat. The maximum of the curve gives the optimal LIS delay. For all three thresholds the same optimal delay is found (as it should). The distribution of the optimal LIS delay for all 2048 channels of one ROB is shown in Figure 4.12b split by the four fibre mats. The outer fibre mats (0 and 3) register the light from the LIS at the same optimal delay. Due to differences in the latency of the Calib C command used to trigger the LIS, the inner two mats, which are controlled by one HalfROB but read out by the other, register the light one bunch crossing earlier or later. This is an unwanted artefact of the routing of the LIS which needs to be accounted for when performing the calibration of the comparator thresholds discussed in Chapter 6.

4.10.3 LIS Threshold Scans

Since the PACIFIC digitises the SiPM signals using a simple set of comparators and not a full ADC, as described in Section 3.6.1, threshold scans under pulsed illumination need to be performed to measure the pulse height spectrum of the SiPM. At each step in a threshold scan, several thousand events are recorded, and the number of events over the threshold is counted while injecting light pulses of ≈15 ns length with the LIS.

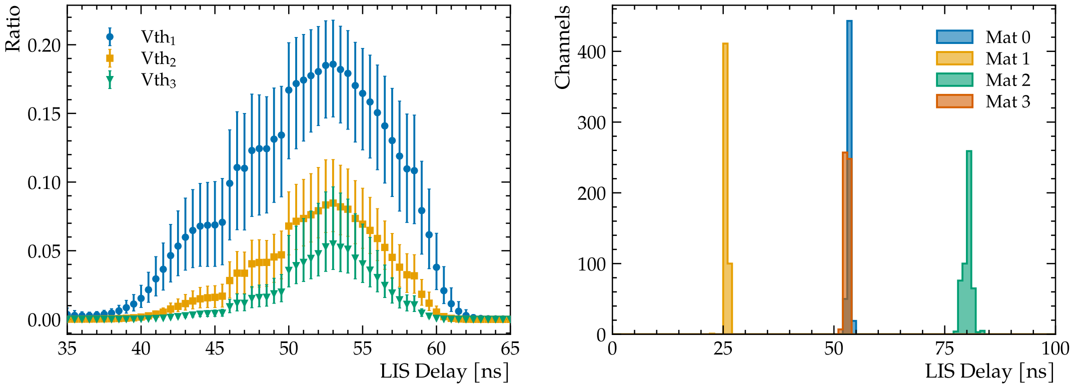

a **b**

Figure 4.12 Left: measurement of the ratio of events over threshold as a function of the LIS delay. Each point is the average over the 512 channels of one fibre mat, with the error bars indicating the standard deviation. Right: distribution of the optimal LIS delays for all 2048 channels of one ROB. Channels belonging to the different mats are shown with different colours.

The resulting curve is called *S-curve*.
Threshold scans with light injection from the LIS are performed using the optimal LIS delays determined in the previous step. As shown in Figure 4.12b multiple threshold scans have to be done at different LIS delay settings to ensure that light pulses are injected at the correct time for all channels. In principle, the previous test also allows for the checking of dead SiPM or PACIFIC channels by verifying that all channels register a significant ratio of events over the threshold. However, performing a threshold scan with light injection allows checking the individual S-curves and exercising the calibration procedure explained in detail in Chapter 6.

S-curve Integral

In order to verify if a channel is broken, a simple but robust metric has been constructed from the S-curves. The so-called *S-curve Integral* measures the area below the S-curve starting from the pedestal position, as shown in Figure 4.13a. If the channel does not register any light, the S-curve Integral measured from the S-curve with (LIS) and without (Pedestal) light injection will be identical. Figure 4.13b shows the S-curve Integral as measured for all 483 328 tested channels. A clear separation between the distributions of the S-curve Integral with and without light injection can be seen, with only three channels showing no difference. These three channels did not register light and were found to be broken already during the SiPM QA. Putting the numbers into relation, this corresponds to a quota of 6.2×10^{-6} broken channels.

LIS Intensity

During the quality assurance of the LIS, the bias and modulation currents of the drivers were tuned to deliver an intensity of about 1.5 pe on average over all 512 channels corresponding to one fibre mat. Assuming a Poisson distribution for the

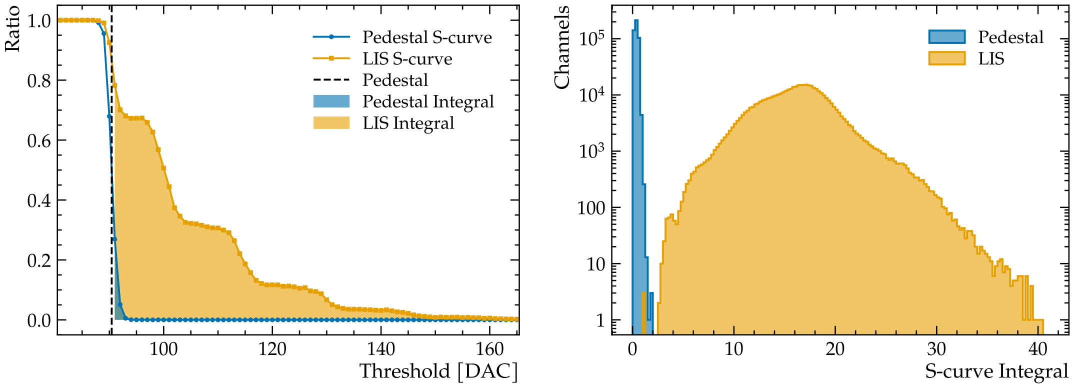

a

b

Figure 4.13 Left: exemplary S-curves with (LIS) and without (Pedestal) light injection for one comparator. The S-curve integral is shown as shaded areas and the pedestal position is indicated as dashed line. Right: distribution of the S-curve integral with (LIS) and without (Pedestal) light injection for all 483 328 tested channels.

detected photons, this intensity is needed to measure 5 pe signals with sufficiently large statistics. Being able to measure the 5 pe signals is vital as the highest of the three PACIFIC comparator thresholds is nominally set to 4.5 pe.

The light intensity is measured using a fit to the S-curve of each channel. The fit model is described in detail in Chapter 6. Figure 4.14a shows the measured light intensity of one light injection bar over the 512 channels corresponding to one fibre mat. Because the optical fibres of the LIS are scratched by hand, the light intensity is not uniform across all channels. This is not an issue as long as all channels see enough light (>1 pe). The light intensity for all 483 328 channels is shown in Figure 4.14b. The average intensity is 1.45 pe with a standard deviation of 0.39 pe. For the majority of the channels the LIS, therefore, delivers enough light for a successful calibration of the comparator thresholds up to 5 pe. However, a small tail extending down to 0.5 pe can be seen with a potentially negative impact on the calibration. If the light intensity is too low events with higher photoelectron amplitudes are less likely, and due to the lower statistics the calibration of these thresholds becomes unreliable or even impossible. This will be further discussed in Chapter 6.

4.11 Commissioning Summary

During the assembly of the SciFi Tracker from 2019 to spring 2022, 12 C-Frames with a total of 236 Readout Boxes have been successfully tested. The remaining 20 ROBs were installed after the installation of the C-Frames in the LHCb experimental cavern.

The commissioning procedure described in this chapter allowed for verifying that the whole electronics chain from SiPMs to the DAQ works as designed. The first threshold scans with light injection from the LIS were also recorded. From these measurements, the average light intensity of the LIS can be determined. A high enough intensity is crucial for successfully calibrating the PACIFIC comparator thresholds. The intensity was within the desired range with an average of 1.45 pe and only three out of 483 328

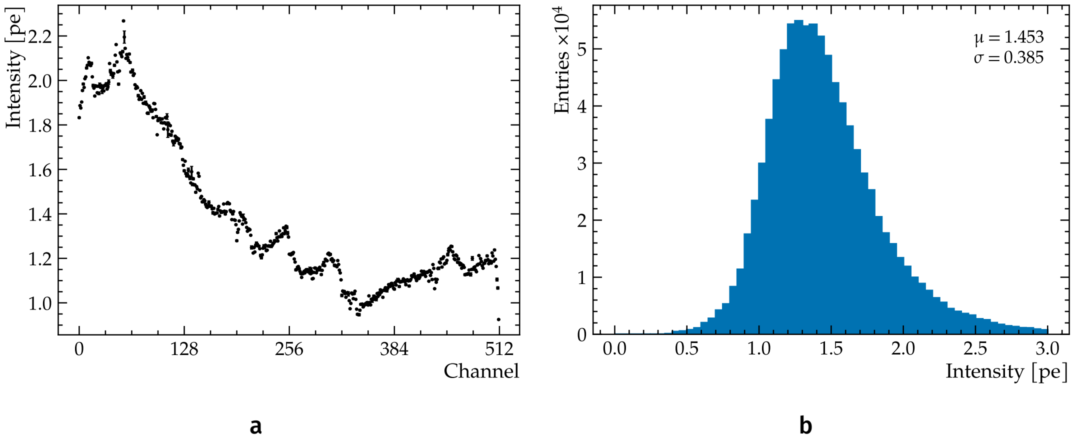

Figure 4.14 Left: LIS light intensity for the 512 channels of one fibre mat. Right: distribution of the light intensity for all 483 328 tested channels.

tested channels were found to be broken

Phase scans of the internal clocks of the frontend electronics were performed and showed that a stable and sufficiently error-free data transmission with a bit error ratio of 5.1×10^{-15} can be achieved using a common set of phase settings across all tested data links.

Some shortcomings, *e.g.* in the design of the power delivery, were identified, and solutions that address the issues were implemented. ROBs that failed to pass a certain test were replaced with spare ones and repaired if possible. In addition to verifying the functionality of the electronics, the commissioning also allowed for gaining first experience in the operation of the detector on a large scale and enabled crucial developments in the control software, various firmwares, and analysis tools.

5 | First Operational Experience in Run 3

The SciFi Tracker was successfully installed in the LHCb experimental cavern before the start of Run 3 in April 2022. With the start of Run 3 also a new period for the SciFi Tracker started. The whole system underwent the same commissioning procedure as presented in the previous chapter, verifying that the installation went as planned and that all frontend electronics work as expected. At the same time the whole detector was integrated into the LHCb control and DAQ systems, to allow for synchronous data taking with all sub-detectors of LHCb.

Using a preliminary calibration of the PACIFIC comparator thresholds determined from the light injection threshold scans performed during the assembly phase of the SciFi Tracker the first data from proton-proton collisions was recorded. Figure 5.1 shows the number of reconstructed clusters in each of the 4096 SiPMs of the SciFi Tracker from a 2022 data-taking run. The 12 layers, with their four quadrants Q0 to Q3, are shown one after each other along the y-axis. The SiPMs of each quadrant are shown along the x-axis going from the beam pipe at $x = 0$ to the outside of the detector. The number of cluster decreases, as expected, from the beam pipe, indicated by the black circle in the center of each layer, towards the outside of the detector. The very central region of the detector is also expected to see less clusters due to the fibre mats being slightly shorter to accommodate the beam pipe. In general the figure shows that the SciFi Tracker works and is able to reconstruct hits from real particles. However, it is also evident that the detector is not optimised yet. Several areas with a significantly larger number of detected clusters can be seen, *e.g.* in Q2 of layer 2. Due to the PACIFIC comparator thresholds settings being too low, in some cases even below the pedestal, channels in these areas produce an excessive amount of noise clusters. In contrast the thresholds in some other parts of the detector are too high resulting in a lower hit detection efficiency and, therefore, lower number of cluster than expected. Additionally, three SiPMs show no clusters (small white rectangles), since the corresponding data links are disabled and excluded from the data taking due to transmission issues.

In parallel to the commissioning of the detector, the SciFi Tracker is regularly included in the global LHCb data taking, enabling the reconstruction of particle tracks and allowing for a first look at the tracking performance. Figure 5.2a shows the distribution of the number of SciFi hits per reconstructed tracks; here *T tracks* are shown, which are reconstructed using only the clusters of the SciFi Tracker. The distribution peaks at value of 10 hits per track with an average of 10.34. In simulation (dashed line), where the hit detection efficiency of the SciFi is assumed to be 99 % and the detector is perfectly aligned, the same distribution peaks at 12 hits per track, *i.e.* each layer provides a hit to the track. Looking at the distribution of the number of hits per track per SciFi layer in Figure 5.2b, it becomes apparent that some layers, such as layer 7, are particularly inefficient. This is predominantly due to high comparator thresholds of (2.5, 3.5, 4.5) pe based on the preliminary calibration results from the assembly of C-Frames, and sub-optimal alignment. The calibration of the thresholds, which will also allow for the nominal lower threshold settings, is the topic of the next chapter.

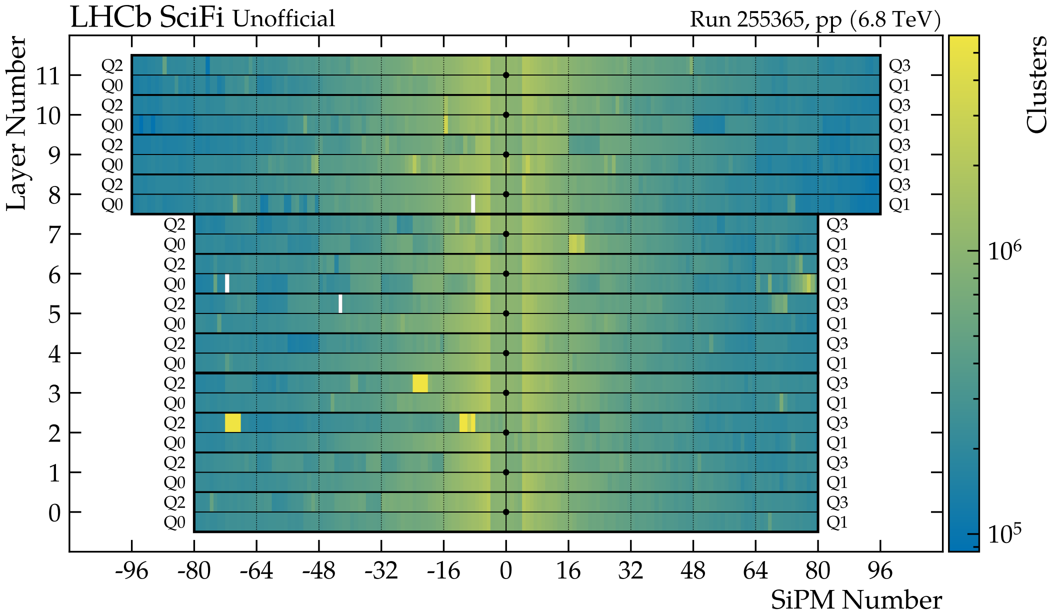

Figure 5.1 Overview of the number of reconstructed clusters per SiPM in the SciFi Tracker during 2022 data taking. Layers are numbered from 0 to 11 and ordered with increasing *z* positions. The SiPMs are numbered from 1 to 96 or −1 to −96 according to their *x* position within the LHCb coordinate system. Positive indices correspond to positive *x* while negative indices correspond to negative *x*.

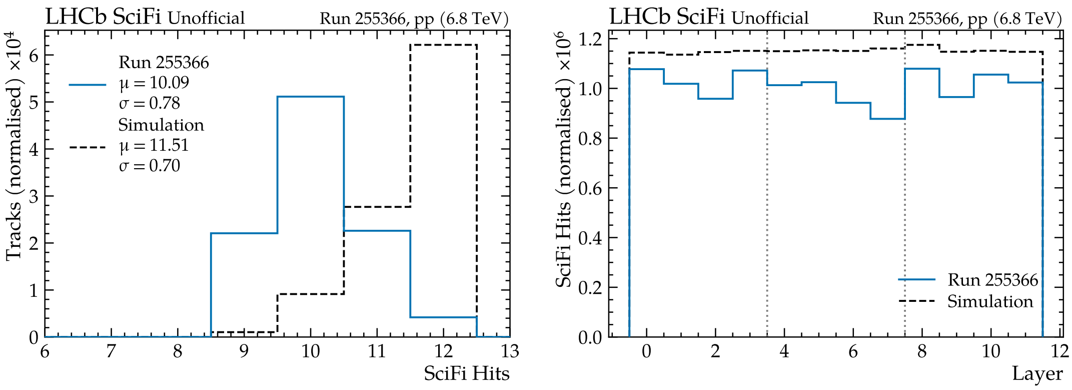

a

b

Figure 5.2 Distributions of the number of SciFi hits per T track for data recorded during Run 3 in 2022 and from simulation. All distributions have been normalised to the same number of tracks. Left: distribution of the number of SciFi hits per track. Right: distribution of the number of SciFi hits per track for each layer. Layers are numbered from 0 to 11 and ordered with increasing *z* positions.

6 Frontend ASIC Calibration

The frontend ASIC of the SciFi Tracker, called PACIFIC, responsible for processing and digitising the SiPM signals, does not implement a full ADC, but rather digitises the SiPM signals with a set of three comparators per channel (see Section 3.6.1). To reconstruct the particle hits and suppress noise, the outputs from multiple PACIFIC channels are combined into clusters in the Cluster FPGA using the algorithm described in Section 3.6.3. As no full amplitude information of the signal in each channel is available for the clustering algorithm, the PACIFIC comparator thresholds must be calibrated to correspond to the equivalent of signals with 1 pe, 2 pe, 3 pe, and so on.
The calibration of the comparator thresholds consists of two parts. First, threshold scans under pulsed illumination from the LIS are performed for each of the three comparators of all 524 288 channels. The data is then fitted based on an analytical description of the SiPM pulse height spectrum. The fit allows for determining the threshold values corresponding to the discrete photoelectron amplitudes and gives access to other parameters, such as the light intensity of the LIS.
A detailed description of the threshold scan procedure is given in Section 6.1, followed by the time alignment of the LIS in Section 6.2. The calibration model and the method are presented in Sections 6.3 and 6.4. Finally the calibration results are discussed in Section 6.6. The correct calibration of the comparator thresholds is essential for the operation of the SciFi Tracker as it directly impacts the detector's hit detection and noise suppression capabilities. The impact of the comparator threshold settings on the hit detection efficiency, single hit resolution, and noise rate is evaluated in Chapter 7.

6.1 Light Injection Threshold Scan

The PACIFIC digitises the SiPM signals using a simple arrangement of three comparators with adjustable thresholds, as described in Section 3.6.1. The thresholds of each of these comparators need to be calibrated with respect to the pulse-height spectrum of the connected SiPM channel. This is done using a threshold scan under pulsed illumination from the LIS. In a threshold scan, the full range of the comparator threshold values is scanned, and at every step, the ratio of events over the threshold for a set number of events is recorded. At the same time, light is injected into the SiPMs using the LIS. The resulting spectrum is shown in Figure 6.1. The black markers show the ratio of events over the threshold for different comparator threshold DAC values, which, due to its shape, is typically called *S-curve*. Each DAC value corresponds to a specific comparator threshold voltage, and higher DAC values also typically give a higher voltage. The underlying pulse-height spectrum, obtained by differentiating the black data points, is shown as a shaded area. It is characterised by discrete peaks originating from the electronic noise or the firing of one or more SiPM pixels, as indicated by the dotted lines. At the beginning of the scan, the threshold is set below the

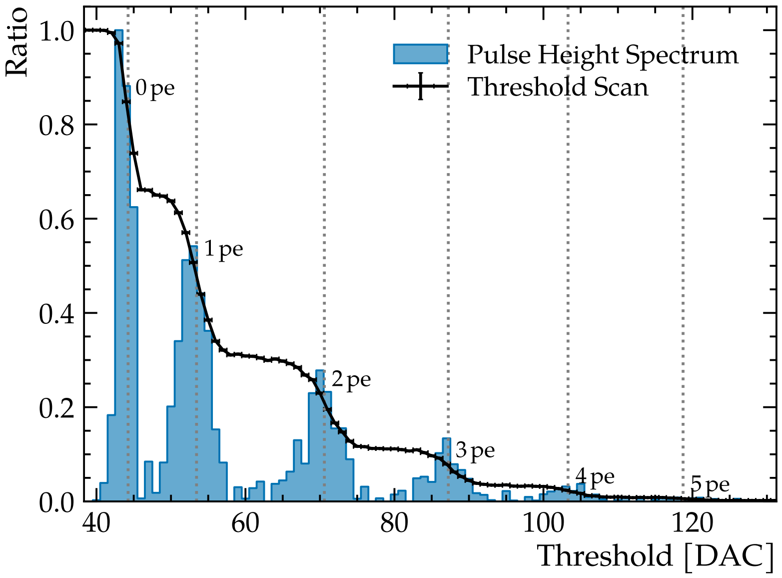

Figure 6.1 LIS threshold scan for one PACIFIC comparator and corresponding pulse-height spectrum. The distance between steps is a measure of the signal amplitude. The amplitude of the first photoelectron (1 pe) is typically lower due to the hysteresis of the comparator.

electronics' noise level, also called *pedestal*. Thus, the ratio of events over the threshold divided by all events is exactly one. Once the threshold is set above the noise level, the ratio is reduced, as some signals fall below the threshold. The ratio continues to fall until the threshold is set between two peaks, leading to the characteristic step-like shape of the S-curve.

6.1.1 Data Taking & Processing

A threshold scan progresses in a number of steps, each consisting of two distinct phases: configuration of the comparator thresholds and recording of events. The thresholds are configured based on a pre-defined sequence that instructs the experimental control system to configure the PACIFIC comparator thresholds with a specific threshold value based on the current step number. This information is stored in special summary XML files which are later used in the analysis of the threshold scans. The threshold scans need to be conducted separately for each of the three comparators of one channel because their output is added together and encoded in a 2-bit number, which loses all information about which comparator has fired. As a result, a total of 768 steps must be scanned to cover the full dynamic range of the 8-bit threshold DAC of all three comparators. During the scanning process of one comparator, the other two are set to the maximum threshold value of 255 DAC. This ensures that the digits sent by the PACIFIC are either `00` or `01`. Higher values are not possible as the other two comparators are functionally disabled by the high threshold setting of 255 DAC.
During the data-taking phase, a fixed number of events is recorded by issuing the Calib C command at a fixed frequency of ≈20 kHz which also simultaneously triggers the LIS pulse. In order to compare the two integrators of one PACIFIC channel, the

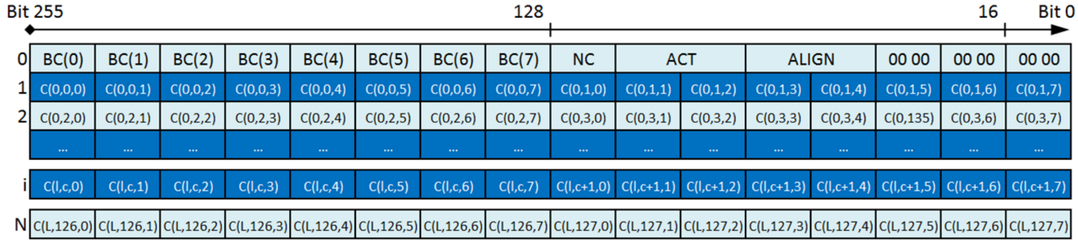

0	BC(0)	BC(1)	BC(2)	BC(3)	BC(4)	BC(5)	BC(6)	BC(7)	NC	ACT		ALIGN		00 00	00 00	00 00
1	C(0,0,0)	C(0,0,1)	C(0,0,2)	C(0,0,3)	C(0,0,4)	C(0,0,5)	C(0,0,6)	C(0,0,7)	C(0,1,0)	C(0,1,1)	C(0,1,2)	C(0,1,3)	C(0,1,4)	C(0,1,5)	C(0,1,6)	C(0,1,7)
2	C(0,2,0)	C(0,2,1)	C(0,2,2)	C(0,2,3)	C(0,2,4)	C(0,2,5)	C(0,2,6)	C(0,2,7)	C(0,3,0)	C(0,3,1)	C(0,3,2)	C(0,3,3)	C(0,3,4)	C(0,135)	C(0,3,6)	C(0,3,7)
	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...
i	C(l,c,0)	C(l,c,1)	C(l,c,2)	C(l,c,3)	C(l,c,4)	C(l,c,5)	C(l,c,6)	C(l,c,7)	C(l,c+1,0)	C(l,c+1,1)	C(l,c+1,2)	C(l,c+1,3)	C(l,c+1,4)	C(l,c+1,5)	C(l,c+1,6)	C(l,c+1,7)
N	C(L,126,0)	C(L,126,1)	C(L,126,2)	C(L,126,3)	C(L,126,4)	C(L,126,5)	C(L,126,6)	C(L,126,7)	C(L,127,0)	C(L,127,1)	C(L,127,2)	C(L,127,3)	C(L,127,4)	C(L,127,5)	C(L,127,6)	C(L,127,7)

Figure 6.2 Illustration of the FTCalibration TELL40 data format. Each block indicates one 16 or 32 bit word. The first row contains the header information and each subsequent row contains the counter values for two channels [88]. The counts *C(l,c,i)* are ordered by the TELL40 link (*l*), SiPM channel (*c*), and the counter index (*i*).

same number of events is recorded in even and odd bunch periods. The frontends are instructed to send non-zero-suppressed data that contains the raw PACIFIC digits for each event. The TELL40s then perform an online counting of the PACIFIC digits using four 16-bit counters[1]. Two sets of these counters are implemented per channel to allow the recording of data for the two integrators of each channel separately by splitting the counting of events into even and odd bunch period. At the end of each data-taking step, the counters are encoded in a special TELL40 calibration data format, called *FTCalibration*, shown in Figure 6.2, and saved to disk.

Once the data from all steps has been written to disk, it undergoes a series of processing steps to generate the S-curves. In the context of this dissertation, the entire processing chain, as illustrated in Figure 6.3, has been integrated within the LHCb software framework [89]. The data processing occurs event by event, starting with the decoding of the FTCalibration raw data. The data format structure is displayed in Figure 6.2. The header, consisting of the first 256 bits, contains amongst others the two 32-bit words, *ACT* and *ALIGN*. These words encode which of the up to 24 TELL40 links took part in the data collection. It is crucial information for mapping the TELL40 data link to its counterpart on the frontend electronics. Each subsequent group of 128 bits contains the counter values for one channel. The first four 16-bit words contain the counts for the even bunch periods, and the last four contain the counts for the odd ones. Each group of four counter values is stored with the information of the step number, channel, and integrator into an *FTCounter* object.
To construct S-curves, the FTCounters from all steps corresponding to one channel, integrator, and comparator need to be combined. It is impractical to keep all 3 145 728 FTCounters from all channels, integrators, and comparators in memory, so the data is temporarily buffered in a file for easy retrieval. Once all events have been decoded and stored in the buffer, the S-curves are created channel by channel using the information stored in summary XML files, which allows the mapping of the step number of an FTCounter to a specific setting of the comparator thresholds. The S-curves are then saved to disk in a simple tuple-like format, where each entry corresponds to one S-curve.

---

[1] Even though only two of the four possible PACIFIC outputs are expected, counters for all four have been implemented in the TELL40 firmware.

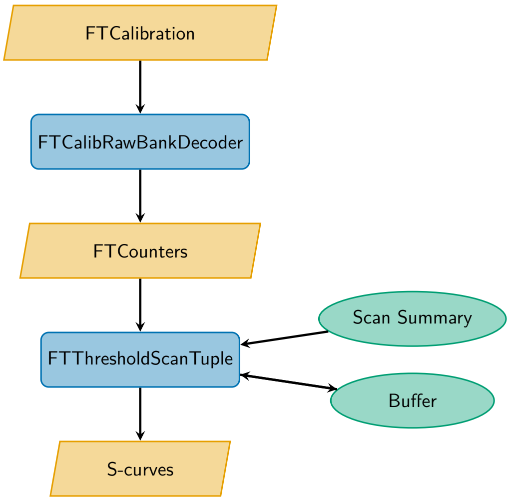

Figure 6.3 Overview of the algorithms used to create S-curves from FTCalibration raw banks. The algorithms are shown as blue rounded rectangles, data as yellow parallelograms, while additional dependencies are shown as green ovals.

6.2 Light Injection Time Alignment

The generation of light pulses by the LIS is triggered by the Calib C command of the TFC system, as previously described in Section 3.6.5. TFC command are issued in a particular bunch clock period by the SODIN board and then distributed to all frontend electronics via the SOL40 boards through ≈400 m long optical fibres. In addition the processing of the Calib C command, generation of the LIS pulse and processing of the signals by the frontend electronics also takes time. All of this introduces a significant latency between the bunch period where the TFC command was issued and when the resulting light pulse is seen by the frontend electronics. This latency needs to be corrected for such that the data is recorded when the LIS pulse is triggered.
However, this only allows for changes in steps of 25 ns, but as previously mentioned in Section 3.6.5, the LIS pulse needs to be injected at the right time within one 25 ns bunch clock period in order to align it with respect to the PACIFIC integration interval. In addition special care needs to be taken to select the correct Calib C phase in the frontend electronics.

6.2.1 Calib C Latency

To determine the latency of the Calib C command, the command is sent at a specific bunch period, and data is recorded using a random trigger to ensure that all bunch periods are uniformly sampled. The PACIFIC comparator thresholds are set to $(1.5, 2.5, 3.5)$ pe and the frontend electronics are configured to send the raw PACIFIC output (NZS data). In the bunch period where the LIS pulse is triggered, the PACIFIC output is expected to be higher on average due to the detection of light. An example of such a measurement can be seen in Figure 6.4a, where the average of the PACIFIC output (mean digit) of 128 channels corresponding to one SiPM is shown as a function of the bunch period number. In this case, the Calib C command was issued in bunch period 92. The following four bunch periods contain no data due to the NZS data requiring four bunch periods to be transmitted from the frontend electronics to the TELL40s, as mentioned in Section 3.6.3. A clear peak can be seen 11 bunch periods later originating from the detection of light from the LIS pulse by the SiPM. Some of the light is also visible one bunch period later, which is likely due to the LIS pulse extending into the next bunch period.

