

DISSERTATION

submitted to the
Combined Faculty of Mathematics, Engineering and Natural Sciences
of the Ruperto Carola University of Heidelberg, Germany
for the degree of
Doctor of Natural Sciences

put forward by
Daniel Alexander Berninghoff
born in Lüdenscheid
Oral examination: 25.05.2022

Commissioning of the Front-End Electronics of the LHCb Scintillating Fibre Tracker

Referees:

Prof. Dr. Ulrich Uwer
Prof. Dr. Klaus Reygers

Abstract

Between 2019 and 2022, the LHCb experiment has undergone a major upgrade. It enables the detector to be operated at an increased luminosity and to be read out at 40 MHz corresponding to the proton bunch crossing rate at the LHC. In the course of the upgrade, the tracking stations downstream of the magnet have been replaced by a tracker made of scintillating fibres. With a fibre diameter of 250 μm , silicon photomultiplier readout and custom front-end electronics, it is well suited for the conditions of the next data taking period.

The front-end electronics of the Scintillating Fibre (SciFi) Tracker follow a modular design and have been optimised for the detector readout at 40 MHz. Before being installed in the upgraded LHCb experiment, every component is thoroughly tested as part of a detailed commissioning procedure. This thesis describes the process of the commissioning and presents the results of a large fraction of the detector. Along with further performance studies, it could be shown that even on large scales the SciFi front-end electronics meets very high quality standards.

Kurzfassung

In den Jahren zwischen 2019 und 2022 wurde das LHCb Experiment einem umfassenden Upgrade unterzogen. Dieses ermöglicht den Betrieb des Detektors bei erhöhter Luminosität, sowie die Auslesung der Daten bei einer Rate von 40 MHz, was der Protonkollisionsrate am LHC entspricht. Im Zuge des Upgrades wurden die Hauptspurkammern hinter dem Magneten mit einem neuen Detektor ersetzt, der auf szintillierenden Fasern basiert. Mit einem Faserdurchmesser von 250 μm und einer Auslese mittels Silizium Photoovervielfachern und maßgeschneiderter Elektronik ist der Detektor gut für die neuen Bedingungen gerüstet.

Die Ausleseelektronik des Scintillating Fibre (SciFi) Trackers ist modular aufgebaut und wurde speziell für die Ausleserate von 40 MHz entwickelt. Vor dem Einbau im LHCb Experiment im Zuge des Upgrades wird jede Komponente in einem detaillierten Inbetriebnahmeverfahren sorgfältig getestet. Diese Prozedur wird im Rahmen dieser Arbeit besprochen, sowie die Testergebnisse eines Großteils des Detektors vorgestellt. Zusammen mit weiterführenden Studien konnte gezeigt werden, dass die Ausleseelektronik des SciFi Trackers auch in großen Stückzahlen sehr hohen Qualitätsanforderungen genügt.

Contents

1	Introduction	1
2	The LHCb Experiment at the LHC	3
2.1	Detector Upgrade	4
2.2	Tracking System	6
2.3	Particle Identification	9
2.4	Trigger and Data Acquisition	11
3	The Scintillating Fibre Tracker	13
3.1	Motivation for the Tracker Upgrade	13
3.2	Performance Requirements	14
3.3	Detector Design and Operating Principle	15
3.4	Scintillating Fibres	16
3.5	Fibre Modules	17
3.6	Silicon Photomultipliers	24
3.7	Front-End Electronics	30
3.8	Infrastructure	31
3.8.1	Electrical Power	32
3.8.2	Cooling	33
3.8.3	Mechanical Support Structure	35
3.9	Tracking Performance	36
4	Front-End Electronics of the SciFi Tracker	39
4.1	PACIFIC Board	39
4.2	PACIFIC ASIC	41
4.2.1	Signal Processing and Digitisation	42
4.2.2	Serialisation	48
4.2.3	Slow Control Interface and Configuration Registers	49
4.2.4	Threshold Scan	50
4.3	Cluster Board	52
4.3.1	Cluster SCA	53
4.3.2	Cluster FPGA	54
4.4	Master Board	58
4.4.1	Master GBTX	59
4.4.2	Master SCA	61
4.4.3	Data GBTX	62
4.5	Light Injection System	63
4.6	Quality Assurance	65

5	Commissioning of the SciFi Front-End Electronics	69
5.1	Complexity of Readout System	70
5.2	Front-End Tester	72
5.3	Commissioning Setup	73
5.3.1	Infrastructure	75
5.3.2	DAQ System	76
5.3.3	Control Software	78
5.3.4	Commissioning Procedure	82
5.4	Communication to Master GBTXs	83
5.5	Configurability of ROBs	84
5.6	Optical Fibre Mapping	88
5.7	Optical Power	89
5.8	Temperature Sensors	92
5.9	Voltage Sensors	96
5.9.1	Low Voltage	96
5.9.2	High Voltage	97
5.10	High Voltage Monitoring Calibration	98
5.11	Bit Error Rate Tests	99
5.12	Clock Timing Scans	103
5.12.1	FPGA Serialiser → Data GBTX (A)	104
5.12.2	FPGA Clustering → Serialiser (B)	105
5.12.3	FPGA Deserialiser → Clustering (C)	106
5.12.4	PACIFIC → FPGA Deserialiser (D)	108
5.12.5	SYNC Pulse → PACIFIC (E)	108
5.12.6	Conclusion	111
5.13	Light Injection System Tests	111
5.13.1	Functional Test and Mapping	112
5.13.2	Delay Scan (Lite)	113
5.13.3	Threshold Scans	115
5.14	Commissioning Summary	119
6	Performance Studies	121
6.1	PACIFIC Pulse Height Digitisation	121
6.1.1	Pedestal Stability	121
6.1.2	Internal Charge Injection	123
6.1.3	Integrator Differences	124
6.1.4	Comparator Differences	128
6.2	Clock Distribution	129
6.3	Dark Count Rate	131
7	Conclusion	133
	Bibliography	135

1 Introduction

The Standard Model (SM) of particle physics [1]–[3] is a well tested and very successful theory that explains the structure of matter and the interactions between the elementary particles. According to the SM, all visible forms of matter are composed of 12 fundamental particles, which are divided into the groups of quarks and leptons. Furthermore, it includes the gauge bosons that are the mediators of the interactions between these particles. One of them is the well-known photon, which is the force carrier of the electromagnetic force. Together with the strong and weak interaction, they represent the three fundamental forces included in the model. Lastly, there is the Higgs boson, which was predicted by Peter Higgs in 1964 [4]. The associated Higgs field plays an important role in the SM, as it explains how the elementary particles gain their masses.

While almost any experimental observation in the field of particle physics is in agreement with the SM predictions, it is not a complete theory. This can already be seen from the fact that the gravitational force, although negligible at the subatomic scale, is not covered by it. Other observations that cannot be explained by the SM are the asymmetry between matter and antimatter in the universe [5], the existence of dark matter [6], and the measurement of neutrino oscillations [7]. It is therefore essential to continue the thorough testing of the SM, as well as theories beyond it.

In 2012, with the observation of the Higgs boson at the Large Hadron Collider (LHC), the last missing elementary particle of the SM was discovered [8], [9]. The LHC near Geneva, Switzerland, is the most powerful particle accelerator in the world. By colliding high-energy protons at enormous intensities, it provides the ideal conditions to probe the SM and search for new physics phenomena.

The LHCb experiment is one of four major experiments that are located at the LHC. It is a dedicated experiment specialised in studying hadrons¹ containing b and c quarks. They are produced in the proton-proton collisions and decay again after a short time into other particles, which can be detected in the experiment. During the first years of operation starting from 2010, the LHCb experiment has recorded the world's largest samples of events containing b and c hadrons [10].

The collected samples allowed for performing important measurements in the field of flavour physics such as the observation of the rare decay $B_s^0 \rightarrow \mu^+ \mu^-$ [11], the determination of the CKM² angle γ [12] or the test of lepton universality [13]. While the latter already revealed a slight tension compared to the SM prediction, the

¹Hadrons are composite particles made of two or more quarks that are bound by the strong interaction. They are further divided into mesons (usually containing two quarks) and baryons (usually three quarks).

²The Cabibbo–Kobayashi–Maskawa (CKM) matrix parametrises the relative weak coupling strengths between the different quark flavours.

significance of the measurement is statistically limited. In fact, many key results at LHCb are limited by the statistical sensitivity of the data samples. Therefore, to be competitive with the uncertainties of the theoretical predictions, the LHCb experiment will collect data at five times the previous luminosities¹ starting from the middle of the year 2022. In order to prepare the detector for the increasing rates, a three-year upgrade period was required during which in particular the tracking system has been improved.

An essential component of the upgraded tracking system is a detector made of scintillating fibres: the LHCb Scintillating Fibre (SciFi) Tracker. With a fibre diameter of 250 μm , it provides sufficient granularity for the higher number of particle tracks in each event. The produced scintillation light is detected by arrays of silicon photomultipliers (SiPMs), whose output signals are processed and digitised by a complex chain of front-end electronics. The electronics is designed and optimised to sample the signals at a rate of 40 MHz, matching the proton bunch collision rate at the LHC. While the initial version of the LHCb experiment was only read out at a fraction of that rate, the so-called trigger-less readout is required to fully exploit the increase in luminosity.

In the context of this thesis, several contributions to the SciFi Tracker, in particular to the front-end electronics, have been made. While supporting various steps in the research and production phase of the different electronic components, the main responsibility was the initial commissioning of the electronics on a large scale alongside the surrounding infrastructure and the upgraded LHCb DAQ system. A major challenge that had to be overcome in this context was the establishment of the 40 MHz readout of the complex detector. In its final configuration, the SciFi Tracker will produce data volumes up to 20 Tbit/s corresponding to about 40 % of the entire LHCb experiment.

As part of the front-end electronics commissioning, a detailed test procedure has been defined, developed and implemented. In the process, every component is thoroughly tested and validated in its final environment before being released for installation in the experiment. This also includes first calibration and tuning steps that are required for the proper operation of the detector.

The following contents will be discussed in the course of this thesis: In Chapter 2, the LHCb experiment is further introduced, followed by the SciFi Tracker in Chapter 3. Afterwards, in Chapter 4, a detailed overview on the SciFi front-end electronics is given, providing a first complete reference of all components involved. In Chapter 5, the commissioning procedure is presented and the results from the commissioning of a large fraction of the detector are shown. Finally, performance studies that were conducted alongside the commissioning are discussed in Chapter 6.

¹At particle colliders, the luminosity \mathcal{L} is a measure for the interaction rate \dot{N} relative to the cross-section σ .

2 The LHCb Experiment at the LHC

This chapter aims at introducing the LHCb experiment at the Large Hadron Collider (LHC). The LHC is the largest circular particle accelerator and collider in the world and is located near Geneva, Switzerland. With a circumference of about 27 km, it is placed in a tunnel between 45 m and 170 m below the surface and is operated by the European Organization for Nuclear Research (CERN) as part of a larger accelerator complex [14].

The LHC is designed to collide two proton beams at centre-of-mass energies up to 14 TeV. The high-energy beams are guided around the accelerator ring by superconducting magnets, which are cooled below 2 K using superfluid helium. The beams consist of bunches spaced 25 ns apart, each containing about 10^{11} protons. They travel in opposite directions in two separate beam pipes and are brought to collision at four different points along the ring, resulting in a maximum bunch crossing rate of 40 MHz and an instantaneous luminosity of up to $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. However, several gaps are foreseen in the filling scheme to allow for reliable dumping of the beams, reducing the average collision rate to about 30 MHz [14].

The four major experiments ATLAS, CMS, ALICE and LHCb are placed around the interaction points to study the collisions. ATLAS [15] and CMS [16] are the largest of the four and are designed as multi-purpose detectors covering a broad physics spectrum. By making use of the high luminosity provided by the LHC, the data collected at ATLAS and CMS led to the discovery of a new particle in 2012 [8], [9], whose properties were later found to be in agreement with the Higgs boson in the Standard Model [4], [17]. ALICE is a general-purpose detector focusing on the physics of strongly interacting matter and the quark-gluon plasma [18]. To allow for this, the LHC is designed to be able to also accelerate and collide heavy (Pb) ions at centre-of-mass energies up to 1.15 PeV [14]. While ALICE is specialised in studying these collisions, the other experiments make use of the data collected during heavy ion runs as well.

This also includes the LHCb experiment that is dedicated to study rare decays of hadrons containing b and c quarks as well as performing precision measurements of CP violation [19]. The experiment is located in a cavern 100 m below the surface at the border between France and Switzerland near the airport of Geneva.

Detector Design

The LHCb detector is designed as a single-arm forward spectrometer covering an angle between 10 mrad and 300 mrad (250 mrad) in the bending (non-bending) plane

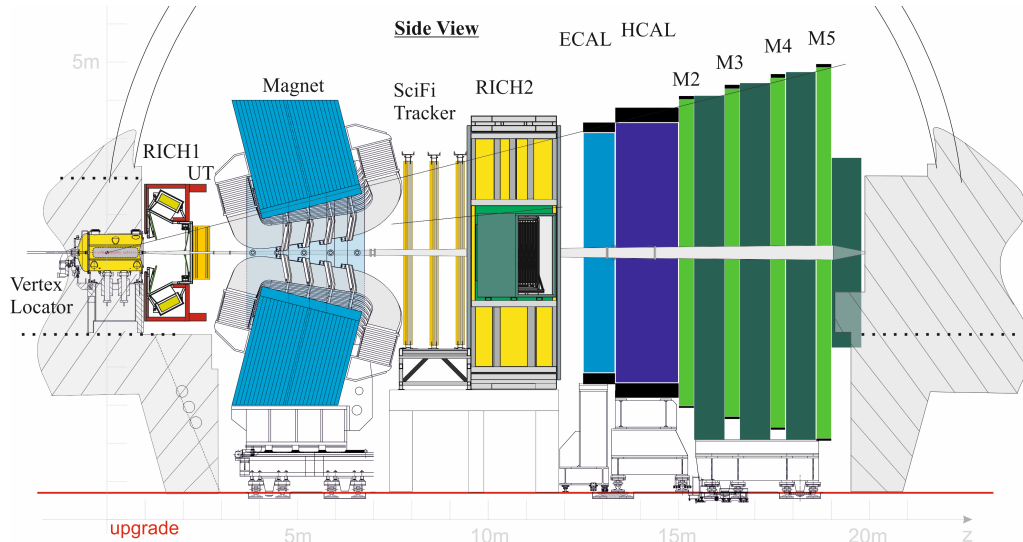


Figure 2.1: Schematic side view of the upgraded LHCb detector. The proton-proton collisions take place inside the Vertex Locator on the left. Image taken from Ref. [20].

of the dipole magnet [19]. This corresponds to a pseudorapidity¹ range $2 < \eta < 5$ and accounts for about 4% of the solid angle. As illustrated in Fig. 2.1, a right-handed coordinate system is adopted at LHCb, whose z -axis runs horizontally in the direction of the beam pipe. The other horizontal axis (x) is the direction in which charged particles are mainly bent when passing through the magnetic field.

Unlike the three general-purpose detectors at the LHC (ATLAS, CMS and ALICE), which are built in a cylindrical shape around the interaction point in order to achieve a large geometric coverage, the special design of the LHCb detector was chosen specifically for its field of application: As shown in the simulations in Fig. 2.2, the $b\bar{b}$ quark pairs that are of particular interest at LHCb are predominantly produced at small opening angles to the beam axis, which is reflected in the geometry of the detector.

Between 2019 and 2022, the detector has undergone a major upgrade in order to boost the future data collection rates and thus significantly reduce the statistical uncertainties of the measurements. In the following, the upgraded LHCb experiment as well as the scope of the upgrade are presented.

2.1 Detector Upgrade

During the first years of operation between 2010 and 2018, the (initial) LHCb detector recorded data corresponding to an integrated luminosity of about 9 fb^{-1} . Towards the end of this period, the detector was operating at a centre-of-mass energy $\sqrt{s} = 13 \text{ TeV}$ and an instantaneous luminosity $\mathcal{L} = 4 \cdot 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$ [22]. It

¹The pseudorapidity η is related to the angle to the beam pipe θ via the equation $\eta = -\ln[\tan(\theta/2)]$.

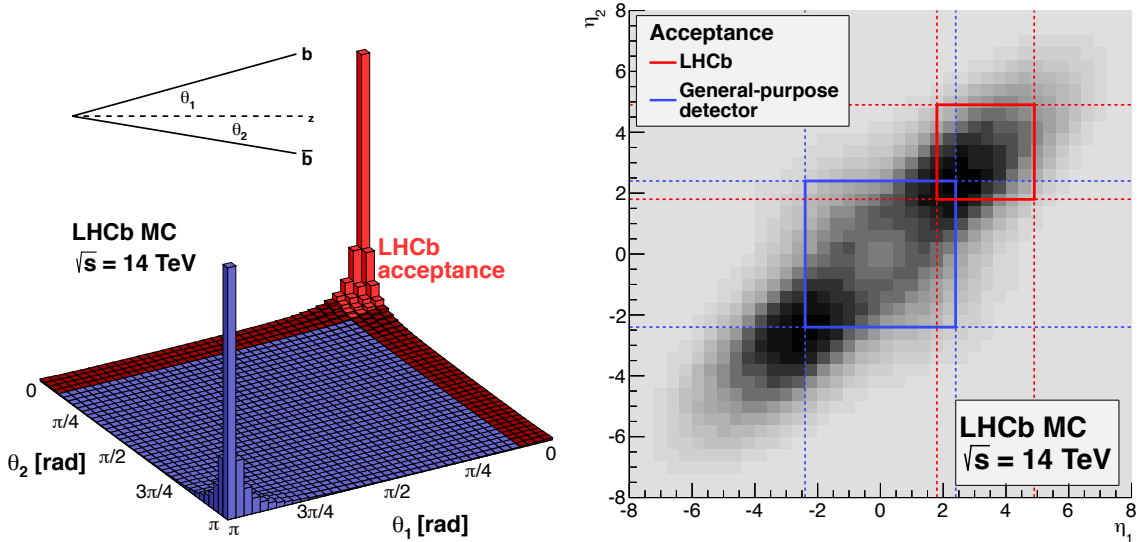


Figure 2.2: Simulated $b\bar{b}$ production angles at centre-of-mass energies $\sqrt{s} = 14$ TeV as a function of the angles $\theta_{1,2}$ (left) and pseudorapidities $\eta_{1,2}$ (right). The LHCb acceptance is highlighted in red and is compared with typical values for a general-purpose detector in case of the pseudorapidities. Images modified from Ref. [21].

should be noted that this value is about two orders of magnitudes below the LHC design luminosity. With the help of a tunable beam focus and separation at its interaction point, LHCb is able to lower the luminosity to an optimal level for the experiment [19]. This procedure is performed in order to limit the number of visible proton-proton interactions per bunch crossing, which is required for a clean and efficient event reconstruction.

Starting from LHC Run 3 in 2022, the LHCb detector will be operated at five times the luminosity compared to before the upgrade. Until the end of Run 4, this corresponds to a collected data volume of 50 fb^{-1} [23]. This results in a significant reduction of statistical uncertainties and allows the experiment to reach its full flavour-physics potential at the LHC. However, several improvements need to be made during the upgrade prior to Run 3 in order to prepare the detector for the larger number of proton-proton interactions and track multiplicities.

This concerns in particular the LHCb tracking system, which is completely replaced during the upgrade between 2019 and 2022. Furthermore, the trigger system undergoes a fundamental change. In the initial version, the readout of the detector was performed at 1 MHz as determined by the first trigger level implemented in hardware [19]. While the so-called L0 trigger provides high efficiencies for dimuon events, it quickly saturates for hadronic channels [23]. To fully profit from the increase in luminosity, it was decided to perform a trigger-less readout of the complete detector at the 40 MHz bunch crossing rate. To achieve this, it is required that the front-end electronics of all sub-systems are revised and replaced.

In the following, an overview of the various sub-systems of the upgraded LHCb




Figure 2.3: Schematic view of one half of the upgraded LHCb Vertex Locator (VELO) consisting of 26 modules (left). Two opposing modules from both halves in the closed position are shown on the right. Image adapted from Ref. [26].

detector is given.

2.2 Tracking System

As previously discussed, the tracking system undergoes major changes during the LHCb upgrade as it is completely replaced by new detectors. The system is responsible for reconstructing the trajectories of charged particles. It consists of three parts that are located closely around the interaction point in order to resolve production and decay vertices, as well as before and after the dipole magnet. By measuring the deflection angle of charged particles caused by the Lorentz force in the magnetic field, this arrangement enables the determination of the particles' momenta.

Vertex Locator (VELO)

The (upgraded) LHCb Vertex Locator (VELO) is a silicon pixel detector that is placed closely around the proton-proton interaction region. It consists of two detector halves equipped with 26 modules each. The modules are arranged along the beam axis and are oriented perpendicular to the beam. The detector consists of a total of about 41M pixels with a size of $55 \times 55 \mu\text{m}^2$ [24]. They are grouped in four silicon hybrid pixel tiles per module, whereas each tile is readout by three VeloPix ASICs. The total output data rate is estimated at 1.6 Tbit/s while reading out the front-end electronics at the full LHC bunch crossing rate of 40 MHz [25].

A special feature of the VELO is the ability to mechanically change the distance between the modules and the beams. In the closed state, the active pixels will be as close as 3.5 mm from the proton beams, allowing precise resolution of the reconstructed tracks and vertices during normal operation of the LHC [25]. On the other hand, this feature offers protection from unstable beam conditions by




Figure 2.4: Schematic layout of the Upstream Tracker (UT) with the four detection layers (left). The detector is composed of four different sensor types: Type A (green), B (yellow), C and D (pink). It is built from vertical staves, which carry the sensors and readout hybrids as shown on the right. Image adapted from Ref. [27].

mechanically increasing the distance to the beams (open state). In order to allow for the two detector halves to be closed completely, the modules are slightly shifted along the beam axis relative to the opposing half. A schematic view of one VELO half as well as two opposing modules in the closed position are shown in Fig. 2.3.

Upstream Tracker (UT)

The Upstream Tracker (UT) is located just in front, i.e. upstream of the dipole magnet. As shown in Fig. 2.4 (left), the detector is arranged in four detection layers that are equipped with silicon strip sensors. Each plane has an active area of approximately 2 m^2 , which covers the full acceptance of the LHCb experiment. In order to obtain the highest possible hit resolution in the horizontal, bending direction of the magnet, the strips in the outer layers are oriented vertically. The resolution in the non-bending direction is achieved by rotating the two inner layers by $\pm 5^\circ$. Depending on the track occupancy and radiation dose, four different sensor types are in use. The majority of sensors (type A) have a strip pitch of $187.5\ \mu\text{m}$, while the pitch of the three sensor types closer to the beam pipe amounts to $93.5\ \mu\text{m}$ [27].

The detector is built from a total of 68 elongated support structures that are referred to as staves. As illustrated in Fig. 2.4 (right), each stave is equipped with 14 to 16 silicon strip sensors. They are mounted from both sides with a slight overlap between neighbouring sensors. Each sensor is accompanied by a readout hybrid carrying up to eight custom front-end chips, called SALT. The radiation-hard chips




Figure 2.5: Expected dose in the x - y plane at $z = 783$ cm after an integrated luminosity of 50 fb^{-1} . The Scintillating Fibre (SciFi) modules of the first tracking station located at this position are shown superimposed. The integrated dose absorbed in the fibres towards the centre peaks at about 35 kGy. Image modified from [20].

are based on 130 nm technology and were developed to sample the signals at the LHC frequency of 40 MHz. The resulting digital signals are sent via flex cables to the near-detector electronics [27].

Scintillating Fibre (SciFi) Tracker

The last component within the LHCb tracking system is the Scintillating Fibre (SciFi) Tracker [20]. It consists of three stations that are located downstream of the magnet. Each station is composed of four detection layers that follow a similar geometry as the UT with the two inner layers being tilted by $\pm 5^\circ$. The active area, which amounts to about 30 m^2 per layer, is built from staggered layers of scintillating fibres with a diameter of $250 \mu\text{m}$. They are read out by arrays of silicon photomultipliers (SiPMs), which are located at the top and bottom of the detector outside of the geometrical acceptance. In the adjacent front-end electronics boxes, the SiPM signals are processed and digitised at the full LHC bunch crossing frequency of 40 MHz. The dimensions of one detector layer are illustrated in Fig. 2.5, along with the expected radiation dose at the position of the first station after an integrated luminosity of 50 fb^{-1} .

The SciFi Tracker and in particular its front-end electronics are the subject of this thesis and are therefore discussed in detail in the following chapters.

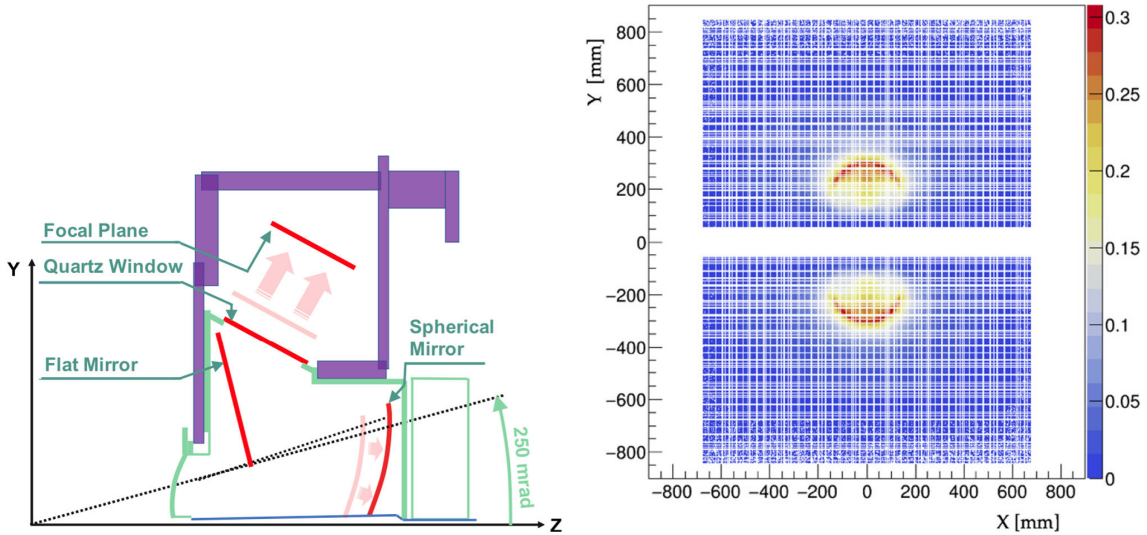


Figure 2.6: Schematic side view of the upper half of RICH1 (left) and simulated single channel occupancy (right). Modifications to the mechanical and optical system in the course of the upgrade are indicated by the red arrows in the side view. Images taken from Ref. [29].

2.3 Particle Identification

The particle identification (PID) capabilities play an important role for any flavour-physics experiment and has been essential for the success of LHCb during the first years of operation [28]. It consists of the Ring Imaging Cherenkov (RICH) system, two calorimeters and the muon stations. Compared to the tracking system, only small changes have been made to the PID system during the LHCb upgrade and mostly affect the RICH detectors.

Ring Imaging Cherenkov (RICH) Detectors

The main purpose of the RICH detectors is the separation of kaons and pions. This is achieved by utilising Cherenkov light, which occurs when a charged particle traverses a medium faster than the speed of light. In combination with the momentum as provided by the tracking system, the emission angle can be used to determine the mass of the particle.

Two separate detectors form the RICH system at LHCb and are referred to as RICH1 and RICH2. As shown in Fig. 2.1 on page 4, RICH1 is placed between the VELO and UT and covers the full angular acceptance of the detector and a momentum range of 10–50 GeV/c. On the other hand, RICH2 is located downstream of the SciFi Tracker and specialises in high momentum particles between 15–100 GeV/c with angles up to 120 mrad [29]. For these purposes, RICH1 uses C_4F_{10} gas as the medium, while CF_4 is used in RICH2 [28].

While the enclosure of the RICH detectors could be reused in the upgrade, some

modifications to the mechanical and optical systems were required in order to expand the focal point and thereby reduce the channel occupancy. The modifications along with the simulated occupancy are shown in Fig. 2.6. Due to the performed changes, the maximum occupancy can be limited to 27% [29].

For both RICH1 and RICH2, commercial multi-anode photomultiplier tubes (MaPMTs) are used for the detection of the Cherenkov light, along with external front-end electronics. The 40 MHz readout is enabled by a custom designed, radiation-hard ASIC in 0.35 μm CMOS technology that is referred to as CLARO [30].

Calorimeters

The upgraded LHCb calorimetry system is composed of two sub-detectors that are located downstream of RICH2: an electromagnetic (ECAL) and a hadronic (HCAL) calorimeter [28]. Combined, they are responsible for the identification of electrons, photons and hadrons, as well as the determination of their energy. Both calorimeters employ a design with alternating layers of absorbers and scintillators. The produced scintillation light is transmitted via wavelength shifting fibres and is detected by photomultiplier tubes (PMTs).

One ECAL module is built from 66 stacked layers of lead absorbers (2 mm thick) and scintillators (4 mm). The module size varies from $121.1 \times 121.1 \text{ mm}^2$ in the outer section of the ECAL, down to $40.4 \times 40.4 \text{ mm}^2$ in the hottest region around the beam pipe. In contrast, the HCAL consists of alternating layers of scintillators (4 mm) and iron absorbers (16 mm), and employs two different module sizes of $131.3 \times 131.3 \text{ mm}^2$ and $262.6 \times 262.6 \text{ mm}^2$ in the inner and outer sections, respectively [19].

No changes to the calorimeter modules have been required during the upgrade since studies have shown that the granularity is sufficient even when being operated at a luminosity of $2 \cdot 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ starting from LHC Run 3. However, new front-end electronics had to be developed to support the readout rate of 40 MHz, as well as having a higher preamplifier sensitivity in order to mitigate the gain degradation of the PMTs caused by increasing levels of radiation [25].

Muon Stations

Similar to the calorimeters and apart from redeveloping the front-end electronics for the trigger-less readout, the muon sub-detector remains mostly untouched during the LHCb upgrade. It consists of four stations downstream of the calorimeters which, for historical reasons, are designated as M2 to M5. They are equipped with multi-wire proportional chambers (MWPCs) in order to reconstruct the particle trajectories. By taking advantage of their penetrating power, a robust identification of muons is possible, which is achieved with the help of interleaved iron walls between the stations. The iron shielding is also the reason that no major changes are required in order to be operated at higher luminosities. One exception to this is the muon station M1 that was installed in front of the calorimeters in the initial LHCb experiment. Since it cannot withstand the huge hit occupancy expected at this location starting

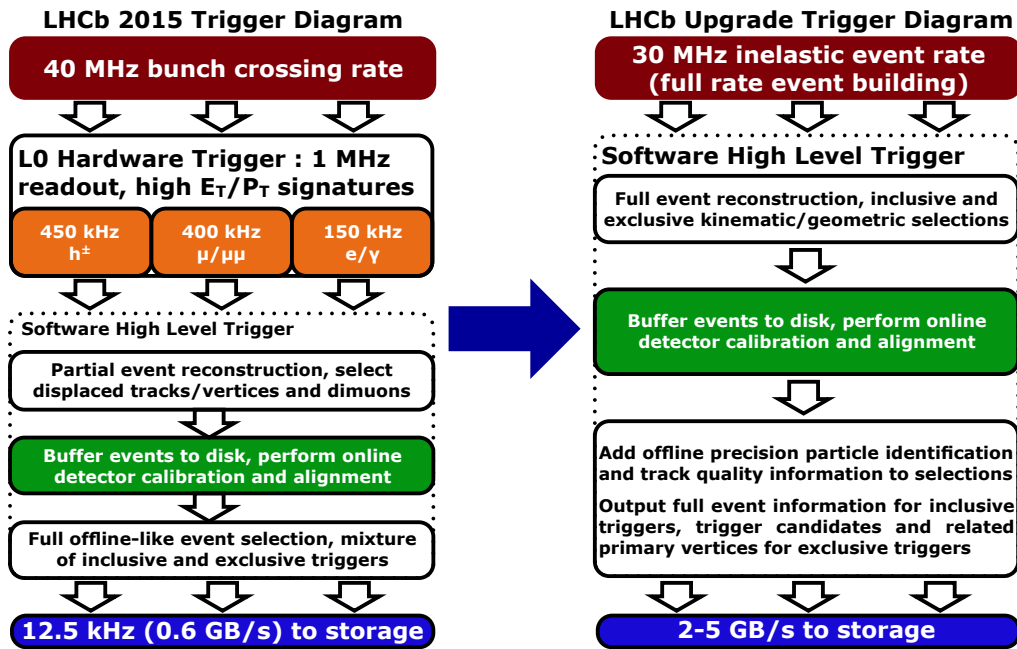


Figure 2.7: Trigger schemes for the initial (left) and upgraded LHCb experiment (right). Images taken from Ref. [31].

from LHC Run 3, it was removed in the course of the upgrade. The associated loss of information is marginal, as M1 was mainly used for the now redundant L0 trigger and was not part of the muon identification algorithms [28].

2.4 Trigger and Data Acquisition

Due to the high bunch crossing rate of 40 MHz at the LHC (30 MHz after excluding empty bunches), it is not feasible to store every event to disk for later access and analysis. In fact, during the first years of operation between 2010 and 2018, the readout of the front-end electronics of the complete detector was throttled to only a fraction of that rate. It was limited to 1 MHz as determined by the first trigger level (L0) implemented in hardware that had only access to the data provided by part of the detector [19].

However, to take full advantage of the fivefold increase in luminosity starting from LHC Run 3, the trigger scheme had to be revised [32]. As illustrated in Fig. 2.7, the reconstruction of the complete event consisting of the data provided by all sub-detectors is now performed at the collision rate of 30 MHz. The full event reconstruction and selection is performed in a trigger level that is fully implemented in software, the so-called Software High Level Trigger. To allow for the 30 MHz readout, the front-end electronics of all sub-detectors had to be renewed as discussed previously. They are designed to be read out via simplex optical links using a common and custom protocol that is referred to as GBT [33]. A total of about

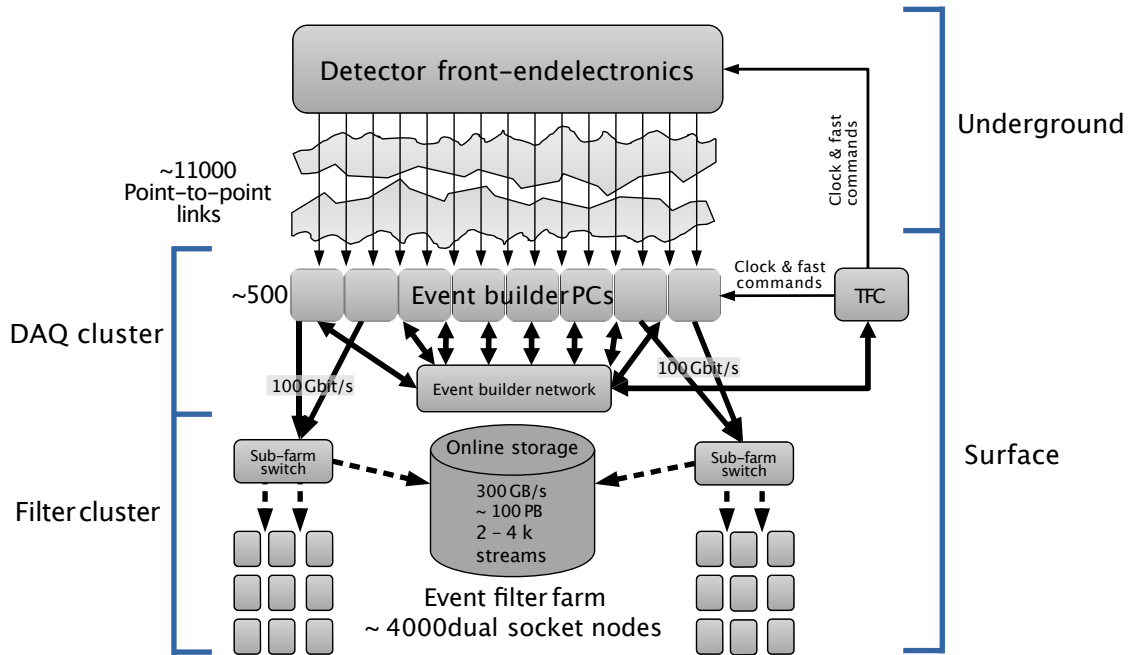


Figure 2.8: Layout of the DAQ system of the upgraded LHCb experiment. Image adapted from Ref. [36].

11 000 optical fibres, each carrying up to 4.48 Gbit/s of user data, transfer the data from the underground cavern to the server farm located on the surface of the LHCb site [34]. The so-called event builder servers are equipped with custom PCI Express readout cards and are responsible for the event reconstruction [35].

The reconstructed event information is used by a second cluster of servers referred to as event filter farm. They have access to a large disk-based buffer and run the selection algorithms in order to filter out the events to be kept for permanent storage [36]. The layout of the upgraded LHCb Data Acquisition (DAQ) system is shown in Fig. 2.8.

The detector front-end electronics and the readout cards are synchronised to the 40 MHz LHC clock by the Timing and Fast Control (TFC) system [37]. The distribution of the information is also done via optical links and the help of dedicated interface cards. An additional 2×2300 bidirectional optical fibres are used for this purpose to enable the communication with the front-end electronics in the LHCb underground cavern [34]. By using the GBT protocol, the bandwidth along with the identical interface cards are shared for the implementation of the Experiment Control System (ECS). While the TFC distributes time-sensitive information to the different components within the DAQ system, the ECS is used for slow-control operations. These include for example the transmission of configuration parameters to the electronics, as well as reading out sensor values or error counters.

Further details about the readout and interface cards enabling the DAQ, TFC and ECS are given in Section 5.3.2.

3 The Scintillating Fibre Tracker

Up to the end of the LHC Run 2 in 2018, two different detector technologies have achieved an excellent tracking performance within the LHCb experiment [38]–[40]. The Scintillating Fibre (SciFi) Tracker is going to replace the tracking stations T1-T3 downstream of the magnet implementing a single technology based on scintillating fibres.

The layout and working principle of the detector is outlined in this chapter, along with the motivation that led to the upgrade of the tracking stations and its requirements. If not stated otherwise, the provided information is taken from the Technical Design Report (TDR) of the upgraded tracking system [20].

3.1 Motivation for the Tracker Upgrade

The Inner (IT) and Outer Tracker (OT), which have formed the downstream tracking stations of the initial LHCb experiment, were designed to provide a high precision estimation of the momentum of charged particles. This was achieved by operating at high hit efficiency as well as high resolution in the bending plane of the magnetic field.

The OT was a gaseous detector covering an area of approximately $6\text{ m} \times 5\text{ m}$ per layer and consisted of 2.4 m long straw tubes with an inner diameter of 4.9 mm [41]. The tubes were filled with a mixture of argon, carbon dioxide and oxygen allowing for fast drift times below 50 ns. The detector modules were composed of two staggered layers of 64 straw tubes each and were arranged in three tracking stations (T1-T3). Each station consisted of four module layers following the so-called x - u - v - x geometry: Modules in the x layers were oriented vertically, while modules in the u and v layers were tilted by $+5^\circ$ and -5° , respectively. Consequently, a total of 24 straw tube layers were arranged along the direction of the beam pipe. The average hit efficiency was estimated to be 99.2% and the position resolution was measured to be close to the design value of $200\text{ }\mu\text{m}$ [38]. Utilising novel real-time calibration methods as described in Ref. [39], the resolution could be further improved to $171\text{ }\mu\text{m}$ during LHC Run 2.

A plus-shaped area of approximately $1.2\text{ m} \times 0.4\text{ m}$ in the centre of each layer was not covered by the OT and instead housed the IT. The IT was a silicon microstrip detector using $320\text{ }\mu\text{m}$ and $410\text{ }\mu\text{m}$ thick single-sided p^+ -on- n sensors with a strip pitch of about $200\text{ }\mu\text{m}$ [42]. Individual silicon sensors were 7.8 cm in width and 11 cm in height and were contained within a total of 12 detector boxes. Each tracking station was composed of four boxes that were located above, below, and on both sides of the beam pipe, thus forming the aforementioned plus-shape. Each detector

box was made up of four layers following the same $x-u-v-x$ geometry as described above for the OT. Despite only covering about 1.3% of the geometrical acceptance, about 20% of all charged particles that were produced close to the interaction point were passing through the active area of the IT. The hit efficiency was estimated to be 99.7% and the resolution was measured to be 58 μm [40].

The initial tracking system was designed to limit the maximum occupancy in the OT to 10% at nominal instantaneous luminosities $\mathcal{L} = 2 \cdot 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$. It has been shown that it is possible to go beyond that design limit and efficiently find tracks at occupancies up to 25% in the hottest regions of the OT [23]. The necessary improvements in the track reconstruction algorithms thereby allowed data collection at instantaneous luminosities $\mathcal{L} = 4 \cdot 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$. However, going beyond $\mathcal{L} = 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ has been found to require detectors of higher granularity. Different scenarios including an increased area of the IT along with shorter OT modules in the central part were considered. In the end, the decision towards a scintillating fibre tracker as a cost-effective, light and uniform detector covering the area of 30 m^2 per layer has been made.

3.2 Performance Requirements

The main objective of the tracking detectors within the LHCb experiment is the accurate determination of the momentum of charged particles. This information can be used to precisely measure the mass and lifetime of decayed particles. In order to achieve that, high hit efficiencies, as well as a good position resolution are necessary. In addition, the material budget has to be kept low to minimise multiple scattering, which is the main limiting factor for the momentum resolution of low momentum particles. Taking these considerations into account, the following requirements for the SciFi Tracker were formulated in the TDR [20]:

- Reaching high hit efficiencies at the level of 99%, while keeping the rate of reconstructed noise clusters low ($< 10\%$ of the recorded hits) at any location.
- The single hit position resolution must be 100 μm or better in the bending plane of the magnet.
- Minimise the effect of multiple scattering by keeping the material budget low. Achieving radiation lengths $X/X_0 < 1\%$ per detection layer is sufficient to not be the limiting factor compared to material upstream of the magnet.
- The readout electronics needs to operate and output data at the full LHC clock speed of 40 MHz. Inefficiencies due to dead time should be minimised by short recovery times.
- The aforementioned requirements should be met over the full lifetime of the upgraded LHCb detector up to an integrated luminosity of 50 fb^{-1} .




Figure 3.1: Schematic side (left) and front (right) view of the SciFi tracking station T3 with a 1.86 m tall human dummy for scale. Image adapted from Ref. [43].

3.3 Detector Design and Operating Principle

The SciFi Tracker fully replaces the OT and IT between the dipole magnet and RICH2 in the LHCb experiment. It is designed to follow the identical layout as its predecessors consisting of three stations T1-T3 with four detection layers each. The layers are arranged in the same $x-u-v-x$ geometry with the two inner layers being tilted by $+5^\circ$ and -5° from the vertical axis. Each layer covers an area of approximately $6\text{ m} \times 5\text{ m}$ and consists of a minimum of 10 fibre modules. The tracking station T3, which is furthest away from the interaction point, is equipped with two additional modules per layer to not limit the geometrical acceptance of the experiment. Figure 3.1 shows the schematic layout of the tracking station T3.

Each detector module is composed of a total of eight fibre mats that are read out by multichannel arrays of silicon photomultipliers (SiPMs) located at the top and bottom of the detector. The fibre mats itself consist of six staggered scintillating fibre layers with $250\ \mu\text{m}$ in diameter.

The operating principle is illustrated in Fig. 3.2. Charged particles that traverse the layers of scintillating fibres dissipate energy as ionisation and excitation. Subsequently, the absorbed energy is transformed into luminescence emission in the scintillation medium [44]. Photons that are emitted along the fibres are trapped inside by total reflections at the fibre boundaries. They follow the course of the fibres until they reach the SiPM channels at the fibre ends.

The amplitude of the resulting electrical signal is proportional to the number of

3 The Scintillating Fibre Tracker




Figure 3.2: Operating principle of the LHCb Scintillating Fibre Tracker. Image adapted from Ref. [20].




Figure 3.3: Longitudinal (top) and cross (bottom) section of a multi-cladding scintillating fibre with exemplary trajectories of trapped photons. Image taken from Ref. [45].

pixels that are hit by photons within each channel. The average light yield, i.e. the number of pixels hit, per traversing particle is in the order of 20 depending on the incident angle and distance from the SiPM array. However, irradiating the fibres with expected doses after 50 fb^{-1} results in a loss of light yield by about 40%. The position of the corresponding hit can be defined as the location of involved SiPM channels weighted by their signal amplitudes.

3.4 Scintillating Fibres

The SciFi Tracker employs scintillating fibres with a round cross section of $250 \mu\text{m}$ in diameter produced by Kuraray¹. Fibres of type SCSF-78MJ were chosen for achieving a fast decay time of 2.4 ns and a large attenuation length of about 3.5 m [46]. A high light yield is achieved by utilising a multi-cladding structure with descending refractive indices towards the surface of the fibre, while also being mechanically more robust. The inner cladding is made of Polymethylmethacrylate (PMMA, $n = 1.49$) and surrounds the Polystyrene (PS, $n = 1.59$) core, while Fluorinated polymer (FP, $n = 1.42$) is used as the outer cladding. Each cladding is about $4 \mu\text{m}$ in thickness [47].

Following Snell's law about the angle dependence of light passing the boundary

¹Kuraray Co., Ltd., Tokyo, Japan

of two isotropic media, there exists a critical angle of incidence

$$\theta_{\text{crit}} = \arcsin \frac{n_2}{n_1} \quad (3.1)$$

above which total internal reflection at the boundary occurs. For the transition between the inner ($n_1 = 1.49$) and outer cladding ($n_2 = 1.42$) the critical angle can be calculated as $\theta_{\text{crit, PMMA} \rightarrow \text{FP}} = 72.4^\circ$. This translates to a maximum opening angle of $2 \times 26.7^\circ$ at the centre of the fibre as illustrated in Fig. 3.3. Assuming an isotropic photon emission, the trapping efficiency is thereby given as 5.3%. Following similar considerations, having a single-cladding fibre with a PS core and PMMA cladding only results in a 3.1% trapping efficiency.

Pure Polystyrene has a relatively poor quantum efficiency and long relaxation time. Therefore, a small amount (about 1% in weight) of p-Terphenyl is added to the PS core. This organic fluorescent dye rapidly (below 1 ns) absorbs the excitation energy of the PS base via a non-radiative dipole-dipole interaction known as Förster Transfer [48]. Subsequently, the absorbed energy is released by the emission of a photon. The fluorescent dye has been chosen for its high quantum efficiency $> 95\%$ and fast decay times below a few ns [20].

A secondary dye is added to the PS core in a very low concentration (below 1‰ in weight). It acts as a wavelength shifter by absorbing the emission of the primary dye and re-emitting photons at a longer wavelength at which re-absorption in the fibre occurs less frequently [49]. Tetraphenyl butadiene (TPB) is used as the secondary dye in the scintillating fibres of type SCSF-78MJ and enables the required large attenuation length. Figure 3.4 shows the resulting photon emission spectrum of the scintillating fibres, which peaks between 450 nm and 500 nm depending on the distance to the point of excitation.

More than 10 000 km of fibres are used for the SciFi Tracker and were delivered to CERN in 48 individual shipments between 2016 and 2018. Every batch of fibres was subject to a detailed test procedure [46]: Samples were taken to evaluate the attenuation length before and after being exposed to X-ray irradiation in order to reveal impurities in the used materials. The ionisation light yield was measured using energy filtered 1 MeV electrons from a Sr-90 radioactive source. In addition to that, for every km of fibre, the quality of the cladding and uniformity of the diameter was verified. Fibre bumps up to 500 μm in diameter were detected and removed by a hot shrinking procedure in a fully automatic manner [51], while larger bumps were cut out manually. The conducted quality assurance measures ensured that only high quality fibres were assembled into the fibre modules as described in the following.

3.5 Fibre Modules

The fibre modules constitute the active part of the SciFi Tracker. In total, 128 of these modules are used and are arranged in 12 layers and 3 stations as described in Section 3.3. Each fibre module is 4.85 m long and 0.52 m in width. Figure 3.5 shows




Figure 3.4: Emission spectrum of a scintillating fibre of type SCSF-78MJ at different distances to the point of excitation. For short distances, the maximum is located at around 450 nm and shifts towards larger wavelengths with increasing distance due to self-absorption of the inserted wavelength shifter. Image taken from Ref. [50].

an exploded view of one fibre module with its key constituents. The core of one module consists of eight fibre mats that are arranged in two rows with four mats lying side by side in each row.

The width of each fibre mat is about 0.13 m and was chosen as a compromise between covering a large area and still being easy to handle. Following similar considerations, the length was chosen such that two mats are required to cover the full module length of 4.85 m. Even though this design introduces a small gap of about 2 mm at the centre of the module between the two rows of fibre mats, it bears the additional advantage that the light yield of a single particle crossing is concentrated into fewer SiPM channels on only one side of the module. Photons that are emitted towards the centre of the module are reflected back to the outside by a mirror foil that is glued to the fibre mat end.

The fibre mats are sandwiched between two half-panels that are made of a carbon fibre, honeycomb structure. The half-panels provide the necessary mechanical stiffness to the modules while keeping the material budget low.

Fibre Mats

The main components of the modules are the fibre mats. Each fibre mat consists of six layers of scintillating fibres as described in Section 3.4 and measures about 2.4 m in length and 0.13 m in width. Figure 3.6 shows the cross section of one fibre mat.

A fibre pitch of 275 μm is chosen to be 10 % larger than the nominal fibre diameter to allow for small production tolerances. Proper fibre positioning is achieved by




Figure 3.5: Exploded view of one fibre module of the SciFi Tracker. Image adapted from Ref. [52].




Figure 3.6: Cross section of one fibre mat used in the SciFi Tracker. Image taken from Ref. [50].

winding the fibre mats on a threaded wheel with a diameter of about 0.82 m [52]. While the scintillating fibres in the first layer are guided by the grooves in the winding wheel, fibres in subsequent layers fall into the gaps of the underlying layer. This results in the staggered arrangement of the layers as can be seen in Fig. 3.6.

Epoxy glue loaded with TiO_2 is added during the winding process to bond the fibres together, while also providing a thin protection coating on the surface. Additionally, the glue gets filled into 2 mm deep holes that are located in the winding wheel with a distance of about 245 mm to each other around the circumference of the wheel. After curing, the glue sticks to the fibre mat resulting in small pins that are used to align the mats with a precision of better than 100 μm in the x -direction during the module production [20]. A close-up view of one alignment pin is shown in Fig. 3.9 on page 21. The machine used for the serial production of the fibre mats is shown in Fig. 3.7.

A reliable winding of the fibre is enabled by applying a constant tension using a




Figure 3.7: Winding machine for the serial production of the fibre mats for the SciFi Tracker. The path of the wound fibre is highlighted in blue. Image taken from Ref. [52].

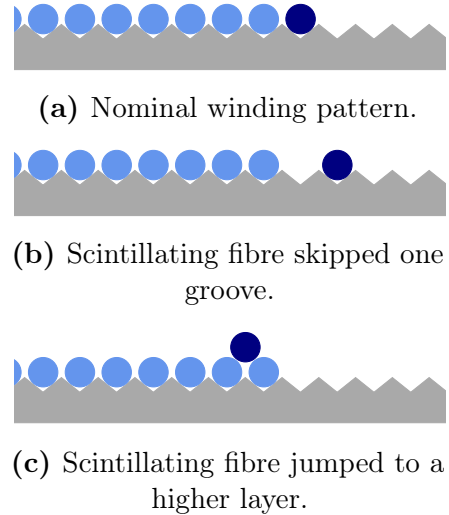


Figure 3.8: Nominal pattern (a) and different types of errors (b,c) that can occur during the winding of a fibre mat. Images adapted from Ref. [52].

dancer roller system. However, different types of errors can occur during the winding of the fibre mats and are illustrated in Fig. 3.8. The winding procedure is monitored with a monochromatic industrial camera that feeds its output to a Convolutional Neural Network (CNN) [53]. This approach allows for a semi-automatic winding procedure where manual intervention is only needed in case of detected defects [50].

After the actual winding is performed, a $25\ \mu\text{m}$ Kapton foil is applied to the accessible side of the fibre mat. The foil acts as an additional protection layer while also shielding the scintillating fibres from external light sources. Afterwards, the fibre mat is cut at a predefined location on the winding wheel perpendicular to the fibre direction and taken off the wheel. In the process, the fibre mat immediately shrinks in the order of 1 cm due to the constant tension applied during the winding. The removed mat is heated up to $40\ ^\circ\text{C}$ to flatten it. Afterwards, the Kapton foil is applied from the remaining side, leaving out small cutouts for the alignment pins.

At this stage, the fibre mat is still slightly larger than the final dimensions. The nominal length of 2424 mm [52] is defined by so-called end pieces made from plastic that are glued to both sides on both ends of the fibre mat. The excess parts of the mat are pre-cut using a saw blade and further processed using a diamond milling head. On one side of the mat, the SiPM arrays will be pressed against the smooth surface that is obtained in this way. On the other side, a mirror foil is applied to reflect back photons that are emitted away from the SiPMs in order to increase the overall light yield. An aluminized Mylar foil is used for that purpose which proved to reach a reflectivity in the order of 80 % [54].

All the steps listed previously were performed in four different locations: TU



Figure 3.9: Close-up view on a wound fibre mat with alignment pin. Image taken from Ref. [50].




Figure 3.10: Finished and laminated fibre mat with plastic end piece. Image taken from Ref. [50].

Dortmund, RWTH Aachen, EPFL in Lausanne and Kurchatov Institute in Moscow. A total of 1024 fibre mats that are needed for the complete detector plus spares were produced in these winding centres. A finished fibre mat is shown in Fig. 3.10 and is already cut to its nominal length. However, when leaving the winding centres, the width is still slightly larger than their nominal value and they feature excess Kapton foil to both edges. This allows for a safe transport to the module assembly centres, where the mats are embedded into the mechanical support structure of the fibre modules.

Mechanical Support Structure

The wound fibre mats that are produced in the winding centres are shipped to two different locations in order to get assembled to full-size detector modules. In these so-called module assembly centres, which are at Heidelberg University and Nikhef in Amsterdam, the fibre mats are embedded into a support structure that provides mechanical stiffness and protection to the modules. Before that can happen, the excess Kapton foil needs to be removed and the fibre mats need to be cut to their nominal width of 130.65 mm.

Afterwards, eight fibre mats are laid out in the final 2×4 arrangement as shown in Fig. 3.5 on page 19 onto a high precision template made from a single aluminium plate. The alignment pins of the mats are facing down and are guided by four longitudinal grooves in the template. The grooves have a centre-to-centre distance of 130.8 mm leaving a gap of 0.15–0.20 mm in-between two neighbouring fibre mats to account for tolerances of the longitudinal cut. The positioning follows the geometry of the SiPM arrays (see Section 3.6) such that no additional loss of acceptance is introduced by the gaps.

In the next step, glue is applied on the surface of the fibre mats that bonds the mats to the 2 cm thick support structure made from a light core material (Nomex

3 The Scintillating Fibre Tracker

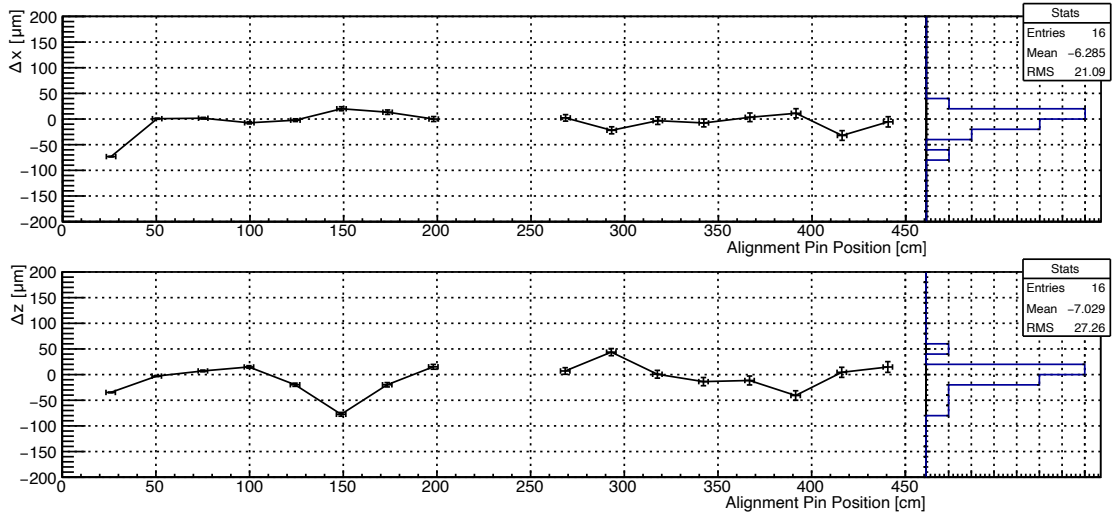


Figure 3.11: Evaluation of the straightness of two facing fibre mats within a fibre module using a test setup consisting of a laser and a beam camera. The deviations from a straight line along the alignment pins are shown in both the x (top) and z (bottom plot) direction within the LHCb coordinate system.

honeycomb). At the ends of the modules, the honeycomb is cut out to provide space for four individual pieces of aluminium, the so-called end plugs. The end plugs act as an anchor point to mount and align the fibre modules in the detector. Additionally, they house the light injection system (see Section 4.5) that is used for calibrating the individual channels. The outer shell of a fibre module consists of a carbon fibre reinforced polymer (CFRP) skin that is 200 μm thick and provides the necessary stiffness to the module.

The partially assembled module is turned over to the other side and placed onto another, flat aluminium template. In this position, the placement and straightness of the fibre mats within the module is verified using a test setup consisting of a laser and a beam camera: The laser is positioned on one side of the module pointing along the centre of two facing fibre mats. The beam camera is successively positioned on top of the total of 16 alignment pins on the two mats and the position of the beam projection is recorded. The camera is mounted on an aluminium block with the same kind of groove as in the aluminium template to allow for an accurate placement on top of the alignment pins. Like this, the straightness of the fibre mats in the module can be evaluated by the deviation of the recorded beam projections to a straight line in both the horizontal and vertical direction with respect to the setup on the assembly table, corresponding to the x and z direction within the LHCb coordinate system. Figure 3.11 shows the exemplary result of two facing fibre mats demonstrating a straightness of better than the required 100 μm along all pins.

After the straightness of the fibre mats is verified, the module is closed from the remaining side using another half-panel consisting of honeycomb core material and a CFRP skin. In order to provide light-tightness to the module, it is finished by




Figure 3.12: Cross section and side view of a fibre module of the SciFi Tracker. The dimensions shown in this image are given in mm. Image taken from Ref. [52].

closing the sides with CFRP sidewalls. Figure 3.12 shows the schematic side view and cross section of a finished fibre module. At this stage, the scintillating fibres are only exposed to the outside between the plastic end pieces of the fibre mats on both sides of the module. They will later be covered by the SiPM arrays (see Section 3.6).

Special modules need to be produced to account for the beam pipe as indicated in Fig. 3.1 on page 15. They feature a rectangular $13\text{ cm} \times 11.5\text{ cm}$ cutout on one side of the module. To allow for this, the two fibre mats on this side of the module are produced with a reduced length. The originally intended circular cutout was found to be difficult to realise due to the mirror foil that needs to be applied to the fibre mat ends [55].

The material budget of the individual components of a SciFi Tracker module are listed in Table 3.1. The aluminium end plugs are not considered because they lie outside the acceptance of the experiment. The total radiation length for one module is $0.984\%X_0$ which translates to $3.936\%X_0$ for one tracking station consisting of four detection layers. This value fulfills the required limit on the material budget as listed in Section 3.2. While the Nomex honeycomb structure accounts for 95% of the module's thickness, it only contributes about 31% to the total radiation length. Note that the total thickness given in Table 3.1 is slightly larger than the nominal value of 41.6 mm due to submersion of the glue into the honeycomb [52].

In total, 128 fibre modules needed for the detector plus spares were produced at the two module assembly centres between 2016 and 2019. Of these, about 20% are special modules with a rectangular cutout on one side to account for the beam pipe in the centre of each of the 12 detection layers.

Table 3.1: Material budget for a scintillating fibre module. Data taken from [52], [56].

Material	Thickness Δz [μm]	Radiation length X_0 [cm]	$\Delta z/X_0$ [%]
Scintillating fibres	1350	33.2	0.407
Kapton foil	2×25	28.6	0.017
Panel assembly glue	4×75	36.1	0.083
Nomex honeycomb	$2 \times 20\,000$	1310.0	0.305
Carbon fibre skin	2×200	23.3	0.172
Total	42 100		0.984

3.6 Silicon Photomultipliers

Silicon photomultipliers (SiPMs) are solid state devices that allow for the detection of individual photons. Combined with their small dimensions and high granularity they meet the needs for a high resolution tracker like the SciFi [20].

Operating Principle

The functionality of solid state photon detectors is based on a p - n junction, which refers to the boundary between n - and p -doped regions inside a semiconductor material. Close to the boundary, electrons from the n -doped region diffuse into the p -doped region and, vice versa, electron holes from the p -doped region diffuse into the n -doped region. By recombination of electrons and holes in the area around the boundary, the so-called depletion zone, the amount of free charge carriers is reduced significantly resulting in a low electrical conductivity. The diffusion continues until the electrical potential that is generated in the process counteracts the motion of the charge carriers.

Forward Bias By applying a positive external voltage to the p -side with respect to the n -side (forward bias) the diffusion process is reversed. When the bias voltage is large enough, the depletion zone becomes very thin and the p - n junction becomes conductive allowing for a (forward) current flow.

Reverse Bias Applying a negative external voltage to the p -side with respect to the n -side (reverse bias) has the opposite effect: The depletion zone expands and the electrical conductivity stays low. Only a small (reverse) current can flow.

A photodiode makes use of operating a p - n junction with a reverse bias voltage applied. Photons with sufficient energy can create electron-hole pairs in the depletion zone via the photoelectric effect. A free electron that is generated by this process is

called (primary) photoelectron (pe). Since the resulting current flow only consists of individual charge carriers, photodiodes are not suitable for low intensity applications like the SciFi Tracker.

Avalanche photodiodes (APD) overcome this issue by operating at higher voltages. The thereby increased electrical field at the p - n junction accelerates the primary photoelectrons up to energies at which they can create electron-hole pairs themselves. This process can repeat itself several times causing an avalanche of charge carriers that is measurable as an electrical current.

When the applied voltage is increased even further, beyond the breakdown voltage V_{BD} , a single primary photoelectron can cause a self sustaining avalanche (Geiger mode). The released charge is independent of the number of primary photoelectrons and results in a permanent current if unquenched. In practice, the avalanche is typically interrupted by inserting a resistor in series to the photodiode. As the current flow increases during the emerging avalanche, the applied voltage is lowered below the breakdown voltage by the so-called quench resistor R_Q such that no more charge carriers can be created.

SiPMs are composed of an array of APD pixels that are operated in Geiger mode (G-APD). They are connected in parallel such that the resulting signal is proportional to the number of pixels with a triggered avalanche. If the probability for multiple photons hitting the same pixel is low, SiPMs allow for the counting of individual photons. This is given when the number of pixels is much greater than the number of incident photons. The amplitude of a signal is typically stated in units of generated primary photoelectrons (X pe), which is equivalent to the number of involved pixels in low light conditions.

Correlated and uncorrelated Noise

SiPMs are subject to different types of noise that affect the operation. One distinguishes between correlated and uncorrelated noise.

Correlated Noise The emergence of correlated noise is, as the name suggests, correlated to the occurrence of an initial signal. A large contribution comes from crosstalk. Crosstalk originates from infrared photons that are emitted in a pixel during an avalanche. These photons can reach neighbouring pixels causing another discharge. Depending on the region where the photon is absorbed, the resulting pulse either happens instantaneously (**direct crosstalk**, within a few 100 ps) or with a significant delay up to 100 ns (**delayed crosstalk**) with respect to the initial signal [57].

Delayed crosstalk must not be confused with **afterpulses**, which have a different formation mechanism. Afterpulses occur due to defects in the silicon lattice that can act as charge traps. A charge carrier created during an avalanche can get caught in these traps. When it is released at a later time, it can trigger another discharge in the same pixel causing a second, delayed pulse distributed over a few 100 ns [57].




Figure 3.13: Electrical circuit of a basic equivalent model for an externally biased G-APD (left). The plot on the right shows the signal pulse after closing (at t_0) and re-opening (at t_{\max}) the switch S . Image inspired from Ref. [58].

Uncorrelated Noise Instead of being excited by an incident photon, an electron-hole pair in the depletion zone can also be generated by thermal excitation. The resulting signal is indistinguishable from a photon induced avalanche. This type of noise can occur randomly and is independent of any initial signal. Because of that, it is also referred to as **dark counts**. Due to the thermal origin, the dark count rate (DCR) strongly depends on temperature. In addition to that, the DCR increases after irradiation.

Various combinations of the different types of noise can occur. For example, it is not uncommon that a dark count arises in conjunction with a crosstalk event. The amplitude of the resulting signal is thereby increased making it more difficult to be identified as noise.

Basic electrical Model

Following Ref. [58], a basic electrical model of a G-APD can be developed to derive the key characteristics of the signal pulse. It consists of the breakdown voltage V_{BD} , capacity C_J and ohmic resistance R_S of the G-APD, as well as the quench resistor R_Q and bias voltage V_{BIAS} . A conceptual switch S is responsible for initiating (S is closed) and terminating (S is open) the avalanche.

Figure 3.13 (left) shows the electrical circuit of the model. In the absence of incident photons and ignoring dark noise, the switch S is open. The capacitor C_J is fully charged by the applied bias voltage V_{BIAS} . No current is flowing. Then, representative for an incident photon triggering an avalanche, the switch S is closed at time t_0 . The capacitor C_J begins discharging through the series resistor R_S as a flowing current

$$I_S(t) \sim 1 - \exp\left(-\frac{t}{R_S C_J}\right) \equiv 1 - \exp\left(-\frac{t}{\tau_S}\right) \quad (3.2)$$

with a time constant τ_S . At t_{\max} , which is in the order of 1 ns after t_0 [58], the maximum current

$$I_{\max} = \frac{V_{\text{BIAS}} - V_{\text{BD}}}{R_Q + R_S} \equiv \frac{\Delta V}{R_Q + R_S} \quad (3.3)$$

is reached. The introduced variable ΔV is called overvoltage. At this time, the applied voltage to the G-APD drops down to approximately V_{BD} , which is not sufficient to sustain the avalanche. The switch S opens and the capacitor C_J is recharged by a current

$$I_Q(t) \sim \exp\left(-\frac{t}{R_Q C_J}\right) \equiv \exp\left(-\frac{t}{\tau_Q}\right) \quad (3.4)$$

with a time constant τ_Q , which is also referred to as the recovery time. The shape of the resulting signal is shown in Fig. 3.13 (right). For $R_Q \gg R_S$ and thereby $\tau_Q \gg \tau_S$, the total charge Q , i.e. area under the signal shape, can be derived as

$$Q = I_{\max} \cdot \tau_Q = \frac{\Delta V}{R_Q + R_S} \cdot R_Q C_J = \Delta V \cdot C_J. \quad (3.5)$$

Since the avalanche is triggered by a single primary electron with elementary charge e , the charge amplification factor – or gain G – is given by

$$G = \frac{Q}{e} = \frac{\Delta V \cdot C_J}{e}. \quad (3.6)$$

Typical operation values for the gain of SiPMs lie between 10^5 and 10^7 [58].

Application at the SciFi Tracker

The SciFi Tracker utilises a total of 524 288 SiPM channels that are grouped in 128-channel arrays. The multichannel arrays referred to as H2017 are custom-made and produced by Hamamatsu¹. One array is 32.54 mm wide and consists of two dies with 64 channels each. The silicon dies are protected with a thin 105 μm epoxy layer [59]. After mounting in the detector, the insensitive area between two adjacent SiPM arrays amounts to 480 μm , while the specifications state the gap between the two dies to be $(220 \pm 50) \mu\text{m}$ [60]. The geometry was chosen such that the width of multiple arrays match the fibre mat and module dimensions. One mat is instrumented with four SiPM arrays, which translates to 16 arrays for a fibre module on each end. Figure 3.14 shows a photo of a H2017 package with a zoomed in view on the individual SiPM channels. The package is mounted on a Kapton flex PCB (Printed Circuit Board) that is about 12 cm long. Two 80-pins, 0.5 mm pitch connectors² are placed on the other end of the flex PCB to provide the SiPM

¹Hamamatsu Photonics K.K., 325-6, Sunayama-cho, Naka-ku, Hamamatsu City, Shizuoka Pref., 430-8587, Japan

²DF12(3.0)-80DP-0.5V by Hirose Electric Co., Ltd.




Figure 3.14: Hamamatsu H2017 128-channel SiPM array with flex cable and connector as used for the SciFi Tracker (left). A zoomed in view on the active area with the gap between the two dies is shown on the right, as well as an illustration of the position of one SiPM channel with respect to the fibre mat. Image adapted from Ref. [59].

connections to the front-end electronics (see Section 3.7).

The SiPM channel pitch is 250 μ m thus matching the diameter of the scintillating fibres. The height is with 1625 μ m about 20 % larger than the fibre mat to allow for slight misalignments during the assembly. One channel is composed of 4×26 pixels each covering an area of 57.5 μ m \times 62.5 μ m. As shown in Fig. 3.15, the large pixel size allows for high photon detection efficiencies up to 44 % at the nominal overvoltage $\Delta V = 3.5$ V [57] for the relevant wavelengths from 400 nm to 600 nm (compare with Fig. 3.4 on page 18). Additionally, it enables the application of isolating trenches to suppress crosstalk. Figure 3.16 shows the crosstalk probability for different overvoltages ΔV , as well as the probabilities for other types of correlated noise. At the nominal overvoltage $\Delta V = 3.5$ V, the total probability for correlated noise amounts to about 7%. It is split equally between the two types of crosstalk: direct and delayed. Afterpulses practically do not play a role for H2017 arrays.




Figure 3.15: Measured photon detection efficiencies (PDE) based on a pulse counting approach for an H2017 SiPM array for different wavelengths λ and overvoltages ΔV . Image taken from Ref. [57].

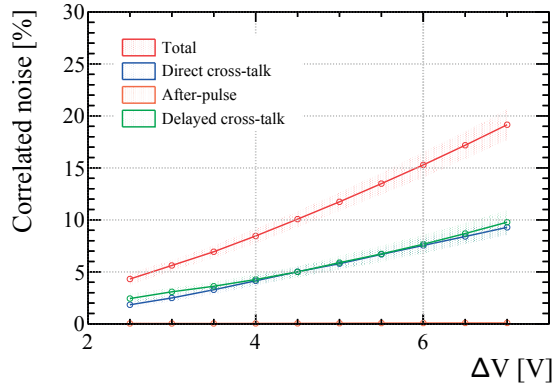


Figure 3.16: Probabilities for different types of correlated noise as determined during the quality assurance of 5000 H2017 SiPMs. Variations in the order of $\pm 1\%$ were observed between the different devices. Image taken from Ref. [61].

The typical gain for an H2017 SiPM array is $G = 4 \cdot 10^6$ at the nominal overvoltage. Figure 3.17 shows the dependence of the gain on the overvoltage for one channel. As derived in Eq. (3.6) on page 27, it follows a linear relation. The recovery time is found to be $\tau_Q = (84.6 \pm 0.2)$ ns. The breakdown voltage V_{BD} ranges from 51.0 V to 52.5 V as determined among 20 devices from the production batch. Within one array, it typically varies by ± 300 mV [59]. When operated at the same voltage around the nominal overvoltage $\Delta V = 3.5$ V, this results in a gain uniformity of better than 10% as required in the TDR [20].

A key parameter for the operation of the SciFi Tracker is the rate at which dark counts occur (DCR), as it is main source of noise for the detector. As described previously, the DCR strongly depends on both temperature and irradiation dose, specifically the 1 MeV neutron-equivalent fluence (n_{eq}/cm^2). Towards the end of the lifetime of the upgraded detector, the expected fluence that the SiPMs will have received is $6 \cdot 10^{11} n_{eq}/\text{cm}^2$ [20]. At that time, the DCR will reach values of several hundred MHz at room temperature, while the DCR of an unirradiated SiPM is in the order of kHz. This drastic increase resulting from lattice deformations is mitigated by cooling of the SiPM arrays. As shown in Fig. 3.18, the DCR can be reduced by 50% with every decrease in temperature by about 10°C . Thereby, at the end of LHC Run 4, manageable DCRs in the order of 10 MHz can be reached at -40°C . Reaching low DCRs of irradiated SiPMs at -40°C was a main criteria for the selection of the best suited technology during the R&D phase of the SciFi Tracker.

The temperature of each array can be monitored with the help of a Pt1000 temperature sensor¹ that is mounted on the backside of the SiPM array. The specific

¹Pt1000 SMD 0603 Class B by Heraeus Nexensos GmbH




Figure 3.17: Dependence of the gain on the overvoltage for one H2017 SiPM channel. The red line shows the result of a linear fit to the data. Image taken from Ref. [57].




Figure 3.18: Dark count rate (DCR) as a function of the temperature for irradiated H2017 SiPMs at $\Delta V = 3.5$ V. The coefficient $T_{1/2}$ that describes the temperature difference at which the DCR is reduced by 1/2 is determined with an exponential fit in the range $[-40, -10]$ °C and indicated by the solid lines. Image adapted from Ref. [61].

sensor was chosen to cover the temperature range of interest between -40 °C and $+30$ °C with a precision of ± 1 °C [59].

3.7 Front-End Electronics

The front-end electronics of the SciFi Tracker are located just outside of the geometrical acceptance on both ends of the fibre modules. They provide the interface between the SiPMs and the control and data-acquisition (DAQ) system. The connection to the SiPMs is established via flex cables as shown in Fig. 3.14 on page 28, while the communication with the back-end electronics happens over optical fibres. The optical connections are divided into unidirectional links for the transmission of the data, and bidirectional control links.

The front-end electronics are grouped into so-called Readout Boxes (ROBs), which are the physical units that are mounted to the fibre modules. In total, the front-end electronics of the SciFi Tracker consists of 256 ROBs with more than half a million channels. A photograph of one ROB with the top cover removed is shown in Fig. 3.19. It consists of two identical halves, which are referred to as HalfROBs. Besides from sharing a common housing and cooling block, the HalfROBs are fully independent from each other in terms of power and control.

The front-end electronics follow a modular design. Each HalfROB consists of one Master Board, four Cluster Boards, and four PACIFIC Boards. The latter are

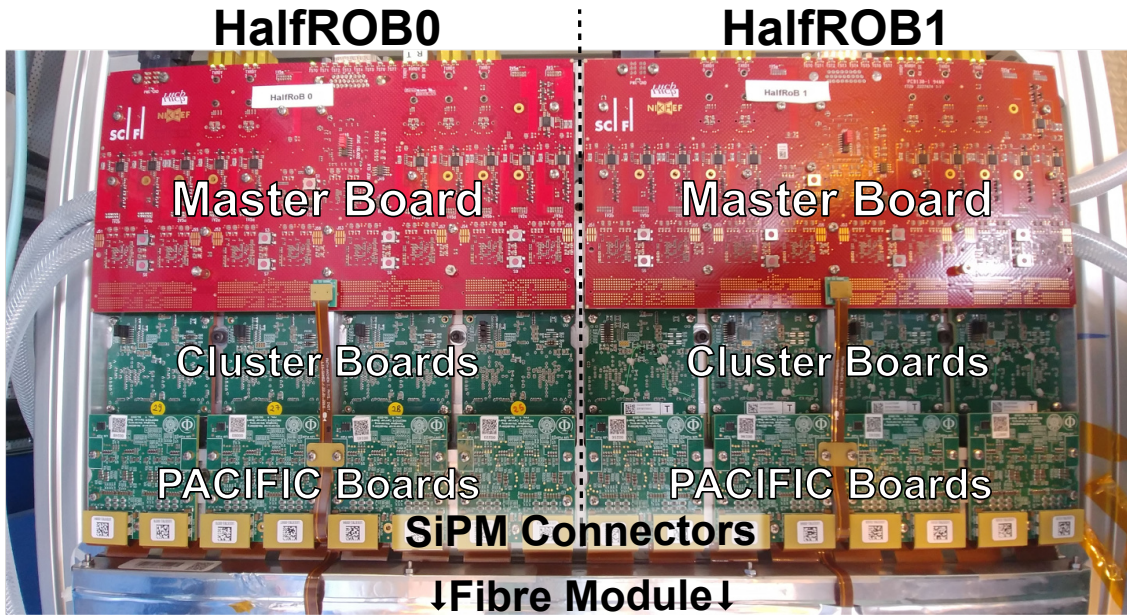


Figure 3.19: Photograph of a SciFi Readout Box (ROB) with the top cover removed. The connection to the SiPM arrays is made via 16 connectors located at the bottom.

the direct interfaces to the SiPMs. Each PACIFIC Board houses four custom 64-channel ASICs (Application Specific Integrated Circuits) that perform the analogue processing and digitisation of the SiPM signals. In the next step, the Cluster Boards perform a hit reconstruction and noise suppression on the digital pattern provided by the PACIFIC Boards with the help of two on-board FPGAs (Field Programmable Gate Arrays). Lastly, the clustered data is encoded and shipped to the DAQ servers via optical transmitters located on the Master Board. In addition, the Master Board is responsible for distributing power, clocks and control commands among all boards within one HalfROB.