The Calib C command and other TFC signals are transmitted to the frontend electronics through control links, with one control link assigned to each HalfROB. When adjusting for the latency of the Calib C command, it is important to consider the mapping of the LIS. As explained in Section 4.10.1, the LIS controlled by one HalfROB injects light into one of the two fibre mats read out by the other HalfROB. Therefore, when adjusting the latency of HalfROB H0, the data recorded for fibre mats 0 and 2 should be used, while for H1, the data from mats 1 and 3 should be used.

Once the latency differences have been adjusted, a cross-check is carried out by recording time alignment events (TAE). This process saves the previous and following two bunch periods in addition to the selected period. Due to bandwidth limitations when using NZS data, this is limited to ± 2 bunch periods and is therefore not suitable for finding larger latencies. Similar to Figure 6.4a, Figure 6.4b shows an example of the average of the PACIFIC digits for the bunch period where the Calib C command was issued, as well as the two before and after. The light from the LIS pulse is observed in the same bunch period where the Calib C command is sent, which indicates that the latency has been correctly adjusted. Figure 6.5 shows the latency of the Calib C command for the entire detector after correcting the latencies for all 512 HalfROBs. For most SiPMs, the LIS pulse is detected during the same bunch period as the Calib C command, resulting in a latency of zero. However, some SiPMs detect light either earlier or later, which is indicated by coloured boxes. The remaining latency differences are caused by interference with other TFC settings, which make it impossible to adjust both HalfROBs of one ROB correctly, or by low light intensity from the LIS which prevents a clear signal from being identified. Unfortunately, this issue could not be investigated further within the time frame of this dissertation, so the calibration of the corresponding PACIFIC channels is also delayed.

6.2.2 Calib C Phase Selection

As discussed in Section 3.6.5, the LIS pulse can become meta-stable if the LIS Start Clock is close to the edge of the Calib C pulse. To avoid this scenario, the Calib C phase can be shifted by 12.5 ns. In order to find the LIS Start Clock delays at which the Calib C phase needs to be shifted, a measurement is performed where the complete delay range is scanned for both phase settings. A set of NZS TAE data is recorded for each delay and phase setting. Figure 6.6 shows the average PACIFIC digits for one SiPM for the different LIS delay settings and the two phase settings. In the central bunch period (delays 0 to 25 ns) the average signal amplitude is rising for both phase settings until a delay of 12.7 ns. Here, the curve for a phase setting of 12.5 ns suddenly drops off. Looking into the next bunch period, it becomes apparent that the pulse is now generated 25 ns later. In order to keep the pulse within the same bunch period, it was decided to set the Calib C phase to 0 ns for most of the delay range and only switch

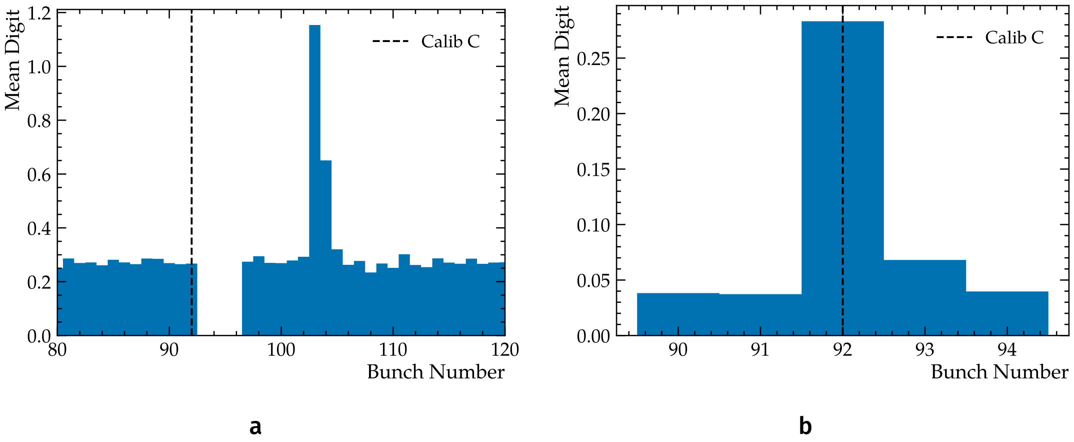

Figure 6.4 Left: exemplary measurement of the average amplitude as a function of the bunch period for one SiPM. The LIS has been triggered in bunch number 300. Right: measurement of the average amplitude after adjustment of the Calib C latencies for two bunch periods before and after the bunch where the LIS is triggered.

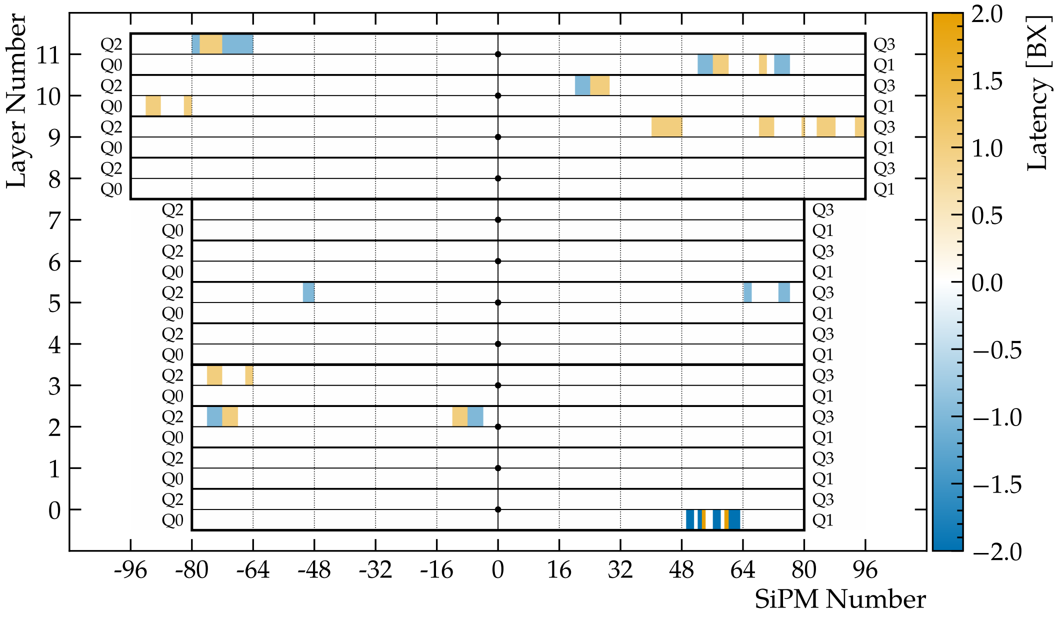

Figure 6.5 Overview of the Calib C latencies after adjustments for each SiPM in the detector where the latency is indicated by the colour. Latency differences are expected to show up in groups of eight SiPMs corresponding to either fibre mats 0 and 2 if controlled by HalfROB H0, or mats 1 and 3 if controlled by HalfROB H1.

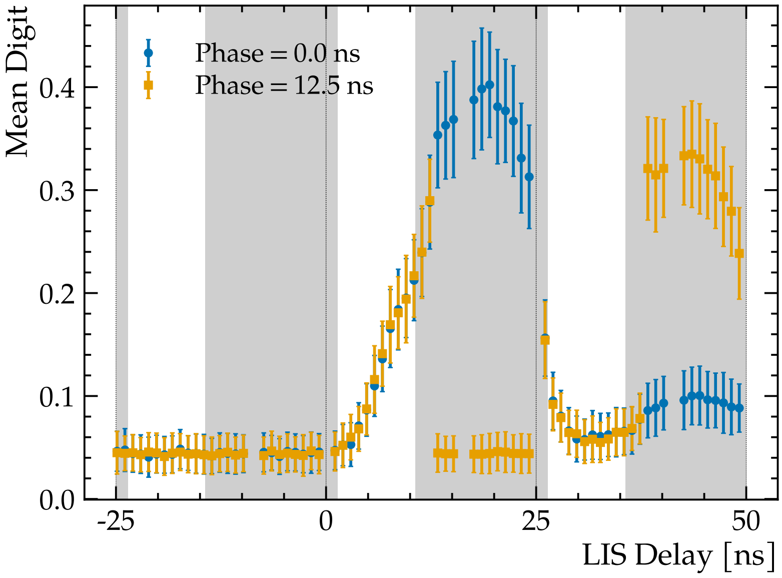

Figure 6.6 Measurement of the amplitude (mean digit) for 128 channel corresponding to one SiPM as a function of the LIS delay using both Calib C phase settings. The error bars show the standard deviation in the amplitude between the channels. The range of LIS delay where the Calib C phase is set to 0 ns is shown as grey shaded area.

the phase within a LIS delay of 1.4 to 10.6 ns as indicated by the grey shaded areas.

6.2.3 LIS Pulse Delay

After the latencies have been adjusted so that the LIS is triggered in the correct bunch period, the next step is to adjust the LIS pulse delay relative to the PACIFIC clock so that the SiPM signals are maximally integrated by the PACIFIC.
As described in Section 3.6.5, the LIS pulse is generated in the Housekeeping FPGA of the Master Board from the combination of the Calib C TFC command and two deskewable clocks that set the start and stop of the pulse. Once the Calib C command is received, the LIS pulse is generated starting from the first rising edge of the Start Clock and lasts until the first rising edge of the Stop Clock signal. By adjusting the Start Clock phase, the LIS pulse delay with respect to the PACIFIC clock can be adjusted. To keep the pulse width the same, the Stop Clock needs to be shifted by the same amount. The delay of the LIS Start Clock needs to be adjusted so that the PACIFIC integrators fully integrate the SiPM signals. Figure 6.7a shows S-curves recorded at different LIS delay settings for an exemplary channel. Only within a few nanoseconds around the optimal delay setting are the characteristic steps clearly visible; after that, the S-curves degrade until only a single transition remains that corresponds to the pedestal of the channel due to less and less of the SiPM pulse being integrated. In order to find the optimal LIS delay, the amplitude per photoelectron called gain is extracted from the S-curve using a fit that is described in detail in Section 6.3. Plotting the gain versus the LIS delay results in a curve as shown in Figure 6.7b, where the gain for all 512 channels corresponding to one fibre mat has been averaged. In this particular example,

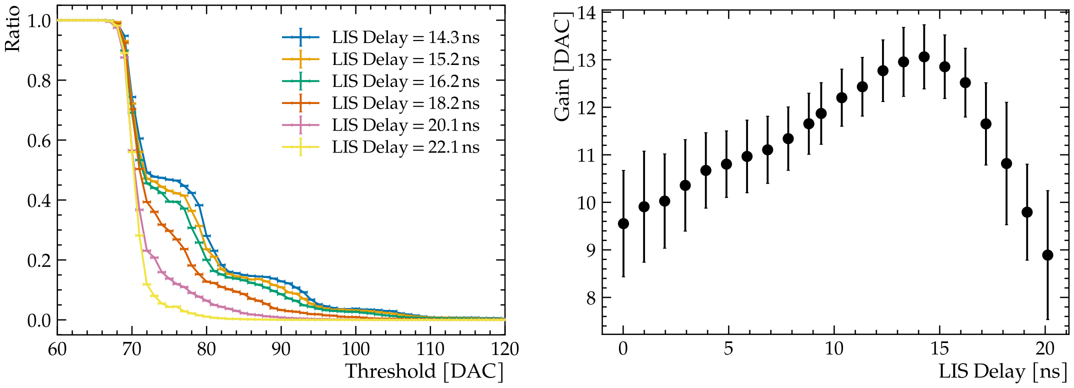

a

b

Figure 6.7 Left: examples of S-curves at different LIS delays for one comparator. Right: average gain as a function of the LIS delay for all 512 channels of one fibre mat.

the maximum gain is at a delay of ≈14.3 ns. For delays larger than ≈21 ns, the S-curves barely exhibit their characteristic steps, and the gain can not be determined reliably. The optimal delay settings for all data links are shown in Figure 6.8. The optimal LIS delay is found in the range of 14 to 15 ns for most data links. The presence of some outliers can be explained by the missing latency adjustment, resulting in the LIS pulse not being triggered in the correct bunch period.

6.3 S-curve Model

In order to set the PACIFIC comparator thresholds to the correct photoelectron amplitudes, the threshold DAC needs to be calibrated with respect to the connected SiPM channel; this is done by recording threshold scans under pulsed illumination from the LIS as presented in Section 6.1 and fitting the entire S-curve to determine the threshold DAC values corresponding to the discrete photoelectron signals. In the following section, an analytical description of the S-curve will be derived.

The pulse height spectrum of an SiPM under pulsed illumination is characterised by a set of separated peaks corresponding to zero, one, or more of the SiPM's pixels firing as shown in Figure 6.1. Every fired pixel in turn corresponds to the detection of a photon. The location of the first peak, called *pedestal*, corresponds to the amplitude of the electronic noise. All further peaks stem from events with one or more detected photons. Since each pixel produces approximately the same signal, the peaks are separated by the same distance called *gain*. Knowing the pedestal p_0 and gain g, the amplitudes p_n for any number of photoelectrons is then given by

$$p_n = p_0 + n \cdot g. \qquad [6.1]$$

Looking closer at the peaks in Figure 6.1, it becomes apparent that the peaks are getting progressively wider. This can easily be understood when remembering that the signal in an SiPM is the analogue sum of all pixels. Therefore, all variations in the signals are also added together. The two notable contributions here are the standard deviation

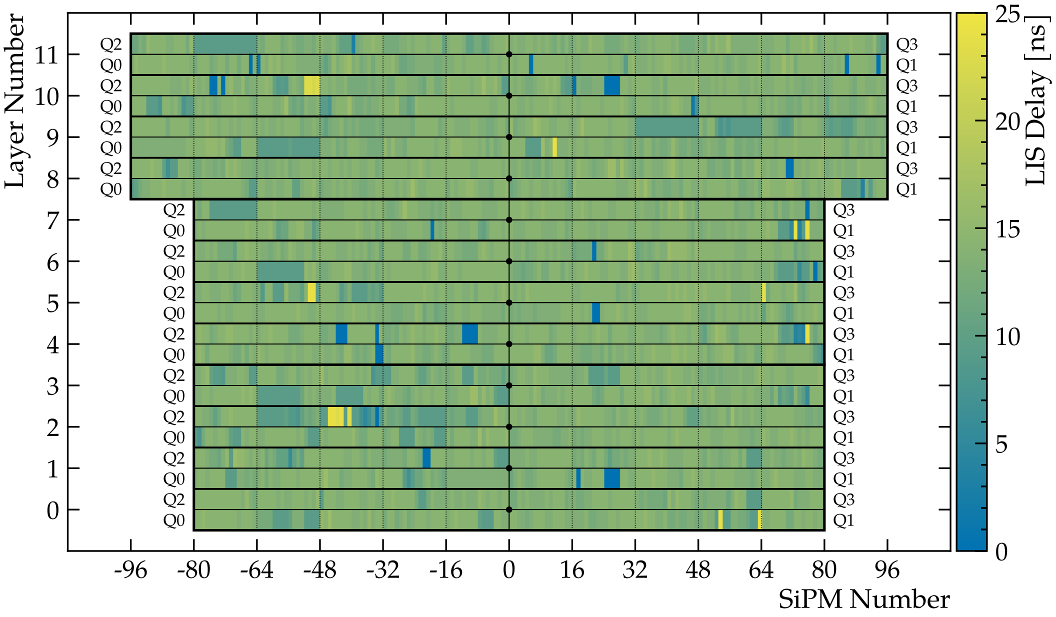

Figure 6.8 Overview of the optimal LIS delay for each SiPM in the detector.

of the electronic noise σ_0, corresponding to the width of the pedestal peak, and the variation in the gain between individual pixels σ_1. The width of the n-th peak is then the sum of squares of the two contributions

$$\sigma_n = \sqrt{\sigma_0^2 + n \cdot \sigma_1^2}. \tag{6.2}$$

Each peak can, to a good approximation, be described by a Gaussian

$$G_n(x) = \frac{1}{\sigma_n \sqrt{2\pi}} \cdot \exp\left(-\frac{(x - p_n)^2}{2\sigma_n^2} \right) \tag{6.3}$$

with mean p_n and width σ_n. The number of events in each peak follows a distribution determined by the photon statistics. For coherent light sources, like a laser, the number of emitted photons follows a Poisson distribution

$$P(n) = \frac{\mu^n e^{-\mu}}{n!} \tag{6.4}$$

with mean and variance μ [90]. This distribution is then sampled with a binomial distribution due to the finite detection efficiency of the SiPM. Since the convolution of a Poisson with a binomial distribution generates a Poisson distribution again, the original nature of the photon statistics is preserved. A complication, however, arises from the fact that SiPMs suffer from crosstalk. As discussed in Section 3.5.3, crosstalk is a form of noise where additional pixels fire correlated to a primary discharge. This effectively shifts the spectrum towards higher mean values. In general, the probability of n pixels firing is given by

$$P(n) = \begin{cases} P_0(0)\,, & n = 0 \\ \sum_{k=1}^{n} P_0(k) \cdot P_k(n)\,, & n > 0 \end{cases} \qquad [6.5]$$

where $P_0(k)$ is the probability distribution of the primary signals and $P_k(n)$ the probability of k triggered pixels given n primaries. The explicit form of the crosstalk probability distribution $P_k(n)$ varies and several models have been proposed in the literature [91–93]. A detailed discussion of a selection of two models is provided in Section 6.3.1.

The full pulse height spectrum can now be described by a sum of Gaussian distributions G_n multiplied by the probability to detect n photoelectrons

$$PH(x) = \sum_n P(n) \cdot G_n(x). \qquad [6.6]$$

Considering now that the PACIFIC digitises the SiPM signals with a set of comparators and not a full ADC, the pulse height spectrum as shown in Figure 6.1 can not be measured directly and a threshold scan as explained in Section 6.1 needs to be performed. For a given threshold, the measurement consists of counting the number of events over the threshold, which in the language of probability distributions is the definition of the complementary cumulative distribution function

$$\overline{F}_X(x) = 1 - F_X(x) = 1 - \int_{-\infty}^{x} f(x)\,\mathrm{d}x. \qquad [6.7]$$

Substituting Equation 6.4 for $f(x)$ the following function for the threshold scan is obtained

$$\begin{aligned} TH(x) &= 1 - \int_{-\infty}^{x} PH(x)\,\mathrm{d}x \\ &= 1 - \int_{-\infty}^{x} \sum_n P(n) \cdot G_n(x)\,\mathrm{d}x \\ &= 1 - \sum_n P(n) \int_{-\infty}^{x} G_n(x)\,\mathrm{d}x \\ &= 1 - \sum_n P(n) \cdot \Phi_n(x) \end{aligned} \qquad [6.8]$$

with Φ the cumulative distribution function of the Gaussian distribution given by

$$\Phi_n(x) = \frac{1}{2}\left[1 + \operatorname{erf}\left(\frac{x - p_n}{\sigma_n \sqrt{2}}\right)\right]. \qquad [6.9]$$

Figure 6.9a shows an example of a threshold scan for one PACIFIC comparator. It is characterised by a series of steps as described by Equation 6.8. Looking at the distance between the steps, it becomes apparent that their location can not be described by the single gain parameter assumed in Equation 6.1. Instead, individual gain parameters g_i are introduced for each of the steps, with g_0 the gain of the first photoelectron, g_1 for

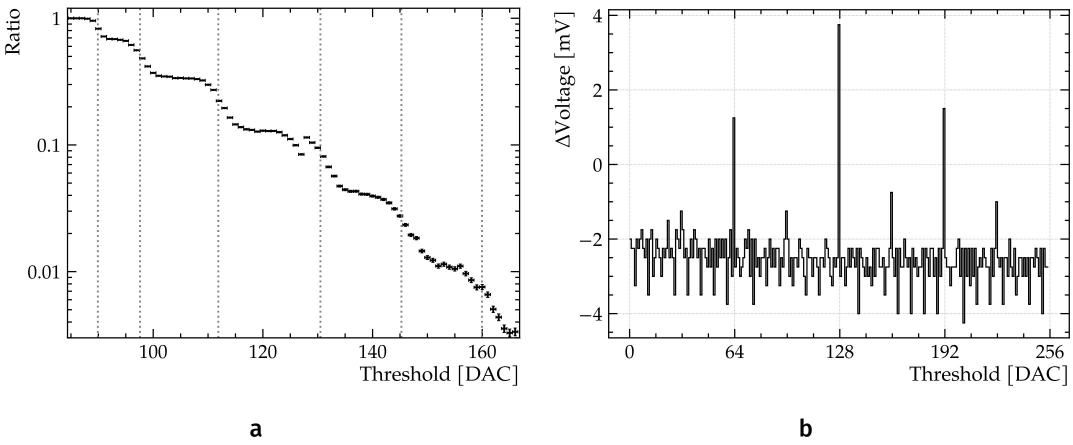

Figure 6.9 Left: threshold scan of one comparator showing the discontinuity at a DAC value of 128. Right: measurement of the threshold voltage step as a function of the threshold DAC for a different comparator during QA of the PACIFIC [94]. The discontinuities occur every 32 DAC values.

the second photoelectron, and so on. Equation 6.1 then becomes

$$p_n = \begin{cases} p_0 & n = 0 \\ p_0 + \sum_{i=0}^{n-1} g_i & n > 0 \end{cases} . \tag{6.10}$$

The gain of the 1 pe signal is typically lower due to the hysteresis of the comparator. The hysteresis protects the comparator from repetitive switching if the signal is close to the threshold. However, it also means that the signals need to be slightly larger than the given threshold voltage to trigger the comparator. The 1 pe signals therefore have a gain that is lower by about 4 DAC given the ≈10 mV hysteresis.

The hysteresis does not impact the higher gain values and should, therefore, be identical. However, the comparator threshold DAC is not entirely linear as is evident from the discontinuity in Figure 6.9a at a DAC value of 128. Figure 6.9b shows a measurement of the threshold voltage step size for each of the threshold DAC values during the QA of the PACIFIC. The step size is approximately -2 to -3 mV over the full range of DAC values, but shows distinct spikes every 32 DAC. These spikes have the effect that different threshold DAC values result in the same threshold voltage, thus leading to the observation of an effectively larger gain in a threshold scan as can be seen in Figure 6.9a around a DAC value of 127.

6.3.1 Evaluation of Crosstalk Models

Several analytical models of crosstalk in SiPMs exist in the literature. Following Reference [93], the models discussed here can all be derived from the same assumptions that crosstalk events are preferably generated in pixels directly adjacent to the pixel where the primary avalanche was triggered. This has led to the adoption of the binomial distribution to describe the number of crosstalk events k given m neighbouring pixels around the primary pixel

$$P(k) = \binom{m}{k} p^k (1-p)^{m-k}, \tag{6.11}$$

with p the probability of triggering a crosstalk event.
Any triggered pixel, either the primary or secondary, can induce crosstalk in n additional neighbouring pixels. This cascading means that Equation 6.11 needs to be applied for each pixel in the sequence. It can be shown that for $m = 1$, the probability distribution to detect k triggered pixels is given by the geometric distribution

$$P(k) = p^{k-1}(1-p) \tag{6.12}$$

with the crosstalk probability p, while for $m \to \infty$ it corresponds to the Borel distribution

$$P(k) = \frac{(\lambda k)^{k-1} \cdot \exp(-\lambda k)}{k!} \tag{6.13}$$

with the mean number of crosstalk events λ. The parameters p and λ of the two distributions are related by

$$1 - p = \exp\left(-\lambda\right). \tag{6.14}$$

The photon statistics $P(n)$ for each model can be derived using Equation 6.5. The compound distribution of the geometric distribution Equation 6.12 and the Poisson distribution is the so-called geometric Poisson distribution [91, 95]

$$P(n) = \begin{cases} e^{-\mu} & n = 0 \\ \sum_{k=1}^{n} \binom{n-1}{k-1} (1-\lambda)^{n-1} \frac{(\mu k)^{k-1} \exp(-\mu k)}{k!} & n > 0 \end{cases}, \tag{6.15}$$

while for the case of the Borel distribution Equation 6.13 the generalised Poisson distribution [92, 96]

$$P(n) = \frac{\mu(\mu + n\lambda)^{n-1} \exp(-\mu - n\lambda)}{n!} \tag{6.16}$$

is derived.
Figure 6.10 shows a comparison of the two models where Equation 6.15 and Equation 6.16 have been evaluated up to $n = 10$ for an average number of primary photoelectrons $\mu = 1.5$ pe and average number of crosstalk events of $\lambda = 0.07$ pe. These values correspond to the expected light intensity of the LIS (Section 4.10.3) and the measured total crosstalk probability Section 3.5.5. The two graphs are almost identical, with only small differences appearing for higher n.

The two crosstalk models presented here have been evaluated using light injection threshold scans obtained during the C-Frame commissioning described in Chapter 4. Exemplary fit results of the entire model based on Equation 6.8 using the different crosstalk models are shown in Figure 6.11a; visually, the fit results are identical. Looking at the distributions of the goodness-of-fit (χ^2/N_{DoF}) for all tested channels, the geometric Poisson model is performing marginally better with an average $\chi^2/N_{DoF} = 9.48$

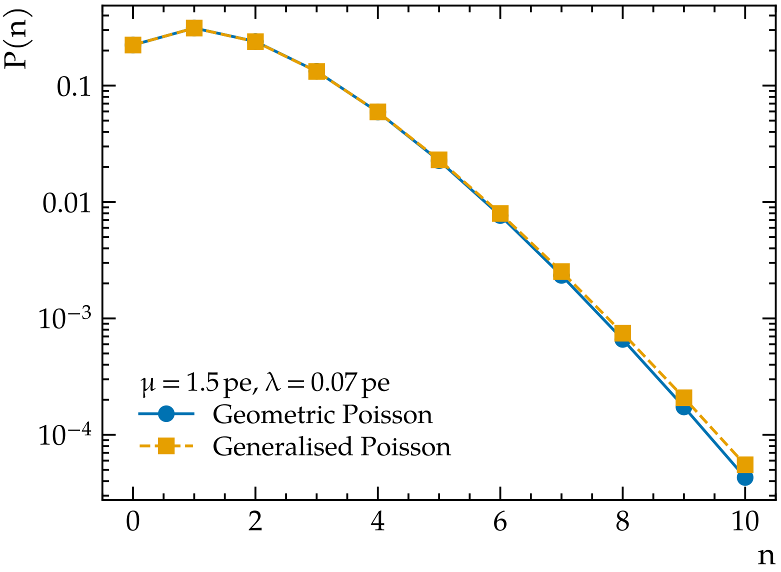

Figure 6.10 Comparison of the two crosstalk models showing the probability to detect *n* photoelectrons.

compared to the 9.56 for the generalised Poisson model, as shown in Figure 6.11b. The distributions for the amount of crosstalk λ and average light intensity μ are shown in Figure 6.12a and Figure 6.12b, respectively. Again, the distributions are almost identical with equal mean values of $\lambda = 0.34$ pe and standard deviations of 0.013 pe. The average light intensity also shows no sensitivity to either model, and the two distributions in Figure 6.12b are almost identical with mean intensities of $\mu = 1.47$. In conclusion, neither of the two crosstalk models is strongly favoured by the data. Looking back at Figure 6.10, this becomes clear. The calibration fits are only performed for $n = 5$, where both models give almost identical probabilities; only for higher values of n can the models be separated. In addition, the overall amount of crosstalk is relatively low with an average value of 0.34 pe, resulting in a crosstalk probability of 3.3 % according to Equation 6.14. The decision was made to use the generalised Poisson model simply due to ease of implementation.

6.4 Calibration Procedure

To facilitate the calibration of all 524 288 channels, a software tool called `ScurveFit`[2] has been developed and implemented in C++, leveraging the threading capabilities of modern multi-core CPUs. For each S-curve, a sequence of configurable algorithms is executed as shown in Figure 6.13. The three major steps are the differentiation of the S-curve and subsequent peak finding, the estimation of the S-curve model's initial parameters, and the final likelihood fit of the S-curve.

---

[2] The source code is available at `https://gitlab.cern.ch/pacific-calibration/scurvefit`.

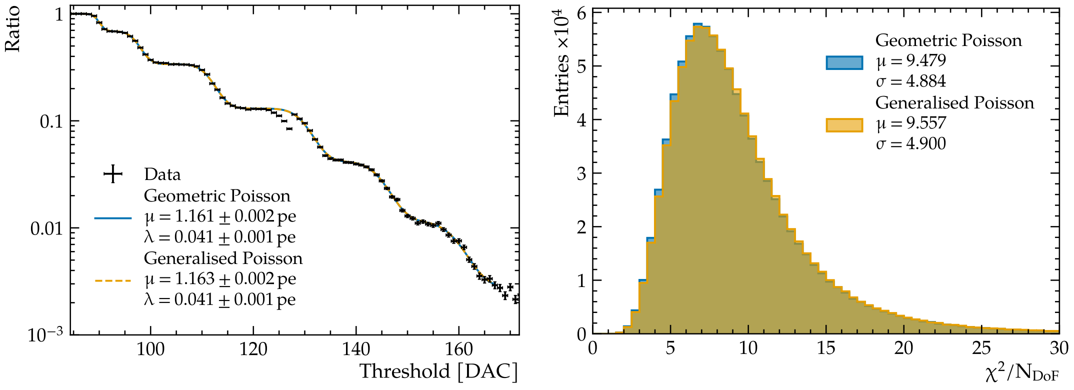

a **b**

Figure 6.11 Left: exemplary threshold scan and fit result for both crosstalk models. Right: distribution of the goodness-of-fit for all channels.

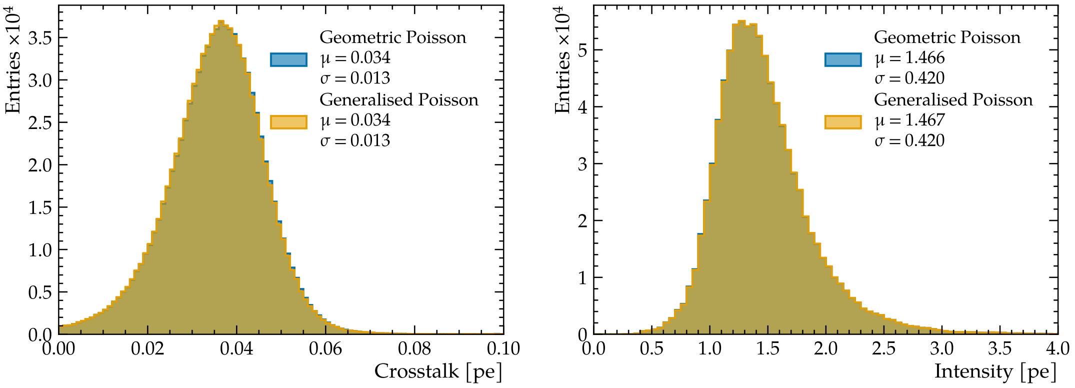

a **b**

Figure 6.12 Left: distribution of the measured crosstalk values for both crosstalk models. Right: distribution if the light intensity for both crosstalk models.

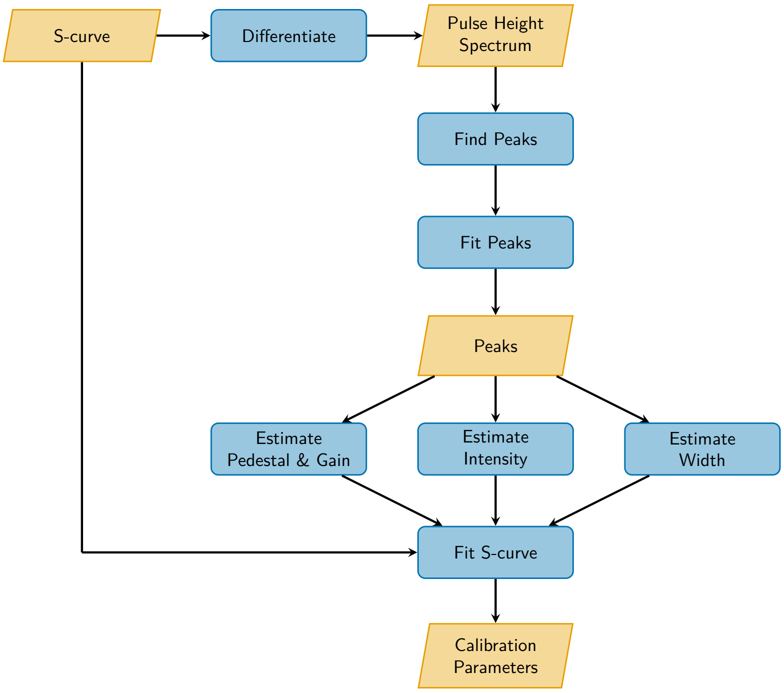

Figure 6.13 Overview of the algorithms and data flow in the calibration software. Algorithms are shown as blue rounded rectangles while input and outputs are shown as yellow parallelograms.

6.4.1 Peak Finding

In the first step, the derivative of the S-curve is calculated to obtain the pulse height spectrum. Following Equation 6.8, the derivative needs to be multiplied by an additional factor -1. An exemplary spectrum is shown in Figure 6.14a. The derivative is given by the central finite differences

$$\Delta f_{central}(x_i) = \frac{1}{2}\left[f(x_{i+1}) - f(x_{i-1})\right], \tag{6.17}$$

apart from the edges where the forward and backward differences

$$\Delta f_{forward}(x_i) = f(x_{i+1}) - f(x_i), \tag{6.18}$$

$$\Delta f_{backward}(x_i) = f(x_i) - f(x_{i-1}) \tag{6.19}$$

are used.

The spectrum shown in Figure 6.14a is then searched for peaks. As discussed in Section 6.3, the distance between the first two peaks is typically lower, and the widths of the peaks increase, making it challenging to find the peaks. The peak finding algorithm is a simplified version of the algorithm described in Reference [97]; by computing the second differences a peak can be identified by a local minimum surrounded by a

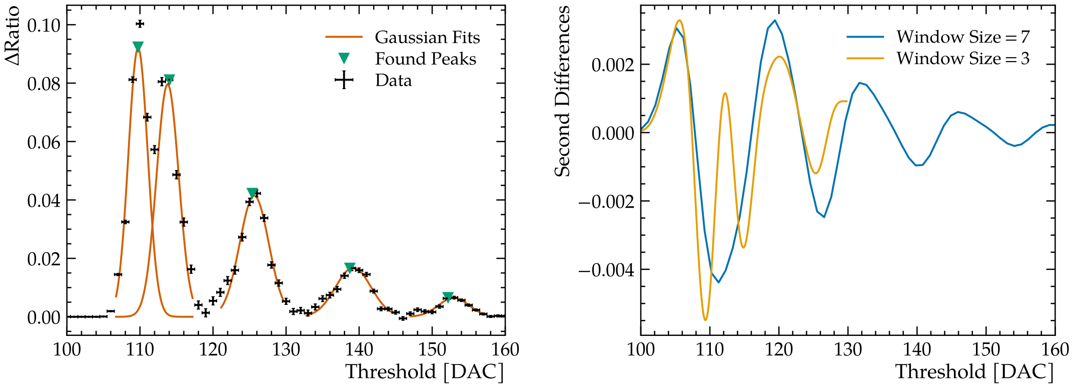

a

b

Figure 6.14 Left: example of the derivative of an S-curve (black markers) with the found peaks (green triangles) and gaussian fits (red line) of each peak. Right: second differences used to identify the peaks on the right for two smoothing window sizes.

local maximum on either side. The second differences are given by

$$
\begin{aligned}
\Delta^2 f(x_i)_{central} &= f(x_{i+1}) - 2 \cdot f(x_i) + f(x_{i-1}), && [6.20]\\
\Delta^2 f(x_i)_{forward} &= f(x_{i+2}) - 2 \cdot f(x_{i+1}) + f(x_i), && [6.21]\\
\Delta^2 f(x_i)_{backward} &= f(x_i) - 2 \cdot f(x_{i-1}) + f(x_{i-2}), && [6.22]
\end{aligned}
$$

where the forward and backward differences are again used at the edges of the histogram.

The second differences are smoothed using a moving average to reduce the influence of statistical fluctuations. The smoothed value of each point x_i denoted as $\tilde{x}_i$ is given by the average of the surrounding values calculated as:

$$\tilde{x}_i = \frac{x_{i-n} + \cdots + x_i + \cdots + x_{i-n}}{2 \cdot n + 1} \qquad [6.23]$$

where $2 \cdot n + 1$ is the smoothing window size. Figure 6.14b shows the smoothed second differences of Figure 6.14a for two window sizes. First, a coarse peak search is carried out using a larger smoothing window size of 7, which is only able to resolve four out of the five peaks. The first two peaks are merged due to the smoothing. Therefore, a second fine peak search with a smaller window size of 3 is performed in the vicinity of the first found peak. Using this smaller window size, the double peak structure can be resolved, as shown in Figure 6.14b by the two minima around 115 DAC, and the peaks can be identified.

The found peak candidates are then verified by fitting a Gaussian distribution to each individual peak. The fit parameters are used in the next step to estimate the initial parameters of the S-curve.

6.4.2 Initial Parameter Estimation

Correctly estimating the initial parameters for the final calibration fit is crucial to ensuring proper convergence and stability. The important parameters are the light intensity, pedestal, gain, and width parameters, all of which are estimated using the peaks from the previous step.

The light intensity is estimated by considering the peak areas A_i given by

$$A_n = \sqrt{2\pi} \cdot \sigma_n \cdot h_n \tag{6.24}$$

with the Gaussian peak width σ_n and height h_n. The areas are expected to follow, to first approximation, a Poissonian distribution

$$A(n) \approx \frac{\mu^n \exp(-\mu)}{n!} \tag{6.25}$$

as discussed in Section 6.3. By fitting a Poissonian to the data as shown in Figure 6.15a the light intensity is directly estimated from the mean value μ. The fit is not expected to give perfect agreement with the data as effects such as crosstalk are not accounted for. However, it reliably provides a good first estimate of the light intensity. The crosstalk is not estimated from the data; instead, an initial value of $\lambda = 0.03\,\mathrm{pe}$ is assumed based on the measurement of the crosstalk during the SiPM QA.

The pedestal and gain values are estimated from the peak positions. As previously discussed, the gain between the pedestal given by the position of the first peak ($n = 0$) and the 1 pe peak is typically lower due to the hysteresis of the PACIFIC comparators. This leads to the observed piecewise linear trend seen in Figure 6.15b that can be described by the following function

$$p(n) = \begin{cases} p_0 + n \cdot g_{low}\,, & 0 \geq n \leq 1 \\ p_1 + n \cdot g_{high}\,, & n > 1 \end{cases} \tag{6.26}$$

with the two slopes g_{low} and g_{high} parametrising the gain and the intercept p_0 the pedestal.

The last missing ingredient for the S-curve model are the widths σ_0 and σ_0. They are estimated from a linear fit to the variance of the peaks given by the square of the peak width as shown in Figure 6.15c. Following Equation 6.2 the variance is given by

$$Var(n) = \sigma_0^2 + n \cdot \sigma_1^2 \tag{6.27}$$

and the parameters σ_0 and σ_0 are thus given by the square roots of the intercept and slope of the linear fit.

6.4.3 S-curve Likelihood Fit

In an S-curve, as shown in Figure 6.15d, each measurement point is the result of a series of trials that have either of two outcomes: the signal is above the threshold or below it. Thus, the number of events over the threshold m can be treated as a binomially distributed random variable. The probability P of finding m signals over the threshold, out of N tested signals, is given by

$$P(m; N, \varepsilon) = \frac{N!}{m!(N-m)!}\varepsilon^m (1-\varepsilon)^{(N-m)}. \qquad [6.28]$$

The goal is to estimate the parameter ε given the measurements of m and N. The log-likelihood function for ε is given by

$$\ln L(\varepsilon) = m \ln \varepsilon + (N-m)\ln(1-\varepsilon) \qquad [6.29]$$

and setting the derivative of $\ln L$ to zero gives

$$\hat{\varepsilon} = \frac{m}{N}, \qquad [6.30]$$

where the hat denotes the estimator of the true parameter ε. The ratio ε_i for each threshold value x_i is parametrised by

$$\varepsilon_i(\vec{\theta}) = 1 - \sum_{n=0}^{N} P(n) \cdot \Phi_n(x_i) \qquad [6.31]$$

with

$$P(n) = \frac{\mu(\mu + n\lambda)^{n-1}\exp(-\mu - n\lambda)}{n!}, \qquad [6.32]$$

$$\Phi_n(x) = \frac{1}{2}\left[1 + \operatorname{erf}\left(\frac{x - p_n}{\sigma_n\sqrt{2}}\right)\right], \qquad [6.33]$$

$$p_n = \begin{cases} p_0 & n = 0 \\ p_0 + \sum_{i=0}^{n-1} g_i & n > 0 \end{cases}, \qquad [6.34]$$

$$\sigma_n = \sqrt{\sigma_0^2 + n \cdot \sigma_1^2}, \qquad [6.35]$$

as discussed in Section 6.3. The goal is then to estimate the parameters $\vec{\theta}$ of the model from the known measurements of m and N. Following Reference [98] the likelihood function for the whole set of measurements m_i and N_i, $i = 1, \ldots, N_{bin}$ is given by their joint probability

$$L(\vec{\theta}) = \prod_{i=1}^{N_{bin}} P(m_i; N_i, \varepsilon(x_i, \vec{\theta})). \qquad [6.36]$$

Taking the logarithm gives the log-likelihood function

$$\ln L\vec{\theta}) = \sum_{i=1}^{N_{bin}} \left[m_i \ln \varepsilon(x_i, \vec{\theta}) + (N_i - m_i)\ln(1 - \varepsilon(x_i, \vec{\theta}))\right]. \qquad [6.37]$$

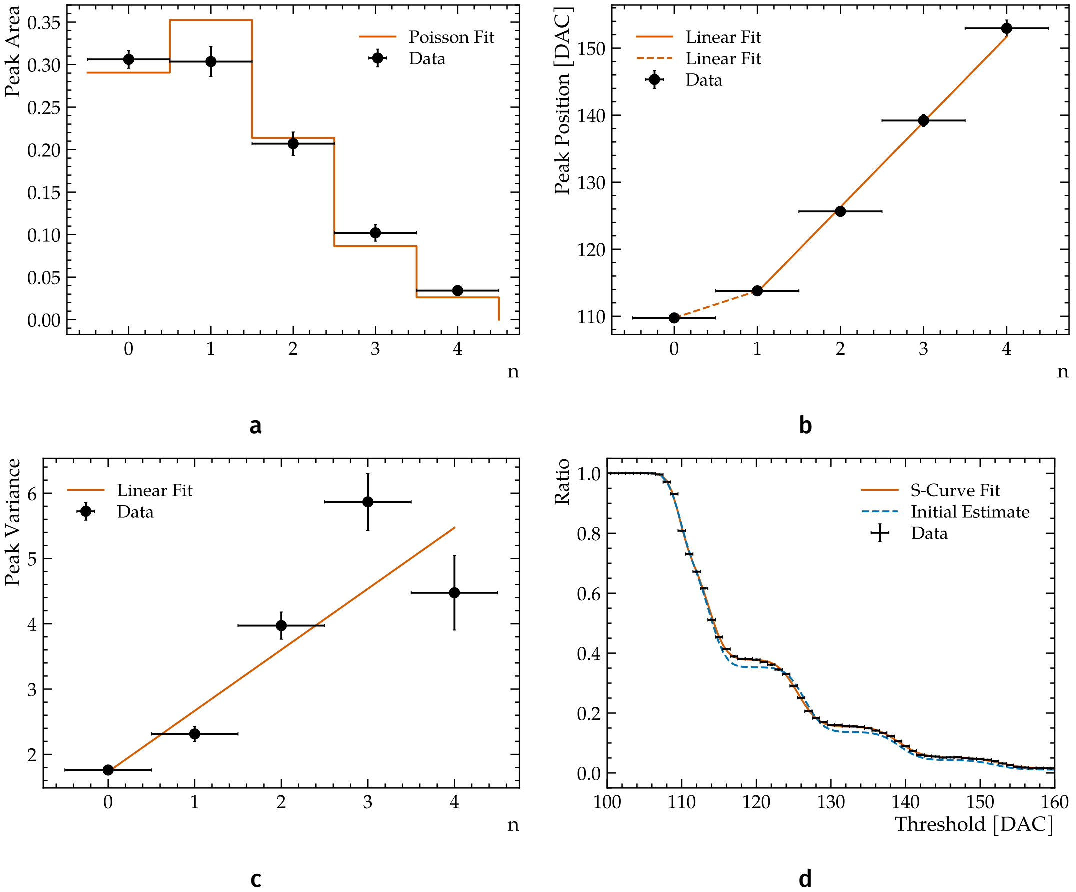

a

b

c

d

Figure 6.15 Examples showing various peak properties as a function of the photoelectron number *n* and fits used to determine the initial parameters as well as the final fit of the S-curve. Top left: peak areas and Poisson fit used to determine the light intensity. Top right: peak positions and linear fits used to determine the pedestal and gain parameters. Bottom left: peak variance and linear fit used to determine the width parameters. Bottom right: example of an S-curve, fit result (solid red line), and initial estimate (dashed blue line).

The likelihood ratio test theorem poses that the *likelihood chi-square* defined by

$$\chi_L^2 = -2\ln\left(\frac{L(\varepsilon)}{L(\hat{\varepsilon})}\right) = 2\left[\ln L(\hat{\varepsilon}) - L(\varepsilon)\right] \qquad [6.38]$$

asymptotically obeys a chi-square distribution [99, 100]. For the binomial log-likelihood function (Equation 6.37) and using $\hat{\varepsilon}(x_i) = \frac{m_i}{N_i}$ the likelihood chi-square becomes

$$\chi_L^2 = 2\sum_{i=1}^{N_{bin}}\left[m_i \ln\left(\frac{N_i \varepsilon(x_i,\vec{\theta})}{m_i}\right) + (N_i - m_i)\ln\left(\frac{N_i(1-\varepsilon(x_i,\vec{\theta}))}{N_i - m_i}\right)\right] \qquad [6.39]$$

The likelihood chi-square function is minimised using the MINUIT[3] library to determine the parameters $\vec{\theta}$.

---

[3] MINUIT is distributed as part of the ROOT framework [101].

6.5 Threshold Calculation

So far, only the S-curve has been discussed, but not the actual calibration of thresholds. The important parameters for calculating the thresholds are the position of the pedestal p_0 and the gain parameters g_i. These parameters are determined from the S-curve fits described in the previous section for the two integrators of each channel. However, the thresholds can not be set separately for each of the integrators and therefore a weighted average

$$p_0 = \frac{w^0 p_0^0 + w^1 p_0^1}{w^0 + w^1} \qquad [6.40]$$

$$g_i = \frac{w^0 g_i^0 + w^1 g_i^1}{w^0 + w^1} \qquad [6.41]$$

is computed with weights $w^n = N_{DoF}/\chi^2$ for the two integrators.
To discriminate between signals with different number of photoelectrons the thresholds are typically set to the half-way point between them. Equation 6.10 therefore becomes

$$p(t) = \begin{cases} p_0 + \frac{1}{2} \cdot g_0 & n = 0 \\ p_0 + \frac{1}{2} \cdot g_n + \sum_{i=0}^{n-1} g_i & n > 0 \end{cases}. \qquad [6.42]$$

As discussed in Section 6.2 the LIS delay is optimised such that the PACIFIC maximally integrates the SiPM signals and, therefore, measures the highest gain. However, the signals due to the detection of particles in the scintillating fibres do not arrive all simultaneously. This is mainly due to the propagation of light in the fibres with a time of 6 ns m^{-1} which considering the 2.5 m long fibres leads to spread of 15 ns in the arrival time of the signals at the SiPMs. Signals that arrive out of time with respect to the optimum for the PACIFIC integrator will not be integrated fully, therefore measuring a lower amplitude. Since the thresholds have been determined at the optimal time, these signals also have a lower chance of passing the comparator thresholds, effectively reducing the hit detection efficiency.
In order to ensure a high detection efficiency along the whole length of the fibres, the thresholds need to be scaled down to account for the loss of signal amplitude due to the different arrival times of the signals. The scaling factor can be derived by considering a simulation of the PACIFIC track and hold output for a one photoelectron signal, as shown in Figure 6.16. Depending on the arrival time of the SiPM pulse, only a certain fraction of the pulse will be integrated. Signals arriving early with respect to the PACIFIC integration window will only have their tail integrated which includes an undershoot due to the shaping of the signal (see Figure 3.12b), and therefore yields a negative amplitude. The more of the pulse is integrated, the higher the amplitude, and the maximum is reached when the aforementioned undershoot is outside the 25 ns integration window. Due to the peaking structure, the output of the track-and-hold stage is strongly dependent on the arrival time of the signals, and within a range of ±3 ns, the amplitude is reduced by 20 %. Considering the range of arrival times of 15 ns as mentioned above, indicated by the black dashed lines in Figure 6.16, the amplitude is reduced by 68 %. This coincides with the scaling factor that needs to be applied.

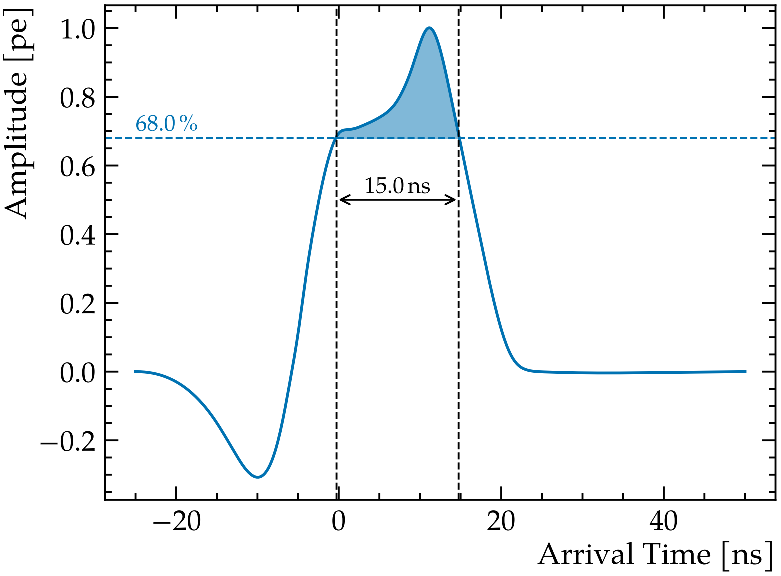

Figure 6.16 Simulation of the PACIFIC track and hold output as a function of the signal arrival time for a 1 pe signal [102].

Introducing the scaling factor s Equation 6.42 becomes

$$p(t) = \begin{cases} p_0 + s \cdot \frac{1}{2} \cdot g_0 & n = 0 \\ p_0 + s \cdot \left(\frac{1}{2} \cdot g_n + \sum_{i=0}^{n-1} g_i \right) & n > 0 \end{cases} . \qquad [6.43]$$

However, it is immediately evident that this scaling factor could trivially be combined with the factor of 1/2, and that separating the two is simply a matter of definition and habit.

6.6 Calibration Results

After the successful time alignment of the LIS presented in Section 6.2 the first full calibration of the entire detector was performed. This means recording threshold scans for the two integrators and three comparators of each of the 524 288 channels and fitting the resulting 3 145 728 S-curves using the model and procedure described in Sections 6.3 and 6.4. The results of the calibration fits are discussed and compared to the initial results obtained from the LIS threshold scans during the detector assembly presented in Section 4.10.

During the C-Frame assembly, only one of the three PACIFIC comparators was scanned, and the SiPMs were operated at ambient temperature. For the calibration of the entire detector, threshold scans for all three comparators have been performed, and the SiPMs were cooled at a set point of −40 °C. In both cases, the SiPM were operated at an overvoltage of 3.5 V above breakdown.

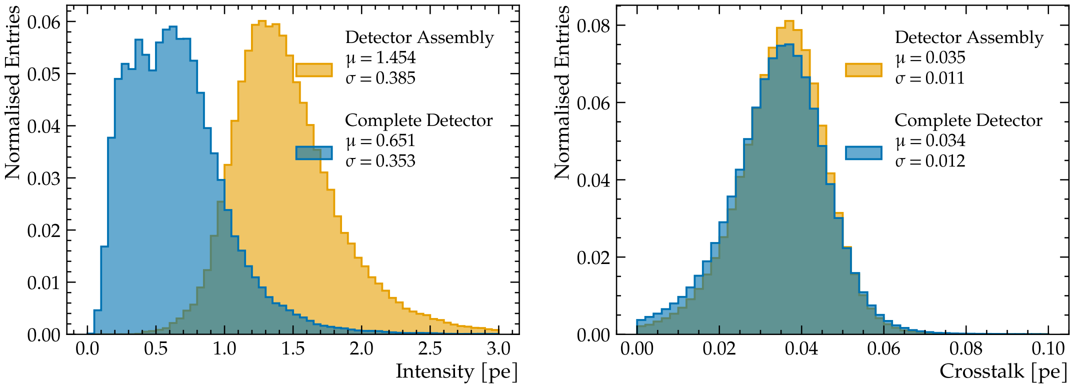

a b

Figure 6.17 Left: Histograms of the LIS intensity determined from light injection threshold scans taken during the assembly of C-Frames and with the complete detector. Right: Histograms of the crosstalk determined from light injection threshold scans taken during assembly and with the final detector.

6.6.1 Intensity

The amount of light the LIS provides is crucial to ensure a reliable calibration of the comparator thresholds. The LIS should provide a high enough light intensity such that, on average, 1.5 pe are detected by the SiPMs. This ensures that the 5 pe signals can be measured with adequate statistics. During the QA of the LIS, the bias and modulation currents were tuned to ensure this intensity. During the C-Frame assembly the light intensity was then measured again using threshold scans as presented in Section 4.10. The LIS intensity is shown in Figure 6.17a for the initial calibration performed during the C-Frame assembly and the first full calibration of the detector. The entries in both histograms have been normalised to the total number of entries since only one comparator was scanned during the assembly phase. For both measurements, the LIS pulse width was set to 15 ns.

Comparing the two histograms, the average LIS intensity has reduced by more than half from about 1.45 pe to only 0.65 pe detected by the SiPMs on average. This is below the specification of the LIS and has a negative impact on the quality of the calibration, especially considering that about 40 % of all channels detect less than 0.5 pe.

It is possible that the low light intensity can be traced back to the optimization of the LIS delay. In principle, the light intensity should not be affected by the LIS delay, and only the gain, the fraction of the SiPM signal being integrated, should change. This is demonstrated in Figure 6.18a, where the gain is plotted against the LIS delay. However, when the intensity is plotted against the LIS delay, as shown in Figure 6.18b, a significant decrease can be seen within the 10 to 25 ns range. This is also where the gain is at its highest, and it is coincidentally where the Calib C phase is set to 0 ns, as indicated in the figures by the grey shaded area. Unfortunately, this was discovered only at a later stage, so further investigations were not possible in the scope of this dissertation. Therefore, the LIS light intensity was left unchanged.

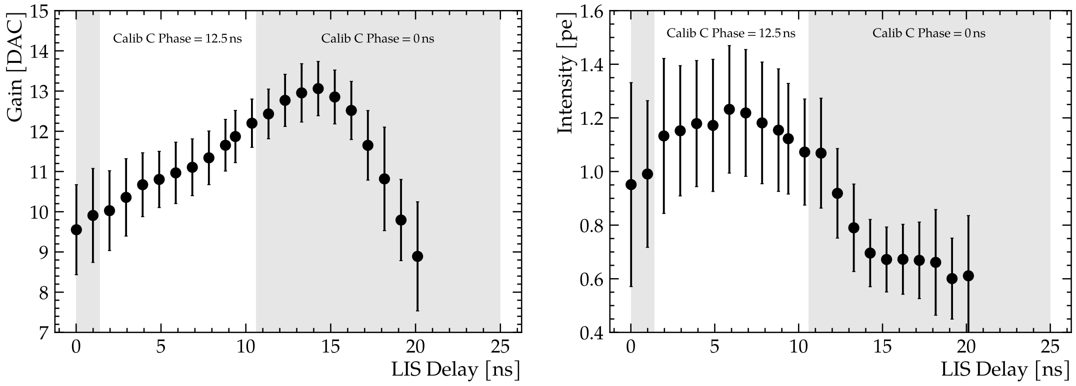

a **b**

Figure 6.18 Average gain (left) and light intensity (right) for the 512 channels corresponding to one fibre mat as a function of the LIS delay. The range of LIS delays where the Calib C phase is set to 0 ns is indicated by the grey shaded area.

6.6.2 Crosstalk

Unlike the light intensity, the crosstalk probability is not expected to impact the calibration of the thresholds. However, it is an essential parameter in the description of the S-curve as discussed in Sections 6.3 and 6.3.1. Crosstalk is mostly determined by the design of the SiPMs and applied over voltage as discussed in Section 3.5.3. For the nominal 3.5 V overvoltage, the expected probability for direct crosstalk is 3.3 % giving an average of 0.034 pe crosstalk events according to Equation 6.14.

The measured crosstalk probabilities are shown in Figure 6.17b again for the detector assembly and the calibration of the whole detector. The distributions are almost identical, with a mean value of about 0.035 pe of crosstalk events.

6.6.3 Pedestal

The integrators are part of the analogue processing in the PACIFIC and integrate the SiPMs signals within each 25 ns bunch period. Without any input signal, the output of the integrator is given by its DC baseline, also referred to as pedestal. Precise knowledge of the pedestal is crucial to determine the pulse height of the signals and to set the comparator thresholds correctly. Noise signals will trigger the comparator if the thresholds are set too close or even below the pedestal.

Figure 6.19a shows the pedestal distributions for the detector assembly and the complete detector. A shift of about 12 DAC is observed. However, it is not expected that the pedestal will be identical between the two measurements as the low-voltage power supply systems used during the assembly of the detector and in the complete detector installed in the cavern are different.

6.6.4 Gain

The pulse height of a detected photoelectron, here called photoelectron gain, is a measure of the signal amplitude. Accurately measuring the gain for each discrete

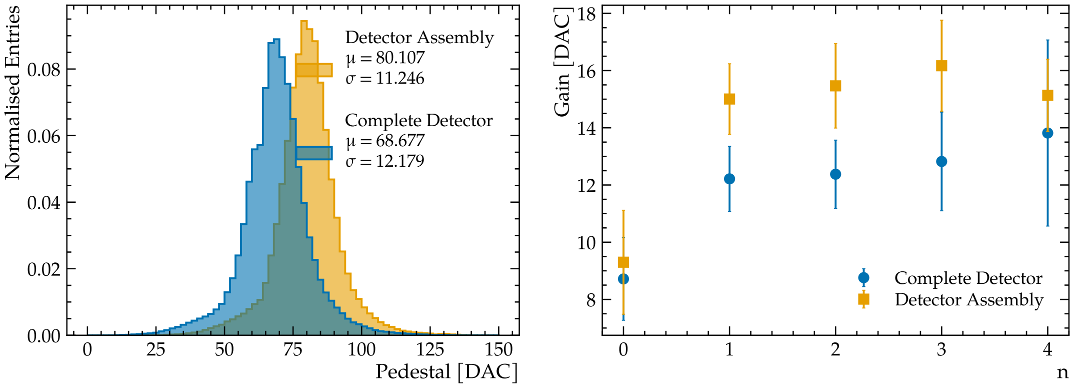

a

b

Figure 6.19 Left: Histograms of the crosstalk determined from light injection threshold scans taken during assembly and with the final detector. Right: Mean (markers) and standard deviation (error bars) of the five gain parameters as measured from light injection threshold scans during assembly and with the complete detector.

photoelectron signal is a key aspect of the threshold calibration. Together with the pedestal it allows to convert an amplitude expressed in photoelectrons into a PACIFIC comparator threshold in DAC values. As discussed in Section 6.3 the gain is given by the distance between the steps in an S-curve. Due to non-linear effects in the comparator threshold DAC, one gain parameter per step is needed.