A detailed description of the various components of the SciFi front-end electronics is given in Chapter 4.

3.8 Infrastructure

The operation of the SciFi Tracker is only enabled with the help of a reliable infrastructure. It includes the provision with electrical power and cooling for the various components. In addition, a mechanical support structure is required that holds the detector modules in place and routes the supply lines to the desired locations. The different parts constituting the infrastructure are described in the following.

3.8.1 Electrical Power

Two different types of power supplies are used for the SciFi Tracker. They can be distinguished by the voltage levels at which they are operated.

Low Voltage

The front-end electronics are powered by MARATON power supplies manufactured by W-IE-NE-R¹. These devices are specifically developed for the operation at the LHC in a magnetic and radiation environment. This is also reflected by the acronym MARATON that stands for **M**agnetism **R**adiation **T**olerant **N**ew power supply system. Each unit offers 12 independent low voltage/high current channels that provide up to 300 W each [62].

The MARATON system consists of three main components: the power box, a primary rectifier, and a remote controller. The primary rectifier converts the 230 V AC (alternating current) mains voltage to a regulated 385 V DC (direct current) voltage that is fed to the power box. In the MARATON power box, DC-DC converters are used to provide up to 12 low voltage floating output voltages. With the help of the remote controller, each channel can be switched off or on independently. In addition, it allows for the monitoring of the output voltages and currents. The adjustment of the output voltage for each channel is possible via a potentiometer that is accessible with a screwdriver at the front of the power box and that can be set up to 8 V. In the same way, the overvoltage protection (OVP) and current limits can be modified.

At the SciFi Tracker, each MARATON channel delivers power for up to two ROBs. To limit the voltage drop in the cables due to high currents, the power box is located near the detector in a harsh environment with high radiation and stray magnetic fields. The connection to the primary rectifier is established with an Amphenol ECTA 133 circular connector, while two 37-pin D-SUB connectors are used for the communication with the remote controller. The system is designed such that the primary rectifier and remote controller are operated in a safe environment under standard industrial conditions up to 120 m apart from the power box. The remote controller is integrated into the LHCb network by a TCP/IP connection. Utilising the Open Platform Communications (OPC) standard, this allows for the control and monitoring of the system within the aforementioned limits from outside the LHCb cavern.

High Voltage

The SiPMs are biased by power supplies manufactured by CAEN². The modules of type A1539BP [63] feature 32 channels that are accessible through two DB25 connectors. The output voltage of each channel can be set between 0 V and 100 V

¹W-IE-NE-R Power Electronics GmbH, Burscheid, Germany

²CAEN S.p.A., Viareggio, Italy

and deliver up to 20 mA. Even though this voltage range is not usually referred to as high voltage (HV), the term will be used throughout this thesis as counterpart to the 8 V LV as provided by the MARATON power supplies. The accuracy of the set voltage is stated as $\pm 0.05\% \pm 50$ mV. The ramp-up and ramp-down rates can be individually adjusted for each channel.

The boards are equipped with both current and voltage protections. The maximum output voltage can be set on a hardware level through a potentiometer that is accessible on the front panel. It applies to all 32 channels. Further restrictions on current and voltage limitations are adjustable in software on a per-channel basis. In addition, it allows for the modification of the trip time after which a channel is turned off when the current limit is exceeded.

At the SciFi Tracker, each HV channel biases four SiPM arrays corresponding to one fibre mat. Intermediary switch panels allow to disable individual SiPM dies to not interfere with neighbouring dies in case of defects. The HV modules are operated in CAEN crates of type SY4527 [64]. Besides from housing up to 16 boards, each crate contains a fan tray section with six fans, and a 600 W primary power supply unit (PSU). Depending on the load, up to three additional booster PSUs can be installed. The crates used for the SciFi Tracker will be expanded with one additional 1200 W booster PSU of type A4533, allowing for the operation of ten A1539BP modules at maximum load. For controlling the HV remotely, each crate includes a CPU and Gigabit Ethernet interface that allows for the integration into the LHCb network.

3.8.2 Cooling

Two different types of cooling circuits are implemented in the SciFi Tracker.

Water Cooling

During operation, the front-end electronics heat up significantly. To limit the temperature to a maximum of 50 °C, a water cooling scheme inspired by the OT is implemented [20]. To allow for an efficient heat transfer, the ROBs are tightly screwed on aluminium cooling blocks. The cooling water flows through a straight copper pipe that traverses the cooling blocks of 5-6 neighbouring ROBs depending on the tracking station. A 180° bend behind the last ROB guides the water to the return pipe that traverses the cooling blocks in reverse order, thereby also contributing to the cooling process.

Based on a heat dissipation close to 100 W per ROB, it was estimated that a flow rate of 2.5 l/min per row of 5-6 ROBs is required. At this rate, the water temperature is expected to increase by 3 °C. It is foreseen to use chilled and demineralised water with a temperature of 16 °C at the inlet [65].

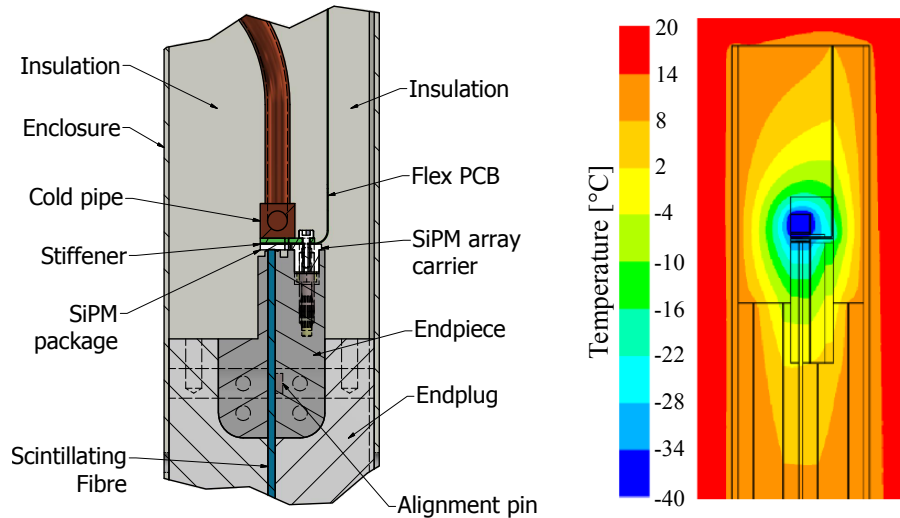


Figure 3.20: Cross-sectional view of a SciFi cold box with its components (left), as well as a thermal simulation (right). The SiPM signals are transmitted through the flex PCB to the front-end electronics located above. Images adapted from [20].

Novec Cooling

As outlined in Section 3.6, the SiPMs need to be cooled to -40°C to mitigate the drastic increase of the dark count rate due to radiation damages. A complex infrastructure is needed to achieve the required cooling performance. The coolant used for that purpose is Novec 649 distributed by 3M¹. The fluid has similar thermo-physical properties as Perfluorohexane C_6F_{14} that has been widely used in previous low-temperature cooling applications at the LHC. Comparable with C_6F_{14} , it is non-combustible, volatile, dielectric and practically non-toxic [20]. However, due to its much lower global warming potential (GWP) around 1, Novec has been chosen as the superior alternative.

The coolant is circulated through a titanium cooling bar to which the SiPMs are glued. The structure is surrounded by an insulating box, which is referred to as cold box. A cross-sectional view of a cold box is shown in Fig. 3.20. Each box contains sixteen SiPMs corresponding to one fibre module and ROB. A pair of inlet and outlet bellows allow for a flexible connection of the individual cold boxes to the cooling circuit. The Novec supply lines as well as the bellows are vacuum insulated.

To avoid condensation and ice formation inside the cold box due to the low temperatures, a system is implemented that constantly flushes the boxes with a dry gas (either dried air or nitrogen). During the commissioning of the first detector parts (see Chapter 5), it was found that condensation can also occur towards the outside. Therefore, a heating system has been developed that covers the surface of the cold boxes, as well as the Novec bellows and dry gas outlets. It consists of Kapton isolated heating wires that are twisted around the mentioned parts. Same as

¹3M, Saint Paul, Minnesota, USA




Figure 3.21: Schematic illustration of a C-Frame of type SX that carries the detector modules, front-end electronics, and the various infrastructure components. The SciFi Tracker consists of 12 of these frames. Image adapted from Ref. [66].

for the front-end electronics, the required heating power is provided by MARATON power supplies.

3.8.3 Mechanical Support Structure

The individual detector modules are carried by C-shaped mounting frames. The complete detector consists of a total of 12 of these so-called C-Frames that are arranged on both sides of the beam pipe. They are mounted on guide rails, which facilitates the installation and also allows for quick access in case of required maintenances.

One C-Frame carries two layers with 5-6 fibre modules each depending on the associated tracking station. While the modules in one layer are oriented vertically (i.e. in y -direction within the LHCb coordinate system), the modules in the second layer are tilted by 5° – the so-called stereo layer. Having two different types of

C-Frames is sufficient to form the typical $x-u-v-x$ geometry (see Section 3.1) in each station on both sides of the beam pipe. They are denoted as C-Frames of type XS and SX, depending on the order of the two layers when having the vertical bar of the frame on the right-hand side.

In addition to the modules, the C-Frames also provide the framework for the various components of the detector infrastructure as described before. They guide the low and high voltage cables to the ROBs, as well as the pipes for the water and Novec cooling. Moreover, cable trays are mounted on the frames to allow for the installation of the optical cables for data transmission and control of the front-end electronics. A schematic representation of a C-Frame with its components is shown in Fig. 3.21.

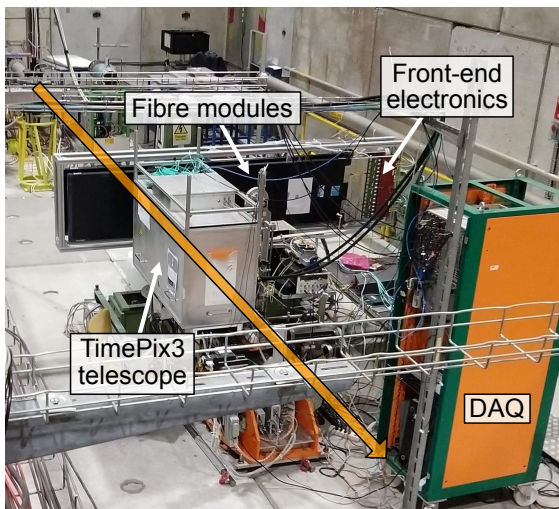
3.9 Tracking Performance

During the preparation of this thesis and throughout the development phase of the SciFi Tracker, multiple test beam campaigns have been conducted in order to evaluate its performance, in particular the tracking capabilities. The most recent one took place during two weeks in July 2018 at the CERN Super Proton Synchrotron (SPS) North Area [67]. The beam in the H8 beamline was set to a width of about 1 cm and provided pions and muons with momenta around 180 GeV/s.


The experimental setup as shown in Fig. 3.22a consisted of two full-width half-length fibre modules with the mirror on one end and 16 SiPM arrays on the other. The SiPM signals were read out by two ROBs that transferred the data via optical links to a mini version of the upgraded LHCb DAQ system [68]. Reference tracks used for the determination of the tracking performance parameters were provided with an uncertainty of approximately 8 μm by a TimePix3 telescope located downstream of the modules [69].

The campaign was a great success as it could be shown that a single hit efficiency beyond 99% can be reached. In addition, as shown in Fig. 3.22b, the resolution in the bending direction (x -axis) within the LHCb coordinate system was measured to be about 65 μm . Thus it could be verified that the corresponding performance requirements as listed in Section 3.2 are fulfilled.

This test beam was the first one during which the full chain of front-end electronics along with the new DAQ system was tested. In addition, all SciFi related components in use were (near-)production versions. Using this setup, it was possible for the first time to read out the 32 optical links at the LHC clock frequency of 40 MHz resulting in a total data rate of about 150 Gbit/s. Getting the two systems to work with each other was not a trivial task and required for extensive preparations, which have been a significant part in the course of this thesis. The software that had to be developed for it also formed the basis for the commissioning of the front-end electronics. More details on this follow in Chapter 5.



(a) Experimental setup



(b) Resolution

Figure 3.22: Test beam setup consisting of two fibre modules along with the corresponding front-end electronics (a) and measured hit resolution in the bending direction (x -axis) within the LHCb coordinate system (b). The particle beam provided in the SPS North Area at CERN is indicated by the orange arrow.

4 Front-End Electronics of the SciFi Tracker

The front-end electronics of the SciFi Tracker follow a modular design. Each HalfROB consists of one Master Board, four Cluster Boards, and four PACIFIC Boards. The latter are the direct interfaces to the SiPMs. Each PACIFIC Board houses four custom 64-channel ASICs (Application Specific Integrated Circuits) that perform the analogue processing and digitisation of the SiPM signals. In the next step, the Cluster Boards perform a hit clustering and noise suppression on the digital pattern provided by the PACIFIC Boards with the help of two on-board FPGAs (Field Programmable Gate Arrays). Lastly, the clustered data is encoded and shipped to the DAQ servers via optical transmitters located on the Master Board. In addition, the Master Board is responsible for distributing power, clocks and control commands among all boards within one HalfROB. Figure 4.1 shows a block diagram of one HalfROB with the indicated data stream from the SiPMs to the optical transmitters.

Following the path of the data, the various components of the SciFi front-end electronics will be discussed in the following sections.

4.1 PACIFIC Board

The PACIFIC Board (PB) is a 14-layer PCB with a nominal thickness of 1.6 mm that measures 65 mm in length and 60 mm in width [71]. It serves as the carrier board for four PACIFIC ASICs, which are responsible for the analogue processing and digitisation of the SiPM signals. Further information on the PACIFIC chip follow in Section 4.2. The layout of both sides of the PACIFIC Board is shown in Fig. 4.2 on page 41.

Besides from housing the PACIFIC ASICs, four 80-pins, 0.5 mm pitch connectors¹ are located on the board. They are used to connect two SiPM flex cables, which hold two female counterparts of these connectors on each cable (see Fig. 3.14 on page 28). On each connector, 64 pins are used for the 64 channels of one SiPM die. The remaining pins are connected to the bias voltage and ground. In addition, for each connector pair, one line goes to the Pt1000 temperature sensor located on the backside of the SiPM arrays. Thereby, a total of 256 SiPM channels are operated by one PACIFIC Board.

The analogue SiPM signals are directed to the PACIFIC ASICs. From there, the digital output is forwarded to a high-speed, 400-pin FMC (FPGA Mezzanine Card)

¹DF12(3.0)-80DS-0.5V by Hirose Electric Co., Ltd.

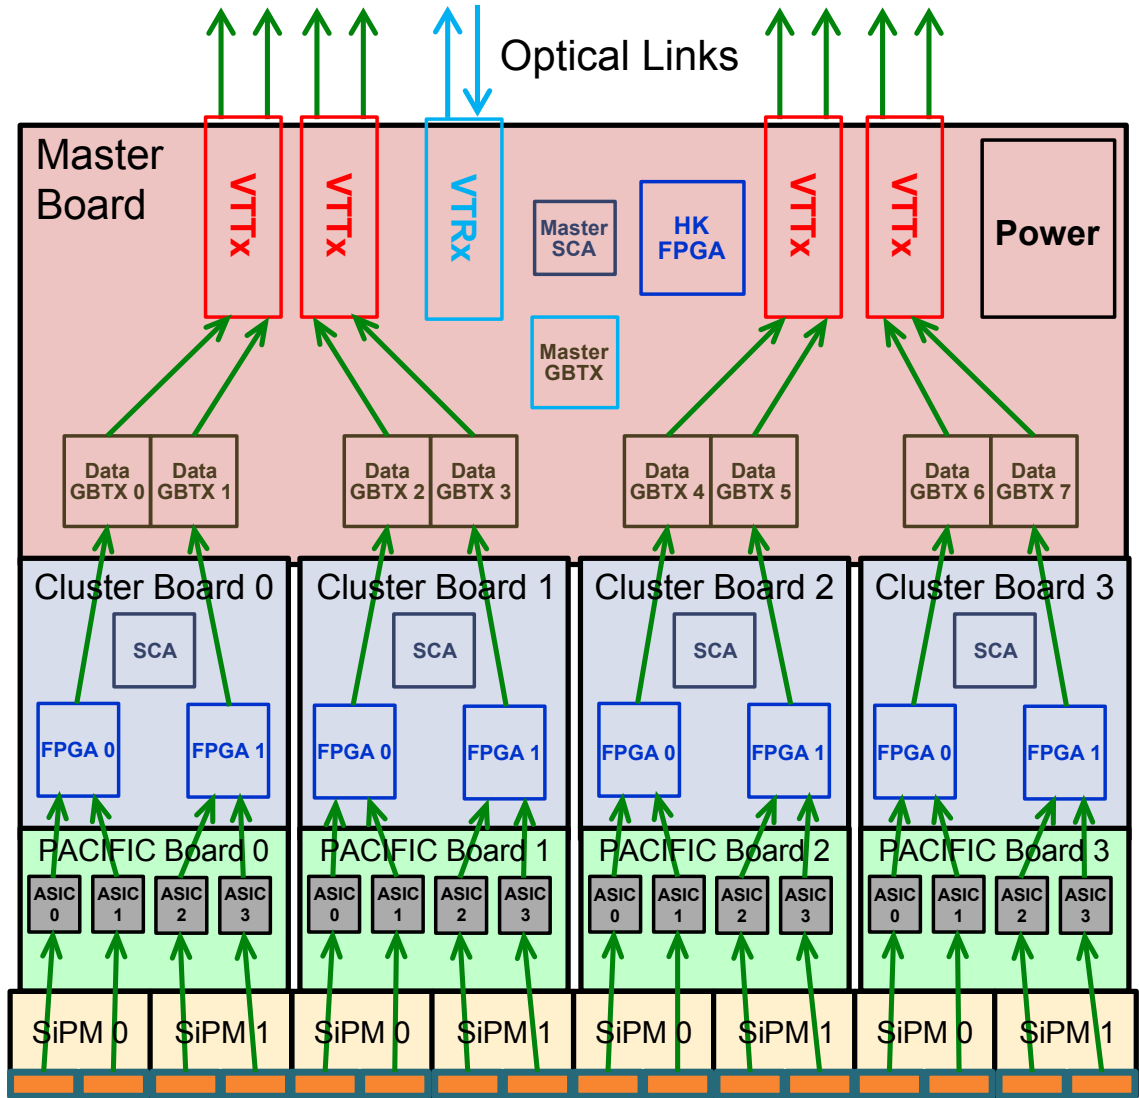


Figure 4.1: Block diagram of a HalfROB with the main components on each board denoted. The data stream is indicated by the green arrows. The blue arrows illustrate the bidirectional optical control link. Image edited from Ref. [70].

connector¹ that allows for the connection to the Cluster Board.

Besides from being responsible for the routing of the signals, the PACIFIC Board contains several monitoring circuits. Two Pt1000 temperature sensors² are located between the four ASICs and the FMC connector on the backside of the board. They are denoted as R10 and R21 on the right side of Fig. 4.2. The readout is performed by an ADC (Analogue-to-Digital Converter) located on the Cluster Board (see Section 4.3). A monitoring of the bias voltages applied to the SiPM dies is enabled

¹ASP-134602-01 (HPC) by Samtec, Inc.

²Pt1000 SMD 0805 Class B by Heraeus Nexensos GmbH

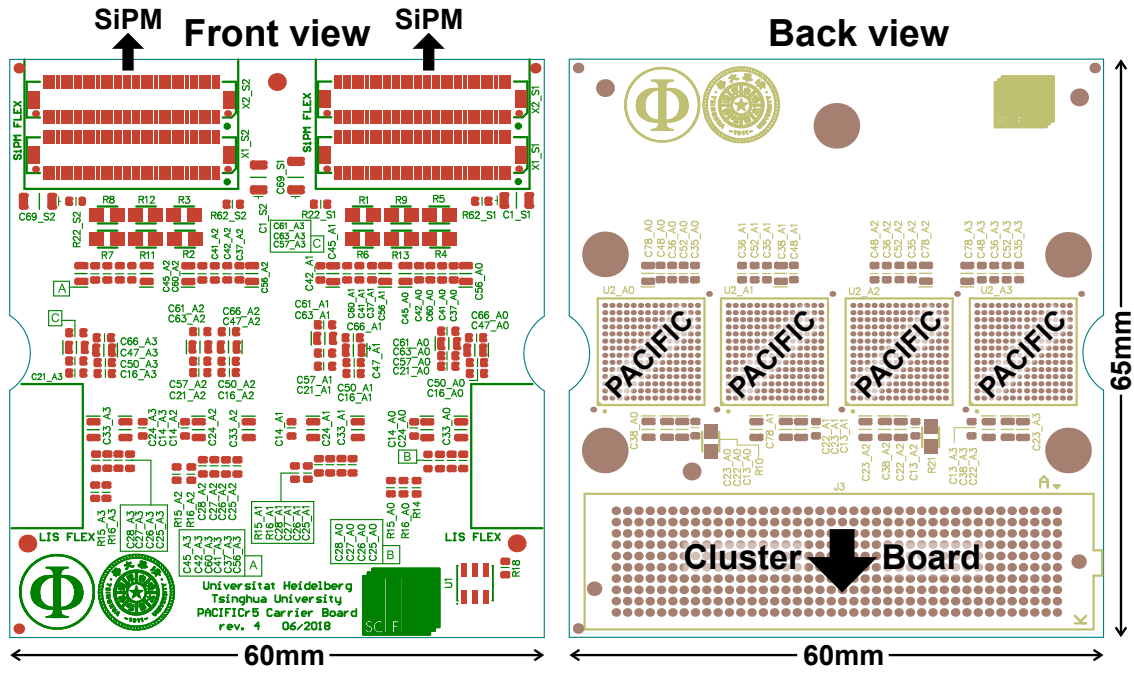


Figure 4.2: Layout of the front (left) and back (right) of a PACIFIC Board. The connectors to the SiPM flex cables are shown on the top left. On the bottom right, the FMC connection points to the Cluster Board are displayed. The four PACIFIC ASICs are soldered on the square areas on the right. Image edited from Ref. [71].

by four voltage divider circuits. The involved resistors are located next to the SiPM connectors and are denoted as R1-R9 and R11-R13 in Fig. 4.2.

Each PACIFIC Board is assigned a unique 48-bit serial number that is stored on a 1-Wire serial ID device¹.

4.2 PACIFIC ASIC

The PACIFIC is a low-power ASIC specifically designed for the SciFi Tracker. This is also reflected by the acronym PACIFIC that stands for low-power AASIC for the Sciⁿtillating Fibre Tracker. The chip allows for the readout of 64 SiPM channels and was developed using TSMC's² 130 nm CMOS technology. It is designed to be radiation tolerant and operate at the LHC clock frequency of 40 MHz with a power consumption of less than 10 mW per channel [72].

As shown in Fig. 4.3, the PACIFIC consists of various logical blocks that are presented in the following.

¹DS2401P+ by Maxim Integrated

²Taiwan Semiconductor Manufacturing Company, Limited, Hsinchu, Taiwan

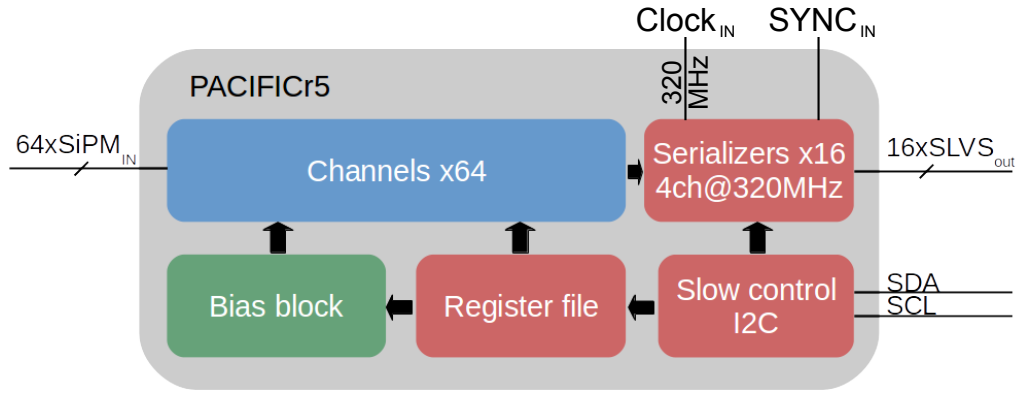


Figure 4.3: Logical structure of the PACIFIC ASIC. Image adapted from Ref. [73].

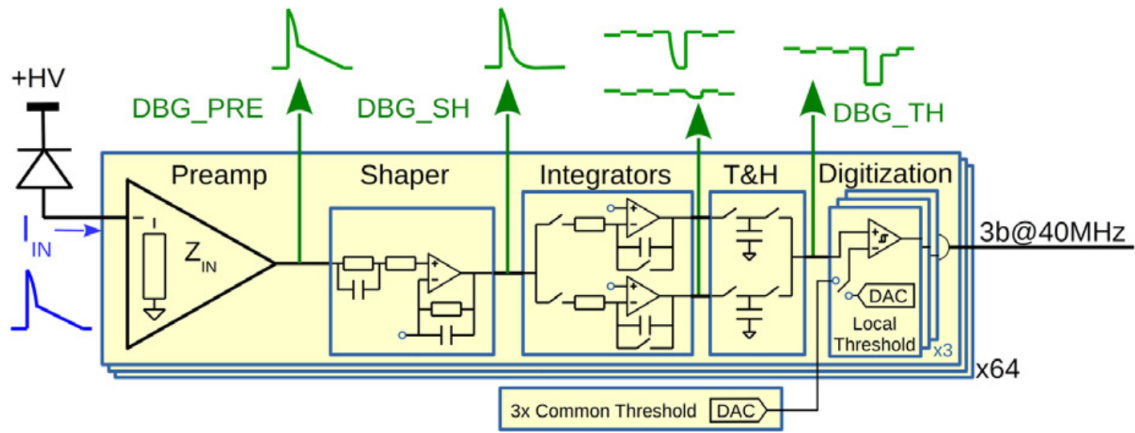


Figure 4.4: Channel block diagram of the PACIFIC ASIC. Each channel consists of five stages in which the analogue SiPM signal is processed and digitised. Image taken from Ref. [72].

4.2.1 Signal Processing and Digitisation

A block diagram of one channel of the PACIFIC ASIC is shown in Fig. 4.4. The analogue processing and digitisation of the SiPM signal can be divided into five successive stages that are discussed in the following.

Pre-amplification

The first stage of the signal processing chain is the pre-amplification, or input stage. It is directly connected to the anode of the SiPM and provides a current mode input with a low impedance and high bandwidth of about $50\ \Omega$ and 250 MHz, respectively [74].

The input stage consists of a current conveyer based on a novel approach using two feedback loops [75]. It allows for the selection of four different gains, as well as the possibility to modify the bias voltage of each individual channel over a range




Figure 4.5: Anode voltage generation circuit within the PACIFIC. Each channel can be connected independently to the output voltages V_0 to V_{15} . Image adapted from Ref. [74].

of 600 mV. This is achieved by the implementation of the voltage divider chain as shown in Fig. 4.5 that produces voltages

$$V_x = \frac{R_x}{R_{\text{tot}}} \cdot 2V_{\text{REF}} = \frac{(x+3)R}{25R} \cdot 2V_{\text{REF}} = \frac{2}{25}V_{\text{REF}}(x+3) \quad (4.1)$$

at the different outputs V_x . With a typical reference voltage $V_{\text{REF}} = 500$ mV, the bias voltage can be fine tuned in 40 mV steps, which corresponds to a change in gain by about 1% around the nominal overvoltage $\Delta V = 3.5$ V as derived in Eq. (3.6) on page 27. Each channel can be connected independently to the output voltages V_0 to V_{15} as determined by four control bits in the channel configuration registers (see Section 4.2.3).

The current conveyer is followed by a closed-loop transimpedance amplifier. It converts the current signal coming from the SiPM to a voltage signal that is used for further processing. As indicated in Fig. 4.4, three extra debug (DBG) pins are included in the design that give access to different stages in the analogue signal processing. An internal multiplexer is used to select the channel to be tested. Figure 4.6 (left) shows the typical debug output of the pre-amplification stage. The signals on the oscilloscope are generated by illuminating the SiPMs with short light pulses (see Section 4.5 for further details). A band structure is visible at early times corresponding to different numbers of pixels with triggered avalanches.

Shaping

As can be seen at the output of the pre-amplification stage in Fig. 4.6 (left), the signals extend over several 25 ns LHC clock cycle periods. This is mainly due to the recovery time of the SiPM pixels as discussed in Section 3.6 and would lead to unwanted spillover into the following bunch crossing intervals. Therefore, a shaper is used within the PACIFIC to suppress the exponential tail of the signal. It is based on a pole-zero cancellation circuit as shown in Fig. 4.7. The impedance between the two nodes is given by

$$Z(\omega) = \left(\frac{1}{R_{\text{PZ1}}} + i\omega C_{\text{PZ1}} \right)^{-1} = \frac{R_{\text{PZ1}}}{1 + i\omega C_{\text{PZ1}} R_{\text{PZ1}}} \quad (4.2)$$




Figure 4.6: PACIFIC pre-amplifier (left) and shaper (right) output. Images adapted from Ref. [73].

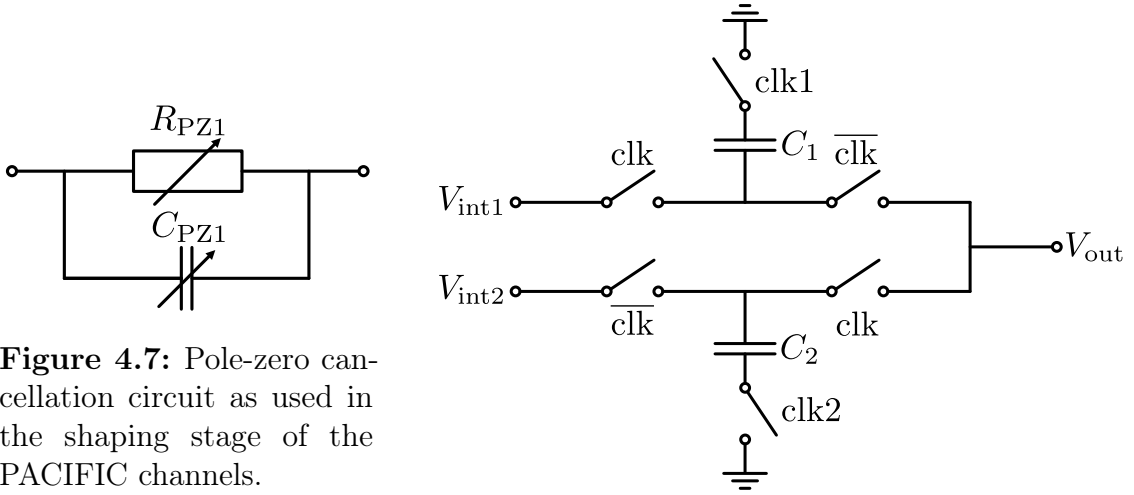


Figure 4.7: Pole-zero cancellation circuit as used in the shaping stage of the PACIFIC channels.

Figure 4.8: Schematics of the track and hold stage of the PACIFIC channels.

which has a pole at

$$\omega = \frac{i}{C_{PZ1}R_{PZ1}} \equiv \frac{i}{\tau_{PZ1}}. \quad (4.3)$$

The time constant τ_{PZ1} thereby determines the fall time of the shaper output. It can be tuned by modifying the resistor R_{PZ1} and capacitor C_{PZ1} within the circuit, which is enabled through dedicated control bits in the configuration registers (see Section 4.2.3).

Typical output signals of the shaping stage as seen on the oscilloscope are shown in Fig. 4.6 (right). Compared to the pre-amplifier output, the fall time is reduced significantly allowing the signal to be fully integrated in the following stage.

Integration

The amplified and shaped signal is integrated in each 40 MHz clock cycle in order to accommodate for variations in the photon arrival time due to the decay and travel time in the scintillating fibres. The implementation uses two gated integrators. They integrate the signals in an interleaved manner at half the system clock frequency (20 MHz). This means that while one integrator is in use, the remaining one is being reset. This structure allows for a reduction of the dead time to a minimum.

It is of crucial importance that the two integrators have similar properties to ensure the same performance in even and odd bunch crossing numbers corresponding to the different integrators. However, due to manufacturing tolerances, especially the DC baselines can show significant differences. The DC baseline refers to the integrator output level in the absence of an input pulse. Since any additional signal is added to it, it is also commonly designated as pedestal. A correction mechanism is foreseen that allows for the adjustment of the reference voltage of each integrator thus effectively shifting the baseline. It can be tuned by means of a 4-bit current output DAC (Digital-to-Analogue Converter) that can feed current in both directions. This enables the reference voltage of each integrator to take 31 different values, which are controlled through dedicated control bits in the configuration registers (see Section 4.2.3). The alignment of the DC baselines of the two integrators is also referred to as trimming and is explained in detail in Section 6.1.3.

Track and Hold

The integrators are followed by a dual passive track and hold stage that stabilises and merges the signals from the two integrators for the subsequent digitisation step. As shown in the schematics in Fig. 4.8, it is based on two capacitors $C_{1,2}$ and several switches that are operating in synchronous with the previous stage. Similar to the integrators, the switching is performed in an interleaved manner using the clock clk and its inverted $\overline{\text{clk}}$ at a frequency of 20 MHz. While integrator 1 is in use, its output V_{int1} is connected to the capacitor C_1 , but disconnected from the common track and hold output V_{out} . At the same time, the capacitor C_2 is connected to V_{out} , but disconnected from integrator 2. After 25 ns, the situation is reversed such that C_2 is connected to the integrator output V_{int2} .

In previous prototypes of the PACIFIC ASIC, the capacitors were directly connected to ground. However, it was found that this design leads to a non-negligible amount of charge sharing in the order of 25% between the two capacitors because they are connected to each other for a brief moment during the switching process. For an unirradiated detector, this resulted in 60% of fake hits in the bunch crossing after the detection of a signal at the same position. The implementation of the track and hold stage was identified as a major contribution to this effect. The redesigned schematic includes two additional switches driven by the clocks clk1 and clk2 as shown in Fig. 4.8. At the moment of the switching, they disconnect the capacitor that is currently not in use for a short period (< 1 ns). In addition, the interleaved




Figure 4.9: Typical output of the track and hold stage of the PACIFIC channels. The distribution on the left represents the pulse height projection of the signals within the brown rectangle in logarithmic scale.

switching procedure as just described is slightly shifted with respect to each other to decouple the two capacitors even further.

The typical output of the track and hold stage is shown in Fig. 4.9. Note that the signals are going in negative direction since the gated integrators are using an inverted structure. The pulse height projection of the signals included in the oscilloscope image illustrates the clear peaking structure corresponding to different numbers of photoelectrons. Unlike the debug output of the pre-amplifier or shaper (Fig. 4.6 on page 44), the output of the track and hold stage strongly depends on the timing of the pulsed light with respect to the system clock. The delay needs to be tuned carefully such that the signals can be fully integrated in the previous stage. At the optimal timing, the distance between the photoelectron peaks is maximised. This aspect is explained in more detail in Section 5.13.

Digitisation

The last stage of the signal processing in the PACIFIC channels is the digitisation stage. The analogue SiPM signals processed in the previous steps are digitised by means of three PMOS comparators that act as a simple non-linear flash ADC [72]. They feature a mean hysteresis of about 10 mV to avoid repetitive switching at the output when the signal level is close to the set threshold [76]. The comparator




Figure 4.10: Relationship between the digital input value and the output voltage of a PACIFIC threshold DAC. In addition, the effect of the slope (left) and offset (right) control is illustrated. Images adapted from Ref. [74].

thresholds are provided by three independent 8-bit current DACs included in each channel. The generated currents are injected into a $\approx 50\text{ k}\Omega$ resistor allowing for voltages between ground and 760 mV [74]. In the default configuration, one threshold DAC step corresponds to about 2.5 mV.

Due to the discrete SiPM signal amplitudes depending on the number of activated pixels (see Section 3.6), having a larger granularity ADC offers little added value, while only consuming a higher bandwidth. On the other hand, the independently adjustable thresholds provide the flexibility to efficiently digitise the signals in preparation for the following processing steps. However, the disadvantage of this design is that appropriate thresholds must be set prior to taking data. The calibration procedure required for this is described in Section 4.2.4.

Besides from the three threshold DACs in each individual channel, the PACIFIC ASIC features an additional set of threshold DACs that are common to all channels. In the following, the latter are referred to as common thresholds, while the individual threshold DACs are called local thresholds. As illustrated in Fig. 4.4 on page 42, a switch is included in each channel that allows for the selection of the different sets of thresholds. During the design phase of the chip, it was intended to primarily use the common thresholds. While this allows for an easier handling, it requires that all 64 channels in each chip have a high uniformity in terms of gain and DC baselines. However, it was found that especially the DC baselines show large variations that can not always be corrected by means of the trimming mechanism described earlier. Therefore, the local threshold DACs became the standard when operating the detector.

A circuit is implemented that allows for the adjustment of the range and granularity of the threshold DACs. The slope can be modified by adding or removing current to the DACs' current reference. A similar mechanism is foreseen at the output of the DACs, just before the resistor that performs the conversion to a voltage. By adding a certain amount of current at this point, the offset can be tuned. The slope and

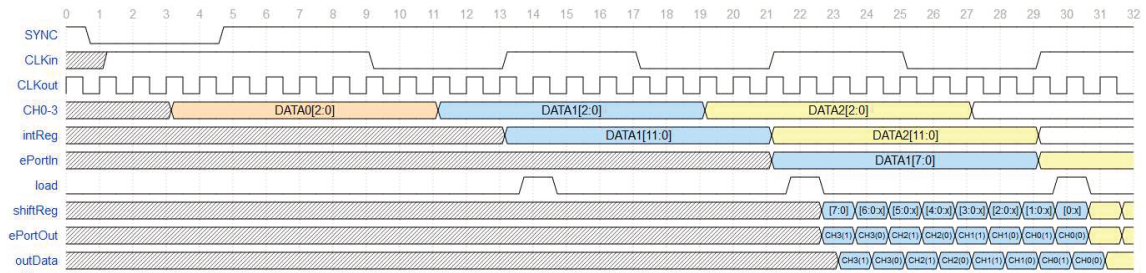


Figure 4.11: Digital timing diagram of the PACIFIC serialiser during normal operation. The x -axis is labelled in units of the 320 MHz clock period (CLKout) that determines the output data stream. Image taken from Ref. [74].

offset adjustment is applied to all threshold DACs (common and local) at the same time. Both variables are controlled by 4 bits in the configuration registers each. An additional bit determines the sign of the slope adjustment.

The relation between the set threshold DAC input value and the output voltage is shown in Fig. 4.10. In addition, it demonstrates the effect of the slope and offset adjustment.

The output of each PACIFIC channel is a 3-bit pattern at 40 MHz corresponding to whether (1) or not (0) the signal amplitude exceeds the thresholds of the three comparators.

4.2.2 Serialisation

A fast serialiser block follows the PACIFIC channels and enables the data transmission to the Cluster Board. The complete serialisation procedure is illustrated in the digital timing diagram in Fig. 4.11. Its operation relies on a 320 MHz master clock (denoted as CLKout in the diagram). Based on this clock, the 40 MHz (CLKin), as well as the two complementary 20 MHz clocks required for the analogue processing and digitisation are generated.

In the first step of the serialisation, the output of the three comparators of each PACIFIC channel is sampled at the rising edge of the 40 MHz clock. The outputs of four neighbouring channels are thereby written into a 12-bit internal register (intReg). To stay within the available bandwidth, a simple compression is performed in which the 3 bits of each channel are arithmetically added to form a 2-bit integer that represents the number of comparators whose threshold levels were exceeded by the signal amplitude. Assuming a strict hierarchy and proper functioning of the comparators, no information is lost in the process. With the help of an internal shift register (shiftReg), the compressed data (ePortIn) is output sequentially over a differential link at the rising edge of the 320 MHz clock (outData) Each chip comprises 16 differential outputs to cover the total of 64 channels, resulting in a data rate of 5.12 Gbit/s.

Furthermore, an additional input signal that is referred to as SYNC is shown in the diagram. As the name suggests, it is responsible for synchronising the total of

Table 4.1: Digital input and output signals of the PACIFIC that interact with the slow control and register blocks.

Signal	Type	Logic	Description
RESET	Input	Active low	Global digital reset
DISACLK	Input	Active low	Disables I ² C clock
LDINIT	Input	Active low	Loads initial register values
ERROR1	Output	Active high	Hamming code detected single bit-flip
ERRORM	Output	Active high	Hamming code detected multiple bit-flips
ERROR1GEN	Input	Active low	Forces single bit-flips for debugging
REFRESH	Input	Active low	Clears error flags ERROR1 & ERRORM

8192 PACIFIC ASICs installed in the SciFi Tracker. It is evaluated at the rising edge of the 320 MHz and triggers a reset of the state machine in the clock generator upon arrival. Thereby, it ensures that the slower 40 MHz and 20 MHz clocks are generated simultaneously in the various chips at a predetermined time.

4.2.3 Slow Control Interface and Configuration Registers

The PACIFIC features a slow control interface that is based on standard 10-bit addressing I²C [77]. Efficient reading and writing of larger amounts of continuous data is enabled by automatic incrementation of the target address without the need to repeat the 2-byte preamble for every transmitted byte of data.

The 10-bit address space is required to allow addressing a total of 340 logical registers. Of these, 336 read/write registers contain the configuration parameters of the chip. The additional four are read-only and allow for the monitoring of internal ADC readings and error detection counters. The registers are logically organised in groups of 8 bits resulting in a total of 2688 control bits of which 2683 are actually in use.

A small fraction of the configuration registers apply to all the 64 channels of the PACIFIC. It includes the parameters for the common bias block and controls, among others, the slope and offset adjustment of the threshold DACs, the pole-zero cancellation parameters, as well as the DAC values of the common thresholds themselves. This functionality is covered by 107 control bits within 14 registers. An additional 80 control bits (10 registers) are used for debugging purposes like controlling an internal channel multiplexer to measure the various stages of the analogue processing chain with an oscilloscope as shown in Section 4.2.1.

By far the largest fraction of the configuration registers determines the behaviour of the individual 64 channels. This includes the SiPM bias voltage offset (4 bits) and trimming of the reference voltages of the two integrators (10 bits). Additionally, the local 8-bit threshold DACs of the three comparators can be selected (1 bit)

and adjusted (3×8 bits). Thereby, 39 control bits are required for each PACIFIC channel, resulting in a total of 2496 bits, which are logically organised in blocks of eight channels.

Protection against erroneous state changes of the control bits due to ionising radiation (single event upset, SEU) is provided by a Hamming encoding algorithm [78]. Each configuration register is split into two 4-bit blocks that are covered by an extended Hamming(7,4) code with an additional parity bit. Thereby, every 4 bits of data are protected by 4 additional parity bits that allow to detect state changes in up to 3 bits. Additionally, it can distinguish between errors in single or multiple bits. While single bit-flips can be corrected on the fly, the configuration has to be rewritten in case of multiple bit-flips. Both types of errors are counted and stored in two of the four read-only registers mentioned before.

A complete list of the configuration registers is given in the PACIFIC data sheet [74]. In addition, the PACIFIC features several digital input and output signals that interact with them. They are summarised in Table 4.1.

4.2.4 Threshold Scan

Since each PACIFIC channel only features a set of three adjustable comparators, the full SiPM photoelectron spectrum as can be seen with a high-resolution ADC can not be resolved. However, the determination of the amplitudes corresponding to different numbers of generated photoelectrons (pe) is required to set appropriate thresholds for the operation of the detector. This process is also referred to as PACIFIC calibration.

Despite the lack of a full ADC, the position of the photoelectron peaks can also be estimated by scanning the threshold DAC counts (hereinafter referred to as DACs) over their entire dynamic range. In each step, several thousand events are recorded and the ratio of events above threshold is determined. At the same time, the SiPM channels are illuminated with pulsed light with the help of the light injection system, which is further described in Section 4.5.

The result of such a threshold scan is an integrated spectrum as shown in Fig. 4.12. Due to the shape of the curve, it is often referred to as S-curve. In order to determine the peak positions of the underlying spectrum, a dedicated calibration tool has been developed [79]. It is based on modeling the spectrum of photoelectron peaks k by a sum of Gaussian functions

$$N(x) = \sum_{k \geq 0} A_k \cdot \frac{1}{\sqrt{2\pi\sigma_k^2}} \cdot \exp\left(-\frac{(x - p_k)^2}{2\sigma_k^2}\right) \quad (4.4)$$

with amplitudes A_k , means p_k and widths σ_k [80]. The width of the first peak σ_0 , corresponding to the channel baseline or pedestal, is determined by the electronic noise, while subsequent peaks are additionally widened due to slight variations σ_1 between pixels

$$\sigma_k^2 = \sigma_0^2 + k \cdot \sigma_1^2. \quad (4.5)$$




Figure 4.12: SiPM photoelectron spectrum and corresponding data (S-curve) as determined by a PACIFIC threshold scan. Image adapted from Ref. [82].

The amplitudes A_k are primarily dependant on the light intensity μ that describes the mean number of photons within a Poisson distribution. However, due to pixel crosstalk λ (see Section 3.6), it is modified to the so-called Generalised Poisson distribution [81]

$$A_k(\mu, \lambda) = \frac{\mu \cdot (\mu + k \cdot \lambda)^{k-1} \cdot \exp(-\mu - k \cdot \lambda)}{k!}. \quad (4.6)$$

When performing the threshold scan as described previously, the course of the S-curve is given by the complementary cumulative distribution function of $N(x)$

$$S(x) = P(X > x) = \int_x^\infty N(x') dx' = 1 - \sum_{k \geq 0} A_k(\mu, \lambda) \cdot \frac{1}{2} \left[1 + \operatorname{erf} \left(\frac{x - p_k}{\sigma_k \sqrt{2}} \right) \right]. \quad (4.7)$$

Ideally, the peak positions p_k have the same distance to each other, which is hereinafter referred to as photoelectron peak separation. It is determined by the gain of the SiPM as well as the analogue processing chain in the PACIFIC. Since the photoelectron peaks become only visible by means of the threshold scans, they are stated in units of threshold DACs in the following. As can be seen in Fig. 4.12, the separation between the peak positions p_k is not constant, which is due to non-linearities in the PACIFIC channels. Therefore, the positions of the first six peaks are determined individually when performing the fits of the S-curves. Along with the light intensity μ , crosstalk probability λ , and the two width parameters σ_0 and σ_1 , a total of 10 free parameters are used in the fit of a single S-curve.

Based on the obtained parameter values, the threshold DAC settings of the three comparators per PACIFIC channel are calibrated. The default operating thresholds are set to be in the middle of the determined positions of the 1st, 2nd, 3rd and 4th photoelectron peak, also referred to as [1.5, 2.5, 3.5] pe. Thereby, in terms of the recorded S-curves, the thresholds are tuned to be in the middle of the plateaus, i.e. the regions where the ratio stays constant over a few DACs. With a typical photoelectron peak separation of 15 DACs and peak width parameters $\sigma_{0,1} \approx 1$ DACs, the width of the 1.5 pe plateau is about 10 DACs, which corresponds to the margin of error in the determination of the threshold. Due to the variations in the SiPM pixels, the peaks are widened with increasing numbers of photoelectrons according to Eq. (4.5) thus reducing the width of the plateaus.

However, the peak separation of 15 DACs is only achieved in the case where the SiPM signals are fully integrated by the PACIFIC channels. For non-optimal timings, the charge accumulated by the PACIFIC integrators is reduced resulting in squashed S-curves in the horizontal direction. On the other hand, the plateau heights are determined by the light intensity, which has to be low enough to be able to locate the pedestal position, but large enough in order to identify the higher order photoelectron peaks. Light intensities corresponding to mean number of detected photons μ between 1 and 2 have been found to be optimal.

Technically the scan procedure is complicated: A separate threshold scan is required for each of the three comparators, resulting in a total of 768 steps per channel. This is because the output of the PACIFIC is encoded in 2 bits per channel (see Section 4.2.2). Thereby, only the number of exceeded comparator threshold levels are kept, while the information about the individual comparator response is lost. Therefore, while scanning one comparator, the thresholds of the remaining two are set to the maximum value (256 DACs). This provides sensitivity to a single comparator, as any change in output can be traced back to the currently scanned comparator.

4.3 Cluster Board

The Cluster Board (CB) forms the bridge between the PACIFIC and Master Board. With 60 mm it matches the same width as the PACIFIC Board, despite being about 35% longer. A close-up photograph of a Cluster Board mounted on the backside of a Master Board is shown in Fig. 4.13. It also displays the black 400-pin FMC connector that is used to establish the connection to the PACIFIC Board. The same connector type is used for mounting on the Master Board.

The main components of the Cluster Board are two FPGAs located on the bottom of the PCB. They perform the hit reconstruction and noise suppression based on the digitised signals provided by the PACIFIC ASICs. The FPGAs, as well as the ASICs on the associated PACIFIC Board, are controlled by a dedicated slow control chip that will be discussed shortly.

Similar to the PACIFIC Board, each Cluster Board is equipped with three Pt1000 sensors to allow the monitoring of the temperatures on the PCB.




Figure 4.13: Close-up view of a Cluster Board mounted on the bottom side of a Master Board. The connection to the PACIFIC Board on the left can be established through the black FMC connector.

4.3.1 Cluster SCA

A CERN GBT-SCA chip [83] is used as the Slow Control Adapter (SCA) on the Cluster Board. It provides the interface to control and monitor the various front-end electronics components. The GigaBit Transceiver (GBT) project was founded to address the challenges of growing beam intensities and radiation doses at the LHC. The GBT-SCA, built in commercial 130 nm CMOS technology, is part of a series of ASICs that were designed in this context to implement a radiation-hard bidirectional 4.8 Gbit/s optical fibre link between the front- and back-end electronics [84].

As an universally applicable chip for the LHCb detector, the GBT-SCA supports various control buses that are commonly used in High Energy Physics (HEP). These include one Serial Peripheral Interface (SPI) master, one JTAG master, 16 independent I²C masters, and 32 general-purpose input/output (GPIO) pins. In addition, it provides 31 analogue inputs that are multiplexed to a pre-calibrated 12-bit ADC covering the range between 0 V and 1 V. Each input features a switchable 100 μ A current source that allows for the operation of externally connected resistance thermometers. One additional input is internally connected to an embedded temperature sensor. Conversely, four 8-bit DACs are available that can provide voltages between 0 V and 1 V on four different analogue output ports [85].

Besides from providing the interface to the various supported control buses, the GBT-SCA features three 8-bit control registers itself. With the help of the 24 available control bits in these registers, the channels that are not in use in a specific application can be disabled to reduce the overall power consumption. The 16 I²C master interfaces can be switched independently. Four additional control bits allow to turn off or on the SPI and JTAG interface, as well as the ADC and GPIO pins, leaving four reserve bits unused. Initially, all control registers start-up with 0x00,

which means that all control buses are deactivated by default. An additional 24-bit read-only register contains a unique ID that is written on internal electronic fuses during the production and testing of the SCA chip.

Application on the Cluster Board

The Cluster (Board) SCA utilises six of the available I²C masters. Since the PACIFIC Board does not have its own SCA, four of these are used for the communication with the I²C slaves of the four associated PACIFIC ASICs. The remaining two channels control the two FPGAs located on the Cluster Board.

The JTAG master interface is connected to both FPGAs in a daisy chained manner according to IEEE 1149.1 Standard [86]. It is thereby possible to re-program the FPGAs, as well as performing boundary scan tests. Alternatively, re-programming of the FPGAs is also possible through the available SPI interface that is connected as a reserve.

Of the available 31 analogue ADC inputs, 11 are in use. They are connected to the Pt1000 temperature sensors on the Cluster Board (3 sensors), the associated PACIFIC Board (2), as well as on the backside of the two connected SiPM arrays (2). Hence, the 100 μ A current source is enabled for these channels to allow for the measurement of the temperature dependent resistance. The remaining four ADC lines are used for monitoring of the bias voltages of the individual SiPM dies.

In contrast to the ADC inputs, almost all 32 available GPIO pins are utilised in the application of the Cluster SCA. With the exception of ERROR1GEN, all digital PACIFIC input and output signals listed in Table 4.1 on page 49 are covered. With four PACIFIC ASICs associated to each Cluster SCA, this results in a total of 24 occupied GPIO pins. Another six lines are divided across both FPGAs for performing different types of resets. The last used pin drives the JTAG reset signal (TRST) leaving only one vacant GPIO pin.

4.3.2 Cluster FPGA

Two radiation tolerant IGLOO2 FPGAs¹ provide the key functionality of the Cluster Board. The flash-based FPGAs are built in 65 nm technology and come in a 484-ball BGA (Ball Grid Array) package for mounting on the bottom side of the board [87].

They perform the hit reconstruction on the digitised SiPM signals received by the PACIFICs, and prepare the data for transmission over the optical links to the DAQ system at the back-end. Both FPGAs are operating based on an identical firmware that can be loaded through the JTAG or SPI interface of the Cluster SCA as described previously. However, a subtle distinction is made at one pin that is pulled to 1.5 V for one FPGA, while being connected to ground for the other. Within the firmware, this information is used to determine the position on the Cluster Board, which is required for locating the origin of the data at the back-end. Further details are given in Section 5.6.

¹IGLOO2 FPGA M2GL090T-FGG484 by Microsemi

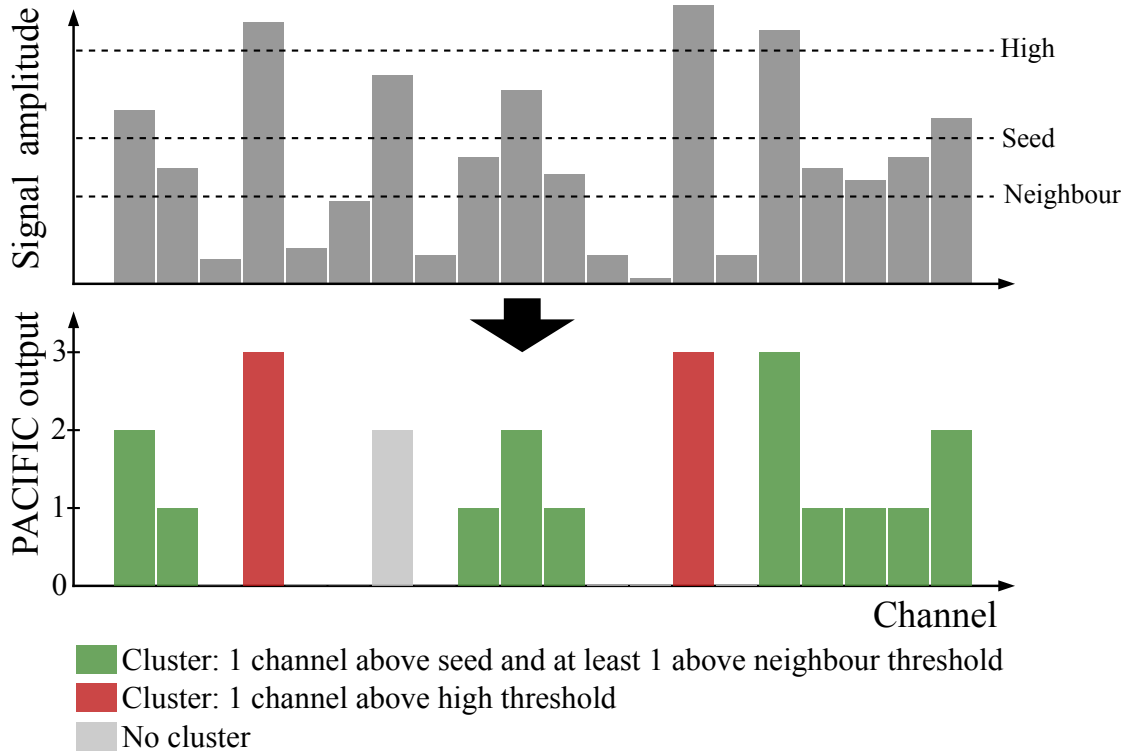


Figure 4.14: Illustration of the clustering algorithm. The height of each rectangle at the top depicts the SiPM signal amplitude in one channel. The 2-bit patterns below represent the outputs of the PACIFIC channels after applying the three comparator thresholds. Five clusters are identified by the algorithm as indicated by the green and red colors depending on the fulfilled condition.

Each FPGA processes the information from one SiPM array (128 channels) that is readout by two PACIFICs (64 channels each). With 16 differential outputs per PACIFIC as discussed in Section 4.2.2, this corresponds to a total of 32 FPGA data input lines.

In the first step, the incoming PACIFIC data stream needs to be deserialised. The block in the FPGA that is responsible for that task operates at a clock frequency of 160 MHz. In order to cope with the input rate of 320 Mbit/s per line, it samples the data at both the rising and falling edge of the clock, which is commonly referred to as Double Data Rate (DDR).

Clustering

The next processing stage is the clustering block that operates at the general clock frequency of 40 MHz. It is used to compress the data before transmission over the optical links by reconstructing hit positions while removing noise and redundant zeros (zero suppression).

The hit finding and clustering algorithm follows a similar logic as the procedure

to form clusters of VELO hits before the LHCb upgrade [88]. It is based on the three comparator thresholds in each PACIFIC channel (see Section 4.2.1). They are named following their function within the clustering procedure:

- **Seed threshold:** The seed threshold is set to a medium level. It marks the starting point for every potential cluster. Depending on the signal amplitude within the channel, as well as the output of the neighbouring channels, it is discarded or becomes an actual cluster.
- **Neighbour threshold:** The neighbour threshold is set to the lowest level of the three. Starting from the channel that exceeds the seed threshold, the cluster algorithm scans the adjacent channels for signal amplitudes above (at least) the neighbour threshold. This process continues in both directions until no more channels exceeding the neighbour threshold are found. At least one adjacent channel above the neighbour threshold is required for a potential cluster to become an actual cluster. All the channels that were collected along the way are considered to be part of the cluster.
- **High threshold:** The high threshold is set to the highest level of the three. If the signal amplitude not only exceeds the seed, but also the high threshold, it is considered an actual cluster. In this case, it is irrelevant whether there are adjacent channels above the neighbour threshold.

The clustering algorithm is illustrated in Fig. 4.14. To achieve a sufficient compression of the data, only the positions of formed clusters are transmitted over the optical links. The position of a cluster \bar{x} is calculated as the weighted arithmetic mean

$$\bar{x} = \frac{\sum_{i=1}^n w_i x_i}{\sum_{i=1}^n w_i} \quad (4.8)$$

with the involved channel indices x_i and corresponding weights w_i that depend on the exceeded threshold values.

However, the weighted mean is only used for clusters that involve up to four channels. Larger clusters are internally split into multiple fragments. Once one fragment contains four channels, a new fragment is created. These types of clusters are flagged and two positions are transmitted: the (unweighted) arithmetic mean of the first and the last cluster fragment. With a mean cluster size of 2.5, up to 10% of the clusters involve four or more channels. Hits of particles that are relevant for physics analyses mostly involve only up to four channels, while larger clusters typically originate from (secondary) particles at high angles, noise, or delta rays. To avoid biases in the position due to these effects, the distinction between large and small clusters is made [89].

Table 4.2: General data format as used in the SciFi Tracker. It consists of a 20-bit header along with a 92-bit payload containing the calculated cluster positions.

	Bit(s)	Description
Header	111:100	Bunch crossing ID. Ranges from 0 to 3653.
	99	Raw data flag. Set to 0 in case of cluster data.
	98	Parity bit of the 19 other header bits.
	97	TFC flag. Set to 1 at arrival of control command(s).
	96:92	Number of formed clusters.
Payload	91:90	Reserved.
	89:81	Position information of the 1st formed cluster (if existing).
	80:72	Position information of the 2nd formed cluster (if existing).

	17:9	Position information of the 9th formed cluster (if existing).
	8:0	Position information of the 10th formed cluster (if existing).

Data Format

After the clustering, the data needs to be formatted and serialised in preparation for the transmission via the optical fibres. The corresponding block in the Cluster FPGA operates at a clock frequency of 80 MHz. Along with the calculated cluster positions, it receives the time-sensitive commands from the control system. The available information is encoded in a 112-bit data frame for each bunch crossing period. Depending on the location of the ROB within the detector, and the mode in which it is operated, different formats are used.

The general data format is shown in Table 4.2. It consists of a 20-bit header and 92-bit payload. Within the header, 12 bits are used to label the bunch crossing ID of the current event. It ranges from 0 to 3653 (incl.) and is required to time-align the incoming data from a total of 4096 optical fibres from the SciFi Tracker at the DAQ servers. The number of formed clusters in the event is encoded as 5 bits. The remaining 3 bits in the header comprises additional meta information and are still subject to change.

The payload contains the calculated positions of the formed clusters. Each cluster is assigned a 9-bit number. The determined position itself is encoded in 8 of the 9 bits. Since each Cluster FPGA processes the output of 128 PACIFIC channels, the position within the SiPM array can thereby be encoded in 0.5-channel steps. For large clusters involving more than four channels, bit 9 is set to 1 and sent along with the 8-bit position (unweighted arithmetic mean) of the last fragment of the cluster. The position of the first fragment of the corresponding large cluster is encoded in the preceding 9 bits.

Since both the header and the payload are of fixed length, the general data format as shown in Table 4.2 is also referred to as Fixed header, Fixed payload (FF) format. It is used in 80% of the ROBs during normal data taking periods. The remaining ROBs are operated in the so-called FV format with a fixed length header but a variable payload size. This concerns 48 ROBs reading out the fibre modules that are directly surrounding the beam pipe. Since these modules are experiencing the highest occupancies, having a payload of variable length allows to circumvent the limitation of ten clusters per event as it is the case for the FF data format. The FV format uses a 32 events deep FIFO in the Cluster FPGA that allows unused space in low occupancy events to be filled with the data of previous events containing a large number of hits [87].

In addition, a modified FF data format exists that is used for debugging and calibration purposes. In these cases, the payload contains the raw 2-bit outputs of the PACIFIC channels instead of the cluster positions. However, the 92-bit payload is not sufficient to hold the raw data of the 128 associated channels. Therefore, it requires multiple 112-bit data frames to transmit a single event, which means that it does not allow for a readout of the detector at the full rate of 40 MHz. This type of data is for instance required for performing threshold scans in order to calibrate the PACIFIC comparators (see Section 4.2.4).

Independent of the format, the data frames are sent to the Master Board via 28 differential outputs. To allow for the transmission of the 4 bits per line and bunch crossing period of 25 ns, the data is serialised at both edges (DDR) of the underlying 80 MHz clock.

4.4 Master Board

The Master Board (MB) presents the interface between the SciFi front-end electronics and the DAQ system at the opposite side of the optical links. The basis of the Master Board consists of a 16-layer PCB with a nominal thickness of 2.15 mm [90]. With a length of 132.5 mm and width of 256.2 mm it directly connects to four adjacent 60 mm wide Cluster-PACIFIC Board pairs [91]. As shown in the exploded view in Fig. 4.15, it houses various components that allow for the data transmission and communication with the back-end, as well as for providing power to the PACIFIC and Cluster Boards.

A total of eight dedicated chips, the so-called Data GBTX ASICs, receive the data from the eight Cluster FPGAs and encode it into the GBT protocol that is used for transmission over the optical links [84]. Another ASIC of the same type, the Master GBTX, operates in a different mode that allows for the bidirectional communication with the back-end electronics. The Master Board features an additional FPGA that is referred to as Housekeeping (HK) FPGA. It enables the operation of the light injection system (see Section 4.5) and is responsible for the monitoring of different Master GBTX status registers.

On each board, 13 FEASTMP DC-DC converters [92] transform the 8 V input




Figure 4.15: Exploded view of a Master Board showing the various mounted components. Image modified from Ref. [91].

voltage provided by the MARATON power supplies (see Section 3.8.1) into four different voltage levels required by the various components, including the ones mounted on the PACIFIC and Cluster boards. In addition, a 26-pin D-SUB connector is installed on the board and connects to the CAEN modules (see Section 3.8.1) for biasing of the 16 associated SiPM dies. Monitoring of the board temperature is enabled by eight NTC thermistors¹ mounted on different positions.

In the following, the key components of the Master Board are discussed in more detail.

4.4.1 Master GBTX

The Master GBTX is essential for the operation of the SciFi front-end electronics as it is responsible for the communication with the back-end [33]. Same as the GBT-SCA covered in Section 4.3.1, the underlying GBTX ASIC is part of the GBT project [84]. In order to establish the bidirectional optical link, it is operated in close conjunction with two other chips of the GBT project.

The downlink is realised by connecting to a GigaBit Trans-Impedance Amplifier (GBTIA) ASIC. The GBTIA amplifies and converts the weak photocurrent provided by the PIN diode that is detecting the incoming light through the optical fibre [93]. The resulting digital differential voltage is received by the input stage of the GBTX.

¹B57301V2472H060 by TDK Electronics




Figure 4.16: Structure of the 120-bit GBT frame as used in the Master GBTX for communication with the back-end electronics. It consists of a 4-bit header, 2-bit internal and external slow control fields (IC and EC), 80 bits of data, and 32 bits used for a Forward Error Correction (FEC) code.

Based on the data stream, a clock and data recovery (CDR) circuit recovers and generates the required 4.8 GHz clock to correctly sample the data. Afterwards, the data is deserialised.

The uplink connection is enabled by a GigaBit Laser Driver (GBLD) ASIC [94]. It receives the serialised data from the GBTX and drives a 850 nm vertical-cavity surface-emitting laser (VCSEL) accordingly. Parameters of the GBLD like bias and modulation currents can be modified via a simplified I²C master within the GBTX. It is solely designed for this purpose and therefore cannot be used for the communication with other I²C devices.

Both the receiving part consisting of the GBTIA ASIC and PIN diode, as well as the transmitting part comprising the GBLD chip and VCSEL are embedded into a single device that is referred to as Versatile Transceiver (VTRx). While the design work of the two associated ASICs was carried out in the framework of the GBT project, the VTRx is part of the Versatile Link project at the LHC [95].

A GBT frame consists of 120 bits that are exchanged in each bunch crossing period (25 ns) resulting in the total data rate of 4.8 Gbit/s. The frame structure is illustrated in Fig. 4.16. Each GBT frame is led by a 4-bit header that is required to synchronise the data stream at the frame level [96]. Repeated recognition of the header results in a valid synchronisation, also referred to as frame-locking. Another 4 bits are reserved for slow control operations, which are split into internal and external control (IC and EC). The 2-bit IC field is used for the internal control of the GBTX itself and can be used to access a total of 366 read/write and 70 read-only registers. The 2-bit EC field is reserved for the implementation of an external slow control application, e.g. a separate GBT-SCA chip. In case of the Master GBTX, this field with an associated bandwidth of 80 Mbit/s is used for the communication with the Master SCA (see Section 4.4.2). The largest fraction of the GBT frame is allocated to an 80-bit data (D) field that can be used flexibly for the transmission of physics data, trigger signals and slow control commands. The unified handling of these three functions is one of the greatest strengths of the GBT architecture. In case of the Master GBTX, 4×2 bits within the data field are used for the communication with the four associated Cluster SCAs (see Section 4.3.1). Another 8×8 bits are allocated for trigger signals going to the eight Cluster FPGAs (see Section 4.3.2) within each HalfROB. An additional 2 bits within the data field are reserved for the control of the light injection system via the HK FPGA (see Section 4.5).

The remaining 32 bits are required for the implementation of a Forward Error Correction (FEC) code. It is based on two interleaved Reed-Solomon RS(15,11) encoded words with 4-bit symbols [97]. With a coding efficiency of 70%, it is capable of correcting bursts of 16 consecutive incorrectly transmitted bits [96]. The FEC is combined with a pseudorandom, self-synchronising scrambling algorithm to achieve a DC-balanced data transmission [98]. DC-balanced signals contain an even distribution of 0s and 1s in order to prevent bit errors that can occur otherwise after a long series of the same bit state (caused for instance by unwanted charging of the coupling capacitor).

The electrical connection of the GBTX chip with the various front-end devices is realised by so-called E-Links. In the standard configuration, one E-Link is built from three differential signal pairs: downlink and uplink data, as well as a differential clock line. However, not all three signal pairs have to be implemented in every application. In addition, the E-Links offer a high flexibility in terms of the used data rate, which can be selected between 80, 160 or 320 Mbit/s resulting in 2, 4 or 8 assigned bits in the GBT frame, respectively. Even higher bandwidths can be achieved by configuring multiple data lines in parallel.

In addition to the E-Links, the GBTX implements a clock manager circuit that provides eight external user clocks with programmable phases and frequencies. The clock frequency can be independently selected between 40, 80, 160 and 320 MHz, while the phase can be varied in 50 ps steps. In case of the Master GBTX, the eight external clocks are set to the system frequency of 40 MHz and are provided as reference clocks to the eight Data GBTX ASICs (see Section 4.4.3).

4.4.2 Master SCA

The Master SCA connects to the 80 Mbit/s E-Link of the Master GBTX that is assigned to the EC field. It is based on the identical GBT-SCA chip employed in the Cluster SCA. General remarks about the GBT-SCA ASIC can be found in Section 4.3.1.

On the Master Board, the SCA utilises ten of the available I²C master interfaces for the control of the eight Data GBTX ASICs, as well as the HK FPGA and light injection system. Re-programming of the HK FPGA is enabled via the SCA's JTAG interface, or alternatively the SPI, which is connected as a reserve. Eight ADC inputs are allocated to monitor the temperature of the PCB using the eight NTC thermistors mounted on the board. An additional ADC line is used to measure the input (low) voltage.

All of the 32 available GPIO pins are in use. They are set up as 15 output and 17 input lines. Four inputs are indicating the connectivity to the four Versatile Twin Transmitter (VTTx) modules (see Section 4.4.3). The remaining input pins are responsible for the monitoring of the power-good signals from the 13 FEASTMP DC-DC converters. The DC-DC converters are arranged in four sections. One section is always enabled, providing power to the essential components that are required for the communication with the back-end: the Master SCA itself, the Master GBTX,




Figure 4.17: Structure of the 120-bit GBT frame in wide mode as used in the Data GBTX for transmission of the physics data. In the wide frame mode, the Forward Error Correction (FEC) is omitted in exchange for additional bandwidth. For compatibility reasons, the allocation of the first 80 bits in the data field is identical to the default GBT frame structure depicted in Fig. 4.16.

and the VTRx module. The remaining three sections can be controlled with the help of three GPIO output lines. Like this, it is possible to independently power on/off the left and right half of the Master Board including the associated Cluster and PACIFIC Boards, as well as the 3.3V DC-DC converter that is powering the Cluster FPGAs. Reset signals going to the Data GBTX ASICs and the HK FPGA (soft, hard and JTAG reset) are driven by 11 output pins and an additional line can be used to disable the light injection system.

4.4.3 Data GBTX

Unlike the Master GBTX that is configured for bidirectional communication with the back-end, the eight Data GBTX ASICs mounted on the Master Board only provide an uplink data stream. Therefore, instead of recovering the clock from the optical link, they receive the 40 MHz reference clock from the Master GBTX. In addition, the internal registers of the Data GBTX are accessed via I²C from the Master SCA.

As discussed in Section 4.3.2, the data transmission from the Cluster FPGA takes place via 28 differential E-Link data lines each operating at 160 Mbit/s. Thereby, each Data GBTX receives a data frame consisting of 112 bits per bunch crossing interval. As this exceeds the available bandwidth in the default GBT frame structure presented in Fig. 4.16, a different format for the optical data transmission to the back-end has to be used. Therefore, the Data GBTX ASICs are operating in the GBT wide frame mode. As illustrated in Fig. 4.17, the available data field is extended to 112 bits in this mode. However, no protection against transfer errors is available as the FEC field is traded for the additional 40% bandwidth.

The Data GBTX marks the last stage of the data chain, before being transmitted over the optical fibres to the back-end. A summary of the complete chain within the SciFi front-end electronics is pictured in Fig. 4.18. The following optical transmission of the data is realised by four Versatile Twin Transmitter (VTTx) modules [99]. The package is identical to the VTRx modules, but, to address the typical need for asymmetric bandwidth in HEP experiments, the optical receiver is replaced with an additional transmitter channel. Thereby, two neighbouring Data GBTX ASICs share one VTTx module.

Five of the eight external user clocks of each Data GBTX are utilised and provided




Figure 4.18: Data stream and corresponding bandwidths between the different components of the SciFi front-end electronics. In total, 4096 instances of the shown chain are present in the detector, resulting in a total data rate of about 20 Tbit/s arriving at the back-end.

to the various stages in the data chain, ranging from 40 MHz as required by the clustering block to 320 MHz used for the operation of the PACIFICS. Their phases have to be carefully adjusted to each other to ensure a stable and reliable data transfer over the complete chain. Further details about this tuning process is given in Section 5.12.

4.5 Light Injection System

The light injection system (LIS) is a critical component for the proper functioning of the SciFi Tracker. It enables the calibration of the individual channels by artificially injecting light into the scintillating fibres close to the SiPMs while performing the PACIFIC threshold scans (see Section 4.2.4).

As shown in Fig. 4.19, the main part of the LIS is located within the end plugs of the fibre modules. It consists of a 670 nm VCSEL and a plastic optical fibre that guides the emitted light towards the fibre mat end piece. The last 13 cm of the optical fibre, matching the width of the fibre mat, is embedded into an aluminium bar that is referred to as light injection bar. Along that length, the fibre is scratched allowing light to escape. Subsequently, the escaped photons travel through the transparent end piece, enter the scintillating fibres, and finally reach the SiPM channels.

The VCSEL is driven by an external GBLD, the same chip that is used in the transmitter channels of the VTRx and VTTx modules. A pair of two GBLDs and two VCSELs is mounted on a 6-layer PCB [100]. Each fibre module is equipped with two of these boards in order to illuminate the four fibre mats contained within. A 28 cm long flex cable is guided along the surface of the fibre module and connects the LIS PCB with the Master Board via a 20-pin low-profile connector¹.

Operational settings of a GBLD pair, like bias and modulation currents, can be accessed from the Master SCA via I²C (see Section 4.4.2). The input signal to the GBLDs, in the following referred to as LIS pulse, is provided by the HK FPGA and determines the delay and duration of the resulting light pulse. The generation of

¹TLH-010-0.50-G-D-A by Samtec, Inc.




Figure 4.19: Schematic illustration of a fully assembled detector module including the light injection system. Image adapted from Ref. [100].

the LIS pulse is based on a dedicated calibration command from the control system. Unlike the slow control commands, the so-called CalibC pulse is a time-sensitive signal that can be issued in a particular bunch crossing period and thus be shifted in steps of 25 ns. As part of the GBT frame, it is received by the Master GBTX and provided to the HK FPGA via a dedicated E-Link.

However, as discussed previously in Section 4.2.1, finer tuning of the LIS pulse other than the 25 ns steps is required in order to align it with respect to the PACIFIC integration intervals. For that purpose, two unused external user clocks with programmable phase and frequency that are available in one of the Data GBTX ASICs are utilised. They are denoted as the start and stop clock and are set to the lowest possible frequency, i.e. to 40 MHz. The generation of the LIS pulse is based on the interplay of the different signals and illustrated in the timing diagram in Fig. 4.20.

Upon arrival of the CalibC pulse, the first rising edge of the start clock defines the start of the LIS pulse. Similarly, the following rising edge of the stop clock marks the end of the resulting pulse. Like this, it is possible to vary the delay of the LIS pulse in very fine steps (50 ps) by shifting the phases of the two clocks equally. In addition, the width of the pulse and thereby the light intensity can be adjusted by only shifting the phase of the stop clock. However, given the underlying clock frequencies of 40 MHz, both the delay and duration of the pulse can only be adjusted within a range of 25 ns. In case of the delay, this is accounted for by issuing the CalibC command in the previous or following bunch crossing period. LIS




Figure 4.20: Timing diagram of the LIS pulse generation in the HK FPGA. The system clock and CalibC pulse are received from the Master GBTX. The start and stop clocks are provided by the third Data GBTX on the Master Board.

pulse widths larger than 25 ns are typically not required but are still possible by instructing the HK FPGA to skip a certain amount of stop clock edges, thereby effectively increasing the width in steps of 25 ns.


In practice, the delay and width can not be varied freely over the entire dynamic range of 25 ns, which can become an issue in particular for the delay, if the best timing happens to be out of reach. The reason for this is that when the clock and CalibC edges are too close to each other (< 1 ns), due to jitter, the order of the signals is not guaranteed. The result is an unstable pulse that shows jumps by about 25 ns in either delay or width. In case of the width, the unstable configurations are simply avoided and compensated by changing the modulation current of the GBLD to tune the light intensity. However, no alternative method is available for changing the delay of the LIS pulse. Therefore, a mechanism was implemented that allows for the clocking of the CalibC signal by the falling edge of the system clock instead of the rising edge, thereby effectively shifting it by 12.5 ns. In principle, this allows for a range of the pulse delay of 37.5 ns, which is still covering the full 25 ns bunch crossing period when avoiding the configurations around the edges.

To validate the LIS pulse generation, the output light pulse was measured with the help of a PIN diode and an oscilloscope using different delays and widths of the generated LIS pulse. Figure 4.21 shows the signal from the PIN diode, as well as the LIS and CalibC pulse, which was used as the trigger. In addition, the measured light pulses with different delay and width configurations are displayed, demonstrating a proper functioning of the LIS and its tuning mechanisms as described previously.