The mean values and standard deviations of the five gain parameters g_n (see Equation 6.1) are shown in Figure 6.19b while the individual distributions are shown in Figure A.1. As discussed in Section 6.3, the 1 pe amplitude g_0 is typically lower by about 4 DAC compared to the other ones due to the hysteresis of the comparators. With the complete detector the gain is about 3 DAC lower compared to the C-Frame assembly, likely due to differences in the overvoltage. The widening of the distributions of g_3 and g_4 for in both measurements is due to the typically large non-linearity at a threshold DAC value of 128 as shown in the S-curve in Figure 6.9a. For the complete detector, g_4 shows a significantly bigger spread indicated by the error bars, with a standard deviation of 3.2 DAC. Due to the low LIS intensity of 0.65 pe, the 5 pe signal, corresponding to g_4 can not be resolved with large enough statistics, and the Likelihood fit produces unreliable results. This becomes apparent when looking at the correlation between the light intensity and the gain for the two data sets shown in Figure 6.20. A clear correlation between the low light intensity and spread of the gain can be observed for the measurement with the complete detector, while no correlation is present in the data from the assembly.

6.6.5 Width

The two width parameters σ_0 and σ_1 describe the electronic noise and the gain variation between the individual pixels, respectively. Similar to the crosstalk parameter, the two noise parameters are not expected to impact the calibration significantly but are an essential ingredient in the correct description of the S-curve. Checking their distributions, therefore, allows for verifying the calibration fits.

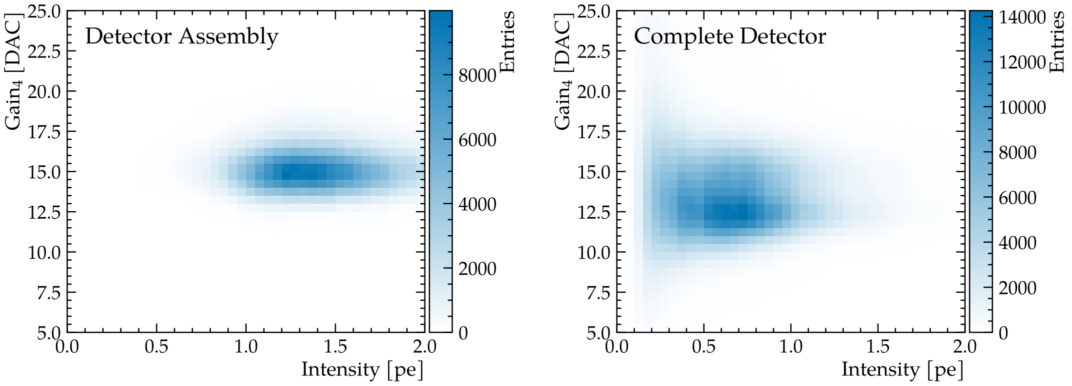

a

b

Figure 6.20 Heat maps showing the correlation of the gain parameter g_4 with the LIS light intensity determined during assembly and in the complete detector.

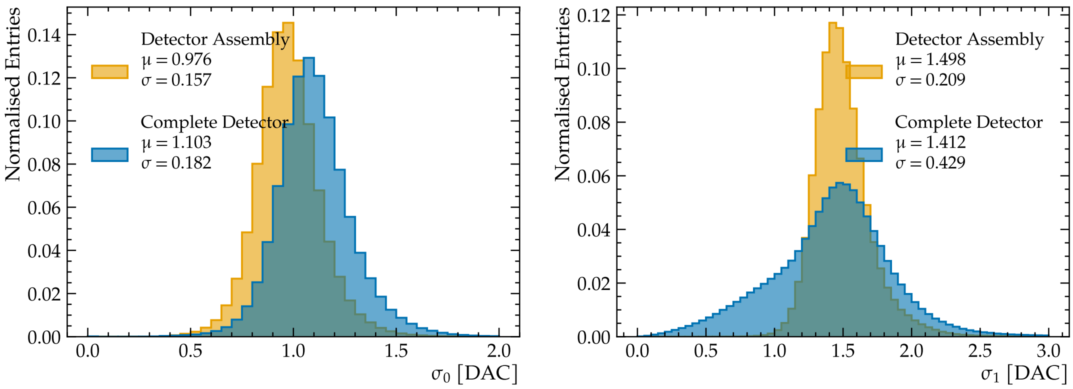

a

b

Figure 6.21 Distributions of the two width parameters σ_0 and σ_1 determined from light injection threshold scans during the C-Frame assembly and in the final detector.

The measured values for both parameters are shown in the two panels of Figure 6.21. The distributions of the electronic noise σ_0 measured during the assembly and with the complete detector have similar shapes with only a shift of about 0.12 DAC between the mean values. The distributions differ vastly for the variation in the width σ_1. While the peak value is almost identical between the two measurements with a value of ≈1.5 DAC, the measurement from the complete detector exhibits a large tail towards smaller values and a slight tail towards higher values. Since the variations in gain between the individual pixels are a property of the SiPMs and are not expected to change when operated at the same overvoltage, the large tails of the distributions hint at some issues in the calibration fits. Considering the correlation of σ_1 with the LIS intensity as shown in Figure 6.22a, a trend can be seen where lower intensities tend to give higher values of σ_1. The large tail towards lower values of σ_1 is present for intensities below 1 pe and leads to fits with a worse fit quality as shown in Figure 6.22b.

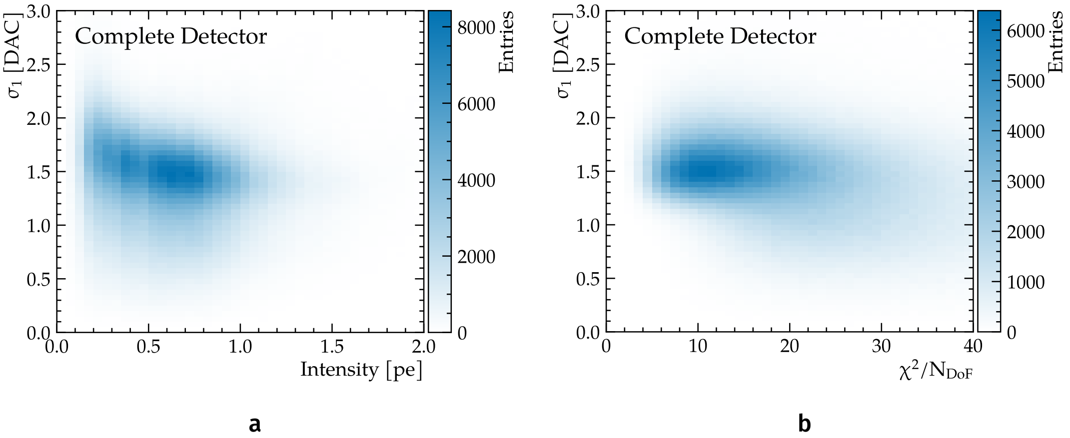

Figure 6.22 Heat maps showing the correlation of the width σ_1 with the light intensity and the fit quality in the complete detector.

6.6.6 Fit Quality

The fit quality (χ^2/N_{DoF}) calculated according to Equation 6.39 is shown in Figure 6.23 for the initial calibration during the C-Frame assembly and the first calibration of the entire detector. With a mean value of 23.2 compared to 11.9, the fit quality is significantly worse. Comparing the two-dimensional distributions of the fit quality versus the light intensity for the measurements during the assembly of the detector and with the complete detector shown in Figure 6.24 no clear correlation between the light intensity and fit quality is observed that would explain the worse fit quality.

6.6.7 Thresholds

Using the pedestal and gain parameters determined from the calibration fits the thresholds are calculated using Equation 6.43. Figure 6.26 shows the threshold DAC values minus the pedestal for all channels of the SciFi Tracker for exemplary threshold setting of $(1.5, 2.5, 3.5)$ pe and a scaling factor of 68 %. The distribution for each of three comparators Vth_1, Vth_2, Vth_3 show two peaks. The large peaks are at DAC values of 10 DAC, 18 DAC, and 26 DAC, respectively, which are the expected values considering the average gain of 12 DAC and a scaling of 68 %. The second smaller peaks are located at higher DAC values. Channels with these threshold DAC values are in regions of the detector that could not be calibrated (see Section 6.2) and therefore use the gain values determined during the assembly of the detector which are on average 3 DAC larger and have no scaling applied. Figure 6.26 shows the threshold DAC values minus the pedestal for one comparator of all channels in the SciFi Tracker. The regions with higher thresholds correspond to the same regions where the LIS latency or LIS delay is not correctly set (see Figures 6.5 and 6.8).

Integrator Differences

The PACIFIC integrates the SiPM signals within each 25 ns bunch period using two interleaved integrators (see Section 3.6.1). While one integrator is working, the other is

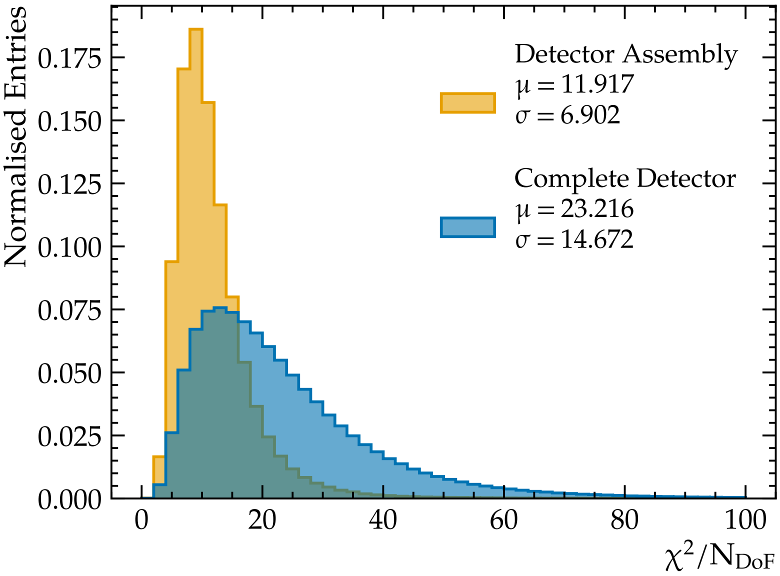

Figure 6.23 Distribution of the fit quality of the calibration fits to light injection threshold scans during the C-Frame assembly and with the complete detector.

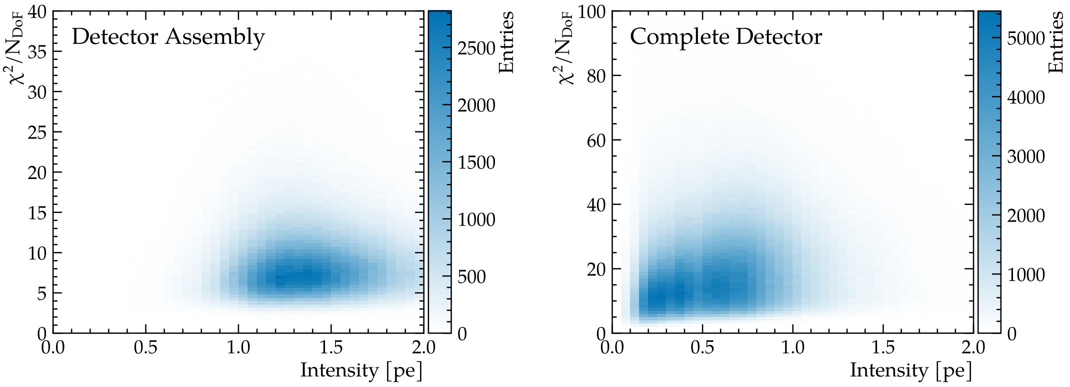

a

b

Figure 6.24 Heat maps showing the correlation of the fit quality and the light intensity determined during the C-Frame assembly and in the complete detector.

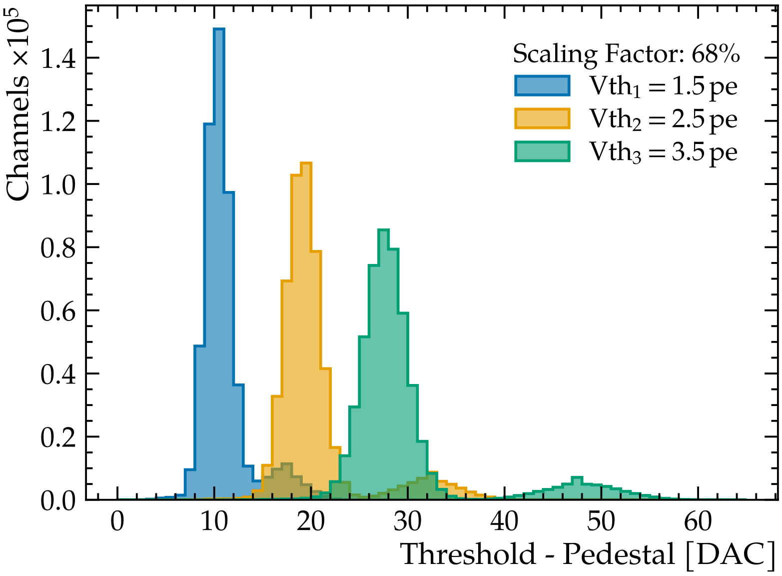

Figure 6.25 Distributions of the threshold DAC values for exemplary threshold setting of $(1.5, 2.5, 3.5)$ pe with a scaling factor of 68 %.

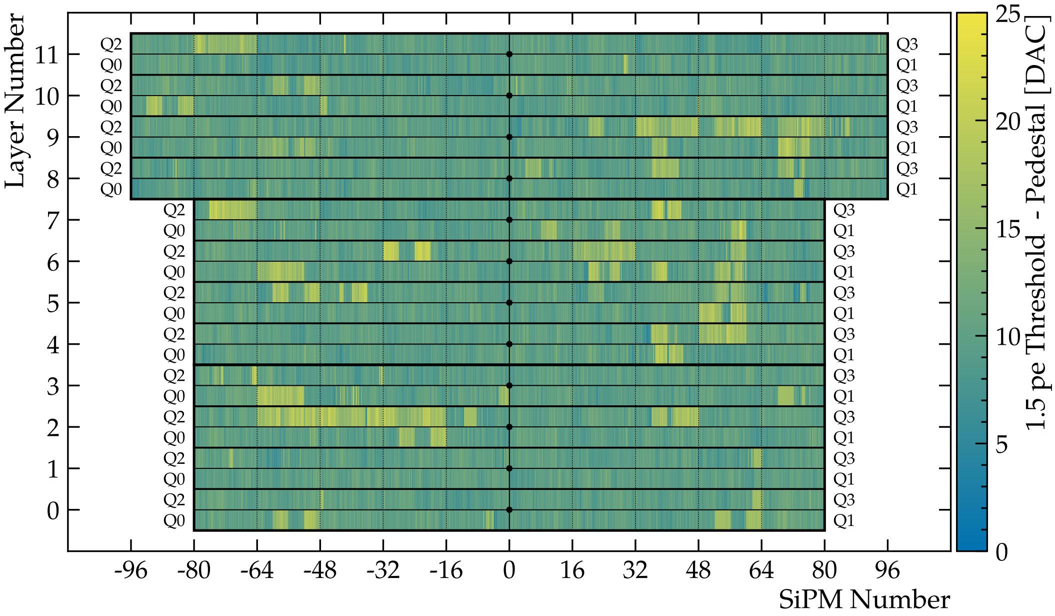

Figure 6.26 Distribution of the 1.5 pe threshold values with a scaling factor of 68 % throughout the detector. Regions with larger thresholds use the calibration constants determined during the assembly of the detector.

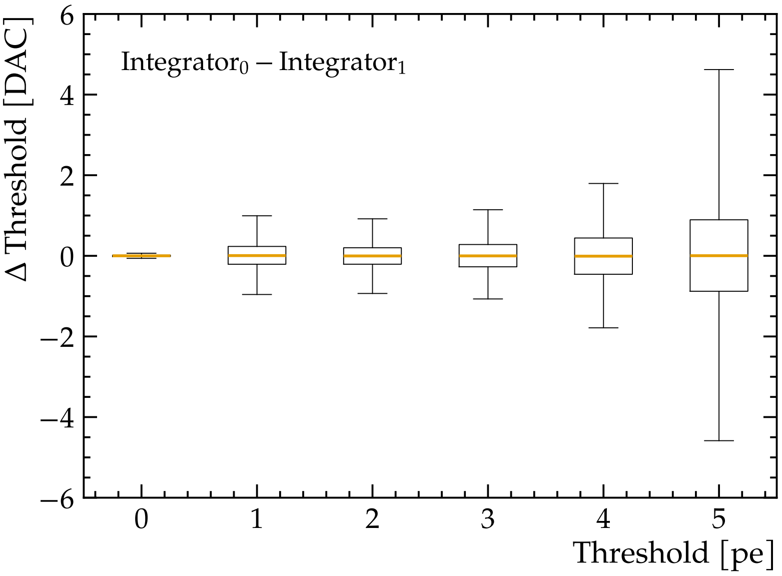

Figure 6.27 Box plot of the differences in the determined photoelectron thresholds between the two integrators of each PACIFIC channel. The box shows the 25th and 75th percentiles, and the orange line indicates the median (50th percentile). The whiskers mark the 5th and 95th percentiles.

being reset. As each of the two integrators within one PACIFIC channel is connected to the same three comparators, a uniform behaviour is important to ensure a consistent performance across all bunch periods.
The difference of the calibrated thresholds corresponding to photoelectron amplitudes p_n between the two integrators are shown in Figure 6.27. The median difference is zero up to the 5 pe threshold, as indicated by the orange line. The boxes show the 25th to 75th percentiles while the whiskers correspond to the 5th and 95th percentiles, respectively. The difference in the pedestal (0 pe threshold) is negligible and far below 1 DAC. The differences increase to higher threshold values to about ±2 DAC. For the 5 pe threshold, the effect of the low LIS intensity becomes apparent again. As already discussed in Section 6.6.4, the low intensity leads to unreliable fit results for the 5 pe threshold, which then also results in large differences of ±5 DAC between the two integrators. Considering an average gain of about 12 DAC this variation is significant. Whether this has an impact on the performance of the SciFI Tracker in terms of hit detection efficiency or noise rates needs to be studied.

Comparator Differences

In contrast to the integrators, differences in the comparators do not directly influence the performance as the thresholds are set and calibrated separately for each comparator. However, having precise knowledge of the correlation between the three comparators of one channel would allow the threshold of all three based on the calibration of one of them. This would reduce the time it takes to perform a threshold scan by a factor of three.

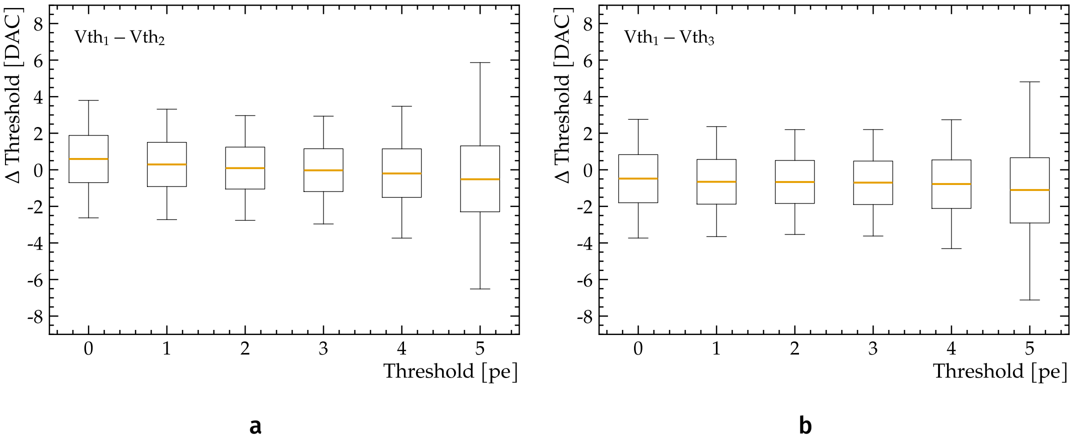

a

b

Figure 6.28 Box plots of the differences in the determined photoelectron thresholds between the three comparators of each PACIFIC channel. The box shows the 25th and 75th percentiles, and the orange line indicates the median (50th percentile). The whiskers mark the 5th and 95th percentiles.

Analogous to the differences between the two integrators, Figure 6.28 shows the differences between comparators Vth_1 and Vth_2, and Vth_1 and Vth_3. A slight bias of 0.5 DAC and -0.5 DAC can be seen for Vth_2 and Vth_3, respectively. Looking at the 5th and 95th percentiles, the overall spread approaches ± 4 DAC. These variations are significant regarding the gain of 12 DAC and illustrate the need to calibrate each comparator individually. The large spread for the 5 pe threshold can again be attributed to the low light intensity and unreliable fit result.

6.7 Calibration Summary

The calibration of the PACIFIC comparator thresholds is integral for the successful operation of the SciFi Tracker. The method presented in this chapter allows for calibrating the threshold DAC of all 1 572 864 comparators using threshold scans under pulsed illumination from the LIS.

For an accurate calibration the injection of the light onto the SiPMs needs to be synchronised with the PACIFIC integration window, such that the resulting SiPM signals are maximally integrated. This has been achieved for 97 % of the detector. Furthermore a sufficient LIS light intensity of 1.5 pe on average is required to resolve signals up 5 pe with sufficient statistics. Due to an unresolved issue with the LIS, the light intensity was only 0.65 pe on average, resulting in an unreliable calibration for about 17 % of all channels for thresholds larger than 4 pe. Rectifying these issues was, unfortunately, not possible within the time frame of this dissertation.

Nonetheless, the comparator thresholds for 417 792 out of 524 288 channels have been successfully calibrated. Using this calibration, the performance of the SciFi Tracker will be evaluated for different threshold settings in the following chapter.

7 | First Detector Performance Studies

Tracking charged particles and reconstructing their momenta with high precision is crucial for the physics programme of LHCb. Located behind the dipole magnet, the SciFi Tracker is responsible for charged particle tracking and provides the best momentum estimate by measuring the deflection of tracks from the VELO and UT detectors inside the magnetic field. The performance of the SciFi Tracker therefore has a direct impact on the tracking and spectrometer performance. To achieve the physics goals of the LHCb experiment, the SciFi Tracker is required to have a high hit detection efficiency of ≈99 %, a spatial resolution better than 100 µm, while at the same time maintaining a low noise cluster rate below 2 MHz per SiPM [40].

The SciFi Tracker reconstructs particle hits and rejects noise by building clusters from the 2-bit amplitude information of each channel already in the frontend electronics. The 2 bits encode how many of the three comparators of a PACIFIC channel were exceeded, and if the thresholds are calibrated, correspond to the number of detected photoelectrons (pe). The cluster-finding algorithm, introduced in Section 3.6.3, only considers channels where the signals exceed the lowest threshold (Vth_1). Adjacent channels with an amplitude above the lowest threshold are grouped, and if the sum of their amplitudes exceeds the sum of the lowest and middle threshold (Vth_2), a cluster is formed. In addition, single channels can form a cluster if their amplitude exceeds the highest threshold (Vth_3). The ability of the SciFi Tracker to detect particle hits and to reject noise from thermal excitations in the SiPMs, therefore, directly depends on the chosen threshold settings and their correct calibration; lower thresholds are expected to result in a higher detection efficiency and higher noise rate while higher threshold will yield a lower detection efficiency and lower noise rate. The hit resolution, on the other hand, is not expected to be strongly dependent on the threshold settings.

Using the calibration constants determined in the previous chapter, three threshold configurations $(1.5, 2.5, 3.5)$ pe, $(1.5, 2.5, 4.5)$ pe, and $(2.5, 3.5, 4.5)$ pe have been computed according to Equation 6.43 for two scaling factors of 68 % and 100 %. The scaling factor has been introduced in Section 6.5 to account for the different arrival times of signals. The resulting six threshold settings, summarised in Table 7.1, are applied to the detector to evaluate the hit detection efficiency, hit resolution, and noise cluster rates. All measurements have been carried out with the SiPMs cooled at a cooling plant set point of −40 °C[1] and operated with 3.5 V overvoltage.

7.1 Hit Detection Efficiency

A high hit detection efficiency is crucial as it directly affects the tracking efficiency and thus the physics performance of the LHCb experiment. The more measurement points are available the better the track reconstruction and subsequent momentum estimate.

---

[1] The temperature at the SiPMs is about 2 °C higher.

Table 7.1 Summary of the threshold settings used to evaluate the performance of the SciFi Tracker.

Vth_1 [pe]	Vth_2 [pe]	Vth_3 [pe]	Scaling Factor [%]
1.5	2.5	3.5	68
1.5	2.5	4.5	68
2.5	3.5	4.5	68
1.5	2.5	3.5	100
1.5	2.5	4.5	100
2.5	3.5	4.5	100

Assuming a hit detection efficiency for the SciFi Tracker of 99 % (97.3 % including acceptance gaps) about 99.6 % of tracks will have ten or more hits in the SciFi Tracker [103].

The hit detection efficiency ε is defined as the number of observed hits in a region of the SciFi Tracker divided by the number of expected hits. The number of expected hits is estimated by considering charged particle tracks from 2023 *PbPb* collision data and extrapolating the tracks to the layer under study. The number of observed hits is estimated from tracks with an associated hit in the layer under study. As illustrated in Figure 7.1, a hit in the SciFi is considered to be a line spanned by the x location of the hit evaluated at $y = 0$ and the orientation of the fibre mat. For each track the distance of closest approach between the lines spanned by the track and the hits in the SciFi is calculated. A track is considered to be *associated* if the distance of closest approach is less than 1 mm. The efficiency at a location $\vec{x}$ is then given by

$$\varepsilon(\vec{x}) = \frac{N^{Tracks}_{associated}(\vec{x})}{N^{Tracks}_{unassociated}(\vec{x}) + N^{Tracks}_{associated}(\vec{x})}. \quad [7.1]$$

During the time of the data taking the VELO was not closed to its nominal aperture but opened to about 49 mm resulting in an asymmetric acceptance for tracks reconstructed using the VELO and SciFi Tracker. For the measurements presented in the following only so-called *T tracks* have been considered. T tracks are reconstructed using exclusively the hits in the SciFi Tracker [104] and are therefore not affected by the limited VELO acceptance. In order not to bias the efficiency measurement, the layer under study is excluded from the track reconstruction. Only well reconstructed tracks with a χ^2/N_{DoF} less than 2, at least one hit in each of the remaining 11 SciFi layer, a momentum larger than 2 GeV, as well as a transverse momentum above 400 MeV have been selected. Furthermore, high-multiplicity events with an estimated number of SciFi hits above 10 000 were removed by a global event cut prior to the track reconstruction and only events with a track-multiplicity between 2 to 100 were considered. All event and track selection criteria are summarised in Table A.1.

The hit detection efficiency is evaluated for each of the six threshold configurations listed in Table 7.1 from the distribution of the measured inefficiencies in each channel

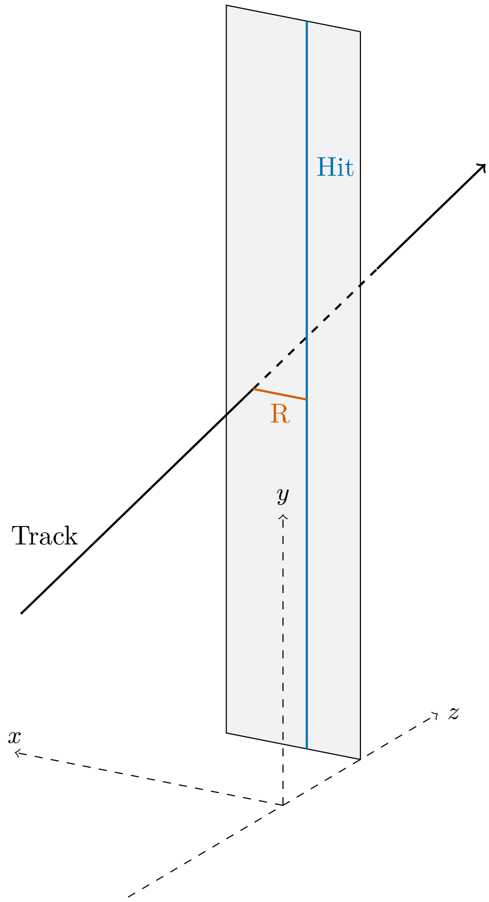

Figure 7.1 Illustration of the association of a track (black line) to a SciFi hit (blue line). A SciFi hit only contains information about the *x* position, a line is constructed from the *x* position and the orientation of the fibre mat. The distance of closest approach (*R*) between the track and the hit is shown in orange.

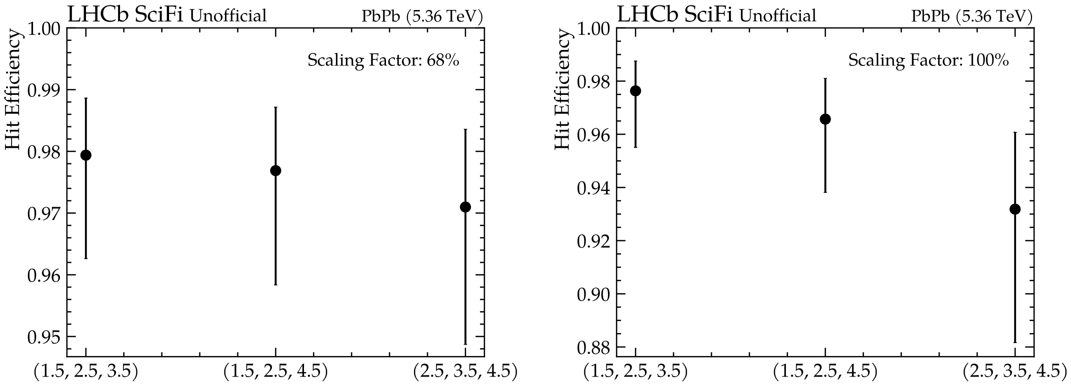

a b

Figure 7.2 Hit efficiency for different threshold configurations and scaling factors wit the markers indicating the median hit efficiency and the errors the 25th and 75th percentiles of a log-normal distribution.

weighted by the number of tracks through the channel. Only channels with at least ten tracks passing through them were considered. Additionally channels close to gaps in the detector have been excluded[2]. Figure 7.2 shows the median hit efficiency as obtained from a log-normal fit [105] to the individual distributions. The error bars indicate the 25th and 75th quantiles of the fitted log-normal distribution. The individual distributions and fit results are given in Figure A.4. Using the lowest thresholds settings of $(1.5, 2.5, 3.5)$ pe and a scaling factor of 68 % the median hit detection efficiency is 98.0 %. As expected, increasing the highest threshold by 1 pe only slightly lowers the efficiency to 97.7 %, while increasing all three thresholds by 1 pe to $(2.5, 3.5, 4.5)$ pe yields an efficiency of 97.0 %. A similar trend is observed when using a scaling factor of 100 %, where the efficiencies are 97.7 %, 96.5 %, and 93.2 %, respectively for the three threshold settings. The efficiency per channel varies by about 3 to 8 % as shown by the size of the error bars. This indicates that some parts of the detector are more efficient than others.

Looking at the median hit efficiency for each of the twelve layers, shown in Figure 7.3a for one of the threshold configurations, two things are immediately noticeable: the central two layers in each station, indicated by the dotted line, have a systematically lower hit detection efficiency and the hit efficiency is decreasing towards the back of the detector. The lower efficiency in the central two layers is reflected by the geometry of the SciFi Tracker, where each station is composed of four layers with the two inner layers rotated by $\pm 5°$. While the vertical layers reach efficiencies of up to $\approx$99 %, the stereo layers have a hit detection efficiency that is about 1 to 1.5 % lower on average. Figure 7.4 shows the average efficiency in dependence of the x and x position of track intercept for a stereo layer (layer 2) and a vertical layer (layer 3). The lower efficiency in the rotated layer is predominantly observed at $y = 0$ with an alternating pattern. Due to the rotation of the stereo layer by 5°, some fibre mats from the upper quadrant reach into the lower one and also the other way around as is illustrated in Figure 7.3b. The

---

[2] The excluded SiPM channels are 0, 1, 62, 63, 64, 65, 126, 127 as they lie at the edge of an acceptance gap.

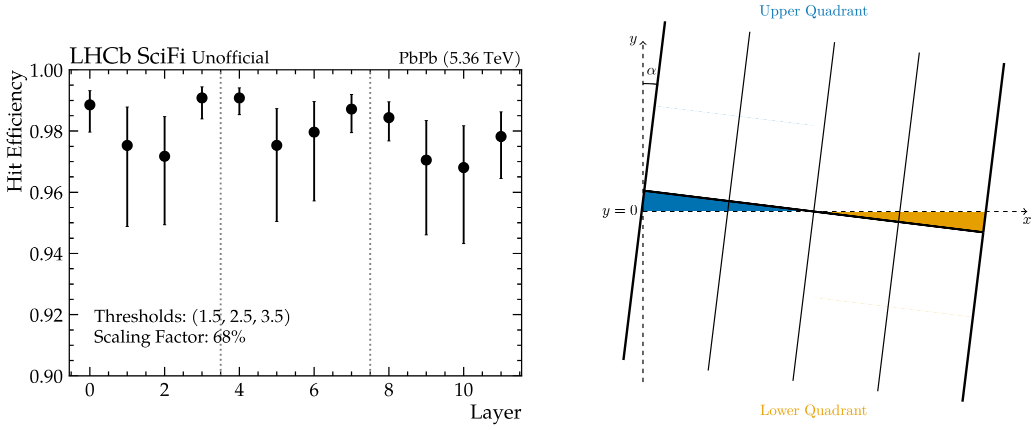

a

b

Figure 7.3 Left: Median hit detection efficiency as a function of the SciFi layer. The error bars indicate the 25th and 75th percentiles of a log-normal distribution. Right: Illustration of one module in a stereo layer of the SciFi Tracker with four fibre mats in the upper quadrant and four mats in the lower quadrant. Two mats from the upper quadrant reaching into the lower one the other way around as shown by the shaded triangles. Adapted from Reference Reference [103].

areas where the fibre mats of either quadrant reach into the other are shown as shaded triangles. As only the hits in the fibre mat intersected by the track are considered, it is possible that due to misalignment the fibre mat belonging to the opposite quarter is picked out and searched for hits; since a hit is unlikely to be found the efficiency will be reduced.

The decrease in efficiency of about 0.5 % towards the last layers can be attributed to a sub-optimal time alignment of the detector relative to the bunch crossings. The processing of the SiPM signals by the PACIFIC is sensitive to their arrival time (see Section 3.6.1). Signals that arrive either to late or to early with respect to the PACIFIC clock will not be integrated fully, resulting in a lower total charge and are therefore less likely to pass the thresholds and create clusters, thus reducing the hit detection efficiency. The time alignment procedure, described in Reference [106], was biased due to the large amount of secondaries present in the *PbPb* collisions, and thus the wrong working points were chosen.

Figure 7.5 shows an overview of the average hit detection efficiency for every SiPM in the SciFi Tracker. The loss in hit efficiency in the stereo layers can be seen, as well as the generally lower efficiency in station T3 (layers 8 to 11). In addition one group of 16 SiPMs corresponding to one module in layer 0, quadrant Q2, have a significantly lower hit efficiency of ≈80 %. Comparing this to Figure 6.26, the threshold settings are not the cause for the loss in efficiency since they are comparable to the neighbouring SiPMs; misalignment is the likely cause for these inefficiencies.

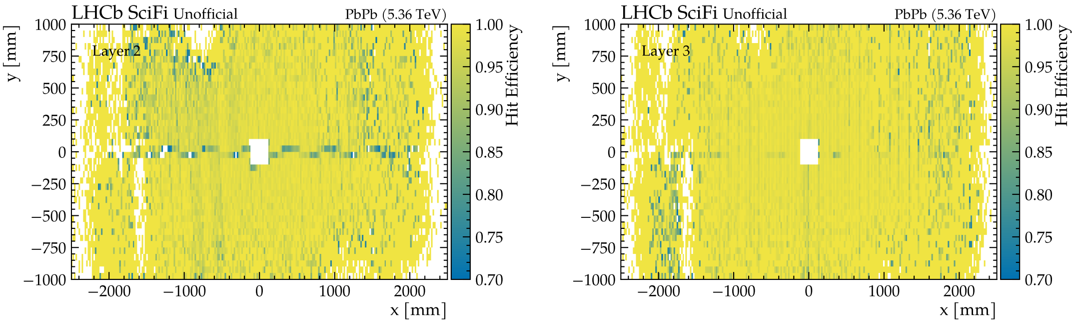

a

b

Figure 7.4 Hit efficiency in dependence of the *x* and *y* position of the track intercept for a stereo layer (left) and a vertical layer (right) using threshold settings (1.5, 2.5, 3.5) pe with 68 % scaling.

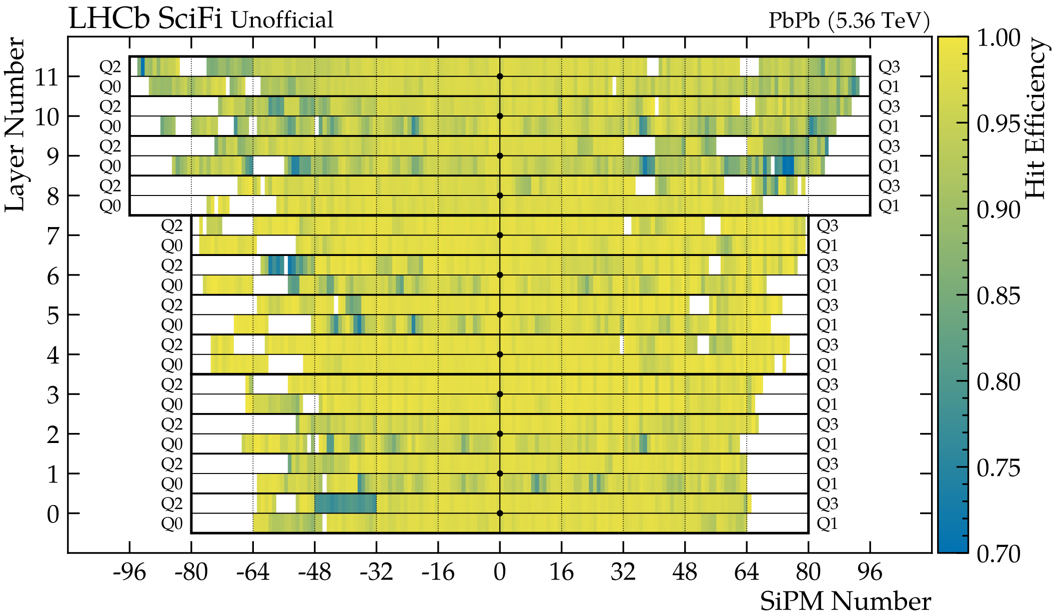

Figure 7.5 Overview of the average hit detection efficiency per SiPM in the detector using threshold settings (1.5, 2.5, 3.5) pe with 68 % scaling. Layers are numbered from 0 to 11 and ordered with increasing *z* positions. The SiPMs are numbered from 1 to 96 or −1 to −96 according to their *x* position within the LHCb coordinate system. Positive indices correspond to positive *x* while negative indices correspond to negative *x*.

7.2 Hit Resolution

The hit resolution of the individual sub-detectors is one factor contributing to the momentum resolution of the LHCb spectrometer. At the location fo the SciFi Tracker the particles have already traversed a significant amount of material, so that the momentum resolution is limited by the multiple scattering for momenta up to about 80 GeV [26]. However, a low hit resolution in the SciFi is still beneficial as it aids the pattern recognition allowing for smaller search windows when searching for candidate tracks [26].

The hit resolution is determined using well reconstructed T tracks from 2023 *PbPb* data. Only tracks with a track-fit $\chi^2/N_{DoF} < 2$, hits in all twelve SciFi layers, and a momentum larger than 10 GeV have been considered. As the performance of the pattern recognition and tracking algorithms tends to decrease for too high occupancies, high-multiplicity events with an estimated number of SciFi hits above 10 000 were removed by a global event cut. The events and track selection criteria are summarised in Table A.2. For a given track the estimated hit position from the track in each SciFi layer is compared with position of the recorded hit. The resulting distribution is shown in Figure 7.6a, where the core of the distribution is fitted by a Gaussian function. A hit resolution of 138 µm given by the width of the distribution is measured. The threshold settings only have a small impact on the hit resolution, with about 1 µm variation going from the lowest to the highest threshold setting, as shown in Figure 7.6b.

As was already discussed in Section 7.1 the alignment of the detector elements is not yet ideal and is therefore also expected to impact the measured residuals. Figure 7.7 shows the mean and width of the unbiased residuals in each of the twelve SciFi layers for one of the tested threshold settings. The mean of the unbiased residual shown in Figure 7.7a shows no clear dependence on the layer and the values are scattered between −50 to 50 µm. The width of the unbiased residuals shown in Figure 7.7b, on the other hand is varying significantly with the layer number. The stereo layers in stations T1 and T3 (layers 1, 2, 9, and 10) exhibit residual widths larger than 160 µm,

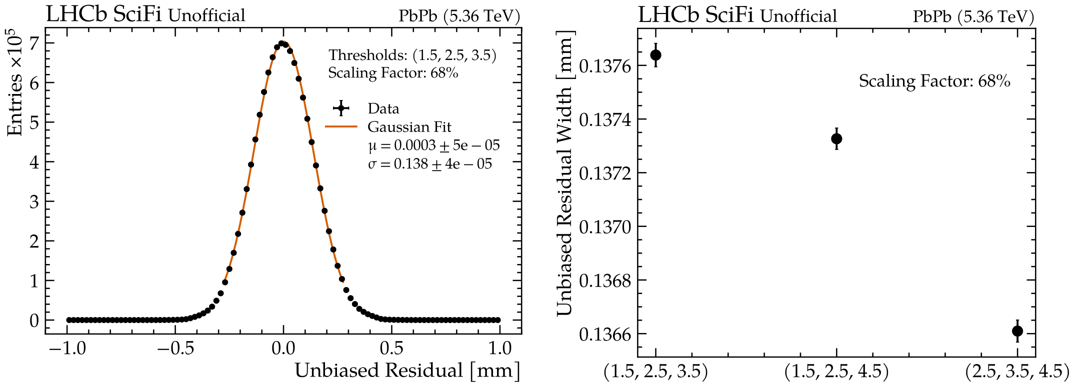

a

b

Figure 7.6 Left: Distribution of the unbiased residuals for threshold of (1.5, 2.5, 3.5) pe with 68 % scaling. A Gaussian is fitted to the core of the distribution. Right: width of the unbiased residual distribution for three different threshold settings.

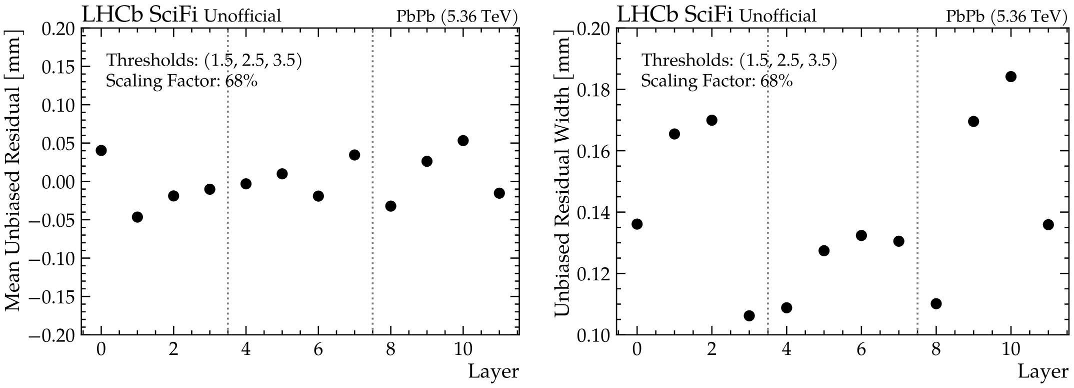

a **b**

Figure 7.7 Mean and width of the unbiased residuals as a function of the SciFi layer for threshold settings of $(1.5, 2.5, 3.5)$ pe with 68 % scaling.

while some of the vertical layers $(3, 4, \text{and } 8)$, are close to the design goal of 100 µm. Figure 7.8 shows the distribution of the width of the unbiased residuals throughout the whole detector, where the residuals have been accumulated always for 512 channels corresponding to one fibre mat. The best hit resolution, in the order of 75 to 100 µm, is measured for the central modules of layers 3, 4, 7, and 8. The hypothesis that the low hit efficiency in the group of 16 SiPMs in layer 0, quadrant Q2 (see Figure 7.5), is due to mis-alignment has to be rejected, as the width of the unbiased residuals is not significantly larger compared to the neighbouring fibre mats.

It is evident that the measurement of the hit resolution contains a series of systematic effects, as shown in Figures 7.7 and 7.8, that require further investigations. However, it should also be noted that the detector reaches the design requirement of providing a single hit resolution of <100 µm.

7.3 Noise Cluster Rates

A limiting factor for the performance of the SciFi Tracker is the rate of noise clusters due to dark counts of the SiPMs. A high noise cluster rate will lead to a higher chance of reconstructing spurious tracks not associated to a real particle. In addition, the limited bandwidth of the readout electronics, which allows to send a maximum of 10 or 16 clusters per SiPM and per event, should not be dominated by noise clusters such that the real clusters will not be transmitted. Based on the available bandwidth, a noise cluster rate of 2 MHz per SiPM or 16 kHz per channel is tolerable [107].

The noise cluster rate is evaluated from data collected during periods where no beam was circulated in the LHC and the SiPMs were cooled at a cooling plant set point of −40 °C. The noise cluster rate for a particular channel x is given by the number of noise hits (N_{noise}) divided by the number of recorded events (N_{events}) and the bunch spacing of 25 ns

$$R(x) = \frac{1}{25} \cdot \frac{N_{noise}(x)}{N_{events}} \text{GHz}. \qquad [7.2]$$

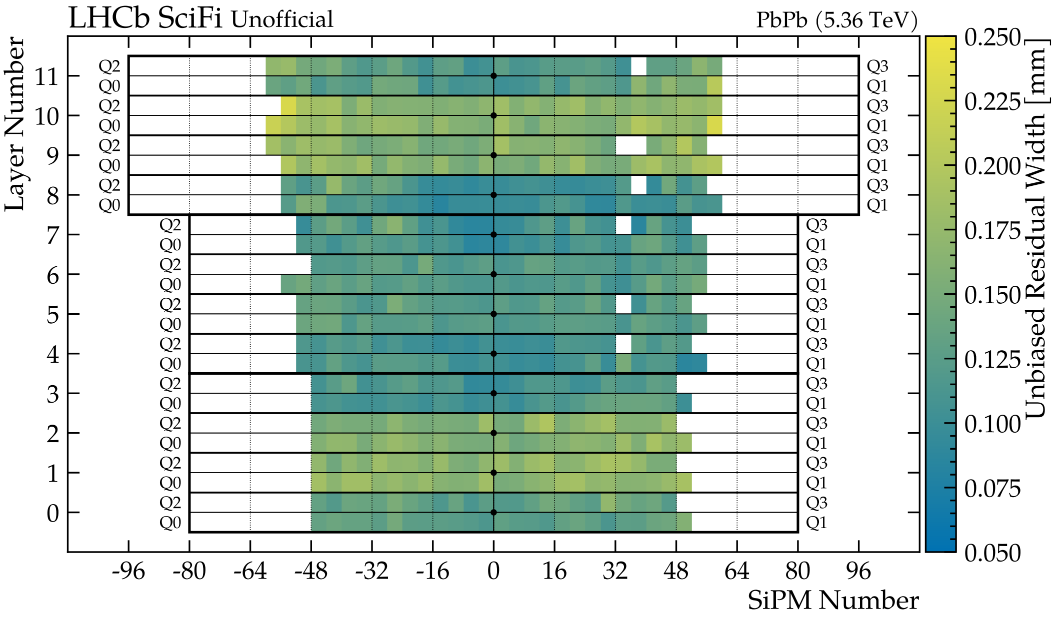

Figure 7.8 Overview of the width of the unbiased residual per fibre mat (four SiPMs) throughout the detector using threshold settings (1.5, 2.5, 3.5) pe with 68 % scaling. Bins with low statistics are shown in white.

Figure 7.9 shows the median noise cluster rates per channel for each of the six threshold configurations. The median, the 25th, and 75 quantiles are obtained from a log-normal fit to the distribution of the noise cluster rates per channel. The distributions and fit results are provided in Figure A.8. With the lowest threshold settings of (1.5, 2.5, 3.5) pe and a scaling factor of 68 % the median noise cluster rate per channel is measured to be ≈500 Hz. Setting the highest threshold to 4.5 pe reduces the noise rate by a factor of two, while increasing all thresholds by 1 pe gives a reduction by factor of ten. A similar trend can be observed for the measurements using a scaling factor of 100 %, however, the noise cluster rates are overall lower by a factor of 15.

It is not expected that the noise rate is uniformly distributed across the whole detector, but rather a gradient from the first to the last layer is expected as the SiPMs located here receive a higher non-ionising dose due to backscattering of low-energy neutrons from the calorimeter system [34]. Figure 7.10 shows the noise cluster rate for every channel in the SciFi Tracker, where a gradual increase in the rates from layer 0 to layer 11 can be seen. It should be noted that the detector so far was exposed to very little radiation corresponding to about 1.55 fb^{-1} of integrated luminosity or about 3 % of the total expected luminosity. In some areas, such as layer 2 quadrant Q2, the rate of noise clusters per channel is close to zero. These areas correspond to exactly the regions where the calibration constants have not yet been updated, as shown previously in Figure 6.26, and the constants determined during the assembly of the detector are still used. The thresholds in these regions are therefore higher by about 5 DAC or 0.4 pe, resulting in significantly lower noise rates of just a few hertz.

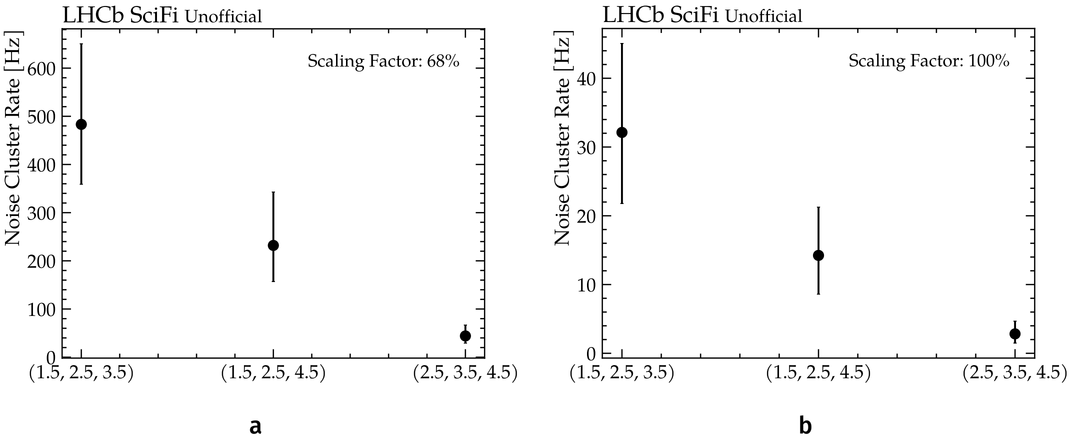

Figure 7.9 Noise cluster rates for different threshold configurations and scaling factors with the markers indicating the median noise cluster rate and the error bars the 25th and 75th percentiles of a log-normal distribution.

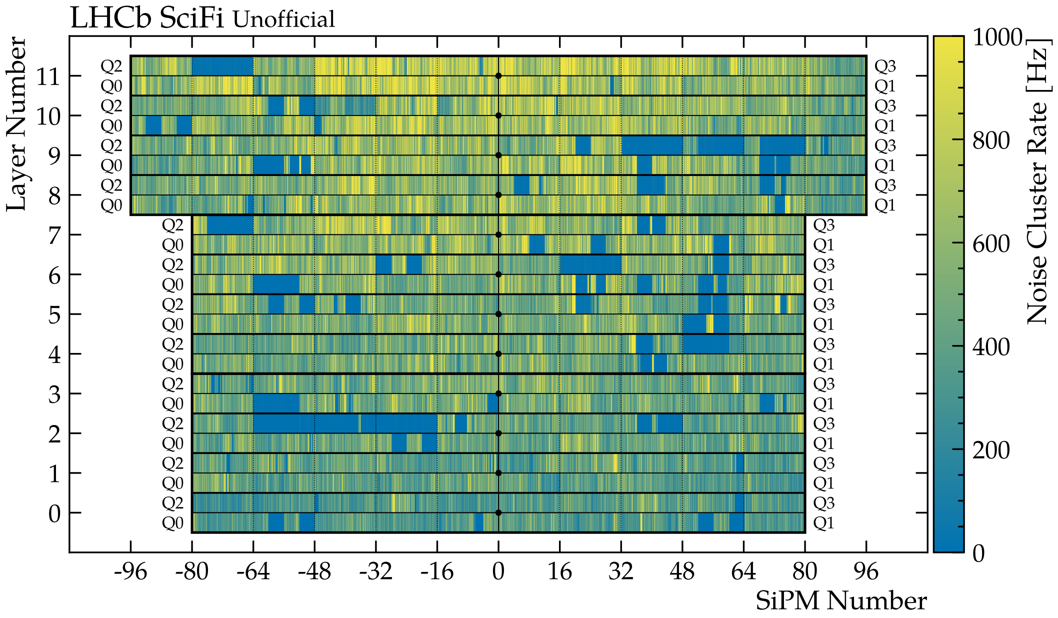

Figure 7.10 Overview of the noise cluster rate per channel using threshold settings $(1.5, 2.5, 3.5)$ with 68 % scaling.

7.4 Noise Digit Rates

Clusters in the SciFi Tracker are constructed from the 2-bit output (*digits*) of the individual PACIFIC channels using the algorithm discussed in Section 3.6.3. These digits have four possible values ranging from 0 to 3 corresponding to the number of fired comparators in a given channel. Considering the noise rates for the different digits, thus allows for investigating whether the thresholds have been set correctly. If a threshold is set too low the noise digit rate is expected to be higher, conversely if it is

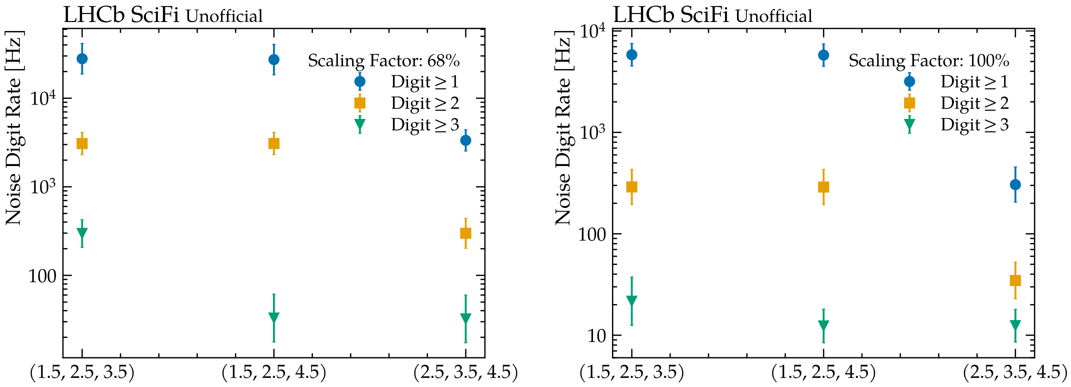

a **b**

Figure 7.11 Noise rates for the raw PACIFIC outputs (digits) for different threshold configurations and scaling factors. The markers indicate the median and the error bars the 25th and 75th percentiles of a log-normal distribution.

set to too high the noise digit rate is lower. Equivalent to Equation 7.2, the noise digit rate per channel x for digits $\geq i, i \in [1, 2, 3]$ is given by

$$R_i(x) = \frac{1}{25} \cdot \frac{N_{noise,i}(x)}{N_{events}} \text{GHz} \qquad [7.3]$$

with $N_{noise,i}(x)$ the rate of digits $\geq i$ in channel x.

Figure 7.11 shows the median noise digits rate per channel for the six threshold settings in Table 7.1. The rates for digits ≥ 1, ≥ 2, and ≥ 3 are shown as different markers and colours. From the lowest to the highest threshold settings the noise rates for all three ADC counts reduces by about a factor of ten for both scaling factors. This illustrates the effectiveness of the cluster-finding algorithm. While the rate of single channel signals with an amplitude above the lowest threshold (digit ≥ 1) is in the order of 10 kHz for a scaling factor of 68 %, the corresponding noise cluster rate is only about 500 Hz (see Figure 7.9).

Channels with threshold settings that are too high or too low can be identified by looking at the distribution of the noise digit rate for digits greater than or equal to one across the entire detector. This is illustrated in Figure 7.12 where a few channels with a noise digit rate exceeding 1 MHz appear as yellow lines, while channels with a lower-than-average noise digit rate appear as blue rectangles.

For example, in layer 8, Q1, SiPM 64 to 80, there is a group of channels with a noise digit rate of approximately 10^3 Hz located next to a group with rates around 10^7 Hz. By comparing this to the threshold values in Figure 6.26, it is evident that the channels with lower noise digit rates have thresholds that are about 5 DAC higher than expected. On the other hand, channels with higher noise digit rates have thresholds that are about 5 to 10 DAC lower than expected. This implies that the thresholds are set below the 1 pe amplitude. Given that the predominant noise in the SciFi Tracker is from single photoelectron dark counts, this leads to a high digit noise rate.