Besides from enabling the calibration of the detector, the LIS also plays an important role for the commissioning of the front-end electronics as it allows to identify malfunctioning channels. Further details about the usage of the LIS in the commissioning is given in Section 5.13.

4.6 Quality Assurance

After production, the individual front-end electronics components undergo different quality assurance (QA) tests to ensure proper functioning. These include standard test procedures of the assembled PCBs like



(a) Overview



(b) Different pulse delays

(c) Different pulse widths

Figure 4.21: Validation of the LIS pulse generation using an oscilloscope and a PIN diode that measures the output light of the VCSEL. The top image (a) shows the three probed signals. The bottom images display the measured PIN diode voltages for different LIS pulse delays (b) and widths (c). The CalibC pulse appears later as it was probed on a separate setup due to difficult accessibility at the LIS itself. Nevertheless, it serves as a reference signal and was used as the trigger for the oscilloscope.

- 3D Automated Optical Inspection (3D AOI)
- Automated X-ray Inspection (AXI)
- Flying Probe Testing (FPT)
- Temperature Cycle Testing (TCT)

that are mainly conducted on-site at the production companies. Additionally, for some parts, extended functional tests are performed at the institutes within the SciFi collaboration [100], [101].

The QA of the PACIFIC Boards is conducted at Tsinghua University and the University of Valencia. Within a custom designed test DAQ developed at Heidelberg

University, each board has to pass a series of tests. Starting with a check of the power consumption, it is ensured that the communication via I²C and the digital input signals are operational. As one of the last steps, the DC baselines of the two integrators per channel are aligned (trimming) and the obtained parameters (trim DACs) are saved along with the ID of the board. Further details about the trimming procedure are given in Section 6.1.3. Prior of being mounted on a PACIFIC Board, each PACIFIC ASIC also undergoes an individual test sequence. Besides from similar checks as performed on the board-level described before, the reference voltages are tuned in the process to match their intended values.

In case of the Cluster Board, the assembled PCBs are examined at LPC in Clermont-Ferrand. Within a dedicated setup, two boards are tested in parallel. The procedure includes the injection and evaluation of test patterns, monitoring the current consumption, as well as testing the functionalities of the GBT-SCA. Furthermore, the programmability of the Cluster FPGAs is ensured.

At RWTH Aachen, the test setup for the LIS components is located. Therein, the operational settings for the bias (I_{bias}) and modulation currents (I_{mod}) of the GBLDs and VCSELs are determined in order to achieve a predefined light intensity. These values are of crucial importance for the calibration of the PACIFIC threshold DACs (see Section 4.2.4) that is required to take high quality physics data. Therefore, a further fine-tuning of I_{bias} and I_{mod} is performed on the LHCb site at CERN during the QA of the fully assembled fibre modules including LIS, cold boxes and SiPMs.

Apart from the tests at the production company, the Master Boards do not undergo individual testing but are examined in conjunction with the remaining boards: In a dedicated setup on the LHCb site, every assembled ROB is thoroughly tested in its final configuration including being readout via the optical links. Further details about the test-stand and procedure are given in Section 5.2.

The test results of the different QA steps are saved in a dedicated database hosted at Heidelberg University [102]. This includes in particular the individual operation parameters required for the later operation of the detector that are obtained during the tests.

5 Commissioning of the SciFi Front-End Electronics

The commissioning of the SciFi front-end electronics is performed along with the assembly and testing of the C-Frames (see Section 3.8.3). It began in 2019 and, after multiple interruptions in the course of the COVID-19 pandemic, is planned to be concluded by spring 2022. The assembly takes place in building 3852 at CERN LHC Point 8, which is the site housing the LHCb experiment and its infrastructure. The location of the building, in the following referred to as (SciFi) assembly hall, is shown on the satellite image of the site in Fig. 5.1.

The assembly and commissioning of the individual C-Frames requires a close collaboration of physicists, engineers and technicians from different institutes. From start to finish, the process for an individual C-Frame takes about four months and consists of the following steps:

- Assembly of the C-shaped mechanical support frame.
- Installation of water cooling pipes and blocks.
- Low and high voltage cables installation and testing.
- Installation of Novec and dry gas lines.
- Laying of optical fibres.
- Mounting of fibre modules.
- Connection of cold boxes to dry gas and Novec bellows.
- Winding of heating wires around the cold boxes.
- Test of vacuum and Novec cooling system.
- Mounting of Readout Boxes (ROBs).
- Connection of power and SiPM flex cables to the ROBs.
- Inspecting, cleaning, and connecting optical fibres to the ROBs.
- **Commissioning of ROBs.**

Many of the steps can be performed in parallel on up to four C-Frames. This is possible by four separate assembly and commissioning slots inside the SciFi assembly hall as shown in Fig. 5.2 on page 71. Each slot, also referred to as C-Cage, is surrounded by a scaffolding that allows work to be conducted on three different levels. However, the available infrastructure only allows for the commissioning of one C-Frame at a time. Due to limited lengths of cables and supply lines it is performed in the two leftmost slots depicted in Fig. 5.2.




Figure 5.1: Location of the SciFi assembly hall, highlighted in blue, at CERN LHC Point 8. The orange path indicates the transport route of the fully assembled and commissioned C-Frames to the access shaft in order to be lowered into the LHCb cavern. The street art on the left shows a projection of the LHCb experiment located 100m underground. Satellite image taken and modified from Ref. [103]

Six further C-Frames can be stored in a dedicated storage cage. Because it is also used for transporting the detector elements to the LHCb cavern (see route in Fig. 5.1) it is referred to as transport and storage cage (TS-Cage). In addition, the TS-Cage allows for flexibly swapping C-Frames between the different slots. This is for instance required for moving an assembled C-Frame to one of the two commissioning slots.

5.1 Complexity of Readout System

Due to the complexity of the system, the operation of the SciFi front-end electronics is challenging. As illustrated in Fig. 5.3, the complete detector consists of many thousand individual components that are installed in a total of 256 ROBs. These include the Master, Cluster and PACIFIC Boards, as well as the various ASICs, FPGAs and other electronic parts that are located on these boards.

The task of the electronics commissioning is the initial operation of this system in conjunction with the surrounding infrastructure and the DAQ system. As mentioned previously, this is performed at the level of individual C-Frames, i.e. in units of $1/12$ of the complete detector. Thereby, it must be ensured that all components are



Figure 5.2: Photograph of the four C-Frame assembly and commissioning slots inside the SciFi assembly hall at night.

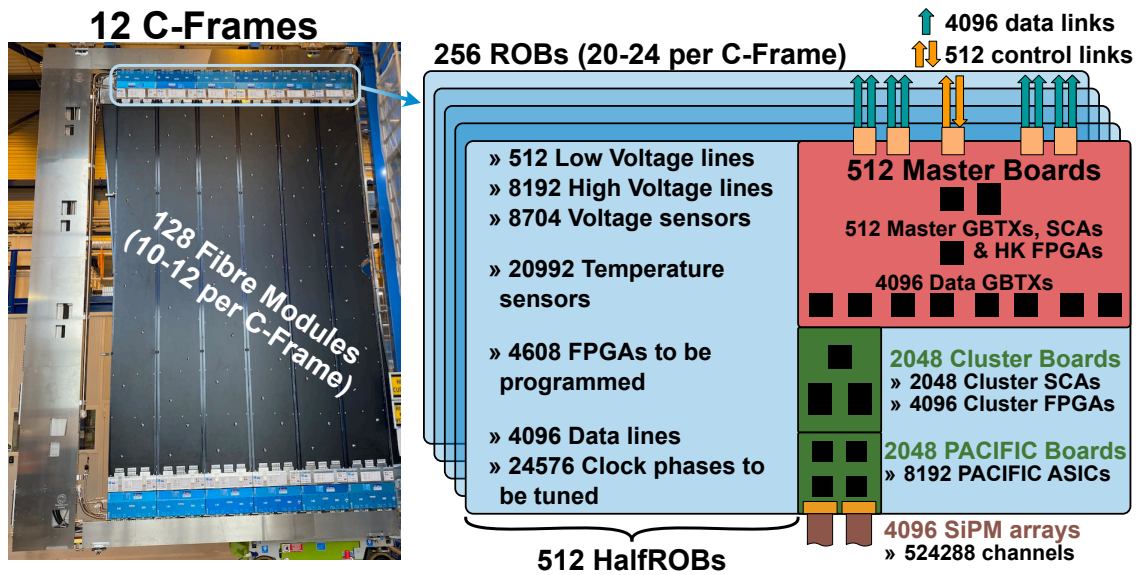


Figure 5.3: Illustration of the complexity of the SciFi front-end electronics. The given values correspond to the number of components installed in the complete detector.

functioning properly and, if not, that the necessary calibration and tuning measures are carried out.

In the scope of this thesis, a detailed commissioning procedure has been defined, developed and implemented. In a series of tests, the correct functioning of the front-end electronics is ensured. In addition, the electronics properties are validated at the system level while being installed on the C-Frames. The complete commissioning procedure is presented in Section 5.3.4 along with the discussion of the results throughout the course of this chapter. A key aspect of it is the verification of a stable data transmission over the total of 4096 data links. Establishing the 40 MHz readout required for that purpose was a major part in the framework of this work.

5.2 Front-End Tester

Before being mounted on a C-Frame, the proper functioning of each individual ROB is verified by a dedicated test stand located in an electronics lab on the LHCb site. The so-called SciFi Front-End Tester (FE-Tester) comprises a signal injection system that was specifically designed for the project. A FE-Tester with a mounted ROB is shown in the photograph in Fig. 5.4, along with a schematic drawing of the complete setup. The FE-Tester itself consists of one control board and eight injector modules with 256 channels each in order to cover the 2048 channels of one ROB under testing. Two of these systems are in operation for the quality control of the SciFi front-end electronics [104].

The connection to the DAQ server is established via a bidirectional optical link. Like the Master Board, the control board features a Master GBTX (see Section 4.4.1) and SCA (Section 4.4.2) to enable the communication. It shares the same back-end electronics as the ROB under testing, as well as an extended version of the control software used in the commissioning, which has been contributed as part of this thesis. Further details about the DAQ system and software are given in the following sections.

The main functionality of the FE-Tester is the injection of calibrated charge pulses into the 2048 PACIFIC channels of the examined ROB. It is thus possible to verify the functioning of each channel by comparing the detected signal amplitude with the expected value, as well as evaluating the noise level. In addition, the FE-Tester allows for the independent charge injection into selected channels, thereby enabling the detection of short-circuited channels. The connection to the PACIFIC channels is established via the same 80-pins, 0.5 mm pitch connectors used for connecting the SiPM arrays (see Section 3.6). However, defects of these high density connectors are to be expected after about 50 mating cycles. Therefore, an interconnected adapter board (connector saver) is utilised to allow for the testing of the larger number of ROBs.

Besides from the main test sequence involving the charge injection from the FE-Tester, each ROB is subject to a series of further checks while being mounted on the test stand. These include the reading of hardware IDs of the individual components

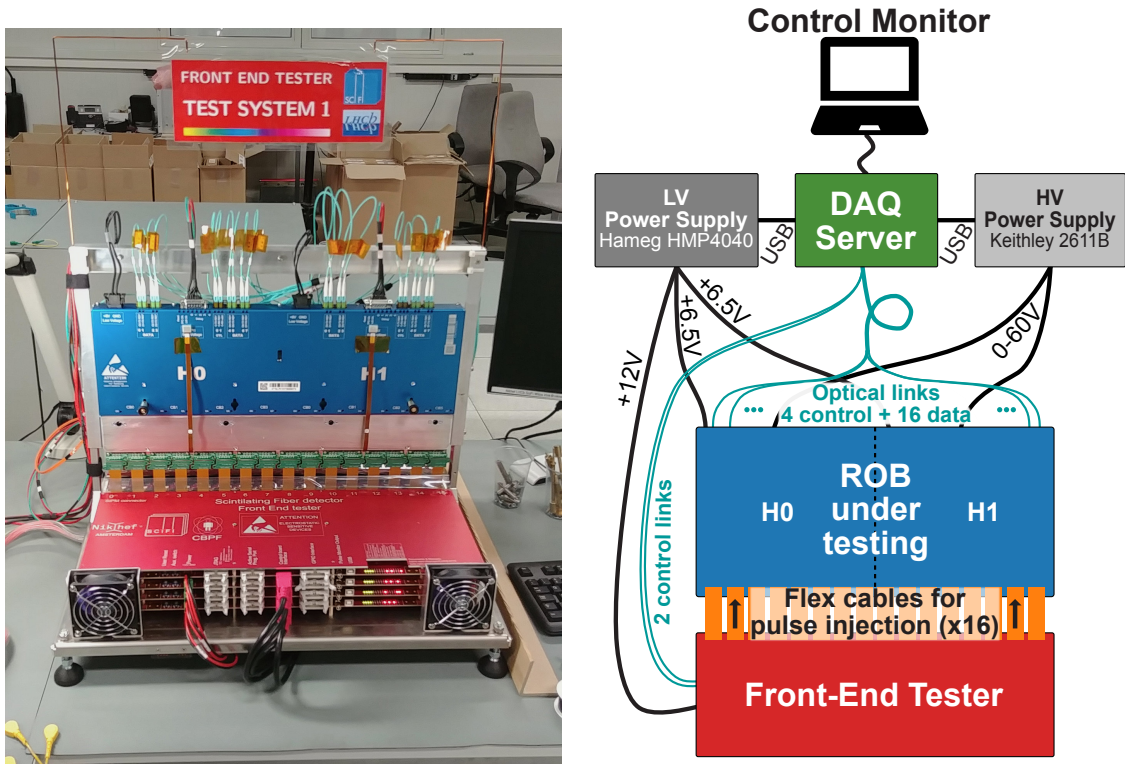


Figure 5.4: Photograph of the SciFi Front-End Tester with a mounted ROB under testing (left). A schematic drawing of the complete setup including the DAQ system and power supplies is shown on the right. Two of these test systems are used to verify the proper functioning of the ROBs before being installed on the C-Frames.

like GBTX and SCA ASICs. By comparing them with the expected values as stored in a database during the assembly of the ROB, it is assured that only intended boards and components are used. Furthermore, the ROBs are prepared for the following commissioning on the C-Frame by programming the FPGAs with later firmware versions, while also testing the programming functionality itself. The commissioning preparations also include a (pre-)calibration of the temperature and voltage monitoring circuits. Further details can be found in Sections 5.8 and 5.10.

After a ROB passed the complete test procedure, it undergoes an optical inspection and cleaning of the VTRx and VTTx modules before being cleared for the commissioning. Figure 5.5 shows a photograph of six ROBs that have been mounted on a C-Frame after passing the tests on the FE-Tester.

5.3 Commissioning Setup

As discussed previously, the commissioning of the SciFi front-end electronics is performed directly following the assembly of each C-Frame. It is carried out in one of two available commissioning slots in the back part of the SciFi assembly hall, close



Figure 5.5: Photograph of six ROBs mounted on the bottom half of a C-Frame and ready to be commissioned.

to the infrastructure needed for that purpose.

Depending on the tracking station, each C-Frame comprises up to 24 ROBs that are commissioned at once¹. All the ROBs have been examined using the two FE-Tester setups prior to being mounted. The aim of the commissioning is to ensure the functionality of the ROBs on the C-Frame within their final configuration of connected power cables, optical fibres, LIS and SiPMs. The connection of the SiPM arrays is an especially delicate operation due to the high density connectors. In the assembly slots, easy access to the different components is provided in case of detected defects during the commissioning. Afterwards, the C-Frames are moved to the TS-Cage and transported as a whole to the LHCb cavern, which has been found to be a gentle process in previous (test) transports.

During the commissioning, the ROBs are already referenced by their location within the final detector. The location identifier (ID) is composed of the station number $T = [1, 2, 3]$, layer within the station $L = [0, 1, 2, 3]$, quadrant within the layer $Q = [0, 1, 2, 3]$, and module within the quadrant $M = [0, 1, \dots, 5]$. Note that the

¹Some commissioning steps cannot be done in parallel for all ROBs, but are performed in two iterations as discussed in the following sections.




Figure 5.6: Schematic view of the SciFi Tracker with the naming convention overlaid. An example ROB and its location identifier is highlighted in red. Distances between the layers (L) and stations (T) are exaggerated for better visibility. The inner layers per station (L1 & L2) are tilted by $\pm 5^\circ$ to obtain the y -coordinate of detected hits.

module number only ranges from 0 to 4 for T1 and T2. The naming convention, as well as the position of an exemplary ROB with location ID T3L2Q3M2, is illustrated in Fig. 5.6. During the commissioning, the location ID of the ROB is often appended with additional information to further localise a component or a problem. As outlined in Section 3.8.3, each C-Frame comprises two layers with two quadrants each. Thus, it is common to also refer to the C-Frames in terms of T, L and Q (e.g. T3L23Q13).

5.3.1 Infrastructure

The infrastructure used for the commissioning represents a subset of the final system. It consists of a MARATON power supply that delivers the low voltage for the front-end boards. With 12 output channels set to 8V, two ROBs are powered by one channel. The SiPM channels are biased by three CAEN modules of type A1539BP inside a SY4527 crate (see Section 3.8.1). Each of the 96 available HV channels power four neighbouring SiPM arrays corresponding to one fibre mat. The front-

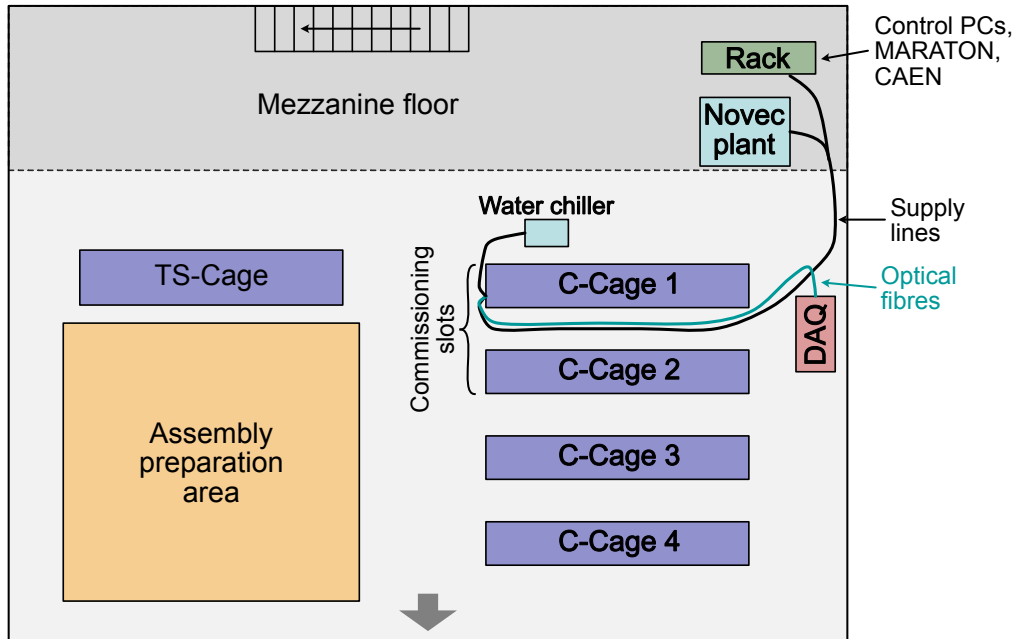


Figure 5.7: Floor plan of the SciFi assembly hall.

end electronics are cooled by a portable water chiller that provides a water flow of 2.5l/min at 20°C to each quadrant. A Novec cooling plant is located next to the LV and HV power supplies on the mezzanine floor in the back part of the assembly hall. During the electronics commissioning of the first C-Frames, the SiPMs were cooled to -40°C . However, due to time constraints, the vacuum and Novec cooling tests were separated and performed prior to the electronics commissioning for the later C-Frames. A floor plan of the assembly hall including the commissioning setup and infrastructure is shown in Fig. 5.7.

5.3.2 DAQ System

The DAQ system used for readout and control of the front-end electronics is located next to the C-Cages. It comprises a 4U dual processor server¹ equipped with two 10-core Intel Xeon Silver 4114 processors. The main components are nine PCI Express cards that act as the interface to the front-end electronics. Since they are specially designed for the trigger-less readout of the LHCb experiment at 40 MHz, they are referred to as PCIe40 cards [35].

The PCIe40 card embeds a powerful FPGA² with 1.15 million logic cells. Its purpose is to enable the flexible implementation of the GBT architecture at the back-end side. By programming the FPGA with different firmwares, it can fulfil three different roles:

¹SuperServer 4029GP-TRT by Super Micro Computer, Inc.

²Arria 10 10AX115 by Altera (now Intel Corporation)

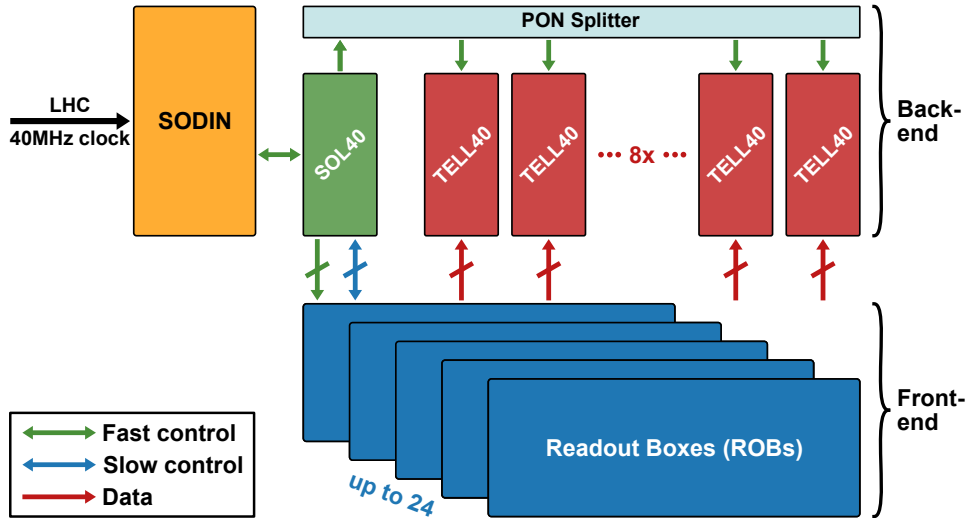


Figure 5.8: Interplay of the back- and front-end electronics in the upgraded LHCb DAQ system as used for the commissioning. During the commissioning, the functionality of the SODIN and SOL40 is combined within a single FPGA.

- The **TELL40** cards receive the data from the front-end electronics and are responsible for merging events belonging to a single bunch crossing.
- The **SOL40s** are the direct interfaces for the communication and control of the front-end boards.
- The readout supervisor **SODIN** is responsible for distributing the 40 MHz LHC clock and fast commands to the other cards.

Each card supports up to 48 bidirectional optical links by means of 4+4 12-channel fibre optics transmitter and receiver modules, called MiniPODs¹. Eight MPO (Multi-fiber Push On) connectors are located on the board to allow for the connection of the optical fibres going to the front-end electronics.

Figure 5.8 illustrates the interplay between the SODIN, SOL40 and TELL40 at the back-end, as well as the front-end electronics. For the SciFi commissioning, eight PCIe40 cards are configured as TELL40s while the remaining one is programmed with a combined SODIN/SOL40 firmware. The latter distributes the clock and fast commands to the TELL40 cards via SFP+ (enhanced small form-factor pluggable) transceivers and a PON (passive optical network) splitter.

Depending on the tracking station, each C-Frame comprises up to 384 data and 48 control links. Thereby, the available PCIe40 cards are just enough for the commissioning of one C-Frame at a time. However, at the time of the commissioning, no TELL40 firmware was available that allowed for the simultaneous operation of all 48 links. Therefore, two separate TELL40 firmwares are in use, each with 24

¹MiniPOD AFBR-811FN3Z (Tx) and AFBR-821FN3Z (Rx) by Avago Technologies (now Broadcom Inc.)

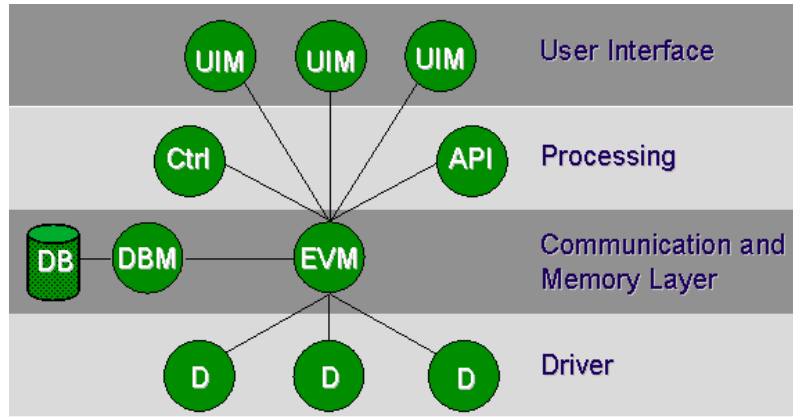


Figure 5.9: Structure of a WinCC OA project. Image taken from Ref. [106].

instantiated data links. The active links are logically split into the lower (quadrants Q0/Q1) and upper (Q2/Q3) half of the C-Frame. As a result, the commissioning steps that require the data taking capabilities need to be performed separately for both halves. Conversely, all ROBs on the C-Frame can be controlled with a single SODIN/SOL40 firmware.

In general, the firmwares for the PCIe40 cards were still experimental and under heavy development especially in the early phase of the commissioning. Therefore, the debugging of the DAQ system was conducted in parallel and several attempts were required in order to obtain a (mostly) stable system. Once a set of functioning firmwares were found, their versions were kept and used throughout the complete commissioning campaign. As a consequence, not the full feature set was available and some workarounds were required in order to achieve the desired results. For compatibility reasons, this also includes the firmwares of the Housekeeping and Cluster FPGAs on the front-end electronics itself.

5.3.3 Control Software

The software to control the SciFi front- and back-end electronics has been developed within the SCADA (Supervisory Control and Data Acquisition) system SIMATIC WinCC OA [105]. A WinCC OA system, also referred to as project, follows a distributed architecture consisting of multiple autonomous processes that are called managers. These can be control managers that are executing program code written in the WinCC OA CTRL language whose syntax shows strong similarities with that of the C programming language. In addition to that, an API (Application Programming Interface) exists that allows to execute external C++ programs in the context of the project. Other examples are user interface managers that are responsible for the visual representation of control panels and the handling of user input. Screenshots of WinCC OA control panels can be found in the following sections.

The typical structure of a WinCC OA project is depicted in Fig. 5.9. While multiple control (Ctrl) and user interface managers (UIM) can be present, a single

event manager (EVM) forms the core of the project. Via an intermediate database manager (DBM), it has access to the project's database (DB). All communication between the various processes, including the drivers (D) that establish the connection to the physical devices, goes through the event manager and the interaction with the database. The objects stored in the database are called datapoints and can follow different structures. These are defined by the datapoint types that can be thought of as (nested) structs within the C programming language and its derivatives. Typically, physical devices like the Master Board with its components are represented as datapoint types.

Besides from the inherent distributed architecture, it is even possible to connect different WinCC OA projects with each other. This is enabled by the distribution managers in each project that allow for the combination of hundreds of systems. Due to its good scalability, WinCC OA is used for controlling all LHC experiments at CERN [107]. However, WinCC OA can not be understood as a ready-to-use control system, but it rather provides the underlying tools to built such a system depending on the individual needs.

JCOP Framework

For that reason, a joint controls project (JCOP) was initiated by the four LHC experiments [108]. Due to similar requirements, it aims at avoiding the duplication of work during the development of new control systems. In addition to including a set of packages with common control solutions, it also provides a framework that allows for the creation and distribution of further JCOP components.

Two of these additional packages are the fwGbt and fwHw components [107]. The fwGbt component integrates the GBT architecture into WinCC OA by providing an interface to the various protocols that the GBT chipset supports (see Sections 4.3.1 and 4.4.1). It is used in close conjunction with the GbtServ that performs the low-level communication with the PCIe40 cards and their firmwares. As such, it forms the bridge between the control system and the hardware level. The fwHw component is designed to model the structure of the front-end electronics in a user-friendly way by hiding the complexity of the underlying WinCC OA datapoint structures. Combined, the fwGbt and fwHw packages allow for the interaction with the different hardware devices without the need to specify parameters like the protocol or address, as these are already defined in the model.

On a macroscopic level, the devices are organised in a hierarchical structure. The tree-like structure consists of interconnected nodes, where each node contains one or multiple children. On the lowest level, representing the leafs of the tree, there are the physical devices. Conversely, the nodes describe logical accumulations of the actual hardware.

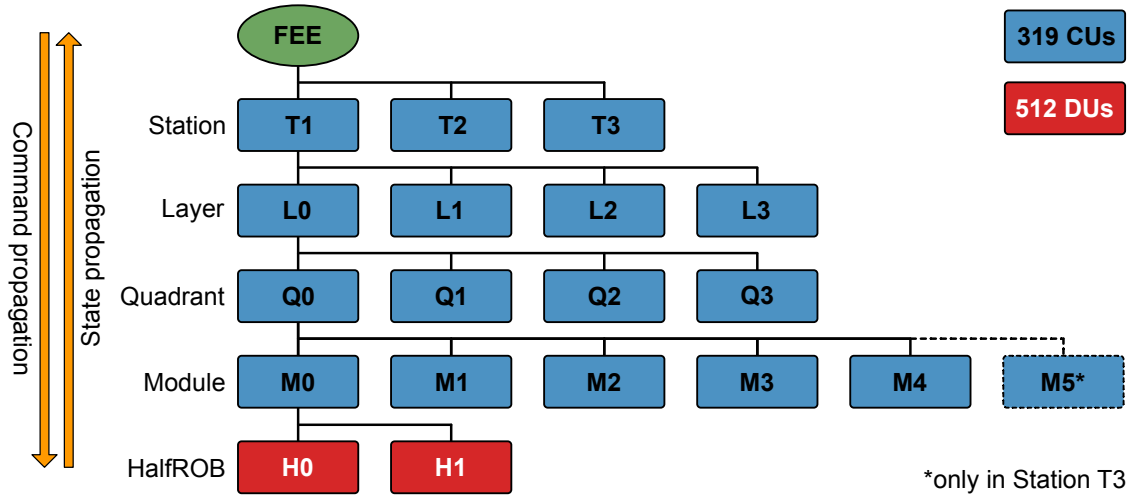


Figure 5.10: Hierarchical structure the SciFi front-end electronics (FEE) as implemented in the control software. It is usually referred to as FSM tree. Image inspired from Ref. [109].

Finite State Machine (FSM)

The creation of the hierarchical structure is enabled by the fwFSM component provided in the JCOP framework [108]. Within the implementation of the package, a node is referred to as control unit (CU) and a leaf as device unit (DU). The naming is inspired by the term FSM (Finite State Machine), since each unit (CU or DU) is assigned a fixed number of states. Transitions between the different states are possible by issuing a command to the unit, while the set of available commands depends on the current state. Since a CU is not a physical device in itself but a logical accumulation, its state depends on the combined states of its children. As a result, when any child is in an erroneous state, the parent CU is in an erroneous state as well. In turn, issuing a command to a CU results in it being forward to its child units. Summarising the above, it can be said that state changes are propagated up the hierarchy (from leaves to the root of the tree) while commands are moving down the hierarchy (root to leaves). For the operator, this allows for a clear overview when controlling and monitoring the detector.

The so-called FSM tree of the SciFi front-end electronics as implemented in the control software is displayed in Fig. 5.10. As can be seen, the hierarchy is strongly oriented towards the naming convention of the ROBs introduced in Section 5.3. Hence, the tracking stations, layers, quadrants and modules each form a new layer of CUs. On the lowest level are the HalfROBs that represent the DUs in the tree. Since each HalfROB is controlled by a separate (bidirectional) control link, it is the logical choice for the smallest unit within the hierarchy. Figure 5.11 shows a screenshot of the Device Editor Navigator (DEN) that is used in WinCC OA to create the tree-like structure.

The possible states and transitions within the FSM of the HalfROBs are illus-

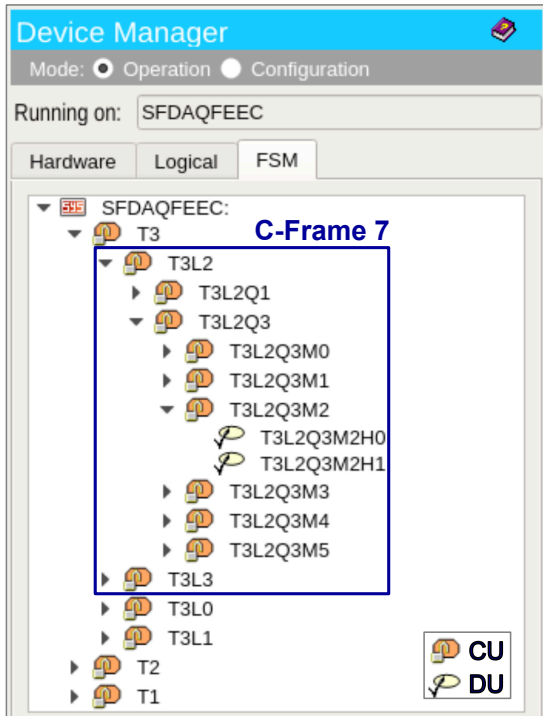


Figure 5.11: Screenshot of the device manager used to create and control the hierarchical structure of the SciFi front-end electronics.

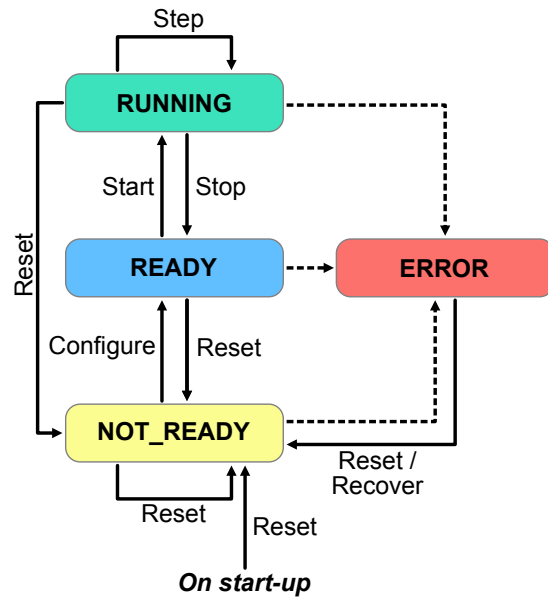


Figure 5.12: Finite State Machine (FSM) inside the device unit (DU) of the SciFi front-end electronics project. Each DU is representing one HalfROB.

trated in Fig. 5.12. On start-up, issuing a Reset command brings the DU into the NOT_READY state in which all the components on the front-end boards are powered. However, a subsequent Configure command is required to get the HalfROB READY for taking data. During a data taking run, the DUs are in the RUNNING state. Any failures during the execution of the different commands result in the transition to ERROR from any of these states. Further details about the different states and commands are given in the following sections.

Based on the aforementioned building blocks, the control system for the C-Frame commissioning has been implemented. Two WinCC OA projects are running on the DAQ server itself and are responsible for the control of the SciFi front- and back-end electronics. Additional projects are installed on supplementary machines in order to regulate the devices providing the infrastructure (see Section 5.3.1). The different projects are integrated into the LHCb network and connected to each other via their distribution managers. Although the system has been developed with the C-Frame commissioning in mind, it also forms the basis for the control system used for the final operation of the detector in the LHCb cavern.

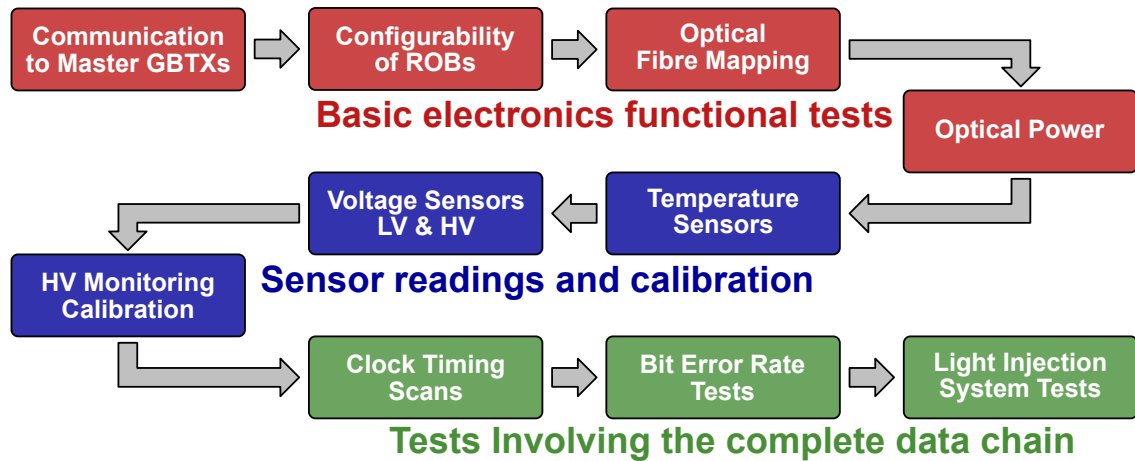


Figure 5.13: Test sequence during the SciFi front-end electronics commissioning.

5.3.4 Commissioning Procedure

The procedure of the SciFi front-end electronics commissioning consists of a number of successive steps. In order to ensure the correct functioning of the ROBs, the test sequence is split into three stages as displayed in Fig. 5.13.

In the first stage, the electronics are tested for their basic functionality. This includes, among other things, checking the communication to the Master GBTX ASICs via the bidirectional optical links, as well as testing the internal communication protocols and registers while configuring the various components within each ROB.

The second stage is about the examination of the readings of the temperature and voltage sensors located on the front-end boards. In case of the HV biasing the SiPM arrays, this includes a recalibration of the monitoring circuit parameters.

The third and last stage consists of extensive tests that involve the data taking functionalities through the complete data chain. These range from injecting test patterns to recording SiPM signals generated by the LIS (see Section 4.5). Given the limitation of the available DAQ system and TELL40 firmwares that only allow to readout half of the data links (see Section 5.3.2), these steps need to be performed twice for each C-Frame.

The complete procedure takes about two working weeks per C-Frame, or longer if defects are found that require a rework or even a replacement of ROBs. However, due to the initial development of tools and the varying situations within the COVID-19 pandemic, the commissioning of the first four C-Frames took significantly more time. Within the scope of this thesis, the test procedure was developed and implemented in software. After overcoming initial challenges and performing the measurements on the first C-Frames, the responsibilities for the further conduction of the different test steps were gradually transferred to other members of the SciFi collaboration.

A list of all C-Frames and the time periods in which they were commissioned is shown in Table 5.1. They are numbered on both sides of the beam pipe from station T3 to T1. Depending on the state of the assembly, the order of commissioning

Table 5.1: Time periods in which the commissioning of the front-end electronics for the 12 C-Frames forming the SciFi Tracker took place.