Though more channels with high noise rates can be observed, investigating all channels with excessive noise rates was unfortunately beyond the scope of this thesis. However,

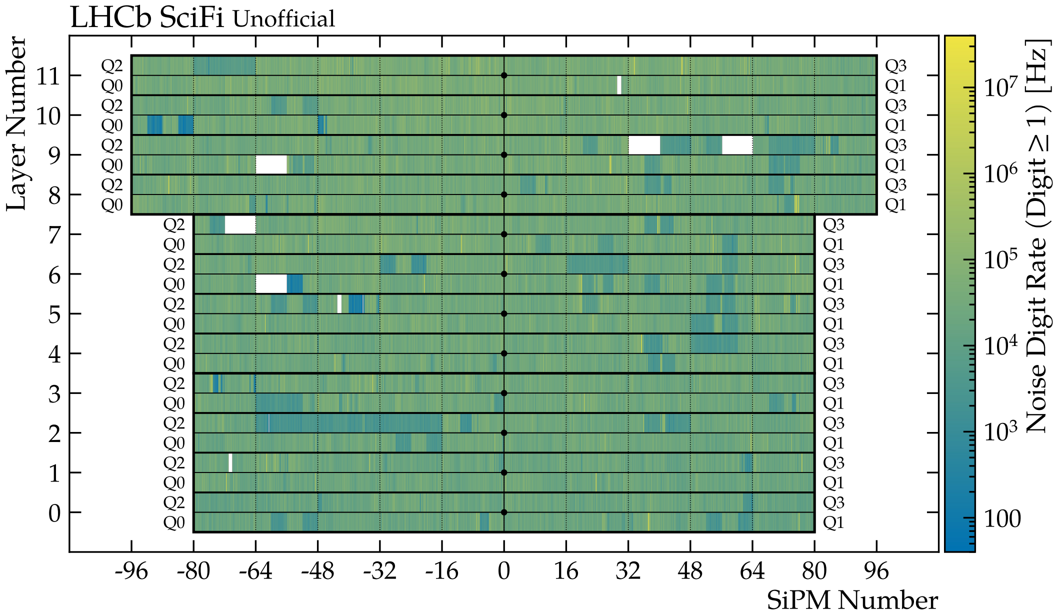

Figure 7.12 Overview of the noise rate of the raw PACIFIC output per channel for digits ≥ 1 using threshold settings $(1.5, 2.5, 3.5)$ with 68 % scaling.

it is vital to investigate such channels and adjust their thresholds accordingly to ensure optimal and uniform performance of the detector.

7.5 Performance Summary

The initial measurements of the hit detection efficiency, single hit resolution, and noise cluster rate of the SciFi Tracker have shown promising results. The hit detection efficiency and single hit resolution have been evaluated using data collected during the *PbPb* data-taking in 2023. Although the measured median hit detection efficiency of 98 % and the median hit resolution of 138 µm across the entire detector do not yet meet the design values, the measurements have shown that with accurately calibrated comparator thresholds, a good alignment of the detector in time relative to the bunch crossings, and a precise spatial alignment, a hit efficiency of over 99 % and a hit resolution of less than 100 µm can be achieved.
Comparing the performance for different threshold settings yields the expected behaviour: lower thresholds result in higher hit efficiency and higher noise cluster rates, whereas higher thresholds give lower efficiency and noise cluster rates. This is summarised in Figure 7.13, where the hit efficiency and noise cluster rates are plotted for all six tested threshold settings. Using 1.5 pe and 2.5 pe for the lowest and middle threshold, a hit efficiency of around 98 % or 97 % is measured for scaling factors of 68 % and 100 %, respectively. Using the highest tested threshold settings of $(2.5, 3.5, 4.5)$ pe efficiencies of 97 % and 93 %, respectively for scaling factors of 68 % and 100 %. The influence of the threshold settings on the single hit resolution is with about 1 µm

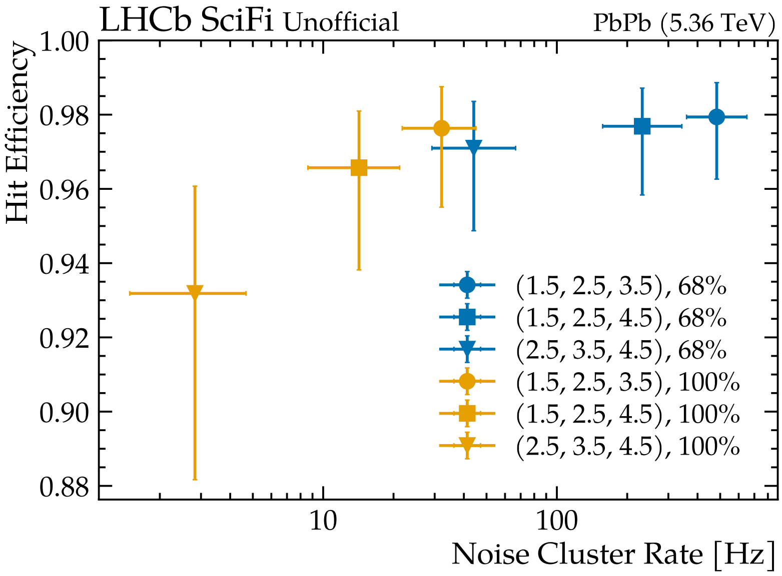

Figure 7.13 Hit detection efficiency versus noise cluster rate per channel for the six evaluated threshold settings. The markers show the median values and the error bars the 25th and 75th percentiles of the corresponding log-normal distributions.

negligible.

Even with the lowest threshold settings of (1.5, 2.5, 3.5) pe, the measured noise cluster rate per channel of 500 Hz is well below the limit of 16 kHz. However, it is important to note that the detector has not yet been irradiated significantly. The noise cluster rates cited here were measured at 1.55 fb^{-1}[108] of integrated luminosity, which only accounts for about 3 % of the total expected integrated luminosity of 50 fb^{-1}. Additionally, it was found that some channels have excessively high noise rates due to low threshold settings. This requires a careful examination of the threshold calibration for these channels.

It is clear that the measurements presented in this study are influenced by several systematic effects such as the alignment of the detector in space and the accuracy of the threshold calibration, among others. These effects should be further investigated to ensure that the detector performs optimally. However, it should be noted that the measurements have also shown that the detector achieves the desired performance with a hit detection efficiency of >99 % and a single hit resolution of <100 µm.

8 Outlook

The successful calibration of the PACIFIC comparator thresholds, as discussed in Chapter 6, has proven to be an essential step towards optimizing the detector. The initial performance measurements, presented in Chapter 7, have shown promising results in hit detection efficiency, single hit resolution, and noise rates. This improvement is reflected in the distribution of clusters throughout the detector, as shown in Figure 8.1 for one data-taking run from the beginning of 2024. Compared to the same distribution from the beginning of Run 3 (see Figure 5.1), the current measured cluster occupancy is considerably smoother, mainly due to the calibration of the comparator thresholds. Only one SiPM with an excessive number of clusters is observed in layer 11; this SiPM was previously excluded and was thus not calibrated. Currently, only four out of 4096 SiPMs are excluded due to transmission issues between the frontend electronics and the DAQ system.

Moreover, since the start of Run 3, a better understanding of the time and spatial alignment of the detector has been achieved, resulting in a much-improved tracking performance. Figure 8.2a displays the number of SciFi hits per track for two runs, 255366 from the beginning of Run 3 in 2022, 289069 from early 2024, and simulation. In run 255366, the distribution peaks at a value of ten hits per track, and only a few tracks have 12 SciFi hits, but it is now completely shifted towards 12 hits per track in run 289069, almost comparable to simulation. This indicates that the detector is more efficient in detecting particle hits overall. The remaining inefficiency can be predominantly attributed to just three layers, as shown in Figure 8.2b, and overall the performance has improved and is now more uniform across all twelve layers compared to the beginning of Run 3.

Looking ahead to the remainder of the data taking in Run 3 of the LHC, it is expected that the SciFi Tracker will reach its design performance. This will be achieved by calibrating the remaining detector channels and improving the spatial alignment of the detector elements.

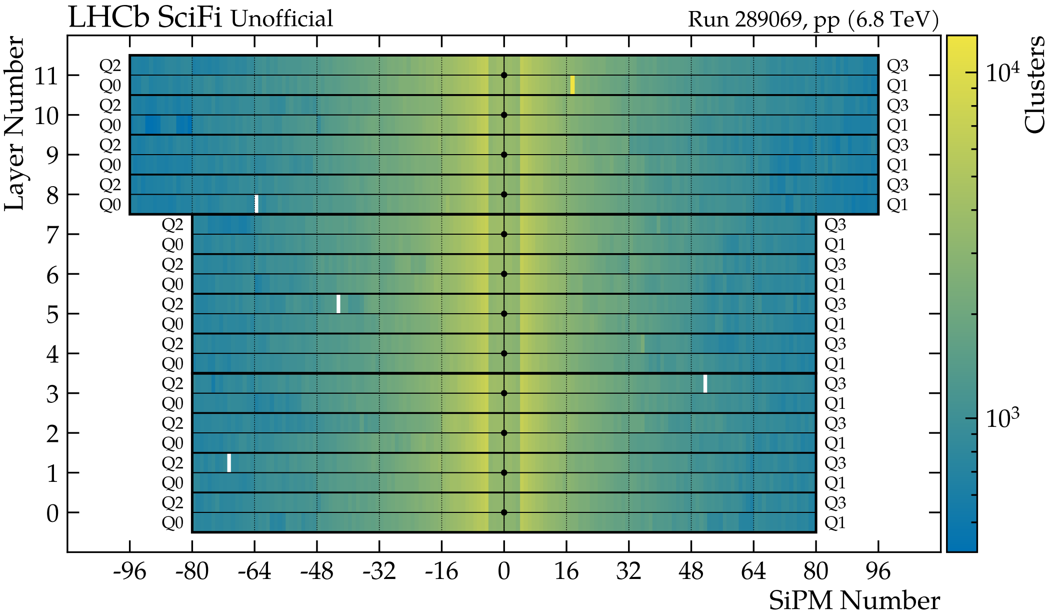

Figure 8.1 Overview of the number of reconstructed clusters per SiPM in the SciFi Tracker during 2024 data taking.

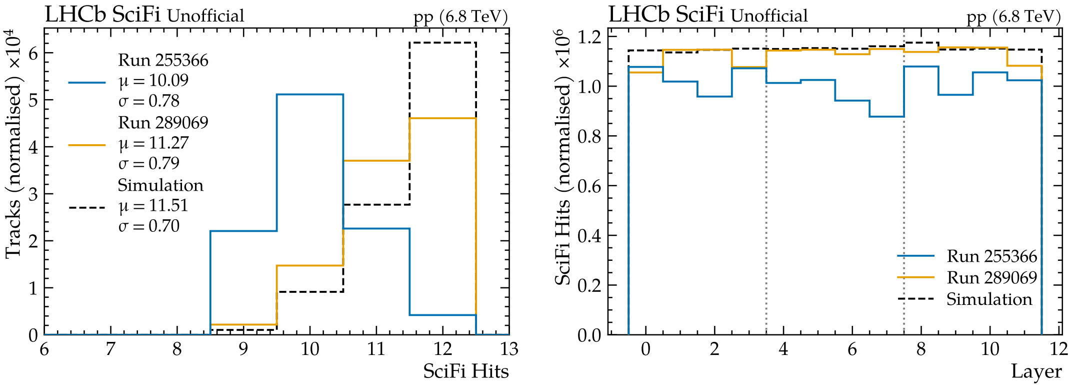

a

b

Figure 8.2 Distributions of the number of SciFi hits per track for data recorded during Run 3 in 2024 and from simulation. All distributions have been normalised to the same number of tracks. Left: distribution of the number of SciFi hits per track. Right: distribution of the number of SciFi hits per track for each layer.

9 | Conclusion

The work presented in this dissertation contributed to the assembly, commissioning, and operation of the LHCb SciFi Tracker and highlights the complexity and challenges that arise in commissioning and operating modern large-scale particle detectors in environments such as the Large Hadron Collider.

The SciFi Tracker uses 250 µm thin scintillating fibres arranged in six-layer fibre mats. The scintillation light is read out by silicon photomultiplier (SiPM) arrays, and the signals are processed and digitised by a custom frontend ASIC, called PACIFIC, at the LHC bunch crossing rate of 40 MHz. The PACIFIC shapes and integrates the SiPM signals every bunch period and digitises the integrated signal using three comparators with adjustable thresholds. Using the combined output of these comparators, an online zero-suppression and cluster-finding is performed by a dedicated FPGA before the clusters are shipped to the data acquisition (DAQ) system.

The SciFi Tracker is part of the more extensive upgrade of the entire LHCb detector for Run 3 of the LHC, which was completed in the spring of 2022. The upgrade enables the experiment to record a data set that is five times larger than the two previous runs combined. This is achieved by operating the detector at a higher instantaneous luminosity, reading out the entire detector at the rate of 40 MHz, and performing a fully software-based event selection. In order to cope with the increase in occupancy, the main tracking stations, previously based on gas straw tubes, were replaced by the SciFi Tracker.

Assembly of the twelve C-Frames that comprise the three stations of the SciFi Tracker progressed from 2019 to 2022 and was accompanied by an extensive commissioning of the frontend electronics. While the individual components that constitute the SciFi Tracker were all previously tested and their performance verified in test beam campaigns, the system was, for the first time, operated on the scale of C-Frames in the scope of this dissertation. After the assembly of each C-Frame, a series of tests was carried out to ensure the correct functionality of the up to 24 Readout Boxes (ROBs), corresponding to 98 304 channels, and their integration into the LHCb data acquisition and experimental control systems. A key challenge was establishing the readout of the detector at a frequency of 40 MHz with sufficiently low transmission errors. This required careful tuning of several clocks along the data path for hundreds of data links per C-Frame, resulting in a stable transmission with an overall bit error ratio of 5.1×10^{-15}. The final test in the sequence consisted of probing the entire signal chain from SiPMs to the DAQ by performing threshold scans of the PACIFIC comparator thresholds under pulsed illumination from the light injection system (LIS). Only three out of more than 480 000 tested channels were found to be broken. Furthermore, these measurements allowed for the first calibration of the comparator thresholds.

After the successful installation of the SciFi Tracker for the beginning of Run 3 in spring 2022, the detector commissioning and integration into the LHCb control and DAQ systems continued. With the preliminary calibration of the comparator thresholds determined during the assembly of the C-frames, the first data from proton-proton collisions was recorded, allowing for the successful reconstruction of particle tracks, albeit not yet with the highest possible efficiency.

Calibrating the comparator thresholds is central to ensure a high tracking performance of the SciFi Tracker. Since the PACIFIC does not feature a full ADC but digitises the SiPM signals using three comparators per channel, no complete amplitude information is available to reconstruct the clusters and suppress noise. The comparator thresholds, therefore, must be calibrated to determine the equivalent photoelectron amplitudes. A method was developed and presented that allows for calibrating the thresholds from threshold scans recorded under pulsed illumination of the LIS. To ensure that the PACIFIC maximally integrates the SiPM signals, the light pulses have to be injected at the correct time. This was achieved for 97 % of the detector and ultimately allowed for the first calibration of the SciFi Tracker. However, due to an unresolved issue with the LIS so far, the delivered light intensity was too low, such that the calibration of about 17 % of all channels proved to be unreliable for signals with 4 pe and higher.

Based on the calibration of the comparator thresholds presented in this dissertation, six different threshold configurations were generated for the entire detector to evaluate the performance of the SciFi Tracker. The hit detection efficiency and single hit resolution were measured using particle tracks recorded during 2023 *PbPb* data taking. The measurements have shown that the detector achieves the design requirements of a hit detection efficiency of >99 % and a single hit resolution of <100 µm while maintaining a sufficiently low noise cluster rate of 500 Hz per channel. However, it was also shown that the measurements are affected by several systematic effects, such as the spatial alignment of the detector and the threshold settings, and that further studies are necessary to achieve a uniform performance across the whole detector.

The calibration of the PACIFIC thresholds was only the first step in optimising the detector performance. With the ongoing effort to optimise the detector, the SciFi Tracker is expected to reach or even surpass the required performance for the remainder of the data-taking period.

Acknowledgements

I want to thank everyone who supported me during my PhD at the Physikalisches Institut Heidelberg and CERN. I am deeply grateful to my supervisor, Ulrich Uwer, for his invaluable guidance throughout my Master's and PhD. I would also like to thank Stephanie Hansmann-Menzemer, with whom I began my academic journey. Additionally, I owe much of my knowledge and experience to Snow and Blake, whose mentorship has been instrumental in my development.

I extend my gratitude to Hans-Christian Schultz-Coulon for serving as my second referee and to Jörg Jäckel and Peter Fischer for being part of my thesis committee.

I want to thank my colleagues from the SciFi team, especially Daniel, Preema, Elisabeth, and Ulisses. Working with them was a pleasure, and their presence truly made the work enjoyable. In addition, I would like to thank Laurent, Nathan, Marian, and Sascha, with whom I have spent great times both on and off the bike.

Furthermore, I would like to thank the entire Heidelberg LHCb group, including all the former and current members I have met since joining, who have contributed to a great work environment.

Finally, my gratitude goes to my friends and family. Without their support, this thesis would not have been possible. In addition to my parents, Eva and Eugen, I want to thank David for our long friendship throughout our studies, and without whom I probably would not have ended up in Heidelberg.

A Appendix

A.1 Gain Distributions

The following figures show the distributions of the five gain parameters $Gain_0$ to $Gain_4$ for as determined from light injection threshold scans during the detector assembly and with the complete detector. The distributions are normalised to the number of entries, since only one of the three comparators was scanned during the detector assembly while for the complete detector all three were scanned.

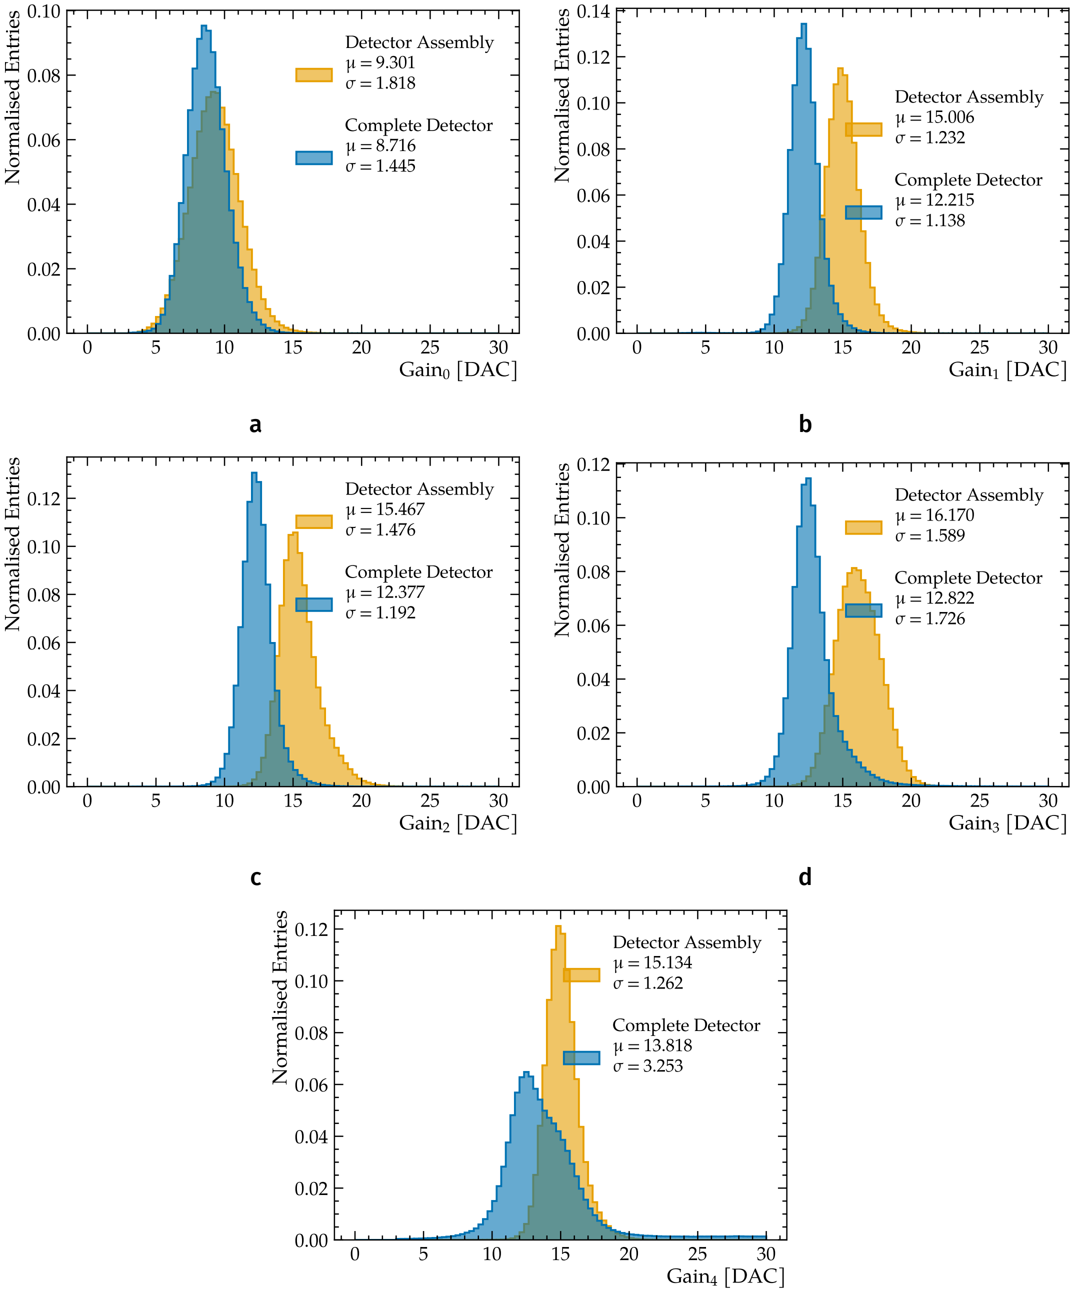

a

b

c

d

e

Figure A.1 Histograms for the five gain parameters Gain$_0$ to Gain$_4$ for the detector assembly and the complete detector.

A.2 Threshold Differences

Threshold scans are recorded separately for the two integrators and the three comparators of each PACIFIC channel. The following figures show the differences in the thresholds determined for the two integrators *Int0* and *Int1* and the differences between the three comparators Vth_1, Vth_2, and Vth_3.

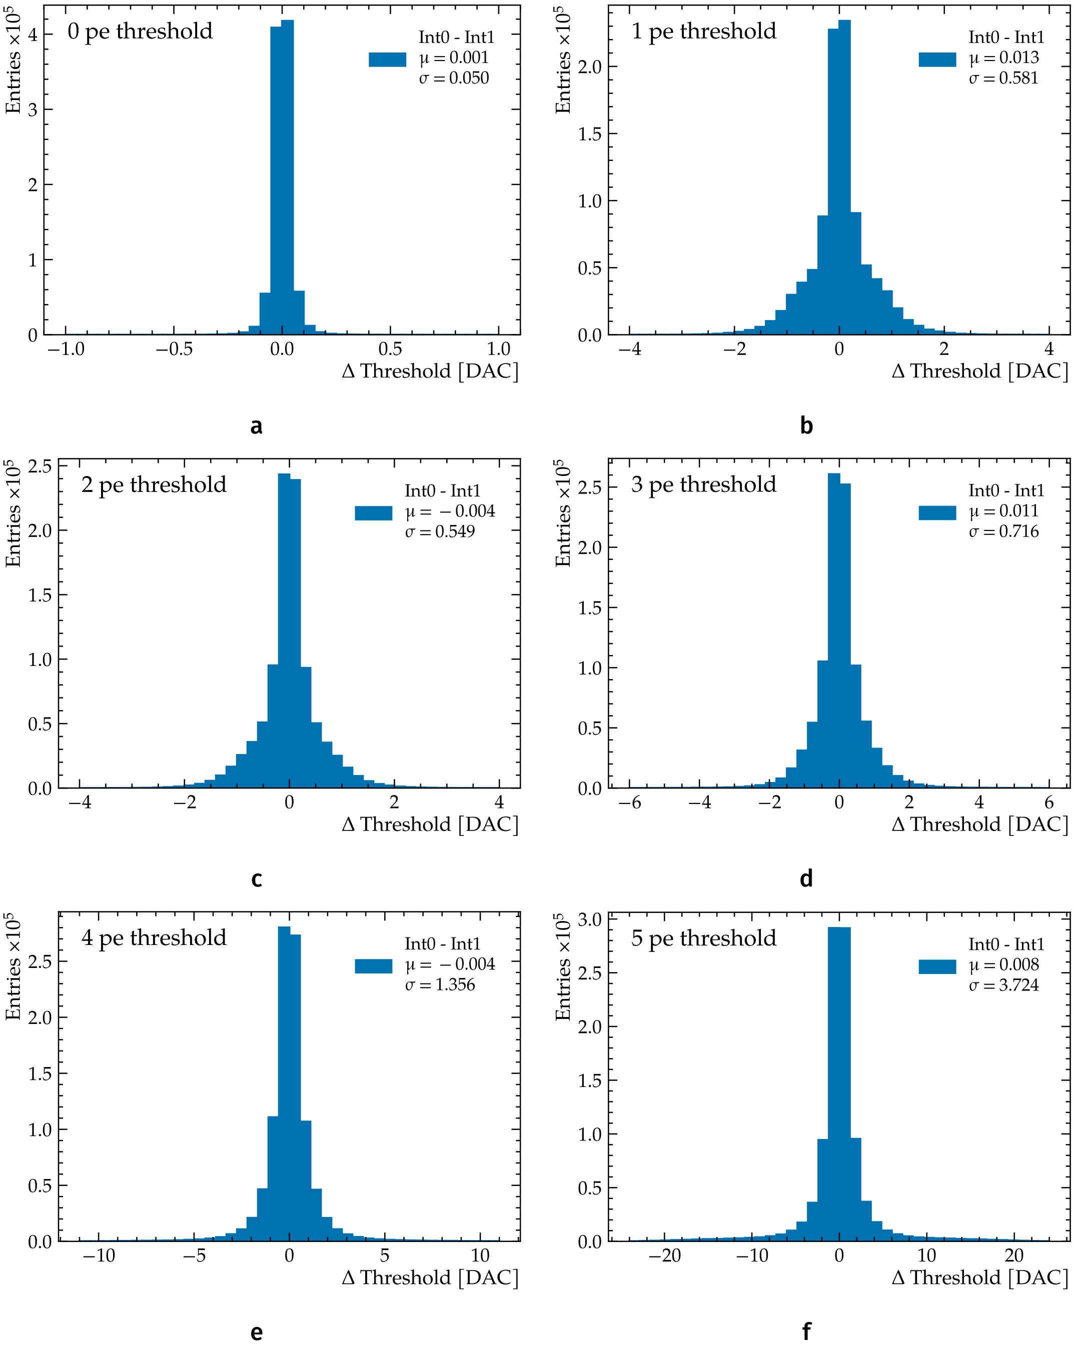
a b c d e f

Figure A.2 Distributions of the differences in the determined thresholds between the two integrators of each channel.

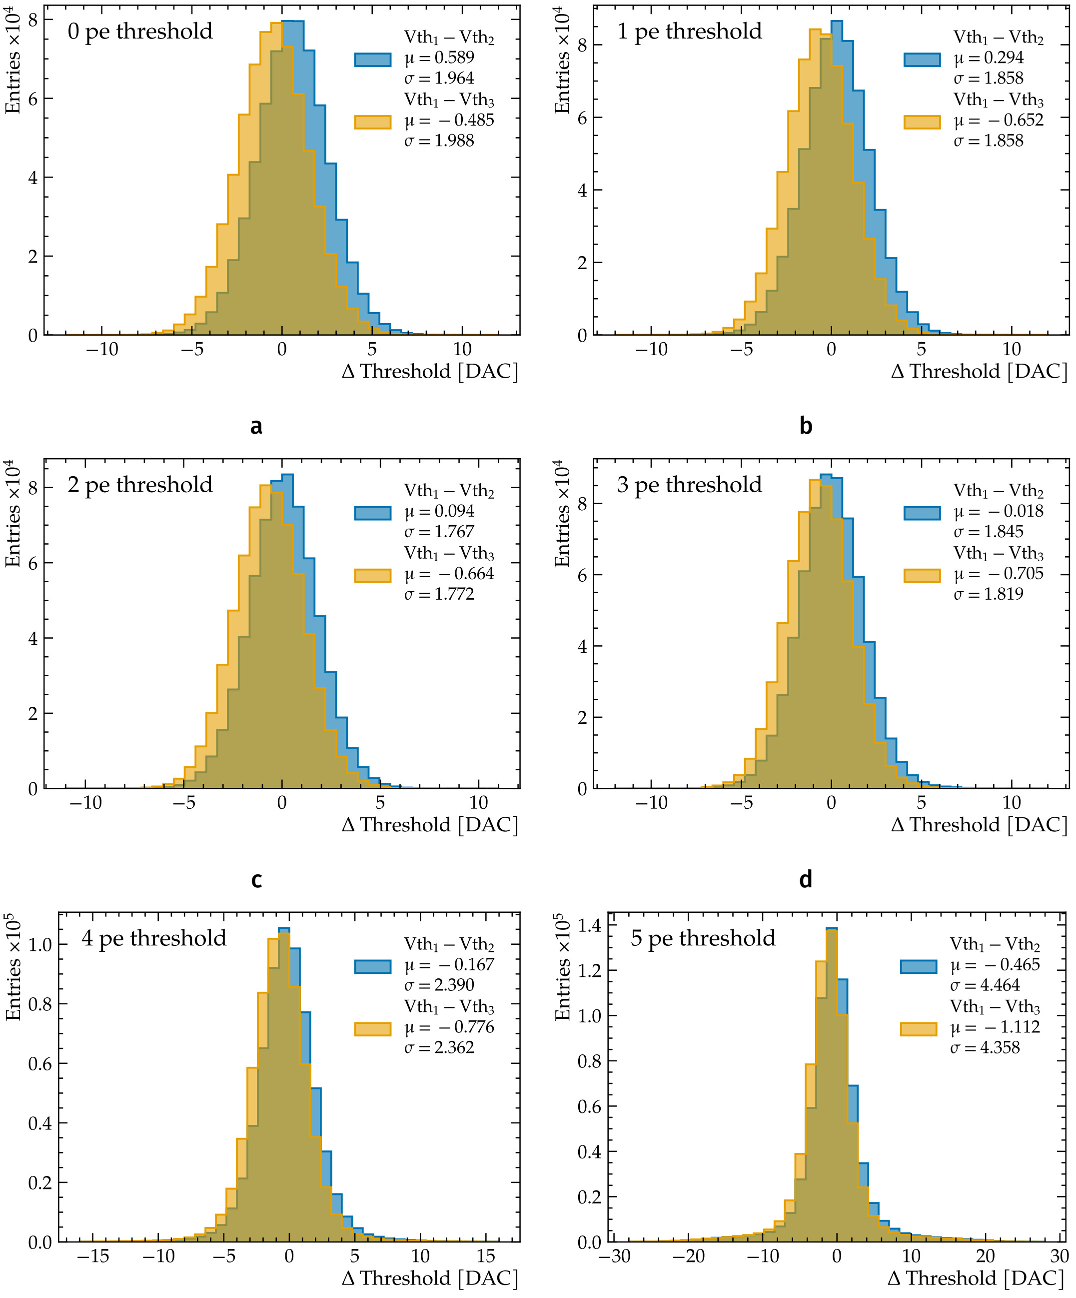

a

b

c

d

e

f

Figure A.3 Distributions of the differences in the determined thresholds between the three comparators of each channel. Comparator *Vth₁* is used as a reference.