C-Frame	Side	Location ID	ROBs	Commissioning	
				Started	Completed
C-Frame 1	C	T3L23Q02	24	25.11.2019	24.02.2020
C-Frame 2	C	T3L01Q02	24	06.11.2020	10.02.2021
C-Frame 3	C	T2L23Q02	20	28.02.2020	21.07.2020
C-Frame 4	C	T2L01Q02	20	04.08.2020	02.11.2020
C-Frame 5	C	T1L23Q02	20	28.05.2021	08.06.2021
C-Frame 6	C	T1L01Q02	20	16.06.2021	01.07.2021
C-Frame 7	A	T3L23Q13	24	13.10.2021	28.10.2021
C-Frame 8	A	T3L01Q13	24	23.09.2021	06.10.2021
C-Frame 9	A	T2L23Q13	20	03.11.2021	18.11.2021
C-Frame 10	A	T2L01Q13	20	Q1 2022*	
C-Frame 11	A	T1L23Q13	20	Q1 2022*	
C-Frame 12	A	T1L01Q13	20	Q1 2022*	

* Expected

roughly followed this numbering. The six C-Frames facing the cryogenic facility (C-side) were commissioned at first in order to be installed in the LHCb cavern prior to the beam pipe. Afterwards, the six C-Frames towards the access point (A-side) followed.

The different test steps are fully described in the following sections, along with the results of the 9 out of 12 C-Frames that were fully commissioned at the time of writing.

5.4 Communication to Master GBTXs

The first step after powering up the ROBs by the MARATON power supply is to check the communication between the front- and back-end electronics via the bidirectional optical links. As outlined in Section 4.4.1, the Master GBTX recovers the clock from the incoming data stream and replies in turn with GBT frames consisting of 120 bits each. As depicted in Fig. 4.16 on page 60, a 4-bit header marks the beginning of each frame and is required for synchronisation on the receiving side. Once the header is recognised 15 times in a row, the data stream is considered as locked. Conversely, missing the header more than 4 times without detecting 8 consecutive valid headers in-between results in a loss of the acquired frame-lock.

For each link of the SOL40 card, there is a register that contains the counted

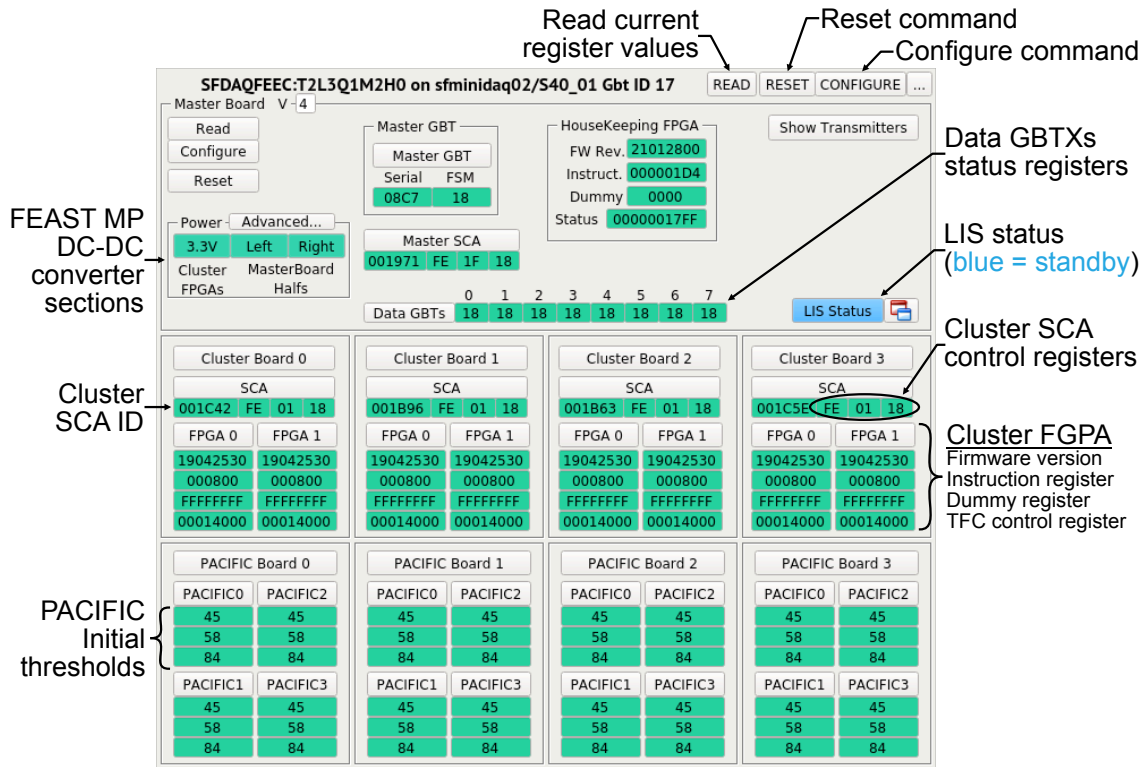


Figure 5.14: WinCC OA control panel showing the readout values of the main registers of one HalfROB. Intended for the use by experts for debugging purposes.

number of clock cycles without frame-lock. When stable connections to the Master GBTX ASICs exists, after powering up the ROB's and resetting the counters, the values should stay at zero. For that purpose, the counters of the 48 SOL40 links are monitored in a dedicated WinCC OA control panel.

During the commissioning of the first 9 C-Frames, corresponding to 392 HalfROBs, it occurred only once that the connection to a Master GBTX could not be established. As a consequence, the affected ROB T3L3Q3M4 on C-Frame 7 was replaced with a spare in order to get repaired.

5.5 Configurability of ROB's

After verifying the connection to the Master GBTXs, the next step is to test the internal communication protocols and control registers of the various components inside the ROB's. These tests are carried out in the process of fully configuring the boards such that they are ready for data taking. As depicted in Fig. 5.12 on page 81, after powering the HalfROBs, getting them to the READY state requires a Reset followed by a Configure command. The constituent actions of these two FSM commands are presented in the following.

Reset Command

The Reset command can be considered as the minimal set of actions that are required to provide power to all components of the ROB. As shown in Fig. 5.12, it can be issued from any state and is used to put the electronics in a well-defined state. The complete procedure consists of the following consecutive steps:

Reset Master Board (MB)

1. Configure the Master GBTX (366 read/write control registers) to set up the E-Links and the Data GBTX reference clocks (see Section 4.4.1)
2. Reset the Master SCA, enable the required control buses and configure directions of the GPIO pins (see Sections 4.3.1 and 4.4.2)
3. Provide power to all components by successively enabling the three DC-DC converter sections through the corresponding GPIO output lines (see Section 4.4.2)
 - a) Left half of the MB including CB0, CB1 and PB0, PB1, as well as the HK FPGA
 - b) Right half of the MB including CB2, CB3 and PB2, PB3
 - c) Cluster FPGAs on all four CBs
4. Reset and configure the Data GBTXs (366 read/write control registers) to provide the external clocks to the Cluster FPGAs and PACIFICs (see Section 4.4.3)

Reset Cluster Boards (CBs)

1. Reset the Cluster SCAs, enable the required control buses and configure directions of the GPIO pins (see Section 4.3.1)
2. Perform device resets of the Cluster FPGAs through the corresponding GPIO output lines (see Sections 4.3.1 and 4.3.2)

Reset PACIFIC Boards (PBs)

1. Reset the four PACIFICs on each PB (see Section 4.2.3)
 - a) Perform a global digital reset (GPIO output signal: RESET)
 - b) Load the initial register values (GPIO output signal: LDINIT)

Each operation is verified by reading back the values from the corresponding control registers. Only if they are matching the previously written values, the procedure continues with the next step. In case of a mismatch, an error message is raised and the HalfROB DU changes to the ERROR state. Otherwise, once all steps have been completed successfully, the FSM state of the DU switches to NOT_READY.

Due to the implementation of the E-Links, all four Cluster and PACIFIC Boards can be handled in parallel. The control software has been developed with that




Figure 5.15: Time course of the input voltage and current consumption of one HalfROB during the processing of the Reset command as measured with an oscilloscope. The slight changes in the input voltage are due to voltage drops in the power cable itself.

functionality in mind and is addressing all boards simultaneously wherever possible. Similarly, on the level of the FSM tree, issuing a command to a CU will be forwarded and processed concurrently by the DUs below. In principle, this allows to reset all HalfROBs on a C-Frame in the same time as an individual one.

Figure 5.15 shows the change of the current consumption and input voltage of one HalfROB during the Reset procedure as measured with an oscilloscope, starting from powering the left half. As can be seen, the complete process takes in the order of a few seconds. The largest increases in the power consumption are related to enabling the three DC-DC converter sections in the beginning, as well as the configuration of the Data GBTXs. The measurement was performed within the commissioning setup as described in Section 5.3. Hence, the probed HalfROB was powered by a MARATON channel operating at 8 V. The slight reduction in the input voltage results from increasing voltage drops in the 30 m LV cable going from the power supply to the C-Frame.

During the commissioning of the first C-Frame, stability issues were encountered that even led to communication losses to the HalfROBs in some cases when issuing the Reset command. It was traced back to be caused by drops of the input voltage that could fall below the level of 6.5 V, which is required for the operation of the FEASTMP DC-DC converters. It was found that an additional capacitance is required to stabilise the input voltage and also to suppress ringing. Therefore, specially designed 10 mF capacitor boards consisting of multiple radiation hard ceramic capacitors were produced [110]. They are connected to each LV splitter

on the C-Frames that distribute the power from one MARATON channel to four neighbouring HalfROBs. Like this, the voltage drops during the Reset procedure are reduced and the input voltage is kept above the limit of 6.5 V as shown in Fig. 5.15.

Configure Command

As depicted in Fig. 5.12, the Configure command can be issued to DUs within the NOT_READY state in which they are found after a successful Reset. The purpose of it is to prepare the front-end electronics with the proper configuration parameters needed for taking high quality data. In contrast to the Reset command that performs the same actions and writes the identical values to the control registers of all HalfROBs, settings individual for each HalfROB are loaded during the Configure command. The values are determined during the production QA tests of the individual components and saved in the SciFi production database (see Section 4.6).

This applies to the parameters I_{bias} and I_{mod} of the LIS, as well as the PACIFIC Boards. After the QA procedure, a configuration file is uploaded to the database for every PACIFIC Board. Each file contains the operational values of the 336 read/write control registers of each of the four assembled PACIFIC ASICs. At this stage, the reference voltage parameters V_{REF} and V_{refDCFB} , as well as the trim DACs are specific to each chip, while the remaining values are common to all. During the Configure command, the configurations are parsed from the files and applied to the devices.

Besides from the control register values, a PACIFIC Board configuration file also contains the serial number of the board. Before applying the configuration, the serial number is read from the board and compared to the value from the file. This is particularly important in the commissioning as it allows to verify that the optical control links are connected to the intended HalfROBs. To confirm the mapping of the data links, an additional step is required that includes the start of a data taking run. In preparation for this, the 10-bit location ID of the HalfROB is written to a register of the HK FPGA during the configuration. Further details are given in Section 5.6.

In the last step of the Configure command, the Master and Cluster SCAs are prepared for the monitoring of the temperatures and voltages. This includes the activation of the ADCs, as well as the switchable 100 μA current sources required for the operation of resistance thermometers.

Completing all steps of the Configure command successfully results in the HalfROB to change to the READY state, otherwise it goes to ERROR. During the commissioning of the first 9 C-Frames, the configuration of a single HalfROB failed because the communication to the LIS GBLDs could not be established. The issue could only be solved by replacing the affected ROB T2L3Q3M0 on C-Frame 9 with a spare.

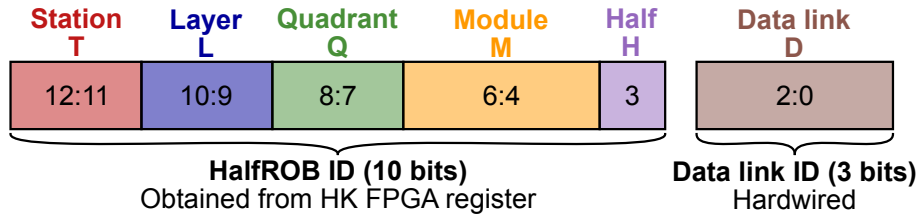


Figure 5.16: Structure of the 13-bit link ID used to uniquely identify the 4096 data links of the SciFi Tracker.

5.6 Optical Fibre Mapping

While the correct cabling of the control links is confirmed in the previous commissioning step when configuring the HalfROBs, further actions are required to verify the mapping of the data links. The process is based on sending a unique identifier (ID) on each link to the TELL40 cards. The structure of the so-called (data) link ID consisting of 13 bits is illustrated in Fig. 5.16. It comprises the 10-bit HalfROB location ID followed by a 3-bit identifier to describe the data link within the HalfROB.

As shown in Fig. 4.1 on page 40, there are eight data links per HalfROB, each one corresponding to one Data GBTX and Cluster FPGA. The 13-bit link ID is constructed in the Cluster FPGA based on the input from 13 individual pins. The 10-bit location ID of the HalfROB is routed from the HK FPGA to all Cluster FPGAs. Its value is determined by the content of the HK FPGA instruction register that is assigned with the correct ID during the configuration of the HalfROB. The remaining 3 bits are composed as follows: Each FMC connector between the Master and Cluster Boards has two dedicated pins that are connected individually to ground or 1.5 V on the Master Board, thereby encoding the location of the Cluster Board within the HalfROB. As already introduced in Section 4.3.2, the same principle is applied on the Cluster Boards themselves to determine the position of the Cluster FPGA on the board.

In order to transfer the information to the TELL40s, a data taking run has to be initiated. However, in the general data format shown in Table 4.2 on page 57, there is not enough bandwidth available for the additional transmission of the 13-bit link ID. Instead, the information is embedded into special GBT frames that are leading every run. The original purpose of these so-called SYNC¹ frames is the re-synchronisation of the bunch crossing IDs in the TELL40s [111]. In addition, it is required for determining the position of the first header in case of the FV data format (see Section 4.3.2). The generation of the SYNC frames by the Cluster FPGAs is triggered by sending the SYNC fast command to the HalfROBs. The content of each SYNC frame is the 12-bit bunch crossing ID followed by a 12-bit programmable SYNC pattern that is set to 0xB4C. The 13 least significant bits (LSBs) are occupied with the link ID. In order to allow the TELL40s to recognise the SYNC pattern,

¹Not to be confused with the SYNC pulse that is sent to the PACIFICs for synchronisation. Accidentally, the same naming has been used even though they are not directly related.

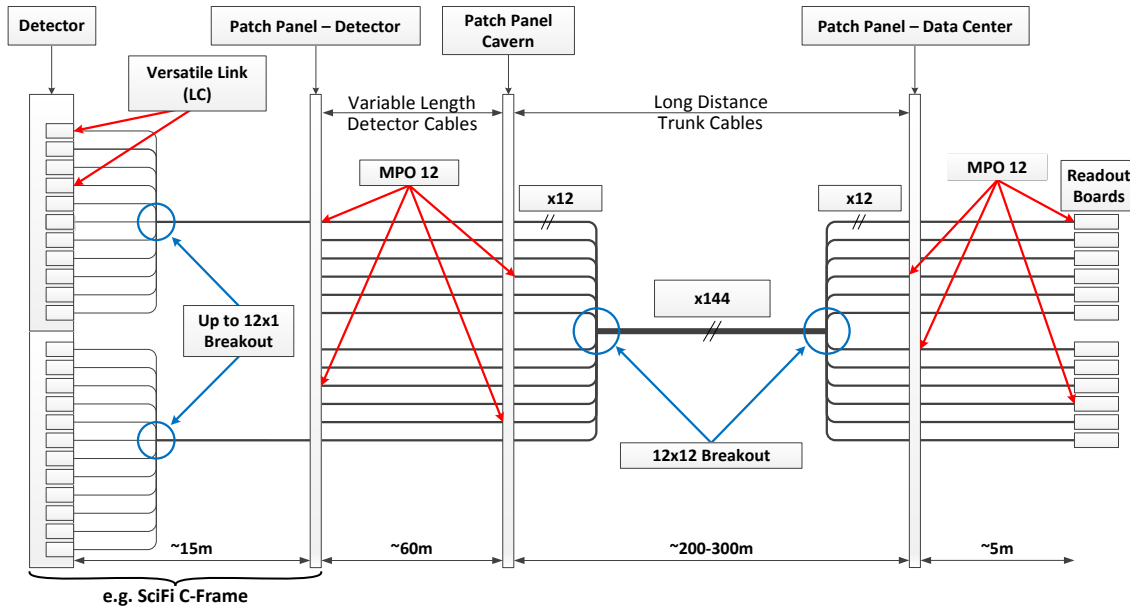


Figure 5.17: Fibre system layout for the readout and control of the front-end electronics of the upgraded LHCb experiment. The complete system will contain about 15 000 individual active fibres [34]. With 4096 data and 2×512 control fibres, the SciFi Tracker contributes more than 30 % to that. Image adapted from Ref. [112].

several consecutive SYNC frames are generated by issuing multiple (typically ten) SYNC commands in a row. By default, this process is performed at the start of every data taking run.

Within the TELL40 firmware, the link ID is captured from the SYNC frames and stored in a dedicated register that can be read out by the control system. By reading out these registers for every TELL40 link after initiating a short run, the Cluster FPGAs on the other end of the optical fibres can be determined and compared to expectations.

Like this, the correct mapping of the optical data links could be confirmed on the first 9 commissioned C-Frames. However, one ROB on C-Frame 8 needed to be replaced due to an issue encountered in this step. Within T3L1Q1M2H1, the two neighbouring data links 6 and 7 were outputting incorrect link IDs. The issue was traced back to a faulty connection between the corresponding Cluster Board 3 and the Master Board which caused the transmitted layer (L) bit from the HK FPGA to be stuck at 0.

5.7 Optical Power

After verifying the correct fibre mapping in the previous commissioning step, a measurement of the optical power is performed. Even though every individual fibre as well as the VTRx and VTTx modules are cleaned and optically inspected prior






Figure 5.18: Results of the optical power measurements as presented in the commissioning report of C-Frame 7: Heatmap of the control links (a) and histogram of all links (b). In addition, an accumulated histogram for the first 9 commissioned C-Frames is shown (c).

to connecting, it has proven to be a useful tool to cross-check for good optical links. The latter are particularly important when operating the detector in the LHCb cavern. As shown in Fig. 5.17, there are several hundred meters of optical fibres and a minimum of 3 breakpoints between the front-end electronics in the cavern and the readout boards in the data center on the surface of the LHCb site [112]. In contrast, the fibre system layout within the commissioning setup only consists of 1 breakpoint (2 for the control links) and 15 m + 60 m of optical fibres.

Within the commissioning, the received power from the optical transmitters of the ROB is measured directly with the MiniPODs that are installed on the PCIe40 cards. For that purpose, the MiniPOD receiver modules feature a Two Wire Serial (TWS) interface to monitor the optical input power, along with other diagnostic information like module temperatures and supply voltages [113]. However, no official specifications regarding the tolerances of these quantities were available. Therefore, a cross-check was performed by remeasuring the optical powers of 96 links with a

calibrated handheld power meter¹. Both sides of the fibre connections were thoroughly cleaned between measurements. By doing so, differences up to 10% between the optical power measurements of the MiniPODs and the handheld power meter were found. This is in the same order as the specified tolerance of ± 0.35 dB (corresponding to about 8%) of the power meter itself [114]. Therefore, it was concluded that the accuracy of the MiniPODs is sufficient for the given application.

For each commissioned C-Frame, an automated report on the optical power readings of the MiniPODs is generated. It includes a histogram of the measured values, as well as heatmaps of the individual control and data fibres. The latter allow for the identification of links that lie outside the acceptable range and require further actions. An exemplary histogram and heatmap from the optical power commissioning report of C-Frame 7 are displayed in Figs. 5.18a and 5.18b.

The optical power must not only be strong enough for a stable transmission within the commissioning setup, but also for the final fibre system layout as shown in Fig. 5.17. Therefore, a limit of $224 \mu\text{W}$ (-6.5 dBm) is applied to all links. The unit decibel-milliwatt (dBm) states the optical power in decibels (dB) relative to 1 mW:

$$P [\text{dBm}] = 10 \cdot \log_{10} (P [\text{mW}]/1 \text{ mW}) . \quad (5.1)$$

It is often used because of the convenience to incorporate losses (in dB) through simple addition.

The limit is derived based on a previous study that evaluated the Versatile Link power budget in the context of the LHCb experiment assuming 400 m of optical fibres and three breakpoints [99]. The relevant values from that study in the upstream direction, i.e. the transmission from the front- to the back-end electronics, are summarised in Table 5.2. The resulting minimal transmission power of -6.5 dBm ($224 \mu\text{W}$) is taken as the limit for the commissioning. Note that this assumes a lossless transmission within the commissioning setup, thus providing an additional safety margin in the order of 1–2 dB.

In practice, a soft limit of $300 \mu\text{W}$ is used during the commissioning. For links below that value, an additional cleaning iteration is carried out. In the process, the transmitter modules as well as the fibre ends are optically inspected and, if required, cleaned again. Despite the additional effort, 9 control and 10 data links still had an optical power reading below $300 \mu\text{W}$ within the commissioning of the first 9 C-Frames, as shown in Fig. 5.18c. Measured against a total of 3528 upstream links, this corresponds to a share of 0.5%. However, the optical power was still above the critical limit of $224 \mu\text{W}$ in all cases and the components were therefore released for operation.

¹MultiFiber™ Pro by Fluke Networks

Table 5.2: Optical power budget of the final fibre system layout of the LHCb experiment in the upstream direction. Values taken from Ref. [99].

Description	Unit	Power (loss)
Max. Rx sensitivity	μW	78
	dBm	-11.1
Fibre attenuation	dB	0.95
Insertion loss	dB	2.25
Mode dispersion	dB	1.3
Fibre radiation	dB	0.1
Min. Tx power	dBm	-6.5
	μW	224

Table 5.3: Temperature sensors in use on the SciFi front-end electronics.

Part	Type	Amount per		Typical Temperatures	Tolerance
		Part	ROB		
Master Board	NTC (4.7 k Ω)	8	16	20 °C to 35 °C	± 3 °C
→ Master SCA	SCA internal	1	2	20 °C to 35 °C	± 2 °C
Cluster Board	Pt1000	3	24	20 °C to 35 °C	± 2 °C
→ Cluster SCA	SCA internal	1	8	20 °C to 35 °C	± 2 °C
PACIFIC Board	Pt1000	2	16	20 °C to 35 °C	± 2 °C
SiPM array	Pt1000	1	16	-40 °C (to 30 °C)	± 2 °C

5.8 Temperature Sensors

As introduced in Chapter 4, the SciFi front-end electronics are equipped with various resistance thermometers that enable the temperature monitoring of the boards and components. An overview of the temperature sensors in use is given in Table 5.3. Ensuring the correct functioning of these sensors is the task of this commissioning step. Moreover, monitoring the temperatures of the water cooled electronics has proven to be a useful tool to detect an incorrect mounting of the ROBs on the C-Frame.

Tolerances

During the commissioning of the first C-Frame, variations by more than 20 °C between neighbouring sensors were observed, which considerably exceeded the expected tolerances. It resulted from the fact that no variations in the 100 μA current,




Figure 5.19: Distribution of the currents as generated by the Cluster SCAs of the 196 ROBs installed on C-Frames 1 to 9. Measured as part of the FE-Tester test procedure. A Gaussian fit is overlaid.




Figure 5.20: Readings of the temperature sensors during the commissioning of the first 9 C-Frames.

which is generated by the GBT-SCAs to operate the temperature sensors (see Section 4.3.1), were taken into account. In order to increase the accuracy of the sensors, the current source of each SCA had to be measured accurately. To achieve this, the connector saver boards of the FE-Tester were redesigned and now feature a high precision $4.7\text{ k}\Omega \pm 0.01\%$ resistor [104]. It connects to the same ADC input line that is normally used for the monitoring of the SiPM temperature via a Pt1000. By knowing the resistance and measuring the voltage drop with the ADC, the generated current by the SCA can be calculated using Ohm's law as

$$I_{\text{src}} = \frac{V_{\text{ADC}}}{R} = \frac{V_{\text{ADC}}}{4.7\text{ k}\Omega \pm 0.01\%}. \quad (5.2)$$

The uncertainty of the voltage measurement is given by the integral nonlinearity (INL) of the ADC that amounts to $\text{INL} < 2\text{ LSB}$ in case of the GBT-SCA [85]. With the 12-bit resolution ADC covering the range between 0 V and 1 V (see Section 4.3.1), this corresponds to an error $\Delta V_{\text{ADC}} = 0.49\text{ mV}$. Compared to the high precision resistor, it is the dominating uncertainty of the current source measurement resulting in an accuracy of $\Delta I_{\text{src}} = \pm 0.1\text{ }\mu\text{A}$.

Using the described method, the current source of all Cluster SCAs installed on the first 9 C-Frames were determined. This was done as part of the test procedure of the ROBs on the two FE-Tester setups. The distribution of the determined values is shown in Fig. 5.19. By performing a Gaussian fit to the data, the mean value and standard deviation are estimated to be about $98\text{ }\mu\text{A}$ and $4\text{ }\mu\text{A}$, respectively.

With the precise knowledge of the current source, the temperature accuracy of the Pt1000 sensors located on the Cluster and PACIFIC Boards, as well as the SiPM arrays, is given by the INL of the ADCs used to measure the voltage drop.

It accounts for an uncertainty of $\pm 1.3^\circ\text{C}$, while the contribution from the current source amounts to $\pm 0.3^\circ\text{C}$. The tolerance of the Class B Pt1000 sensors in use add another $\pm 0.3^\circ\text{C}$ resulting in the total tolerance of about $\pm 2^\circ\text{C}$ for the corresponding temperature readings of the sensors. Note that the uncertainties were simply added up linearly in that case since the underlying distributions are not known. However, since the sources of errors are production tolerances and integral nonlinearities, the total uncertainty of $\pm 2^\circ\text{C}$ can be understood as a tolerance as well, meaning that the true value is in any case within the given interval.

In case of the NTC thermistors located on the Master Boards, the current source is still the limiting factor of the precision. Since the generated current by the Master SCAs cannot be determined in the same way, one has to assume a spread of about $\pm 10\ \mu\text{A}$ that covers 99% of the distribution shown in Fig. 5.19. However, due to a much steeper characteristic curve compared to the Pt1000 sensors, the resulting uncertainty on the temperature measurement only amounts to about 2.4°C . The INL can be neglected in this case, while another 0.7°C has to be added that takes the 3% tolerance of the used NTC thermistors into account.

No further details about the internal SCA temperature sensors were available. However, a similar spread as for the Pt1000 sensors was observed. Therefore, the same tolerance of $\pm 2^\circ\text{C}$ is assigned to both the Master and Cluster SCA's internal sensors.

Temperature Readings

As part of the front-end electronics commissioning, the values of all 82 temperature sensors located on each ROB are recorded. Similar to the measurement of the optical power, an automated report is generated that includes summary plots and highlights malfunctioning sensors. A sensor is considered faulty if it returns an unrealistic reading, meaning a value below the cooling water temperature of 20°C or more than 15°C above.

During the commissioning of the first 9 C-Frames, only 1 of the total of 16 072 temperature sensors was found to be defective. The sensor concerned is a Pt1000 located on the backside of an SiPM array that reads temperatures beyond 80°C . Using a multimeter, it was found that the resistance deviated by about 20% from the nominal value. It was decided to keep the affected SiPM array on the C-Frame and instead mask the temperature readings from it within the control software.

Figure 5.20 shows the temperature readings of the remaining 16 071 sensors, with the cooling water temperature set to 20° . As can be seen, the sensors on the PACIFIC Boards return the largest values. This is because they receive a weaker cooling capacity since they are furthest away from the cooling block that is located just below the Master Boards. In addition, the PACIFICs draw a significant amount of power.

Note that the Novec cooling plant was not used in parallel to the electronics commissioning of the C-Frames. Therefore, the temperature of the SiPM arrays essentially follows the temperature in the assembly hall resulting in large variations




Figure 5.21: Temperature readings of the front-end electronics located on T3L2Q2. With the help of this plot as presented in the automated commissioning report, a mounting deficiency of the ROB T3L2Q2M4 was identified.

throughout the year depending on the time of the commissioning. Due to the water cooling, these seasonal effects are only faintly noticeable at the front-end electronics within the ROBs.

Mounting Check

In addition to the functional tests of the individual sensors, the monitoring and comparison of the temperature readings revealed several deficiencies in the mounting of the ROBs during the commissioning of the first 9 C-Frames. Figure 5.21 shows an exemplary plot of the automated commissioning report that uncovered such an issue on C-Frame 1. It turned out that the significantly higher temperature readings of the ROB T3L2Q2M4 compared to the neighbours were caused by a pinched cable between the ROB and the cooling block, which prevented a good mechanical and thermal contact between the two. Apart from this case, three further ROBs with insufficient thermal couplings to the cooling blocks due to loosely tightened screws were identified and corrected during the commissioning of subsequent C-Frames.

Another type of problem that can be revealed in this way are faulty connections of the SiPM flex cables. Plugging in these high density connectors is a delicate operation that occasionally does not succeed perfectly when connecting up to 384 SiPMs per C-Frame. Depending on the severity, this either leads to a poor conductivity or no electrical connection at all. As a result, unrealistic temperatures are measured by the Pt1000 on the backside of the SiPM array.

This type of error occurred mainly with the first 4 C-Frames. During that time, 10 partially loose SiPM flex cables were detected and corrected in this way. Thereby, valuable time could be saved in the later, more time consuming commissioning steps in which these issues would have appeared as dead channels.




Figure 5.22: Low voltage monitoring circuit. Resistance values taken from Ref. [91].




Figure 5.23: Monitored input voltages of the HalfROBs mounted on C-Frames 1 to 9.

5.9 Voltage Sensors

Following the functional tests of the temperature sensors, the voltage monitoring circuits are examined in this commissioning step. As introduced in Chapter 4, each HalfROB is equipped with one circuit to monitor the low voltage (LV), and 16 circuits to monitor the high voltage (HV) for each SiPM die.

5.9.1 Low Voltage

The monitoring of the HalfROB input voltage is based on the voltage divider as shown in Fig. 5.22, which is located on the Master Board. Using this circuit, the input voltage is given by

$$V_{\text{in}} = \frac{R_1 + R_2}{R_2} V_{\text{mon}} = 21 V_{\text{mon}} \quad \text{with} \quad R_2 = \left(\frac{1}{R_{21}} + \frac{1}{R_{22}} \right)^{-1} = 500 \, \Omega, \quad (5.3)$$

where the monitoring voltage $V_{\text{mon}} \in [0, 1] \text{ V}$ is measured with the ADC of the Master SCA. With the two (effective) resistors R_1 and R_2 in series, each with a tolerance of $\pm 1\%$, only a rough estimation with an accuracy of $\pm 2\%$ is possible¹, which corresponds to about 140 mV for typical input voltages.

The value of the external supply voltage is critical for the correct functioning of the ROBs since the FEASTMP DC-DC converters require an input voltage $> 6.5 \text{ V}$ for a stable operation. During the commissioning, it is verified that the value is above that limit. Figure 5.23 shows the monitored input voltages of the HalfROBs

¹Due to the unknown distribution of the production tolerances, they are again simply added up linearly to estimate the resulting tolerance.




Figure 5.24: High voltage monitoring circuit. Resistance values taken from Ref. [71].




Figure 5.25: Monitored bias voltages applied to the SiPM dies of C-Frames 1 to 9.

mounted on C-Frames 1 to 9 with a median value of 7.1 V. The difference to the set voltage of 8 V at the MARATON power supply is due to voltage drops in the power cables. Significantly lower values would be an indication for unknown resistances in the power line, i.e. due to improper connections of plugs and terminals.

As outlined in Section 5.3.1, one MARATON channel generally powers two neighbouring ROBs that draw a typical current of 20 A, i.e. 5 A per HalfROB. However, since there are only five ROBs per quadrant on the C-Frames belonging to stations T1 and T2, the outer ROBs (M4) on these C-Frames do not share a MARATON channel with a neighbour. Consequently, the voltage drops in the cables are halved for these ROBs, resulting in larger values of the monitored input voltages around 7.6 V.

During the commissioning, it was found that one MARATON channel was outputting a lower voltage than the nominal 8 V. As a result, the HalfROBs that were powered by this channel were monitoring input voltages below 6.8 V, as can be identified as a small peak on the left in Fig. 5.23. However, they were still above the limit of 6.5 V and operating stably throughout the commissioning, along with all the other ROBs mounted C-Frames 1 to 9.

5.9.2 High Voltage

The HV monitoring circuits are located on the PACIFIC Boards. For each HV line powering one SiPM die consisting of 64 channels, a voltage divider as shown in Fig. 5.24 is implemented. To reduce the current consumption due to the monitoring circuits, large resistors were selected. The conversion between the applied bias

voltage and the monitored voltage as measured with the ADC of the Cluster SCAs is given by

$$V_{\text{bias}} = \frac{R_{11} + R_{12} + R_2}{R_2} V_{\text{mon}}. \quad (5.4)$$

Despite the low tolerance resistors ($\pm 0.1\%$) in use, it was decided to perform a calibration of the monitoring circuits when testing the ROBs on the FE-Tester setups. As discussed in Section 3.6, the SiPM gain depends critically on the applied voltage which is why a precise monitoring is required in this case. To perform the calibration, bias voltages in the range from 5 V to 70 V are applied in steps of 1 V using a high precision power supply. By recording the monitored voltage for each bias voltage and performing a linear fit, the conversion factor is determined for each individual HV line. In addition to the slope, which reflects the conversion factor in Eq. (5.4), the intercept is also determined by the fit to account for small (\sim mV) variations due to grounding fluctuations [104].

During the front-end electronics commissioning on the C-Frames, an initial check of the HV monitoring is carried out by applying 50 V to the SiPMs using the CAEN power supply modules. Figure 5.25 shows the monitored bias voltages as calculated using the conversion factors determined during the HV calibration on the FE-Testers. As can be seen, the distribution is not centered around the nominal voltage and shows a relatively large spread. Using a multimeter, sample checks were performed and confirmed that the applied voltage is indeed 50 V and only deviates from it by up to 50 mV in accordance with the specification of the CAEN modules [63]. Therefore, it was concluded that the HV calibration parameters cannot be applied directly from the FE-Tester to the C-Frame commissioning setup and a re-calibration had to be performed. Although the exact reasons for the deviations could not be clarified, differences in the grounding are assumed to play an important role.

At this stage, the monitored bias voltages are considered okay if they differ by less than ± 1 V from the applied voltage. During the commissioning of the first 9 C-Frames, it occurred only once that the measured value was outside of this range: The monitoring circuit of one SiPM die on C-Frame 7 reported a voltage of 0 V. It was found that the SiPM die was indeed not biased, which was due to a broken HV line within the associated ROB. The issue was resolved by replacing the ROB.

5.10 High Voltage Monitoring Calibration

As mentioned in the previous section, it was found that the calibration of the HV monitoring circuits had to be repeated during the commissioning in order to achieve acceptable results. For this purpose, using the CAEN power supply, bias voltages between 20 V and 56 V are scanned in a non-equidistant manner with an emphasis on values around the expected operational voltage (> 50 V). Afterwards, a linear fit is performed in order to calibrate the conversion factor between the monitored and bias voltage as given in Eq. (5.4). An example is shown in Fig. 5.26.

This procedure is applied to all monitoring circuits during the front-end electronics





Figure 5.26: Linear fit to calibrate the bias voltage monitoring with the determined slope (p_1) and intercept (p_0) overlaid. **Figure 5.27:** Monitored bias voltages applied to the SiPMs after the HV calibration of C-Frames 1 to 9.

commissioning. Subsequently, the check from the previous commissioning step is repeated: A bias voltage of 50 V is applied to all SiPMs and compared against the monitored voltages. Figure 5.27 shows the monitored values after the HV calibration of the first 9 commissioned C-Frames. In contrast to Fig. 5.25, the obtained distribution is now centered around the applied voltage and varies with a standard deviation $\sigma = 19$ mV about 5 times less than before the calibration. The maximum deviation of the total of 6272 HV monitoring circuits after the calibration of C-Frames 1 to 9 was found to be ± 100 mV.

5.11 Bit Error Rate Tests

Even though the Bit Error Rate (BER) tests of the optical data transmission are not scheduled as the subsequent commissioning step, they form the basis for the clock timing scans as presented in Section 5.12. Therefore, the BER tests are discussed here at first.

The BER is defined as the number of bit errors divided by the total number of transmitted bits for a data link. Following the operating specifications of the Versatile Link, the SciFi Tracker aims at achieving error rates below 10^{-12} [99]. This includes the electrical data transfer within the front-end electronics, as well as via the optical links. The basis for the latter is a sufficient optical power as verified in a previous commissioning step (see Section 5.7).

An error-free transmission of the data within the ROBs requires the precise coordination of various functional blocks. The careful tuning of the underlying clocks to allow for this is described in Section 5.12. Figure 5.28 illustrates the interplay

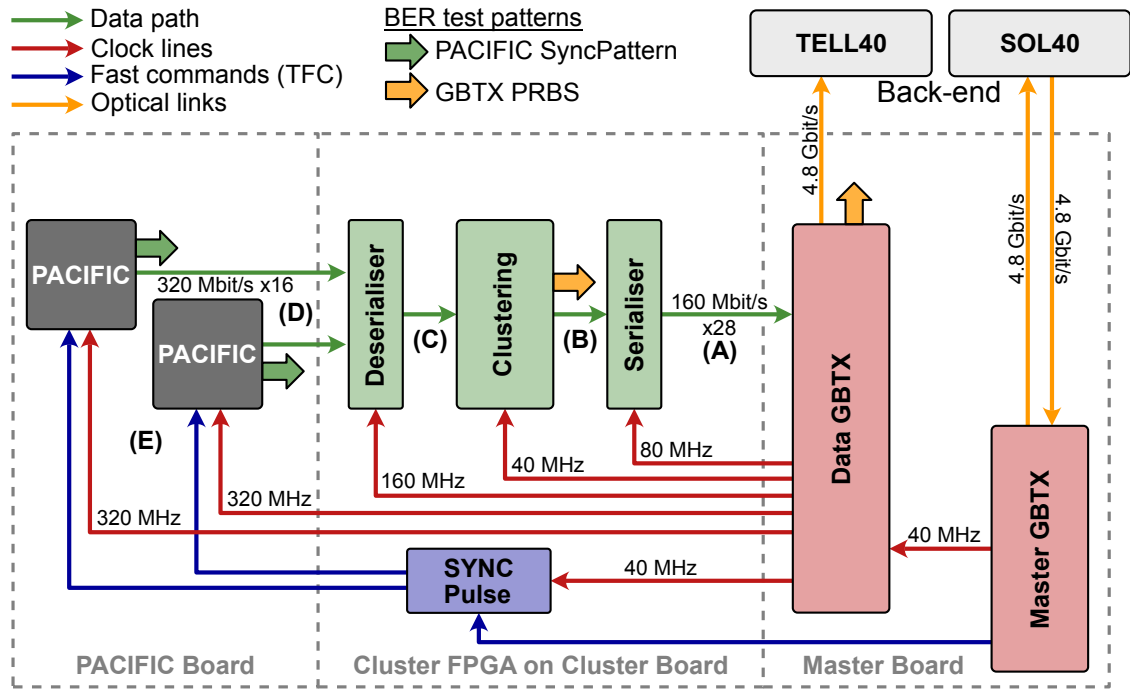


Figure 5.28: Clocking scheme of the SciFi front-end electronics. The error-free data transmission relies on the clock domain crossings labelled as (A)–(E). BER test patterns can be generated at three different points along the data path as indicated by the thick arrows.

between the different components: The SiPM signals digitised by the PACIFIC ASICs are transferred at 320 Mbit/s to the Cluster FPGA. In a first stage inside the FPGA, the serial data is deserialised in order to perform the clustering in the subsequent step. Afterwards, the clustered data is again serialised in order to be transferred to the Data GBTX, which prepares the GBT frames to be sent via the optical link to the back-end.