A.3 Hit Detection Efficiency

The following section contains a summary of the event and track selection criteria for the hit detection efficiency and all log-normal fits used to determine the median efficiency and the 25th and 75th percentiles.

Table A.1 Event and track selection criteria for the hit detection efficiency measurements.

Variable	Selection
Global event cut	< 10000
Number of tracks	$\geq 2, < 100$
Chi square per d.o.f.	< 2
Momentum	$> 2\,\text{GeV}$
Transverse momentum	$> 400\,\text{MeV}$
Number of hits on track	$= 11$

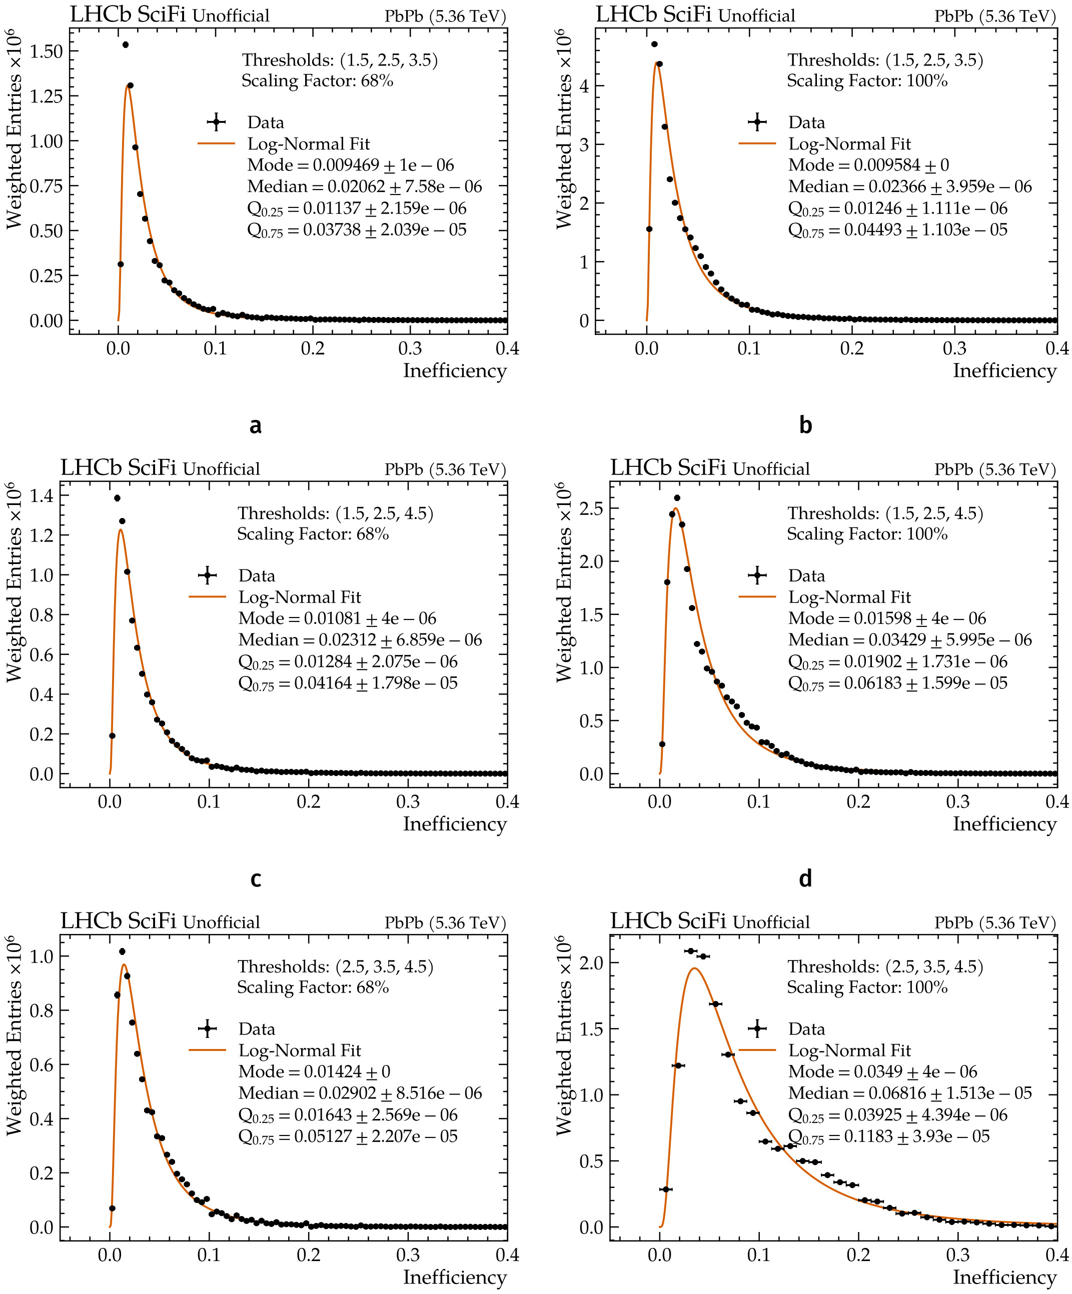

a

b

c

d

e

f

Figure A.4 Hit inefficiency distributions and log-normal fits for the six different threshold configurations.

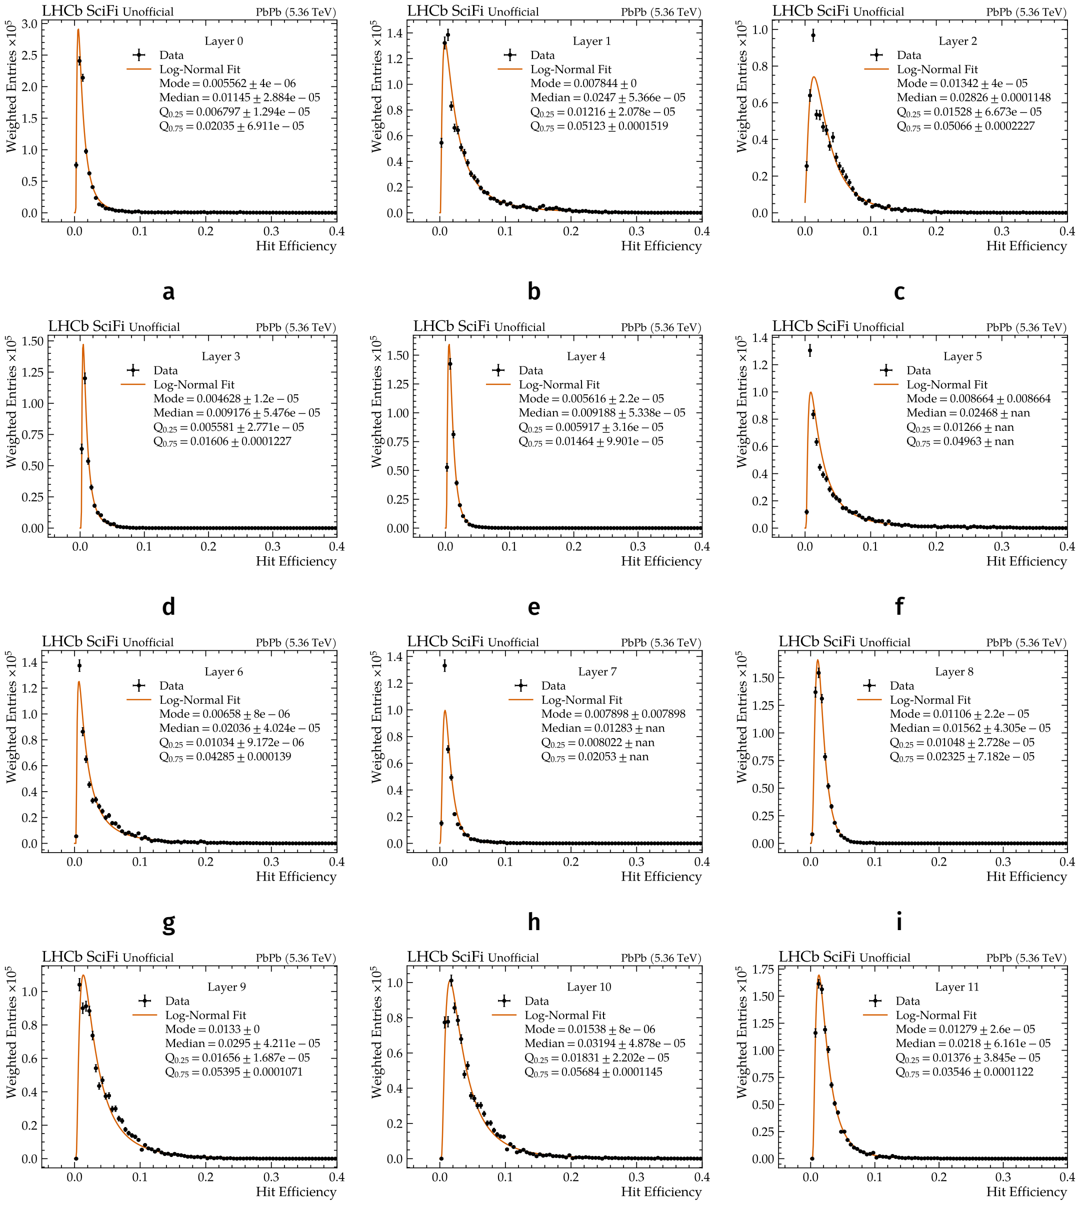

a b c

d e f

g h i

j k l

Figure A.5 Hit inefficiency distributions and log-normal fits for each of the 12 SciFi Tracker layers. The data was recorded using threshold settings $(1.5, 2.5, 3.5)$ pe with 68 % scaling.

A.4 Single Hit Resolution

The following section contains a summary of the event and track selection criteria for the single hit resolution measurements and all Gaussian fits used to determine the mean and width of the residuals.

Table A.2 Event and track selection criteria for the single hit resolution measurements.

Variable	Selection
Global event cut	< 10000
Chi square per d.o.f.	< 2
Momentum	$> 10\,\mathrm{GeV}$
Number of hits on track	$= 12$

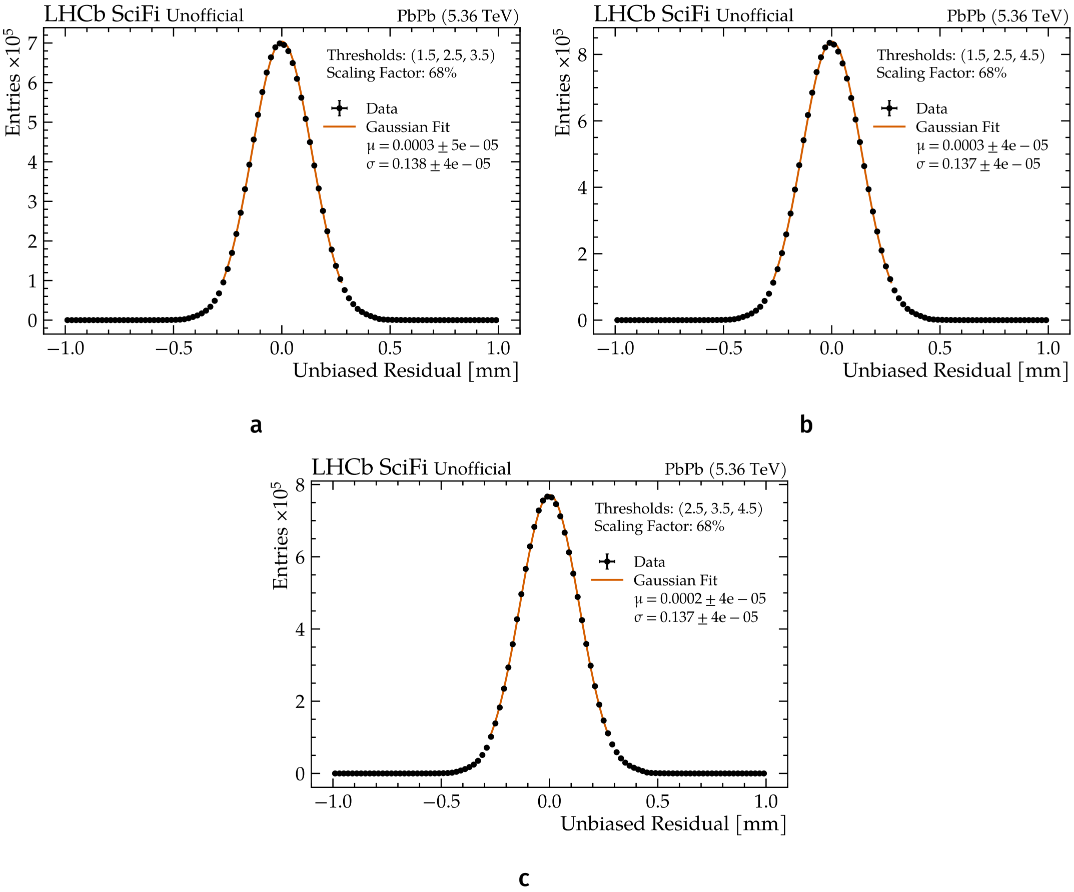

Figure A.6 Unbiased hit residual distributions and Gaussian fits for three different threshold configurations.

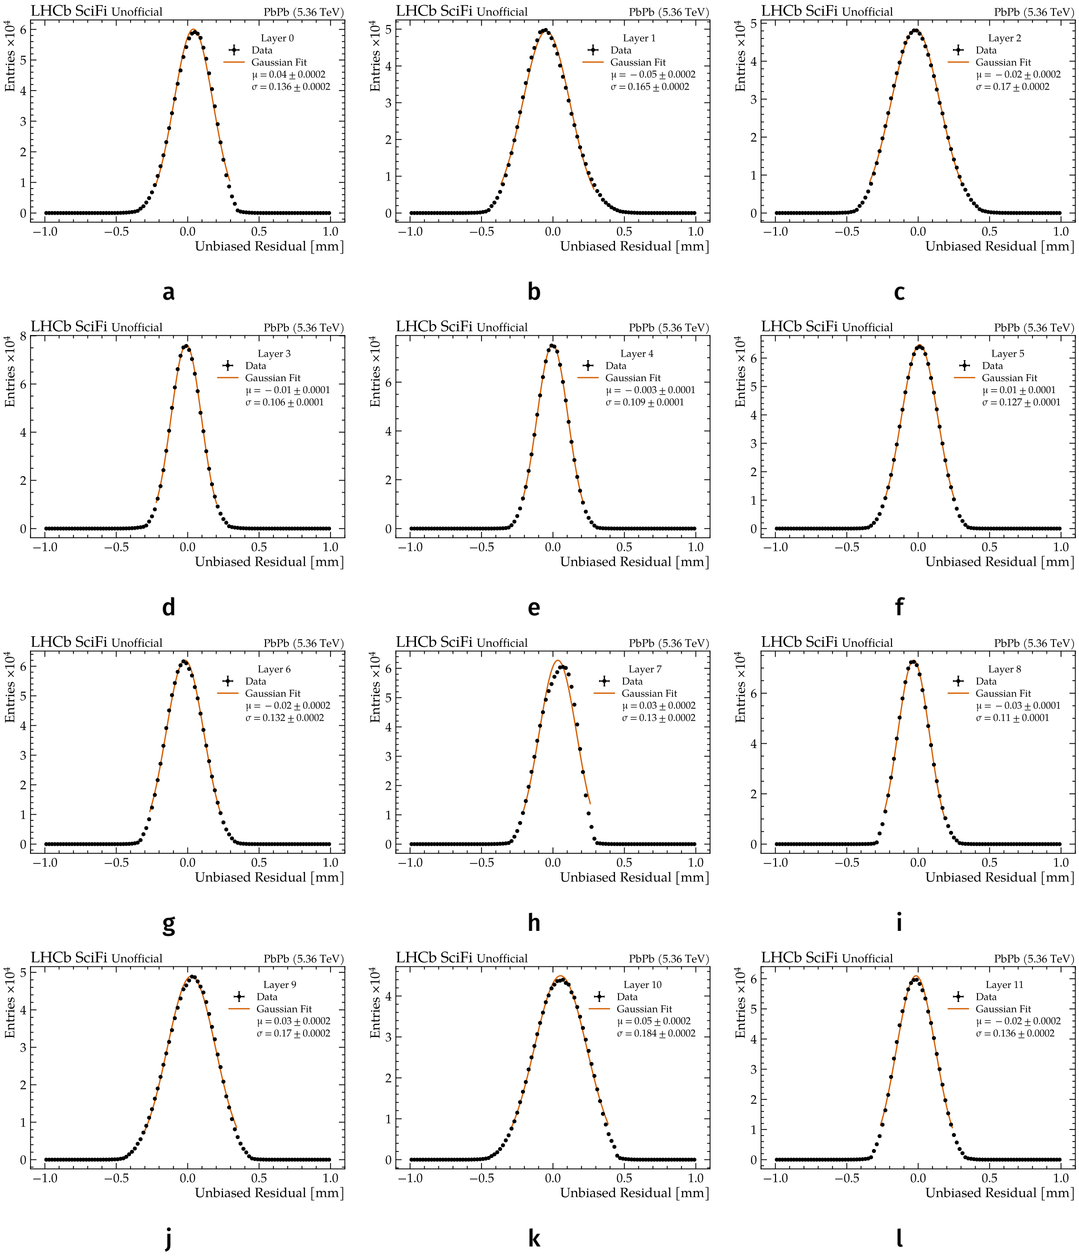

a b c d e f g h i j k l

Figure A.7 Unbiased hit residual distributions and Gaussian fits for each of the 12 SciFi Tracker layers. The data was recorded using threshold settings $(1.5, 2.5, 3.5)$ pe with 68 % scaling.

A.5 Noise Cluster Rates

In this section the distributions of the measured noise cluster rates per channel are shown including the log-normal that are used to determine the median and the 25th and 75th percentiles.

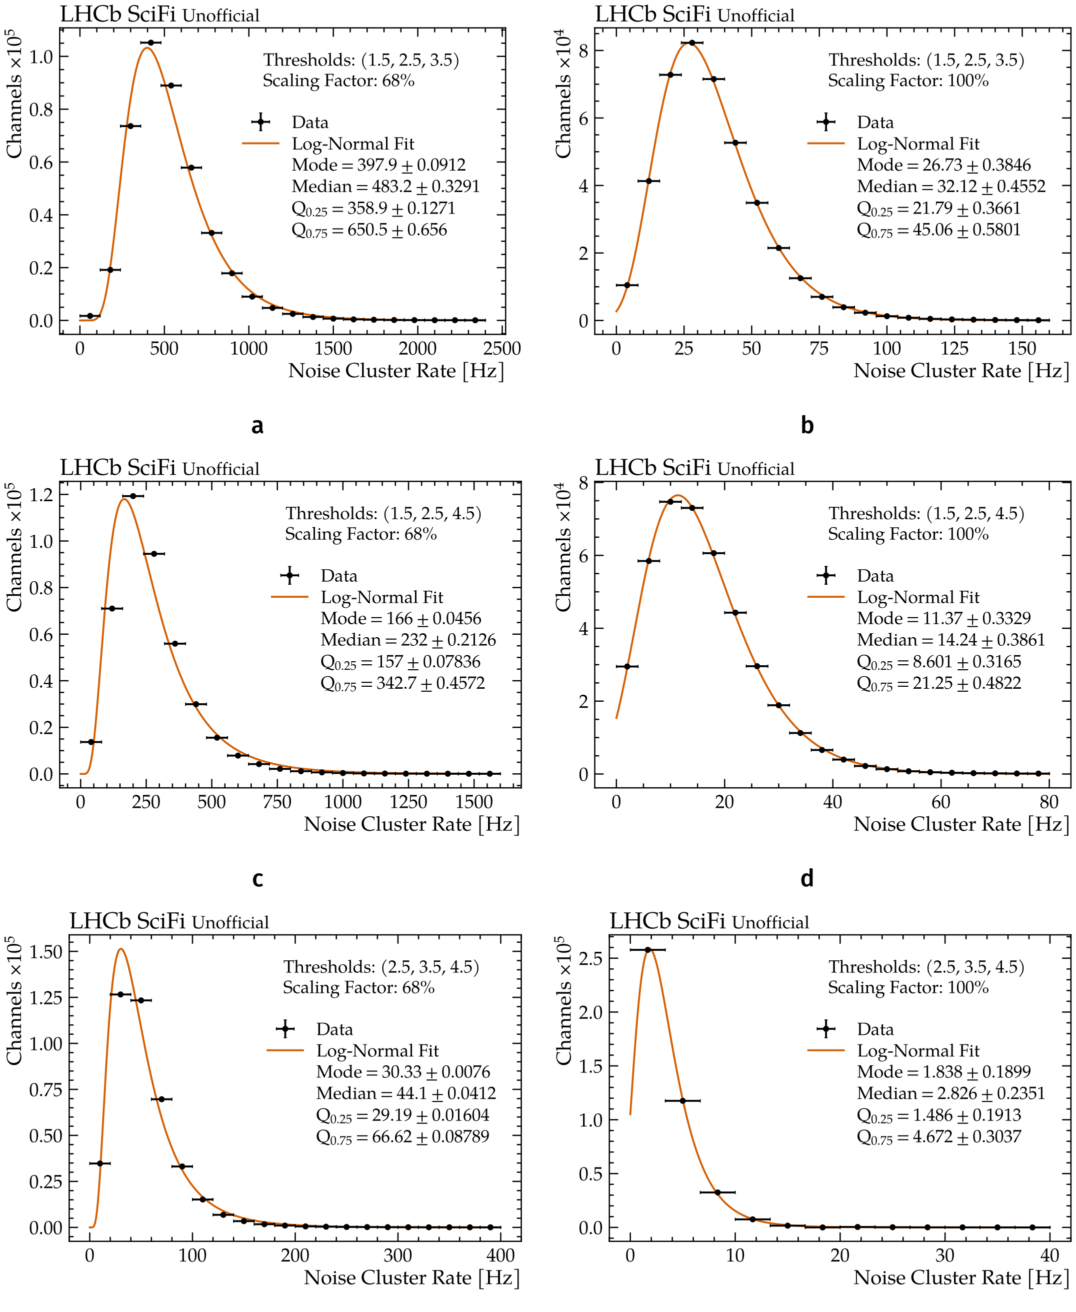

a b
c d
e f

Figure A.8 Noise cluster rate histograms and log-normal fits for the six tested threshold settings.

A.6 Noise Digit Rates

In this section the distributions of the measured noise digit rates per channel are shown including the log-normal that are used to determine the median and the 25th and 75th percentiles. The distributions for digits ≥ 1, ≥ 2, and ≥ 3 are shown separately.

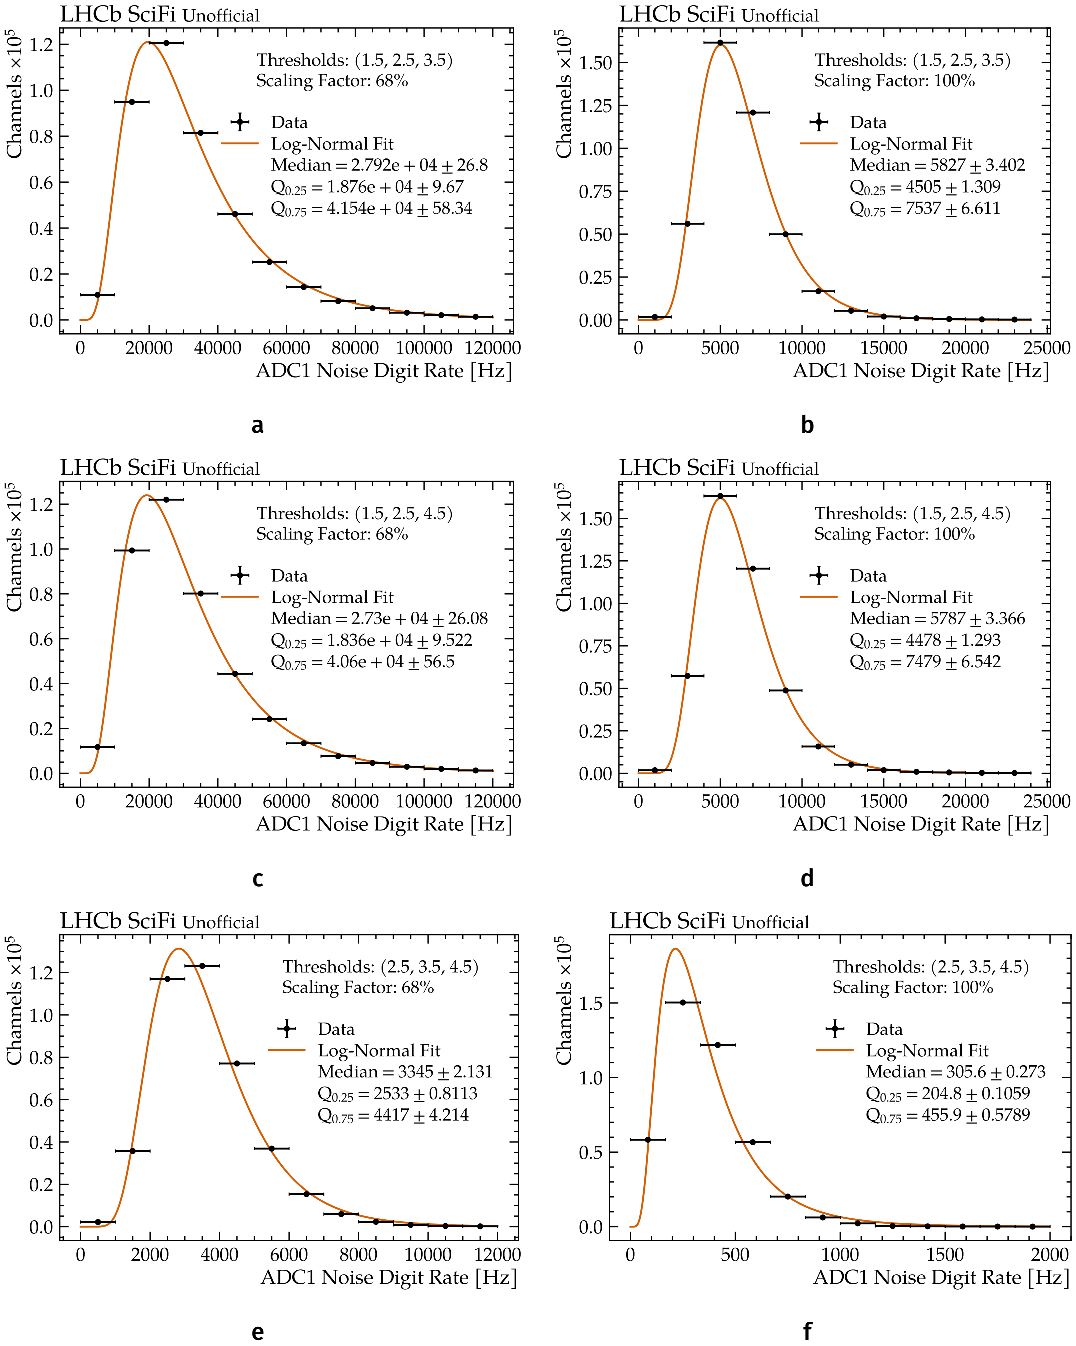

a b c d e f

Figure A.9 Noise digit rate distributions and log-normal fits for digits ≥ 1 for the six tested threshold settings.