In addition to the normal operating mode as just described, Fig. 5.28 shows two types of test patterns that can be generated at three different points along the path of the data. These patterns are used to determine the BER of each data link and are discussed in detail in the following.

GBTX PRBS

The GBTX PRBS is one of two BER test patterns used in the SciFi front-end electronics. As the name suggests, it consists of a pseudorandom binary sequence (PRBS) that is generated by the GBTX ASIC. It is implemented by a 7-bit linear feedback shift register that produces a PRBS7 sequence [96]. The register value is copied 16 times to fill the 112 bits of payload of the GBT frame.

Due to the pseudorandom nature of the pattern, an appropriate decoder at the back-end can predict the sequence and compare it against the received data. It




Figure 5.29: PACIFIC SyncPattern generation. The 4 input bits are concatenated with itself in an inverted manner to avoid that phase shifts by 12.5 ns remain undetected.

counts the numbers of mismatches, i.e. the bit errors in the transmission, and stores them in a register to be read via the control software. A separate 8-bit counter exists for each 7-bit pattern. Like this, the BER can be determined.

In addition to being generated by the GBTX ASIC, the identical PRBS generator has been implemented in the Cluster FPGA firmware. The advantage is that a larger part of the data chain can be tested, since the transmission from the Cluster FPGA to the Data GBTX is also included in this case. This part of the data chain is particularly relevant as it involves a plug connection between the Cluster and Master Board, which may be susceptible to a bad electrical contact. Since the same GBTX PRBS sequence is used, no additional resources are required at the back-end.

Starting from the commissioning of C-Frame 4, a BER test utilising the GBTX PRBS generated by the Cluster FPGAs is an integral part of the commissioning procedure. The importance of these checks were duly noted in the course of the operation and commissioning. Two overnight runs with a duration of at least 9 h per run are performed. Due to DAQ firmware limitations (see Section 5.3.2), two runs need to be performed to cover all data links of each C-Frame. Prior to C-Frame 4, these measurements were only done occasionally and with shorter durations.

In total, $6.1 \cdot 10^{17}$ bits were probed during the tests up to the commissioning of C-Frame 9. In the process, 9 errors were found resulting in a BER of $1.5 \cdot 10^{-17}$, which is well below the target rate of 10^{-12} .

PACIFIC SyncPattern

The PACIFIC SyncPattern is the second type of test patterns that are implemented in the SciFi front-end electronics. As shown in Fig. 5.28, it is generated by the PACIFIC ASICs and from there follows the complete data chain. The exact pattern is determined by a group of four pads that are driven by the Cluster FPGA. Thereby, it can be chosen from a total of 16 different patterns. In order to produce the 8-bit word per PACIFIC output line and bunch crossing, the 4-bit input is concatenated with itself in reverse order as illustrated in Fig. 5.29.

A custom TELL40 firmware block has been developed that analyses the received patterns at 40 MHz and counts the number of detected bit errors. For each PACIFIC

channel, i.e. two output bits, there is a separate 8-bit counter that can be read out by the control system. However, this process requires that the PACIFIC output is passed through by the Cluster FPGA in its raw format. Due to bandwidth limitations, the raw data format used for this purpose has four modes, each containing the 2-bit outputs of 32 PACIFIC channels. Therefore, to probe the data transmission of all 128 channels associated with one Cluster FPGA, four separate BER tests utilising the PACIFIC SyncPattern need to be performed.

Starting from the commissioning of C-Frame 5, the procedure has been standardised to record data for 3 h for each group of 32 channels. In addition, the tests are performed with two complementary PACIFIC SyncPatterns

$$\begin{aligned} 0x5 &\rightarrow 0b1010\ 0101 && \text{and} \\ 0xA &\rightarrow 0b0101\ 1010. \end{aligned}$$

Thereby, a total of 24 h are required to cover all four raw data modes and both patterns. Due to the DAQ firmware limitations used for the commissioning in the assembly hall (see Section 5.3.2), the process needs to be performed separately for both the upper and lower half of each C-Frame.

Up to the commissioning of C-Frame 9, $7.7 \cdot 10^{17}$ bits were probed in this way. The number of observed errors is with about 2500 significantly larger than the BER tests utilising the GBTX PRBS generated by the Cluster FPGAs. However, the nature of the errors, which mostly occurred in all 32 observed channels per link, suggests that they are related to the generation of the PACIFIC SyncPattern itself, rather than the actual data chain. The exact cause, however, could not yet be clarified. Despite this effect, the resulting BER of $3.2 \cdot 10^{-15}$ is still well below the target rate of 10^{-12} .

Loss of Frame-Lock in Control Links

During the BER tests, which typically run for several hours at once, the optical links are also monitored for losses of the acquired frame-locks. As described in Section 5.4, this happens when the GBT frame header is not recognised correctly for a few times. Since all communication to the ROBs takes place via the optical control links, it is one of the most critical failures that can occur.

During the SciFi front-end electronics commissioning, temporary losses of frame-locks have been observed occasionally with some ROBs after a few hours of operation. It later turned out that there is a general issue with the used VTRx modules that may occur with some units due to manufacturing problems related to insufficient epoxy curing [115]. Of the total of 392 commissioned HalfROBs mounted on the C-Frames 1 to 9, this type of failure was observed 6 times, corresponding to a rate of affected control links of 1.5 %. The affected ROBs were replaced with spares for further evaluation and repair.




Figure 5.30: Characteristics of an eye diagram. Image adapted from Ref. [116].

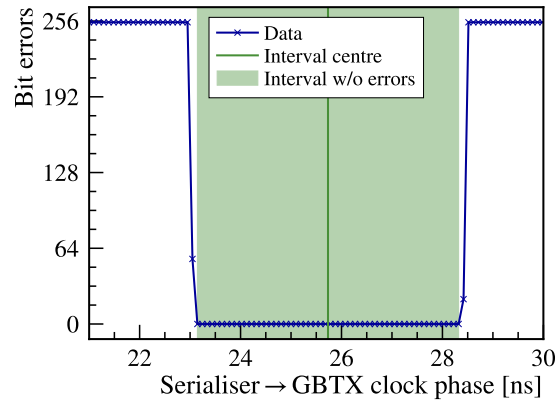


Figure 5.31: Determined bit errors of one PRBS7 word depending on the phase between the Cluster FPGA Serialiser and the Data GBTX.

5.12 Clock Timing Scans

In order to achieve low transmission error rates as presented in Section 5.11, a careful tuning of the relative phases of various clocks is necessary. Within each data link, a total of six clocks need to be adjusted that affect the internal data transmission within the front-end electronics. The contributing clocks are shown in Fig. 5.28 and range from the 320 MHz (3.125 ns clock period) PACIFIC clocks to the clustering blocks of the FPGAs that are operating at 40 MHz (25 ns). Note that the 40 MHz clock that the Data GBTX receives from the Master GBTX does not have an impact on the internal data transmission, since it serves as the common reference for the generation of the remaining clocks. However, it will become relevant during beam data taking as it allows to time align the data taking and processing to the LHC collisions.

The aim of the clock phase tuning is to achieve the highest possible stability in the data transfer between the different processing blocks. This is accomplished when the incoming bit stream from one block to the next is sampled in the centre of the bit period. Figure 5.30 illustrates this by means of an eye diagram, which is the superimposed representation of all possible digital waveforms. As the sampling time approaches the edges of the eye diagram, bit errors start to occur due to jitter and incorrect sampling of the data.

In practice, the tuning of the involved clocks consists in varying their phases one by one in both directions until bit errors occur. Like this, it is possible to determine the bit period and thus the optimal sampling time, which is located in the centre. The detection of bit errors is enabled by using the BER test patterns introduced in Section 5.11.

One difficulty is that an isolated tuning of the individual phases is generally not possible. When using the GBTX PRBS as generated by the Cluster FPGA, both

the 40 MHz clustering and 80 MHz serialiser clock need to be correctly adjusted in order to allow for a stable data transmission. On the other hand, when injecting the PACIFIC SyncPattern, all six clock phases must be set properly. Since scanning all possible combinations in the six-dimensional space is not feasible, the practical approach is to start from a (semi-)stable configuration and vary the phases based on this. Thereby, starting from the processing blocks most upstream of the data chain, i.e. going from right (A) to left (E) in Fig. 5.28, the bit periods and thus the optimal settings can be determined one after another.

During the start of the commissioning, the procedure has not been fully implemented in the control software. This is because it requires a complex interaction between the front- and back-end electronics in a way that was not foreseen by the centrally provided DAQ system. Therefore, the results presented in the following do not include the first 2 commissioned C-Frames 1 and 3. However, the remaining 7 C-Frames covered here contain 152 ROBs, corresponding to 60% of the complete detector.

5.12.1 FPGA Serialiser → Data GBTX (A)

The first clock domain crossing to be optimised is between the Cluster FPGA Serialiser and the Data GBTX - labelled as (A) in Fig. 5.28. This is achieved by shifting the phase of the Serialiser clock in steps of about 100 ps, while the (internal) clock of the GBTX remains unchanged. As shown in Fig. 5.28, the GBTX PRBS pattern generated by the Cluster FPGA can be used for this purpose. Utilising this pattern, the number of bit errors can be determined in units of 7 bits, which are referred to as PRBS7 words in the following. Thereby, with 112 bits of payload in the GBT wide frame mode, 16 separate counter values are available for every data link.

An exemplary result of an FPGA Serialiser clock timing scan is given in Fig. 5.31. It shows the number of detected bit errors depending on the clock phase for one PRBS7 word within one data link. For every phase configuration, 500k GBT frames at 40 MHz are recorded. Within a central interval, no errors are found. At the edges, the number of errors increases and quickly reaches the limit of 255, which is due to the underlying 8-bit error counters.

An important thing to note is that during the scan of the phase between the Serialiser and Data GBTX, the phase between the Clustering and the Serialiser clocks is kept constant. This is achieved by shifting the Clustering clock in parallel to the Serialiser and is necessary because the GBTX PRBS pattern is generated at the level of the Clustering block. Therefore, varying the timing of the Serialiser alone does not provide useful information since it affects both the upstream and downstream clock domain crossings in an interfering manner.

The error-free intervals for an exemplary ROB including 16 Cluster FPGA-Data GBTX pairs and thus 16 corresponding clock domain crossings are shown in Fig. 5.32. Local variations due to differences in the routing lengths and signal qualities can be identified. Furthermore, the overlap of the intervals is indicated, which represents the common error-free interval of all 152 ROBs for which this measurement has been




Figure 5.32: For an exemplary ROB, the error-free intervals in the phase between the Cluster FPGA Serialiser and the Data GBTX are shown. One ROB comprises 16 individual data links (i.e. 16 different clock domain crossings) with 16 PRBS7 words each, resulting in the total of 256 words and corresponding error counters.

performed. In practice, this means that all data links can be operated stably with the identical clock phase configuration between the Cluster FPGA Serialiser and the Data GBTX. By setting the phase in the centre of the 3.2 ns wide common interval, the margin of error towards both edges amounts to at least 1.6 ns for all links.

Figure 5.33 shows summary distributions of the interval centres and widths. Note that the absolute values of the interval centres have no further significance, but only the relative differences are relevant. As can be seen, the centres lie well within the interval overlap of all 152 included ROBs. On the other hand, the interval widths are a measure of the signal quality. The distribution shows a narrow peak around the mean value of 5.24 ns, which is about 16 % smaller than the theoretical limit of 6.25 ns. In addition, a shoulder-like bump towards lower values can be identified that indicates electrical lines with slightly lower signal quality. The theoretical limit is calculated based on the transmission speed of 160 Mbit/s per lane between the Cluster FPGA and Data GBTX. Thereby, under ideal conditions, the signal levels can be sampled accurately over the complete bit period of 6.25 ns. However, as illustrated in Fig. 5.30, due to finite fall and rise times as well as jitter, the measured error-free interval widths will always be below this value.

5.12.2 FPGA Clustering → Serialiser (B)

After tuning the phase between the Cluster FPGA Serialiser and the Data GBTX in the previous step by setting it to be in the centre of the interval overlap, a scan of the Clustering clock is performed. Thereby, the stable phase intervals between the FPGA Serialiser and Clustering blocks can be determined, corresponding to the




Figure 5.33: Error-free intervals in the phase between the Cluster FPGA Serialiser and Data GBTX clock. The distributions of the interval centres (left) and widths (right) include data from 152 ROBs as obtained during the front-end electronics commissioning. The theoretical limit of 6.25 ns on the width results from the transmission speed of 160 Mbit/s per lane.

clock domain crossing (B) in Fig. 5.28. Again, the GBTX PRBS test pattern as generated by the Clustering block is used.

The bit-error-free intervals for an exemplary ROB are shown in Fig. 5.34. A very regular pattern can be recognised that repeats itself every 16 PRBS7 words, i.e. in units of individual data links and Cluster FPGAs. This is due to the fact that the clock domain crossing examined in this step happens entirely within the FPGA itself. Therefore, the observable pattern completely depends on the firmware and the internal signal generations and routings in the FPGA. The interval overlap of all 152 examined ROBs provides with a width of 9.0 ns plenty of margin of error when setting the common phase to lie in the centre of it.

Summary histograms of the interval centres and widths are shown in Fig. 5.35. Due to the regularity of the intervals, the distributions differ greatly for the different PRBS7 words, i.e. different parts in the resulting GBT data frame. This applies in particular to the interval centres. Nevertheless, they lie well within the interval overlap as indicated in Fig. 5.34. Based on the internal data transmission inside the FPGA, the theoretical limit on the interval width amounts to 12.5 ns.

5.12.3 FPGA Deserialiser → Clustering (C)

A similar situation exists for the clock domain crossing (C) between the Deserialiser and Clustering blocks that happens entirely within the Cluster FPGA. But unlike the previous step, the PACIFIC SyncPattern has to be used, since the GBTX PRBS pattern is generated downstream of this part of the data chain. As discussed in Section 5.11, the corresponding error detection logics in the TELL40s provide separate 8-bit error counters for the 2-bit output data of each PACIFIC channel.




Figure 5.34: Error-free intervals in the phase between the Cluster FPGA Clustering and Serialiser blocks for an exemplary ROB. The data within each individual Cluster FPGA and thus clock domain crossing is composed of 16 PRBS7 words.




Figure 5.35: Error-free intervals in the phase between the Cluster FPGA Clustering and Serialiser blocks. The distributions of the interval centres (left) and widths (right) are shown.

Thereby, the total amount of error counters is a factor of 8 larger compared to the GBTX PRBS test pattern. In order to record the data from all 128 channels per link, 4 separate runs need to be performed.

To scan the desired phase, the FPGA Deserialiser clock is varied, while the Clustering clock is kept at the common value as determined in the previous step. The clocks of the SYNC Pulse and the PACIFICs need to be shifted in synchronous to avoid interferences with the phases further upstream in the data chain. Figure 5.36a shows the error-free phase intervals for two neighbouring exemplary Cluster Boards. Similar to the previous scan (B), a regular pattern can be recognised that repeats every 128 channels, i.e. in units of individual Cluster FPGAs.

Projecting the interval centres onto the y-axis for all 152 examined ROBs results in the distribution as shown in Fig. 5.37a on page 110. In addition, the interval widths including the theoretical limit are displayed. Due to the significantly increased number of error counters (i.e. channels) per data link, a breakdown into different groups of channels as shown in Fig. 5.35 has been omitted in this case.


5.12.4 PACIFIC → FPGA Deserialiser (D)

The last part of the data chain is between the PACIFIC ASICs and the Cluster FPGA Deserialiser. For this, the data is serialised in the PACIFICs and transmitted to the FPGAs on the Cluster Boards at rates of 320 Mbit/s per lane. To scan the phase between the sender and receiver, the clocks of both PACIFICs within each data link are varied, while keeping the Deserialiser clock phase at the previously determined common setting. The SYNC Pulse clock is shifted in synchronous with the PACIFIC clocks in order to maintain the same phase between them.


Because the signals need to be transferred across multiple PCBs with length variations in the signal lanes, the resulting error-free intervals in the phase between the sender and receiver show relatively large variations as displayed in Fig. 5.36b. This is true on the microscopic level, as well as in blocks of 64 channels corresponding to different PACIFIC ASICs. Due to the high transmission rates, the overlap of all error-free intervals of the examined 152 ROBs only amounts to 1.2 ns. However, as presented in Section 5.11, sufficiently low bit error rates can still be achieved by setting the phase of all PACIFIC clocks to a common setting in the centre of the overlap. A summary of the interval centres and widths is given in the form of histograms in Fig. 5.37b on page 110. As can be seen, the measured widths are only slightly below the theoretical limit of 3.125 ns indicating a good signal quality.

5.12.5 SYNC Pulse → PACIFIC (E)


Even though the SYNC Pulse is not part of the data chain itself, it is still crucial for the transmission of the data since it is responsible for the synchronisation of the PACIFIC ASICs. As described in Section 4.2.2, the signal is evaluated at the rising edge of the 320 MHz PACIFIC clock. Therefore, it is expected that the SYNC Pulse clock can be varied by about 3.125 ns before the signal is evaluated at the previous



(a) Cluster FPGA Deserialiser → Clustering

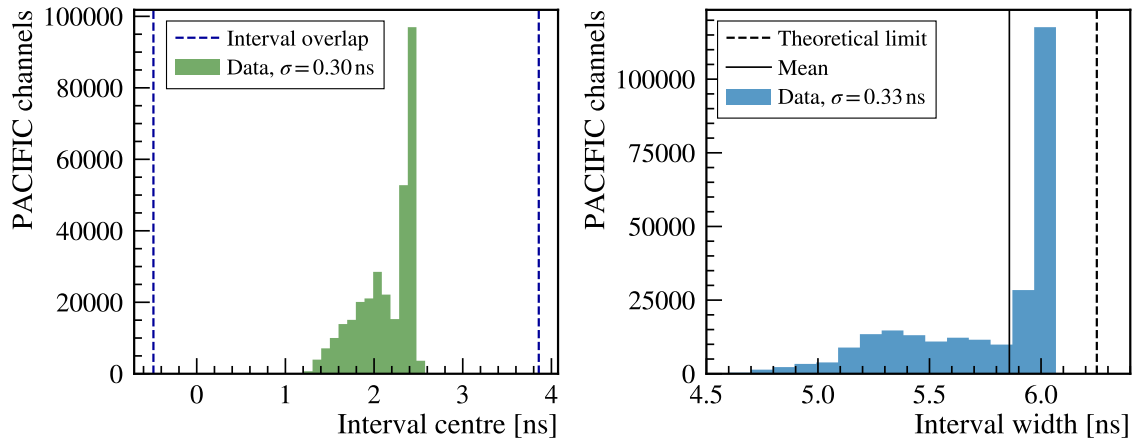


(b) PACIFIC → Cluster FPGA Deserialiser




(c) SYNC Pulse → PACIFIC

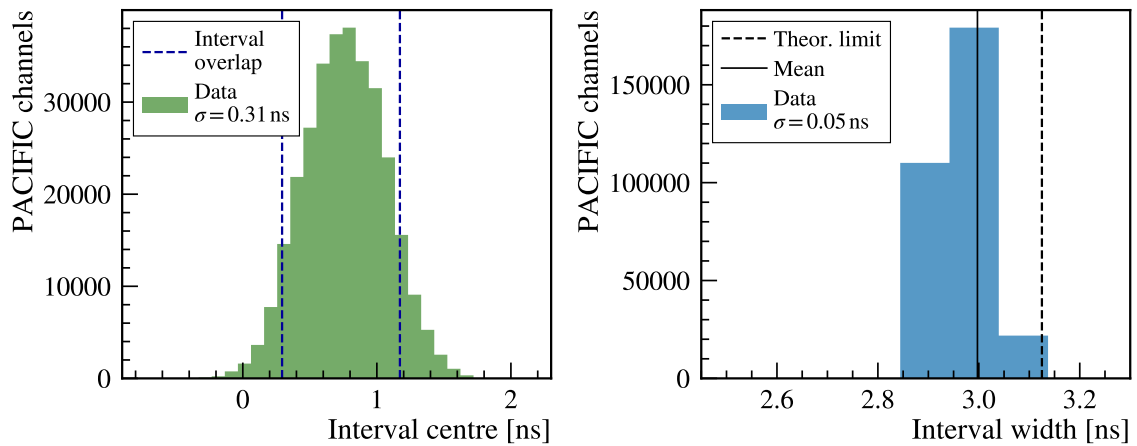
Figure 5.36: Error-free intervals in the different phases along the data chain for two exemplary neighbouring Cluster-PACIFIC Board pairs utilising the PACIFIC SyncPattern. The individual clock domain crossings are indicated by the vertical dashed lines. Within each data link, one Cluster FPGA processes the data from two PACIFIC ASICs with 64 channels each.



(a) Cluster FPGA Deserialiser → Clustering



(b) PACIFIC → Cluster FPGA Deserialiser



(c) SYNC Pulse → PACIFIC

Figure 5.37: Error-free intervals in different phases along the data chain utilising the PACIFIC SyncPattern. The distributions of the interval centres (left) and widths (right) are shown.

Table 5.4: Parameters of the error-free phase intervals as determined by the clock timing scans. At the centre of the interval overlap (second column from the right), the SciFi front-end electronics can be operated stably with the same clock phase configuration for every ROB. The numbers are given in ns.

Clock Phase	Int. Centres		Int. Widths		Int. Overlap	
	Mean μ	Std. σ	Mean μ	Std. σ	Centre	Width
Serialiser \rightarrow Data GBTX	25.32	0.26	5.24	0.25	25.44	3.22
Clustering \rightarrow Serialiser	30.17	0.51	10.71	0.81	29.61	9.03
Deserialiser \rightarrow Clustering	2.13	0.30	5.86	0.33	1.68	4.35
PACIFIC \rightarrow Deserialiser	13.29	0.25	2.96	0.08	13.13	1.17
SYNC Pulse \rightarrow PACIFC	0.77	0.31	3.00	0.05	0.73	0.88

Std. $\hat{=}$ Standard deviation Int. $\hat{=}$ Interval

or next rising edge causing the output data to be shifted by one PACIFIC clock cycle. This can be confirmed by looking at the distribution of the error-free interval widths as shown in Fig. 5.37. Since the SYNC Pulse is generated by the Cluster FPGA, the interval centres show a similar spread as the timing scan between the PACIFIC and the FPGA Deserialiser.

However, as shown in the exemplary view of the intervals for two neighbouring Cluster-PACIFIC Board pairs in Fig. 5.36c on page 109, the variations occur predominantly in blocks of 64 channels, i.e. in units of different PACIFIC ASICs. This is due to the fact that each PACIFIC receives and processes a single SYNC Pulse. The interval overlap of all 152 examined ROBs is with 0.9 ns in the same range as for the phase between the PACIFIC and FPGA Deserialiser clocks.

5.12.6 Conclusion

A summary of the obtained parameters during the clock timing scans as previously presented is given in Table 5.4. As given by the centre of the overlap region of the error-free intervals, the ROBs can all be operated at the identical clock phase configuration while still achieving low bit error rates (see Section 5.11). The smallest margin of error is present in the phase between the SYNC Pulse and PACIFIC clocks and amounts to 0.44 ns in both direction (0.88 ns in total).

5.13 Light Injection System Tests

In the last step of the SciFi front-end electronics commissioning procedure, the complete data chain is examined. While the digital data transfer has already been tested and tuned during the BER tests (see Section 5.11) and clock timing scans (see Section 5.12), this part also includes the analogue signal generation and processing




Figure 5.38: SciFi ROB cover with the routing of the LIS GBLDs to the fibre mats indicated. Image modified from Ref. [117].

in the SiPM and PACIFIC channels. In the absence of ionising radiation in the assembly hall, these tests are performed with the help of the light injection system (LIS). As described in detail in Section 4.5, the LIS allows to inject light pulses into the scintillating fibres towards the end of the fibre mats near the SiPM arrays. Thereby, it is possible to test the full data chain and identify channels that do not register any signals from the incident photons. This can either be caused by broken SiPM or PACIFIC channels, or due to an insufficient electrical connection between the two.

The tests involving the LIS are subdivided into three steps that are discussed in the following.

5.13.1 Functional Test and Mapping

As a first step, the functioning of the LIS itself is verified. This is done only with the help of the SiPM arrays that are biased with the nominal overvoltage $\Delta V = 3.5 \text{ V}$. By monitoring the current drawn from the CAEN modules with and without injected light pulses, the correct functioning of the LIS can be confirmed. Typical dark currents are in the order of $100 \mu\text{A}$ per HV channel, corresponding to four adjacent SiPM arrays. It increases to about $200 \mu\text{A}$ when illuminating the SiPMs with 15 ns wide LIS pulses at a rate of 20 kHz. However, large variations in the light intensity can be observed that depend on the determined operating settings of the GigaBit Laser Drivers (GBLDs) during the QA of the fully assembled fibre modules.

For further tuning of the light intensity, the mapping between the GBLDs and light bars must be known for each module. As illustrated in Fig. 5.38, the light bars mounted on fibre mats 0 and 2 are controlled by HalfROB H0, while mats 1 and 3 are illuminated by the GBLDs connected to H1. However, due to the routing of the LIS fibres within the end plugs of the modules, it is not clear at first which GBLD

of each half connects to which of the two possible mats¹. In order to create the corresponding mapping, the current draw is monitored as described previously while enabling the individual GBLDs one by one.


5.13.2 Delay Scan (Lite)

After verifying the functioning of the LIS itself and determining the mapping, the actual test consists of checking that every channel correctly detects the incident photons. For this, it is necessary that the light pulses are triggered at the right time with respect to the PACIFIC clock such that the resulting SiPM signals are fully integrated by the PACIFIC integrators. For an exemplary channel, Fig. 5.39a shows the S-curves for different LIS pulse delays as obtained by scanning the PACIFIC comparator thresholds. As can be seen, the characteristic step structure as described in Section 4.2.4 is only visible within a few nanoseconds around the optimal delay, which amounts to 29 ns in this case. On the other hand, when moving away from the optimum, the S-curves degrade until eventually only a single transition remains corresponding to the pedestal of the channel.


In principle, recording the S-curves by scanning the PACIFIC comparator thresholds for different delays in order to find the best delay is only needed once. However, the limitations of the available DAQ firmwares in combination with the design of the LIS requires for a delay scan for essentially every C-Frame commissioning. On the one hand, the SOL40 control links have different latencies in the order of ± 25 ns relative to each other. In addition, these are not fixed but may vary after every reset. By itself, this is unproblematic as long as each control link corresponding to a HalfROB remains an isolated system, which is the case for many applications in the context of the front-end electronics commissioning. However, due to the design of the LIS, the two halves of one ROB are closely entangled. This is because the two inner fibre mats of one module are illuminated by light bars that are controlled by the neighbouring HalfROB as illustrated in Fig. 5.38. Therefore, while the outer mats are read out by the same HalfROB that controls the light bars, the optimal LIS delays of the inner mats are shifted in opposite directions depending on the relative latency between the two halves.

Doing a full delay scan as shown in Fig. 5.39a is a time-consuming procedure that is not feasible to do for every C-Frame. With the limited firmwares and the associated software, scanning the 256 threshold DACs for a fixed delay already takes about 20 min. For this reason, a simplified version is performed that is referred to as delay scan lite. It is based on varying the LIS delay while keeping the three comparator thresholds fixed. Although in principle, any threshold values above the pedestal and within a few photoelectron (pe) amplitudes are suitable, they are set to the default settings of [1.5, 2.5, 3.5] pe. However, since the precise threshold calibration as described in Section 4.2.4 is not yet possible without knowing the optimal delay, the values are estimated assuming a typical distance of 15 DACs

¹Accidentally, this was not specified during the production of the modules.



(a) Full delay scan



(b) Delay scan lite

Figure 5.39: Threshold scans of one channel for different delays of the light injection pulse (a). The vertical dashed lines indicate the thresholds of the three PACIFIC comparators that are used for the delay scan lite (b).

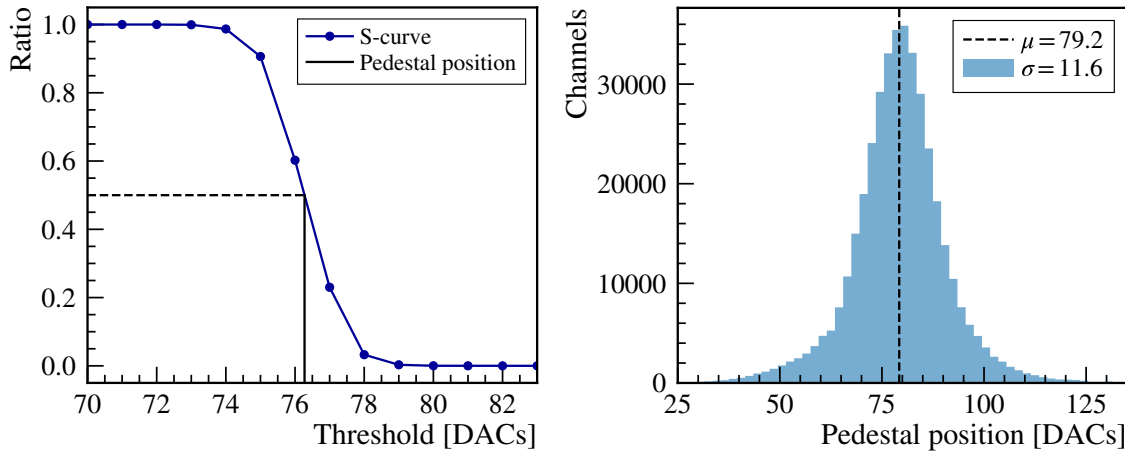


Figure 5.40: Exemplary S-curve of one channel without the injection of any signal showing the pedestal transition (left). The position of the pedestal is given by the point on the linear interpolated curve that matches the ratio of 0.5. The distribution on the right shows the determined pedestal positions of the 400k channels within C-Frames 1 to 9.

between photoelectron peaks.

The required position of the pedestal is determined by doing a single threshold scan without injecting light pulses. A zoomed in view of a pedestal transition is shown in Fig. 5.40 (left). As illustrated, the position of the pedestal is calculated by taking the point on the linear interpolated curve that matches the ratio of events above threshold of 0.5. Figure 5.40 (right) shows the pedestal positions of the 400k channels that are obtained in that way during the commissioning of C-Frames 1 to 9.

The threshold values calculated based on an assumed photoelectron peak separation of 15 DACs and the determined pedestal position as just described are indicated by the vertical dashed lines in the example in Fig. 5.39a. As can be seen, they are slightly off compared to the intended positions in the middle of the S-curve plateaus. However, as noted earlier, the exact values are not relevant for this particular case. Figure 5.39b shows the result of the subsequent delay scan lite for the exemplary channel. The optimal delay is given by the maxima of the curves.

In the same way, the best delays are determined for all channels of each C-Frame. Figure 5.41 shows an exemplary result for the 2048 channels of one ROB, split by the four read out fibre mats. As expected, the two outer mats 0 and 3 show very similar optimal delays, while the values of the inner mats 1 and 2 are shifted in opposite directions.

5.13.3 Threshold Scans

In principle, achieving ratios of events above threshold that are significantly larger than 0 for all channels at any delay already indicates that the channels are correctly




Figure 5.41: Distribution of the optimal LIS delay for an exemplary ROB based on the ratio of events above the lowest threshold corresponding to about 1.5 pe.

detecting the incident photons. However, as an additional test, full threshold scans are performed at selected LIS delays based on the results from the delay scan lite. For the exemplary ROB shown in Fig. 5.41, threshold scans at three different delays are required. Typically, a handful of different settings need to be chosen to cover all channels during the commissioning of one C-Frame.

For each channel, the S-curve at the best timing is evaluated in order to identify dead channels. The property that is used for this purpose is the integral under the S-curve, starting from the position of the pedestal. This is illustrated in Fig. 5.42. As shown in Fig. 5.39a, the area under the S-curve is rising with LIS delays closer to the optimum. This is because the charge accumulated by the PACIFIC integrator increases while the light intensity, as determined by the settings of the GBLD, is constant.

An example of a type of failure that can be identified in this commissioning step is shown in Fig. 5.43. As can be seen, the depicted S-curves do not follow a monotonous course that is to be expected due to the nature of the S-curves as introduced in Section 4.2.4. Instead, they reveal several jumps along the threshold scan. Since these occur in a similar way for both PACIFIC integrators, the error likely lies in the relation between the threshold DACs and the resulting comparator voltage level. Due to the generation of the thresholds by the three 8-bit current DACs per channel (see Section 4.2.1), some jumps in the S-curves are to be expected. However, these mostly occur at multiples of 64 DACs and predominantly at 128 DACs, but not to the extend as shown in Fig. 5.43. Therefore, these types of channels are classified as broken during the commissioning.

In terms of dead channels, i.e. channels that are unable to detect the incident photons, the area under the S-curve as introduced previously is considered. Figure 5.44 on page 118 shows the S-curve integral over the course of the 2048 channels of one fibre module. A clear separation between the threshold scans under injected light pulses, and pedestal-only runs can be identified, indicating that the channels




Figure 5.42: Illustration of the S-curve integral for one channel. The S-curves with illumination of the SiPM (LIS), as well as without any signal (Pedestal) are shown.




Figure 5.43: Example of a broken PACIFIC channel showing a non-monotonous S-curve with several jumps.

are working properly. The course of the LIS curve is mainly determined by the luminosity along the 13 cm light bars. Since the fibres inside the light bars are scratched manually, each one has a unique light profile.

In summary, the distributions as shown in Fig. 5.45 are obtained. They show the S-curve integrals for all 401 408 channels that were examined during the commissioning of C-Frames 1 to 9. A total of 7 channels exhibit a LIS integral below 5, which has been defined as the lower limit in the course of the commissioning. However, after closer examination, only 3 of these turned out to be dead channels, while the remaining 4 suffered from a low light intensity. The cause of the malfunctions was found in the associated 3 SiPM channels, which were already marked as dead during the QA of the assembled SiPM arrays inside the cold boxes.

In addition to the 3 dead channels, 7 channels were identified that show similar jumps in the S-curves as depicted in Fig. 5.43. In relation to the total of 401 408 examined channels, this corresponds to a quota of $2.5 \cdot 10^{-5}$ malfunctioning channels.

Outlook

Regarding the threshold scans that will need to be performed for the precise calibration of the threshold DACs during normal operation, it is expected that no preceding delay scans (lite) will be required every time. However, this requires that the control links are provided with fixed latencies by the SOL40s. Nevertheless, it will be indispensable to once determine the optimal LIS pulse delay by means of the presented methods.




Figure 5.44: S-curve integral along the 2048 channels of one exemplary fibre module. The profile is shown without injecting any signal (pedestal-only), as well as under illumination with pulsed light (LIS).




Figure 5.45: Distributions of the S-curve integrals for the 400k channels as examined during the commissioning of C-Frames 1 to 9.

5.14 Commissioning Summary

In this chapter, the commissioning of the front-end electronics of a large fraction of the SciFi Tracker has been presented. Due to the complexity of the system (see Section 5.1), it has been a challenging task. This concerns in particular the commissioning of the early C-Frames that took significantly longer because many tools had to be developed alongside the commissioning procedure itself. It was the first time that the electronics were operated on such a large scale and at the system level along with the surrounding infrastructure.

After overcoming the initial challenges, including the establishment of the 40 MHz readout in conjunction with a near-final version of the upgraded LHCb DAQ system (see Section 5.3.2), the commissioning procedure steadily became a smooth and routine operation. Nevertheless, there were also a few critical failures that could only be resolved by replacing the affected ROBs. However, due to the fact that the commissioning takes place in a dedicated hall on the surface of the LHCb site, the replacement operations were well possible.

Overall, it could be verified that the SciFi front-end electronics generally works flawlessly and meets very high quality requirements. This was made possible by means of a strict quality assurance (QA) before the installation of the various electronics components (see Section 4.6), as well as before mounting the individual ROBs on the C-Frames (Section 5.2). Among other things, it could be demonstrated that, after careful tuning of the underlying clock phases, a stable data transmission way below the required bit error rate of 10^{-12} is achievable on a large scale consisting of a few thousand optical data links. In addition, only a handful of channels were found to be malfunctioning out of a total of 400k detector channels tested.

6 Performance Studies

The performance of the SciFi front-end electronics crucially depends on the digitisation of the SiPM signals as well as the error-free transmission of the data to the DAQ system. Moreover, it relies on experiencing low noise levels in order to avoid the occurrence of false hits.

Therefore, in addition to the regular commissioning steps as discussed in Chapter 5, further studies in these areas were performed on individual C-Frames. These tests were mainly conducted on C-Frame 2 that was commissioned at the beginning of 2021 (see Table 5.1). During that time, the assembly of the next C-Frame was on hold due to restrictions in the course of the COVID-19 pandemic that did not allow the required international personal to travel from abroad. As a result, C-Frame 2 was available for these performance studies at the system level within the commissioning setup (see Section 5.3) for some time. Due to the limitations in the available DAQ system (see Section 5.3.2), the presented measurements mostly include the data from only half of the C-Frame (i.e. one quadrant). However, this still corresponds to 12 ROBs with a total of 25k channels, i.e. to about 5% of the final detector.

6.1 PACIFIC Pulse Height Digitisation

The digitisation of the SiPM pulses is a key aspect for the operation of the detector as it is the basis for the reconstruction of hit clusters. A deep understanding of the analogue processing and digitisation in the PACIFIC ASICs is essential for an accurate calibration of the threshold DACs (see Section 4.2.4), which is required to achieve high hit efficiencies along with good noise rejection.