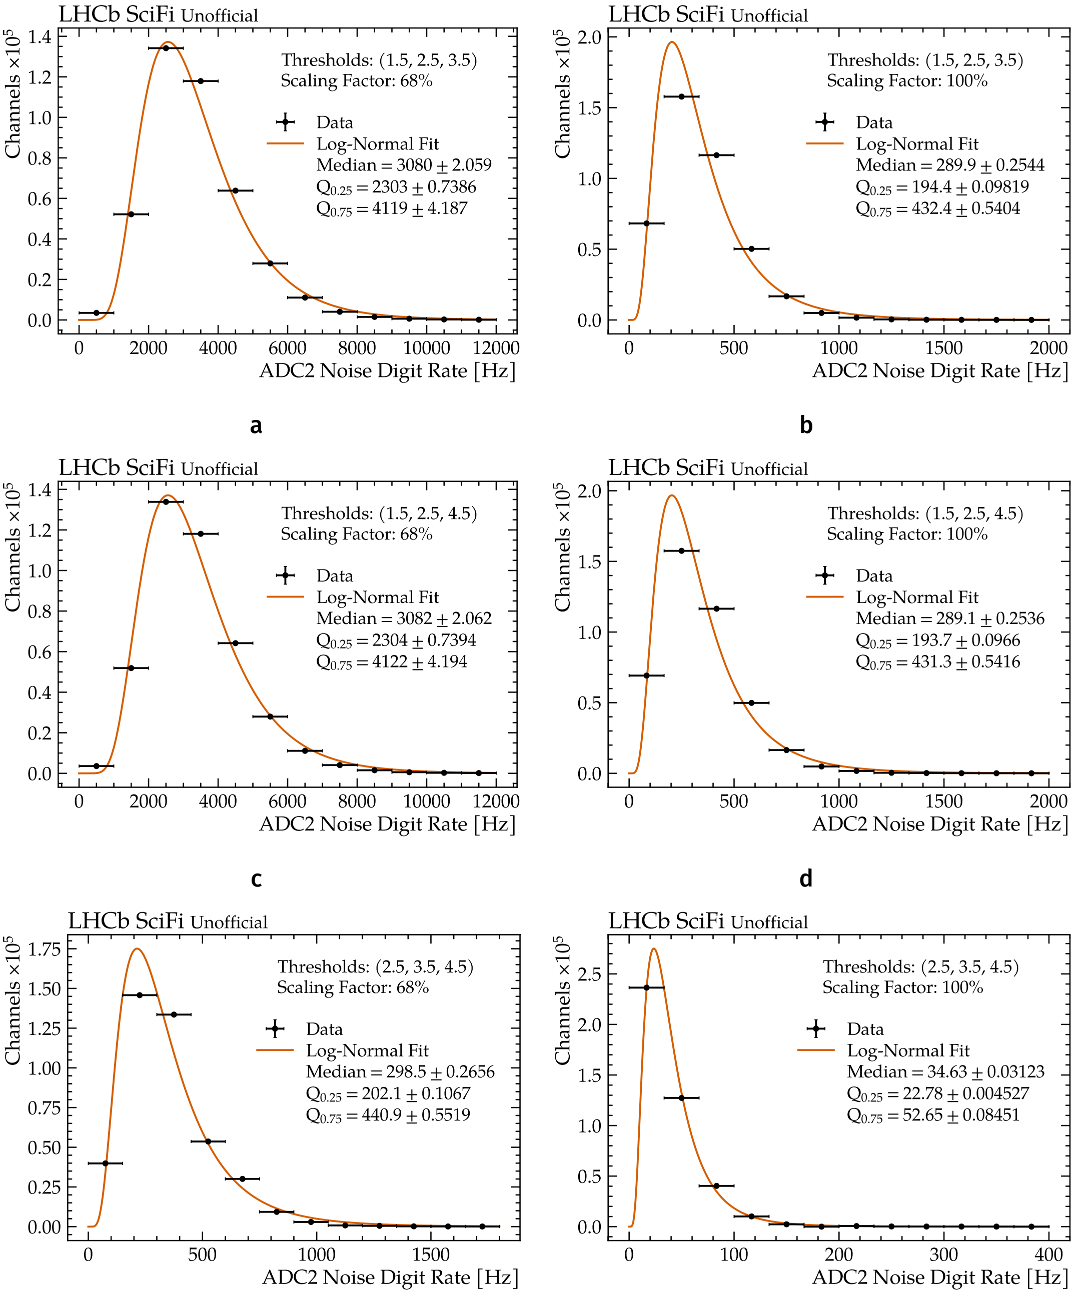

a

b

c

d

e

f

Figure A.10 Noise digit rate distributions and log-normal fits for digits ≥ 2 for the six tested threshold settings.

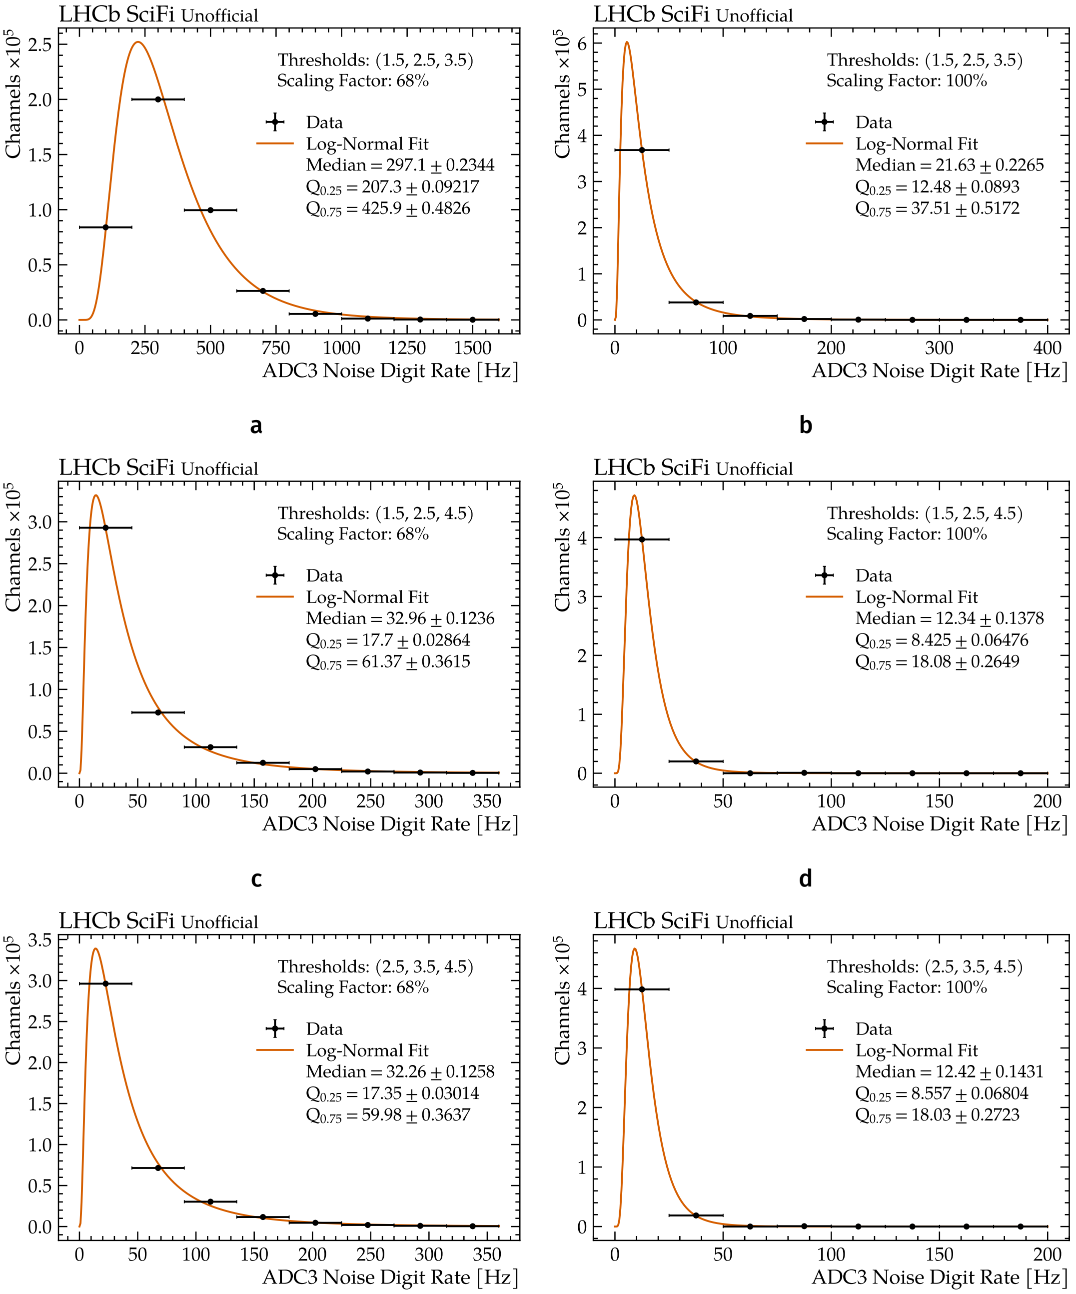

a

b

c

d

e

f

Figure A.11 Noise digit rate distributions and log-normal fits for digits ≥ 3 for the six tested threshold settings.

Bibliography

[1] Sheldon L. Glashow. 'Partial-Symmetries of Weak Interactions'. In: *Nuclear Physics* 22.4 (Feb. 1961), pp. 579–588.

[2] A. Salam and J.C. Ward. 'Electromagnetic and Weak Interactions'. In: *Physics Letters* 13.2 (Nov. 1964), pp. 168–171.

[3] Steven Weinberg. 'A Model of Leptons'. In: *Physical Review Letters* 19.21 (20th Nov. 1967), pp. 1264–1266.

[4] F. Englert and R. Brout. 'Broken Symmetry and the Mass of Gauge Vector Mesons'. In: *Physical Review Letters* 13.9 (31st Aug. 1964), pp. 321–323.

[5] Peter W. Higgs. 'Broken Symmetries and the Masses of Gauge Bosons'. In: *Physical Review Letters* 13.16 (19th Oct. 1964), pp. 508–509.

[6] H. Fritzsch, M. Gell-Mann and H. Leutwyler. 'Advantages of the Color Octet Gluon Picture'. In: *Physics Letters B* 47.4 (Nov. 1973), pp. 365–368.

[7] William Crookes. 'Certain Properties of the Emanations of Radium'. In: *Chemical News* 87 (1903), p. 241.

[8] M. Blau and H. Wambacher. 'Disintegration Processes by Cosmic Rays with the Simultaneous Emission of Several Heavy Particles'. In: *Nature* 140.3544 (Oct. 1937), pp. 585–585.

[9] Charles Thomson Rees Wilson. 'On an Expansion Apparatus for Making Visible the Tracks of Ionising Particles in Gases and Some Results Obtained by Its Use'. In: *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* 87.595 (19th Sept. 1912), pp. 277–292.

[10] Donald A. Glaser. 'Some Effects of Ionizing Radiation on the Formation of Bubbles in Liquids'. In: *Physical Review* 87.4 (15th Aug. 1952), pp. 665–665.

[11] Carl D. Anderson. 'The Positive Electron'. In: *Physical Review* 43.6 (15th Mar. 1933), pp. 491–494.

[12] C. M. G. Lattes et al. 'Processes Involving Charged Mesons'. In: *Nature* 159.4047 (May 1947), pp. 694–697.

[13] Ernest Rutherford and Hans Geiger. 'An Electrical Method of Counting the Number of α-Particles from Radio-Active Substances'. In: *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* 81.546 (27th Aug. 1908), pp. 141–161.

[14] H. Iams and B. Salzberg. 'The Secondary Emission Phototube'. In: *Proceedings of the IRE* 23.1 (Jan. 1935), pp. 55–64.

[15] G. Charpak et al. 'The Use of Multiwire Proportional Counters to Select and Localize Charged Particles'. In: *Nuclear Instruments and Methods* 62.3 (July 1968), pp. 262–268.

[16] B. Hyams et al. 'A Silicon Counter Telescope to Study Short-Lived Particles in High-Energy Hadronic Interactions'. In: *Nuclear Instruments and Methods in Physics Research* 205.1-2 (Jan. 1983), pp. 99–105.

[17] D. P. Barber et al. 'Discovery of Three-Jet Events and a Test of Quantum Chromodynamics at PETRA'. In: *Physical Review Letters* 43.12 (17th Sept. 1979), pp. 830–833.

[18] Ch. Berger et al. 'Evidence for Gluon Bremsstrahlung in E+e- Annihilations at High Energies'. In: *Physics Letters B* 86.3-4 (Oct. 1979), pp. 418–425.

[19] R. Brandelik et al. 'Evidence for a Spin-1 Gluon in Three-Jet Events'. In: *Physics Letters B* 97.3-4 (Dec. 1980), pp. 453–458.

[20] G Sajot et al. 'Experimental Observation of Isolated Large Transverse Energy Electrons with Associated Missing Energy at $\sqrt{s} = 540\,\text{GeV}$'. In: (1983).

[21] P. Bagnaia et al. 'Evidence for Z0→e+e- at the CERN p Collider'. In: *Physics Letters B* 129.1-2 (Sept. 1983), pp. 130–140.

[22] G. Arnison et al. 'Experimental Observation of Lepton Pairs of Invariant Mass around 95 GeV/C2 at the CERN SPS Collider'. In: *Physics Letters B* 126.5 (July 1983), pp. 398–410.

[23] Shanzhen Chen et al. 'Heavy Flavour Physics and CP Violation at LHCb: A Ten-Year Review'. In: *Frontiers of Physics* 18.44601 (Aug. 2023).

[24] Bo Fang. 'Recent LHCb Results on Heavy Hadron Spectroscopy'. In: *Nuclear and Particle Physics Proceedings* (Nov. 2022), pp. 66–72.

[25] Christoph Michael Langenbruch. 'Rare Decays. Blois 2023: 34th Rencontres de Blois on "Particle Physics and Cosmology"'. 2023. URL: `https://cds.cern.ch/record/2859123` (visited on 09/04/2024).

[26] LHCb Collaboration. *Letter of Intent for the LHCb Upgrade*. Geneva: CERN, 2011. URL: `https://cds.cern.ch/record/1333091`.

[27] LHCb collaboration et al. 'Implications of LHCb Measurements and Future Prospects'. In: *The European Physical Journal C* 73.4 (Apr. 2013), p. 2373.

[28] The LHCb Collaboration et al. 'The LHCb Detector at the LHC'. In: *Journal of Instrumentation* 3.08 (Aug. 2008), S08005.

[29] LHCb collaboration et al. 'LHCb Detector Performance'. In: *International Journal of Modern Physics A* 30.07 (10th Mar. 2015), p. 80.

[30] Lyndon Evans and Philip Bryant. 'LHC Machine'. In: *Journal of Instrumentation* 3.08 (Aug. 2008), S08001.

[31] Stephane Fartoukh et al. *LHC Configuration and Operational Scenario for Run 3*. 2021. URL: `https://cds.cern.ch/record/2790409` (visited on 25/04/2024). preprint.

[32] *LHC Long Term Schedule*. URL: `https://lhc-commissioning.web.cern.ch/schedule/LHC-long-term.htm` (visited on 01/02/2024).

[33] F. Follin and D. Jacquet. ‘Implementation and Experience with Luminosity Levelling with Offset Beam’. In: (2014). URL: https://cds.cern.ch/record/1955354.

[34] LHCb collaboration et al. *The LHCb Upgrade I*. 17th May 2023. URL: http://arxiv.org/abs/2305.10515 (visited on 20/06/2023). preprint.

[35] Christoph Hasse. *Simple Sketch of the LHC*. 27th Nov. 2023. URL: https://github.com/hassec/LHC_Sketch (visited on 09/04/2024).

[36] Christian Elsasser. *$\bar{b}b$ Production Angle Plots*. URL: https://lhcb.web.cern.ch/speakersbureau/html/bb_productionangles.html (visited on 09/04/2024).

[37] LHCb Collaboration et al. ‘LHCb Magnet: Technical Design Report’. In: Technical Design Report. LHCb (2000). URL: https://cds.cern.ch/record/424338 (visited on 09/04/2024).

[38] LHCb Collaboration. ‘LHCb VELO Upgrade Technical Design Report’. In: CERN-LHCC-2013-021 (4th Nov. 2013). URL: https://cds.cern.ch/record/1624070 (visited on 10/10/2022).

[39] T. Poikela et al. ‘VeloPix: The Pixel ASIC for the LHCb Upgrade’. In: *Journal of Instrumentation* 10.01 (30th Jan. 2015), p. C01057.

[40] LHCb Collaboration. ‘LHCb Tracker Upgrade Technical Design Report’. In: CERN-LHCC-2014-001 (21st Feb. 2014). URL: https://cds.cern.ch/record/1647400 (visited on 10/10/2022).

[41] Carlos Abellan Beteta et al. ‘The SALT—Readout ASIC for Silicon Strip Sensors of Upstream Tracker in the Upgraded LHCb Experiment’. In: *Sensors* 22.1 (24th Dec. 2021), p. 107.

[42] LHCb Collaboration. ‘LHCb PID Upgrade Technical Design Report’. In: CERN-LHCC-2013-022 (4th Nov. 2013). URL: https://cds.cern.ch/record/1624074 (visited on 10/10/2022).

[43] M. Baszczyk et al. ‘CLARO: An ASIC for High Rate Single Photon Counting with Multi-Anode Photomultipliers’. In: *Journal of Instrumentation* 12.08 (29th Aug. 2017), P08019.

[44] R. Aaij et al. ‘A Comprehensive Real-Time Analysis Model at the LHCb Experiment’. In: *Journal of Instrumentation* 14.04 (15th Apr. 2019), P04006.

[45] R. Aaij et al. ‘Tesla: An Application for Real-Time Data Analysis in High Energy Physics’. In: *Computer Physics Communications* 208 (Nov. 2016), pp. 35–42.

[46] S. Borghi. ‘Novel Real-Time Alignment and Calibration of the LHCb Detector and Its Performance’. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 845 (Feb. 2017), pp. 560–564.

[47] J.P. Cachemiche et al. ‘The PCIe-based Readout System for the LHCb Experiment’. In: *Journal of Instrumentation* 11.02 (18th Feb. 2016), P02013.

[48] F Alessio et al. ‘The LHCb Run Control’. In: *Journal of Physics: Conference Series* 219.2 (1st Apr. 2010).

[49] D.R. Myers. 'The LHC Experiments' Joint Controls Project, JCOP'. In: International Conference on Accelerator and Large Experimental Physics Control Systems. Trieste, Italy, 1999.

[50] R. C. Ruchti. 'The Use of Scintillating Fibers for Charged-Particle Tracking'. In: *Annual Review of Nuclear and Particle Science* 46.1 (Dec. 1996), pp. 281–319.

[51] P R Barbosa-Marinho et al. 'LHCb Inner Tracker: Technical Design Report'. In: Technical Design Report. LHCb (2002). URL: `https://cds.cern.ch/record/582793` (visited on 10/04/2024).

[52] P R Barbosa-Marinho et al. 'LHCb Outer Tracker: Technical Design Report'. In: Technical Design Report. LHCb (2001). URL: `https://cds.cern.ch/record/519146` (visited on 10/04/2024).

[53] LHCb Outer Tracker group et al. 'Performance of the LHCb Outer Tracker'. In: *Journal of Instrumentation* 9.01 (7th Jan. 2014), P01002.

[54] O. Borshchev et al. 'Development of a New Class of Scintillating Fibres with Very Short Decay Time and High Light Yield'. In: *Journal of Instrumentation* 12.05 (17th May 2017), P05013.

[55] Hermann Kolanoski and Norbert Wermes. *Particle Detectors: Fundamentals and Applications*. First edition. Oxford ; New York, NY: Oxford University Press, 2020. 927 pp.

[56] A.B.R. Cavalcante et al. 'Refining and Testing 12,000 Km of Scintillating Plastic Fibre for the LHCb SciFi Tracker'. In: *Journal of Instrumentation* 13.10 (19th Oct. 2018), P10025.

[57] Christian Joram et al. *LHCb Scintillating Fibre Tracker Engineering Design Review Report: Fibres, Mats and Modules*. 2015. URL: `https://cds.cern.ch/record/2004811`.

[58] P Buzhan et al. 'An Advanced Study of Silicon Photomultiplier'. In: *ICFA Instrumentation Bulletin* 23 (2001), pp. 28–41.

[59] P. Buzhan et al. 'Silicon Photomultiplier and Its Possible Applications'. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 504.1-3 (May 2003), pp. 48–52.

[60] Hamamatsu Photonics. *Si APD Technical Note*. URL: `https://www.hamamatsu.com/content/dam/hamamatsu-photonics/sites/documents/99_SALES_LIBRARY/ssd/si-apd_kapd9007e.pdf` (visited on 31/03/2024).

[61] Hamamatsu Photonics. *What Is an SiPM and How Does It Work?* URL: `https://hub.hamamatsu.com/us/en/technical-notes/mppc-sipms/what-is-an-SiPM-and-how-does-it-work.html` (visited on 16/02/2024).

[62] Olivier Göran Girard. 'Development of the Scintillating Fibre Tracker Technology for the LHCb Upgrade and the LHC Beam Profile Monitoring System'. Lausanne: EPFL, 2018. 140 pp.

[63] Axel Kuonen. 'Development and Characterisation of Silicon Photomultiplier Multichannel Arrays for the Readout of a Large Scale Scintillating Fibre Tracker'. Lausanne: EPFL, 2018. 163 pp.

[64] Olivier Girard et al. *Characterisation of Silicon Photomultipliers Based on Statistical Analysis of Pulse-Shape and Time Distributions*. 17th Aug. 2018. URL: `http://arxiv.org/abs/1808.05775` (visited on 11/10/2022). preprint.

[65] S Baron et al. 'The GBT Project'. In: *Proceedings of the Topical Workshop on Electronics for Particle Physics* (2009), pp. 342, 346. URL: `https://cds.cern.ch/record/1235836`.

[66] L Amaral et al. 'The Versatile Link, a Common Project for Super-LHC'. In: *Journal of Instrumentation* 4.12 (3rd Dec. 2009), P12003.

[67] *Read Out Box - SciFi General - Nikhef Projects on Embedded Systems (ET & CT)*. URL: `https://redmine.nikhef.nl/et/documents/167` (visited on 18/04/2024).

[68] Albert Comerma Montells and Jose Mazorra De Cos. *PACIFICr5 ASIC Data Sheet*. Geneva: CERN, 2021. URL: `https://cds.cern.ch/record/2766810`.

[69] Albert Comerma Montells. 'Development of a Multichannel Integrated Circuit for Silicon Photo-Multiplier Arrays Readout'. PhD thesis. Universitat de Barcelona, 9th Jan. 2014.

[70] Albert Comerma. 'PACIFIC: The Readout ASIC for the LHCb Scintillating Fibre Tracker'. In: *2017 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*. 2017 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC). Atlanta, GA: IEEE, Oct. 2017, pp. 1–3.

[71] Mirco Deckenhoff. *Signal Shape and Time of Light Propagation in Scintillating Fibre SCSF-78MJ from Kuraray*. Geneva, 2014. URL: `https://cds.cern.ch/record/1662537` (visited on 21/02/2024).

[72] José Mazorra de Cos et al. 'PACIFIC: SiPM readout ASIC for LHCb upgrade'. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 912 (2018). New Developments In Photodetection 2017, pp. 354–358.

[73] Walter Blum, Werner Riegler and Luigi Rolandi. *Particle Detection with Drift Chambers*. 2. ed. Particle Acceleration and Detection. Berlin Heidelberg: Springer, 2008. 448 pp.

[74] Sebastiana Giani et al. *Digitization of SiPM Signals for the LHCb Upgrade SciFi Tracker*. Geneva, 2014. URL: `https://cds.cern.ch/record/1662549` (visited on 25/02/2024).

[75] Wilco Vink. *LHCb SciFi, Clusterization Board Specifications and Manual*, p. 31. URL: `https://redmine.nikhef.nl/et/projects/clusterization_board`.

[76] A. Caratelli et al. 'The GBT-SCA, a Radiation Tolerant ASIC for Detector Control and Monitoring Applications in HEP Experiments'. In: *Journal of Instrumentation* 10.03 (20th Mar. 2015), p. C03034.

[77] Sevda Esen et al. *Clustering and Rawbank Decoding for the SciFi Detector*. Geneva: CERN, 2018. URL: `https://cds.cern.ch/record/2630154`.

[78] Giacomo Ripamonti et al. ‘Regulated Resonant Switched-Capacitor Point-of-Load Converter Architecture and Modeling’. In: *IEEE Transactions on Power Electronics* 36.4 (Apr. 2021), pp. 4815–4827.

[79] Giovanni Mazza et al. ‘High-Speed, Radiation-Tolerant Laser Drivers in 0.13µm CMOS Technology for HEP Applications’. In: *IEEE Transactions on Nuclear Science* 61.6 (Dec. 2014), pp. 3653–3659.

[80] Hans Verkooijen and Wilco Vink. *Master Board, Technical Manual*, p. 71. URL: `https://edms.cern.ch/document/2154019/1`.

[81] Waclaw Karpinski. ‘Light Injection System Performance and Production’. SciFi Front-End Electronics Production Readiness Review. 23rd Jan. 2018. URL: `https://indico.cern.ch/event/690897/` (visited on 11/04/2024).

[82] W-IE-NE-R Power Electronics GmbH. *MARATON Power Supply System Techical Manual*. URL: `https://file.wiener-d.com/documentation/MARATON/WIENER_MARATON_HighRad_Manual_A2.pdf`.

[83] CAEN S.p.A. *A1539B Power Supply Board*. URL: `https://www.caen.it/products/a1539b/` (visited on 11/04/2024).

[84] Christian Joram, C. Insa and A. Sciuccati. ‘Requirements for the LHCb SciFi Water Cooling System’. In: EDMS 1919838 v.2 (2018). URL: `https://edms.cern.ch/ui/#!master/navigator/document?D:100097508:100097508:subDocs` (visited on 11/04/2024).

[85] Daniel Alexander Berninghoff. ‘Commissioning of the Front-End Electronics of the LHCb Scintillating Fibre Tracker’. Heidelberg University Library, 2022.

[86] Paulo Moreira, Jorgen Christiansen and Ken Wyllie. *GBTx Manual*. URL: `https://espace.cern.ch/GBT-Project/GBTX`.

[87] Francois Vasey. *The Versatile Link Application Note*. URL: `https://edms.cern.ch/document/1309737/1` (visited on 28/02/2024).

[88] Luigi Del Buono et al. ‘LHCb Upgrade SciFi Tracker TELL40 Data Processing’. In: (). URL: `https://edms.cern.ch/document/1904563/5`.

[89] LHCb Collaboration. *LHCb GitLab*. GitLab. 16th May 2024. URL: `https://gitlab.cern.ch/lhcb` (visited on 16/05/2024).

[90] Mark Fox. *Quantum Optics: An Introduction*. Oxford: Oxford University Press, 2006.

[91] S. Vinogradov et al. ‘Probability Distribution and Noise Factor of Solid State Photomultiplier Signals with Cross-Talk and Afterpulsing’. In: *2009 IEEE Nuclear Science Symposium Conference Record (NSS/MIC)*. 2009 IEEE Nuclear Science Symposium Conference Record (NSS/MIC). Oct. 2009, pp. 1496–1500.

[92] S. Vinogradov. ‘Analytical Models of Probability Distribution and Excess Noise Factor of Solid State Photomultiplier Signals with Crosstalk’. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 695 (Dec. 2012), pp. 247–251.

[93] L. Gallego et al. ‘Modeling Crosstalk in Silicon Photomultipliers’. In: *Journal of Instrumentation* 8.05 (May 2013), P05010.

[94] *SciFi Production Database*. URL: https://scifi.physi.uni-heidelberg.de/db/prod/login.php (visited on 09/05/2024).

[95] Craig C. Sherbrooke. 'Discrete Compound Poisson Processes and Tables of the Geometric Poisson Distribution'. In: *Naval Research Logistics Quarterly* 15.2 (June 1968), pp. 189–203.

[96] P. C. Consul and G. C. Jain. 'A Generalization of the Poisson Distribution'. In: *Technometrics* 15.4 (Nov. 1973), pp. 791–799.

[97] M.A. Mariscotti. 'A Method for Automatic Identification of Peaks in the Presence of Background and Its Application to Spectrum Analysis'. In: *Nuclear Instruments and Methods* 50.2 (May 1967), pp. 309–320.

[98] Glen Cowan and RHUL Physics. *Error Analysis for Efficiency*, p. 9. URL: https://www.pp.rhul.ac.uk/~cowan/stat/notes/efferr.pdf.

[99] Steve Baker and Robert D. Cousins. 'Clarification of the Use of CHI-square and Likelihood Functions in Fits to Histograms'. In: *Nuclear Instruments and Methods in Physics Research* 221.2 (Apr. 1984), pp. 437–442.

[100] Paul G. Hoel and Paul Gerhard Hoel. *Introduction to Mathematical Statistics*. 4. ed. Wiley Series in Probability and Mathematical Statistics. New York, NY: Wiley, 1971. 409 pp.

[101] Fons Rademakers et al. *Root-Project/Root: V6.20/06*. Version v6-20-06. 14th June 2020.

[102] *PACIFIC5 Integrator Response*. URL: https://edms.cern.ch/ui/#!master/navigator/document?D:100455263:100455263:subDocs (visited on 09/05/2024).

[103] Andre Günther. 'Track Reconstruction Development and Commissioning for LHCb's Run 3 Real-time Analysis Trigger'. Heidelberg U., 2023.

[104] Salvatore Aiola et al. 'HybridSeeding: A Standalone Track Reconstruction Algorithm for Scintillating Fibre Tracker at LHCb'. In: *Computer Physics Communications* 260 (Mar. 2021), p. 107713.

[105] Robert D. Cousins. *Probability Density Functions for Positive Nuisance Parameters*. Oct. 2008. URL: http://www.physics.ucla.edu/~cousins/stats/cousins_lognormal_prior.pdf (visited on 03/04/2024). preprint.

[106] LHCb Collaboration. *SciFi Fine Time Alignment*. 2023. URL: https://cds.cern.ch/record/2868138/.

[107] Olivier Callot. *Acceptable Noise Rate in the SiPM of the Fibre Tracker*. Geneva: CERN, 2012. URL: https://cds.cern.ch/record/1496111.

[108] *LHCb Operations Plots*. URL: https://lbgroups.cern.ch/online/OperationsPlots/ (visited on 07/05/2024).