6.1.1 Pedestal Stability

The PACIFIC integrators play an important role in the analogue processing. They integrate the signals within each period of 25 ns corresponding to the proton bunch crossing intervals. Without any input signal, the output level is determined by the DC baseline of the integrator. It is also referred to as pedestal since any additional signal is added to it. Knowledge about the pedestal level is essential to determine the pulse height, i.e. amplitude, of a given signal. In addition, it is required that it remains stable after being once determined by means of a PACIFIC threshold scan.

The stability of the pedestal first came into focus during the last test beam campaign that took place in the summer of 2018 at the CERN SPS (see Section 3.9). In the course of the two test weeks, occasional shifts of the pedestal positions of up to 40 threshold DACs were noticed. Compared to a typical photoelectron peak

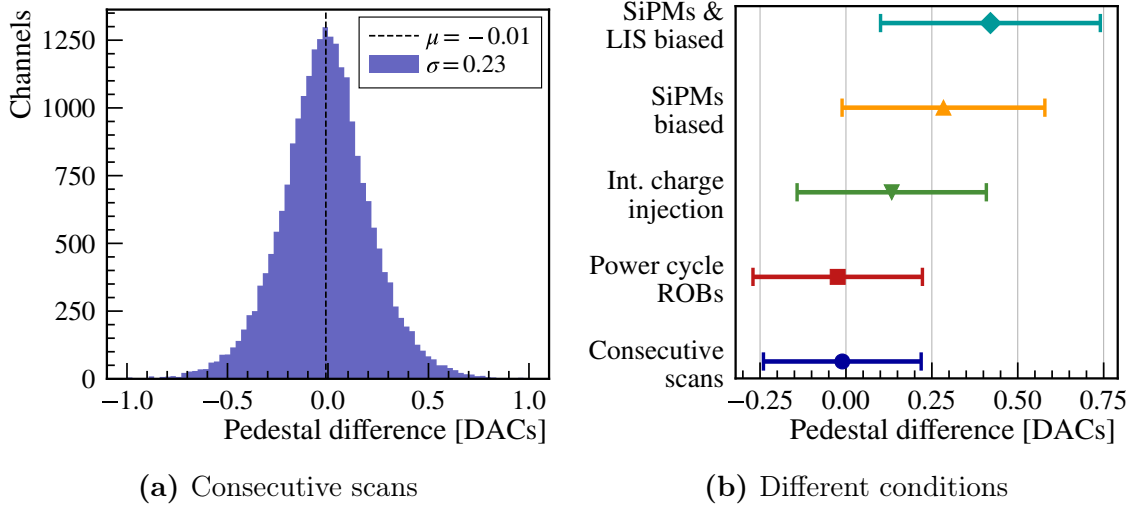


Figure 6.1: Variations of the determined pedestal positions between two consecutive threshold scans (a) and under further conditions (b) for the 25k channels of one C-Frame half. The reference scan has been performed with the HV and LIS turned off. In the right figure (b), the mean values μ and standard deviations σ of the underlying distributions are indicated by the markers and error bars, respectively.

separation of about 15 DACs, this far exceeded the expected variations. Possible explanations are ranging from insufficient grounding to accidental changes in the PACIFIC configurations. However, up to now, the effect seen in the test beam could not be reproduced or understood. Among the extensive studies that have been performed in this context are several measurements within the SciFi commissioning setup that are presented in the following.

As a basis, one has to consider the general variations of the determined pedestal position within each channel. Between two consecutive threshold scans, the pedestals show deviations as displayed in Fig. 6.1a for the 25k channels. The standard deviation of the distribution $\sigma = 0.23$ DACs can be interpreted as the overall uncertainty on the pedestal positions, which are determined according to the illustration in Fig. 5.40 on page 115.

By performing the threshold scans under different conditions, an attempt was made to reproduce similar pedestal shifts as observed during the 2018 test beam campaign. Examples showing the largest displacements that could be achieved are given in Fig. 6.1b. While power cycling the front-end electronics virtually made no difference compared to two consecutive threshold scans, small biases could be produced in other cases. When injecting a signal with the help of the internal charge injection mechanism (see following Section 6.1.2), the pedestals are systematically shifted by 0.13 DACs. Larger displacements are achieved by biasing the SiPM arrays (0.28 DACs) and additionally enabling the LIS, although without injecting light pulses (0.43 DACs). However, even these differences are negligible compared to the typical photoelectron peak separation of 15 DACs and especially compared to the




Figure 6.2: Typical S-curve when injecting a fixed charge 50% of the time using the internal charge injection mechanism for an exemplary PACIFIC channel.




Figure 6.3: Internal charge injection amplitudes depending on the delay for an exemplary PACIFIC channel. The amplitudes of the three possible charge values are shown.

40 DACs test beam variations. Therefore, there is no indication that the studied C-Frame suffers from the problems observed in the test beam.

6.1.2 Internal Charge Injection

An important feature for the characterisation of the PACIFICs is the internal charge injection, which is used to simulate the input of SiPM pulses with fixed amplitudes. It consists of two small on-channel capacitors (nominal 630 fF and 1 pF) that can be discharged through two relatively large resistors (12 k Ω and 9 k Ω , respectively). The discharges can be triggered independently via two digital input signals that are controlled by two fast control commands, which are referred to as CalibA and CalibB. Thereby, three different charge values can be injected into the pre-amplification stage of each channel (see Section 4.2.1). The capacitance and resistance value pairs were selected based on simulations that showed good approximations to actual SiPM signal shapes. Using the slow control of the PACIFIC, any combination of channels can be selected for internal charge injection [74].

The internal charge injection is an important debugging tool as it allows to inject signals into the PACIFIC channels without the need of additional hardware. It is used in particular to examine the per-channel differences between the PACIFIC integrators and comparators as presented in the following Sections 6.1.3 and 6.1.4.

Figure 6.2 shows a typical S-curve in this context. It is generated by triggering the discharge of the 1 pF capacitor (CalibB) 50% of the time. As a result, two transitions are visible corresponding to the pedestal and charge signal levels. Thereby, the charge amplitude is given by the difference of the two transitions as illustrated in Fig. 6.2.




Figure 6.4: Amplitude distributions of the three charge values available through the internal charge injection mechanism. The mean values are indicated by the dashed lines.

Similar to injected light pulses (see Section 5.13.2), the total charge accumulated by the PACIFIC integrators strongly depends on the injection delay with respect to the internal clock. Figure 6.3 shows the dependence for the three possible charge values of an exemplary PACIFIC channel. Within a 40 MHz clock cycle, the delay can be chosen to take eight different values in steps of 3.125 ns, as controlled by a corresponding register of the Cluster FPGA that is responsible for issuing the digital discharge signals to the PACIFICs. For 80% of the examined channels, the maximum amplitude is reached at a delay of 9.375 ns, with the remaining ones peaking at 12.5 ns. However, as can be seen in Fig. 6.3, the charge difference between the two central points is relatively small and amounts to about 5% on average for all channels.

At the overall best delay of 9.375 ns, the internal charge amplitudes for the 25k channels are distributed as shown in Fig. 6.4. With mean values between 23.3 and 55.3 DACs for the smallest and largest available charge values, the internal charge injection mechanism covers a similar range as the SiPM signals of 1-4 photoelectrons. With the default operating thresholds [1.5, 2.5, 3.5] pe lying in the similar range, it is therefore well suited for debugging of the detector. Since no additional hardware is required, it can also be used in the final detector between normal data taking runs and without manual intervention.

6.1.3 Integrator Differences

The internal charge injection mechanism is mainly used to characterise the differences between the two integrators of each PACIFIC channel. This is an important quantity for the operation of the detector since both integrators share the same set of three comparators. Therefore, a uniform behaviour is required to ensure consistent

performance of the two integrators and thus for even and odd events.

Due to the structure of the PACIFIC comparators, determining the integrator differences requires the use of the internal charge injection. The reasoning behind that is illustrated in Fig. 6.5. To avoid repetitive switching of the comparator output when the signal level is close to the set threshold, the PACIFIC comparators feature a design with a 10 mV hysteresis. As a result, there exist two slightly different voltage levels causing the comparator output to change. Going above V_{\uparrow} changes its state from $0 \rightarrow 1$, while going below V_{\downarrow} changes its state from $1 \rightarrow 0$. The difference $V_{\uparrow} - V_{\downarrow}$ is referred to as hysteresis.

As a consequence, the pedestals might overlap perfectly when performing a threshold scan, even though the baselines might differ by up to 10 mV (see bottom left part of Fig. 6.5). In the default configuration of the PACIFIC ASICs, a 10 mV difference at the comparator input corresponds to about 4 DACs in the threshold scan (see Section 4.2.1). To determine the integrator differences with a higher precision than that, a signal needs to be injected whose amplitude exceeds the level of the hysteresis in order to resolve the two S-curves (bottom right part of Fig. 6.5). The internal charge injection mechanism is well suited for this as it does not require any additional hardware. As shown in Fig. 6.4, any of the three available charge levels by far exceed the hysteresis and are in principle suitable for the task. However, to minimise the effect of any potential and interfering effects like timing differences between the two integrators, the smallest possible charge amplitude (CalibA) is used.

Trimming

A uniform behaviour of the PACIFIC integrators to ensure consistent performance is achieved through the trimming, which refers to the process of aligning the two integrator baselines per channel. As discussed in Section 4.2.1, the trim DAC and thus the reference voltage of the integrators can take 31 different values. On average, incrementing the trim DAC by 1 shifts the baseline by about 7 mV (using the default operating settings). When performing a threshold scan, these shifts of the baseline effectively move the S-curves in the horizontal direction by about 2.8 threshold DACs per trim DAC. The covered range by the trimming is illustrated in Fig. 6.6 for the two integrators of one PACIFIC channel. The goal of the trimming is to align both integrator baselines such that the S-curves are overlapping. This is the prerequisite for registering the same pulse height for the same SiPM signal and thus achieve a consistent response for the even and odd integrators.

The trimming is performed as part of the test procedure during the quality assurance (QA) of the PACIFIC Boards (see Section 4.6). In the first step, three threshold scans are performed with the trim DACs set to the minimum, maximum and central (default) values. As discussed previously, this happens in each case under injection of the smallest charge (CalibA) with the help of the internal charge injection mechanism. The covered range for each channel and integrator is determined with the three initial threshold scans. Especially the overlap between the two integrators is

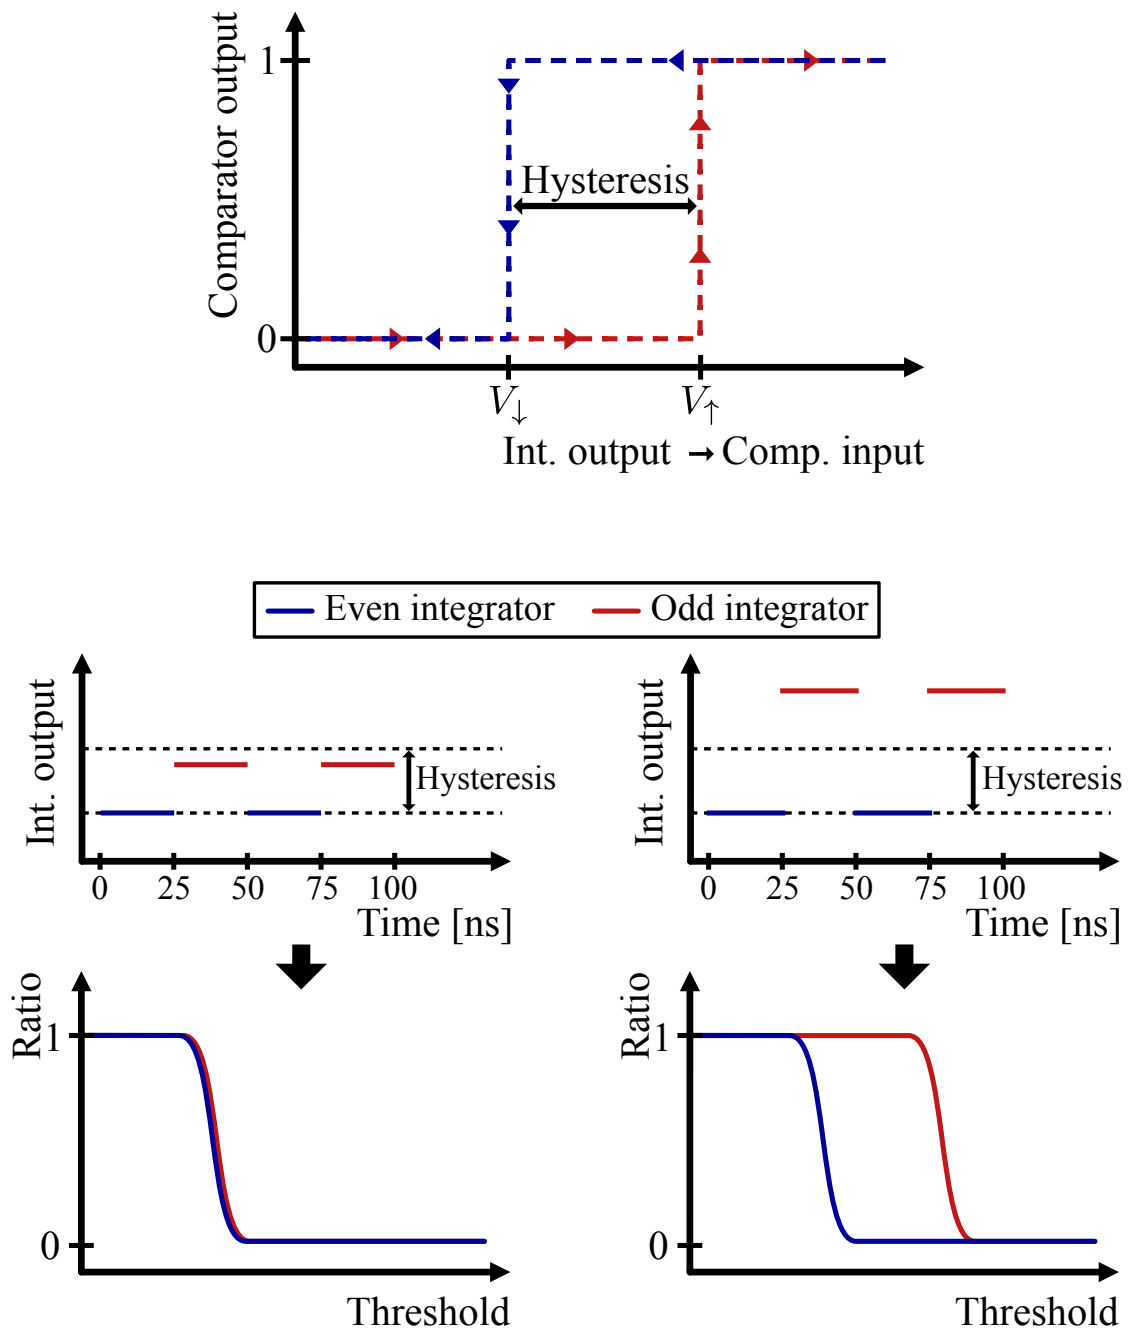


Figure 6.5: Illustration of the switching process of the PACIFIC comparator output depending on the input voltage in the presence of a hysteresis (top). Shown below are the consequences that arise from it during a threshold scan for the cases where the even and odd integrator levels are within the hysteresis (left) and exceeding it (right). Note: The output of the integrator (Int.) is fed to the input of the comparator (Comp.).

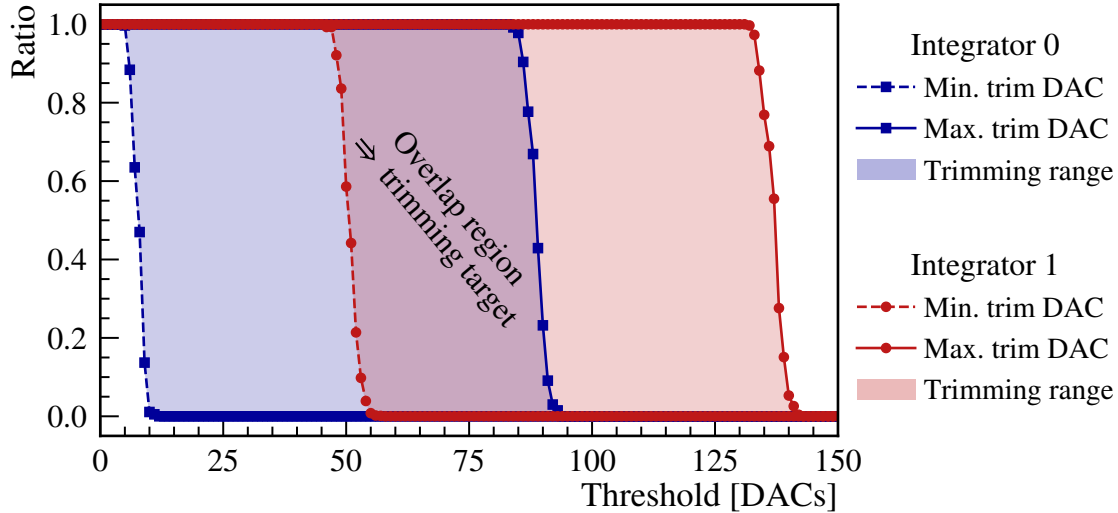


Figure 6.6: Amplitude distributions of the three charge values available through the internal charge injection mechanism. The mean values are indicated by the dashed lines.

of interest as it defines the available range where the two integrator baselines can be aligned (see Fig. 6.6).

Within that range, a target value is defined to which the two integrators are trimmed. To calculate the required change of the trim DAC value, the difference of the current S-curve transition to the target value (in threshold DACs) is divided by 2.8, i.e. the average change in threshold DACs per trim DAC. To achieve a good agreement of the two integrator baselines, this process is repeated multiple times. After each iteration, another threshold scan is performed to evaluate the overlap of the S-curves. The trimming procedure is considered successful when the S-curves of the two integrators are aligned within 2.5 threshold DACs for all channels of the tested PACIFIC Board. This is the level of accuracy that can be achieved given the resolution of the trim DACs. In comparison to a typical photoelectron peak separation of 15 threshold DACs, this is an acceptable level of accuracy.

Figure 6.7a shows the integrator differences of the 192 PACIFIC Boards that are mounted on C-Frame 2 using the trimming constants as determined during the QA procedure for the boards. By design, the variations are within 2.5 threshold DACs when verifying the trimming within the QA stand. However, it was found that this result is not directly applicable to the case where the identical boards are installed in ROBs and mounted on the C-Frame. Instead, as can be seen, the trimming distribution is smeared out and shows a significantly larger spread.

In order to mitigate the effect, an attempt was made to redo the trimming of the PACIFIC integrators within the commissioning setup. As shown in the corresponding box plot in Fig. 6.7b, after one iteration of the retrimming, the same performance can be achieved as before in the QA stand. After the second iteration, this is even slightly exceeded. Beyond that, no further improvements can be observed.

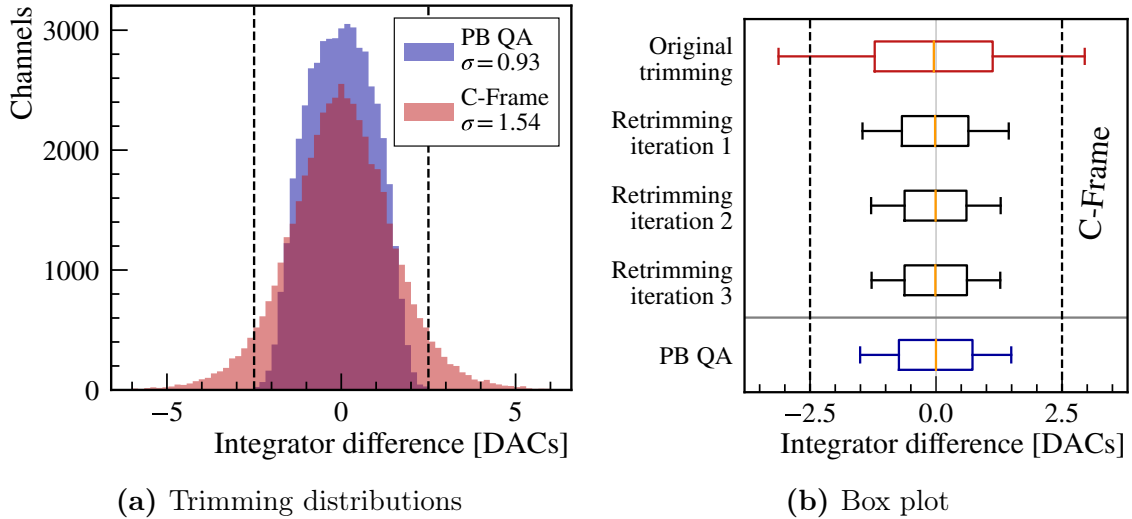


Figure 6.7: Integrator differences of 192 PACIFIC Boards (PBs) after the trimming as performed and measured during the PB QA and after being installed on C-Frame 2 using the identical trimming constants (a). As illustrated in the box plot (b), one retrimming iteration is required on the C-Frame to achieve the same performance as before in the QA stand. Additional trimmings only change the observed differences slightly. The box boundaries indicate the 25th and 75th percentiles, in addition to the median (50th percentile) as the orange lines inside the boxes. The whiskers mark the 5th and 95th percentiles. The dashed lines indicate the region in which the trimming is considered successful.

The tight schedule of the commissioning (see Chapter 5) did not allow for the retrimming to be performed for each C-Frame during the commissioning itself. However, as shown here, it is needed in order to restore or even exceed the intended quality. Therefore, to ensure consistent performance between the two PACIFIC integrators, a retrimming of all channels will be necessary after the installation of the complete detector in the LHCb cavern and before starting to take data.

6.1.4 Comparator Differences

Another fundamental property of the PACIFIC channels is the uniformity of the three comparators. Although, in contrast to the integrator differences, it does not directly influence the performance, knowing the level of agreement can have an impact on the daily operation of the detector. For instance, if a strong correlation between the three comparators is observed, only scanning the thresholds of one comparator would be required and could be applied to the remaining two (compare with Section 4.2.4).

The measurement is carried out analogously to the assessment of the integrator differences as presented previously in Section 6.1.3. By using the internal charge injection mechanism, it is defined by the differences in the determined charge transition




Figure 6.8: Differences in the determined charge transitions between the different PACIFIC comparators. The underlying data is taken from the 50k channels on C-Frame 2 using the internal charge injection mechanism. The differences to the (reference) comparator 1 in each channel are displayed. The mean values μ and standard deviations σ of the distributions are indicated on the right.

positions (see illustration in Fig. 6.2) between the different comparators. The results as obtained from the 50k channels on C-Frame 2 are shown in Fig. 6.8. Compared to comparator 1, which is used as reference, a similar spread of $\sigma = 1.6$ DACs can be observed for both remaining comparators. In addition, systematic biases up to 0.9 DACs between comparator 3 and 1 is visible.

While the slight biases are negligible compared to the overall spread, the resulting differences are at the edge of being relevant compared to a typical photoelectron peak separation of 15 DACs. For the majority of channels, scanning one of the three comparators and assuming a similar behaviour for the remaining two can be sufficient for some applications. However, for optimal performance, a full threshold scan of all comparators is advisable for the calibration of the PACIFIC thresholds prior to taking physics data.

6.2 Clock Distribution

In view of partially small margins of errors when operating the front-end electronics using a common set of clock timing settings as presented in Section 5.12, concerns about the stability of the clock phase configurations arose. This is to be investigated by means of a series of clock timing scans:

1. Initial scans (reference)
2. Consecutive scans with identical conditions
3. After power cycling the ROBs

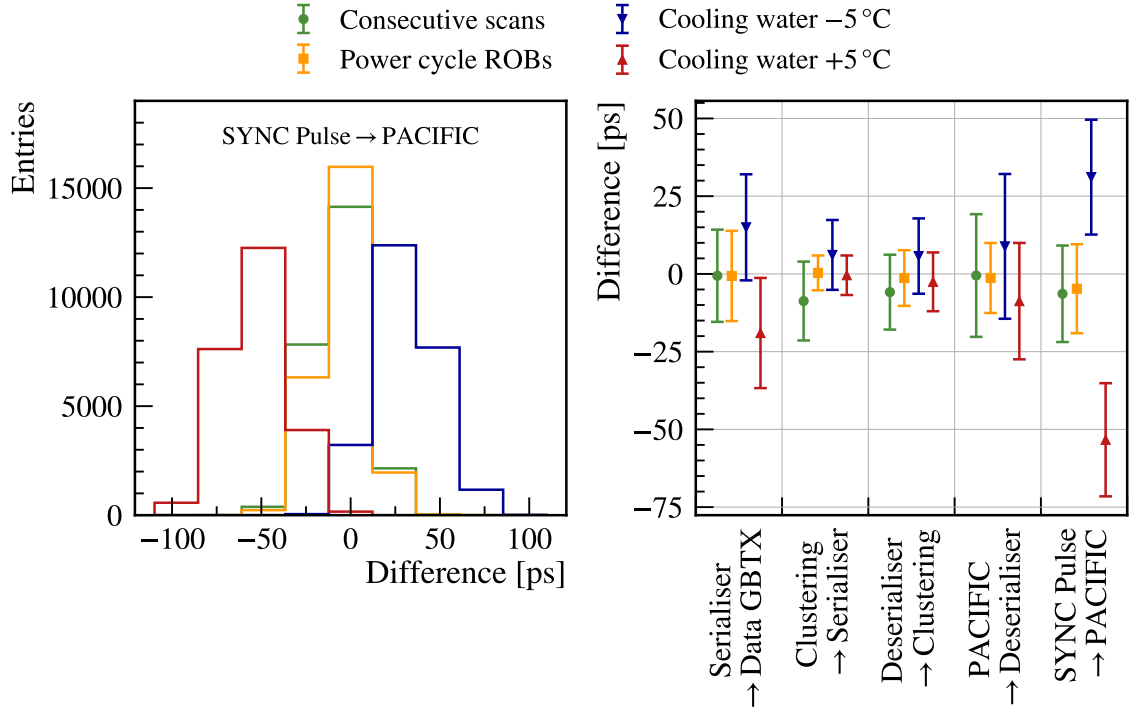


Figure 6.9: Differences in the optimal clock phases between the SYNC Pulse and PACIFIC clock for different run conditions compared to a reference (left). Measured by the centres of the error-free intervals. The right plot summarises the mean and standard deviation of the underlying distributions for the remaining clock phases within the SciFi front-end electronics as indicated by the markers and error bars, respectively.

4. Change cooling water temperature by -5°C ($T = 15^{\circ}\text{C}$)
5. Change cooling water temperature by $+5^{\circ}\text{C}$ ($T = 25^{\circ}\text{C}$)

The clock scans were carried out on the electronics located on the bottom half of C-Frame 2 (12 ROBs, 25k channels). It is expected that especially the variations in the cooling water and therefore the temperature of the electronics has an effect due to slight changes in the electrical conductivity. Note that when changing the cooling water temperature by $\pm 5^{\circ}\text{C}$, the average temperatures of the boards changed by about $\pm 4^{\circ}\text{C}$.

The clock phases are scanned in the smallest available step size of 48.8 ps in order to achieve the highest possible level of accuracy. The determined error-free clock phase interval are each compared with the initial scans that serve as a reference. In general, differences to the reference scans are found to be in the order of the step size resolution and therefore subject to relatively large uncertainties. However, significant differences are observed in the phase between the SYNC Pulse and PACIFIC clock for the different cooling water temperature conditions. This is shown in Fig. 6.9 (left) that displays the phase differences between the centres of the error-free intervals.

While performing a consecutive run with identical conditions as well as power cycling the ROBs only results in minor variations around 0, systematic biases can be observed between different cooling water temperatures. On average, a shift by about 80 ps is visible between the two scans taken with the water temperature set to 15 °C and 25 °C, which can be explained by the length of the corresponding clock and signal lines. As shown in Fig. 6.9 (right), a similar trend can be observed for the other clock phases, although not to the same extent.

In summary, any of the observed variations lie well below the margins of errors as found in Section 5.12 when operating the front-ends at a common set of clock timing settings. As discussed, this also applies to rather unrealistic changes of the cooling water temperature. After all, the temperature of the cooling water, as well as the ambient temperature, is more strictly controlled in the LHCb cavern as it is the case in the assembly hall on the surface.

6.3 Dark Count Rate

It has been discussed in Section 3.6 that the dark count rate (DCR) contributes significantly to the detector noise once the SiPM arrays are exposed to radiation. To mitigate this, the detector is designed to allow for the cooling of the SiPMs down to $-40\text{ }^{\circ}\text{C}$ (see Section 3.8.2). While the radiation dose cannot be varied within the commissioning setup consisting of new hardware, it can be used to study the influence of temperature and the applied overvoltage (ΔV) on the DCR using the complete readout chain.

Because of the processing of the SiPM signals inside the PACIFIC ASICs, the total DCR is reduced compared to the DCR of the SiPM arrays alone. This is due to the integration stage that accumulates the signals within each bunch crossing of 25 ns. Thereby, a dark count event may fall below the applied threshold if it arrives between two integration windows. In the following, the thresholds of the PACIFIC comparators are set to be just (0.5 pe) above the determined pedestal positions. For instance, with a typical photoelectron peak separation of 15 DACs at $\Delta V = 3.5\text{ V}$, this corresponds to 7.5 DACs. For each condition, a total of 60M events are recorded.

Due to the thermal origin of the dark counts, the DCR strongly depends on the temperature of the SiPMs. To verify this, a series of measurements were carried out at different temperatures: $-40\text{ }^{\circ}\text{C}$, $-10\text{ }^{\circ}\text{C}$ and $20\text{ }^{\circ}\text{C}$. Figure 6.10a shows the determined rates along the 128 channels of one exemplary SiPM array. In addition, the box plot in Fig. 6.10b summarizes the results from the 192 arrays located on the bottom half of C-Frame 2, revealing an exponential relationship between the temperature and DCR. Due to the fact that unirradiated SiPMs are in use at this stage, the measured rates even at room temperature are still comparatively low and are in line with the expected values of a few kHz.

Analogous to this, measurements were also carried out at different overvoltages: 2.5 V, 3.5 V (nominal) and 4.5 V. However, due to the overall low rates no significant difference could be found between the three settings.

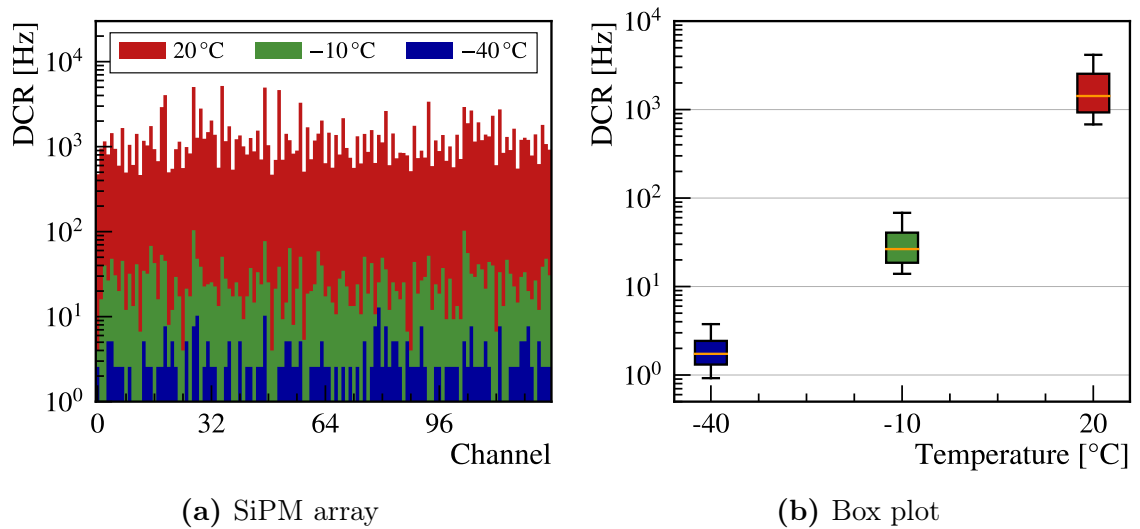


Figure 6.10: Dark count rate (DCR) along the 128 channels of one SiPM array for different SiPM temperatures (a). The box plot (b) shows the measured rates for the 192 SiPMs located on the bottom half of C-Frame 2. The box boundaries indicate the 25th and 75th percentiles, in addition to the median (50th percentile) as the orange lines inside the boxes. The whiskers mark the 5th and 95th percentiles.

7 Conclusion

The LHCb experiment at the LHC is a specialised detector that probes the Standard Model of particle physics and theories beyond it by performing precision measurements of CP violation and studying rare decays in the b and c quark sector. During the first years of operation between 2010 and 2018, huge samples of processes involving these quarks were collected and enabled the publication of key results in the field of flavour physics. Although the detector was already operating at twice the design luminosity towards the end of this period, many measurements are still limited by the statistical sensitivity of the data samples. To overcome this, LHCb has undergone a major upgrade in preparation for another fivefold increase in instantaneous luminosity starting from the next data taking period in the course of 2022. The complete tracking system has been replaced in the process, along with the front-end electronics of all sub-detectors. The latter was required in order to enable a trigger-less readout of the complete detector at the LHC bunch crossing rate of 40 MHz, which will in particular benefit analyses with hadronic final states.

The Scintillating Fibre (SciFi) Tracker is an essential component in the upgraded tracking system of the LHCb experiment. With a fibre diameter of 250 μm , sufficient granularity is provided for the increasing track multiplicities. In addition, the SciFi front-end electronics have been developed and optimised to enable the readout of the detector at 40 MHz.

During the three-year upgrade period, the 12 C-Frames that make up the SciFi Tracker have been assembled. While the overall (tracking) performance was already verified in previous test beam campaigns for individual fibre modules and Readout Boxes (ROBs), the following challenge was to ensure that the same quality could be maintained for the entire detector. In case of the front-end electronics, this means that a total of 256 ROBs need to be thoroughly examined in their final environment.

In the scope of this thesis, the complex system was put into operation for the first time at the level of individual C-Frames. It consists of up to 24 ROBs that need to be operated in conjunction with the surrounding infrastructure and the upgraded LHCb DAQ system. A key challenge in this context was to establish the readout of the detector at 40 MHz. With several hundred optical data links per C-Frame, data rates of up to 2 Tbit/s are generated and need to be processed. To ensure an error-free transmission along the path of the data, this requires a careful tuning of half a dozen clock phases per connection.

After overcoming the initial challenges, a detailed test sequence has been defined, developed and implemented. The so-called commissioning procedure ranges from a series of basic functional tests to taking measurements that involve the complete data chain. It is performed on each individual C-Frame before it is released for installation

7 Conclusion

in the LHCb cavern. After a slow start and despite multiple interruptions due to the COVID-19 pandemic, the commissioning quickly became a routine operation.

The results from the first 9 commissioned C-Frames corresponding to 196 ROBs and 400k channels were presented. Critical failures occurred in 5% of the cases and required the replacement of the corresponding ROBs. Since the commissioning takes place in a dedicated hall on the LHCb site, this operation is significantly easier and safer than performing it after the installation of the detector in the cavern. The most common error concerned the communication between the control system and the front-end electronics via the optical links. At a total of 6 ROBs, temporary losses of the acquired frame-locks were observed. These typically only last for a few milliseconds, but can potentially lead to unwanted resets of the front-end electronics.

Apart from the critical errors, minor issues such as broken sensors were encountered that did not require further interventions. This also includes a total of 3 dead and 7 malfunctioning detector channels, which were identified in the course of the commissioning. However, in view of more than 400k channels, this translates to a rate of functional channels well above 99.99%.

The performance studies that have been conducted alongside the commissioning confirmed the high quality of the front-end electronics. It was demonstrated that they behave consistently and stably even with slight variations in conditions. However, it was found that at least one additional retrimming iteration of the internal integrator baselines is required after the installation of the C-Frames in order to achieve a uniform performance between the two PACIFIC integrators in each channel.

Overall, the commissioning of the SciFi front-end electronics has been a great success and demonstrated its flawless functionality on a large scale. No unsolvable issues have been encountered along the way allowing the detector to be installed in the LHCb cavern and ready to take data. The commissioning procedure has been developed using the same software and tools as will be used in the final system. In addition, it has been conducted in conjunction with the same hardware that implements the control and DAQ system in the upgraded LHCb experiment. Thereby, a lot of experience has been gained that will be invaluable for the operation of the detector in the course of the next years.

Bibliography

- [1] S. L. Glashow, “Partial-symmetries of weak interactions”, *Nucl. Phys.*, vol. 22, pp. 579–588, 1961, ISSN: 0029-5582. DOI: 10.1016/0029-5582(61)90469-2.
- [2] A. Salam and J. Ward, “Electromagnetic and weak interactions”, *Phys. Lett.*, vol. 13, pp. 168–171, 1964, ISSN: 0031-9163. DOI: 10.1016/0031-9163(64)90711-5.
- [3] S. Weinberg, “A Model of Leptons”, *Phys. Rev. Lett.*, vol. 19, pp. 1264–1266, 21 Nov. 1967. DOI: 10.1103/PhysRevLett.19.1264.
- [4] P. Higgs, “Broken symmetries, massless particles and gauge fields”, *Physics Letters*, vol. 12, no. 2, pp. 132–133, 1964, ISSN: 0031-9163. DOI: [https://doi.org/10.1016/0031-9163\(64\)91136-9](https://doi.org/10.1016/0031-9163(64)91136-9).
- [5] A. D. Sakharov, “Violation of CP invariance, C asymmetry, and baryon asymmetry of the universe”, *Soviet Physics Uspekhi*, vol. 34, no. 5, pp. 392–393, May 1991. DOI: 10.1070/pu1991v034n05abeh002497.
- [6] G. Bertone, D. Hooper and J. Silk, “Particle dark matter: Evidence, candidates and constraints”, *Physics Reports*, vol. 405, no. 5, pp. 279–390, 2005, ISSN: 0370-1573. DOI: <https://doi.org/10.1016/j.physrep.2004.08.031>.
- [7] T. Kajita, “Atmospheric neutrino results from super-kamiokande and kamiokande — evidence for ν_μ oscillations”, *Nuclear Physics B - Proceedings Supplements*, vol. 77, no. 1, pp. 123–132, 1999, ISSN: 0920-5632. DOI: [https://doi.org/10.1016/S0920-5632\(99\)00407-7](https://doi.org/10.1016/S0920-5632(99)00407-7).
- [8] ATLAS Collaboration, “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”, *Physics Letters B*, vol. 716, no. 1, pp. 1–29, 2012, ISSN: 0370-2693. DOI: <https://doi.org/10.1016/j.physletb.2012.08.020>.
- [9] CMS Collaboration, “Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC”, *Physics Letters B*, vol. 716, no. 1, pp. 30–61, 2012, ISSN: 0370-2693. DOI: <https://doi.org/10.1016/j.physletb.2012.08.021>.
- [10] LHCb Collaboration, “LHCb detector performance”, *International Journal of Modern Physics A*, vol. 30, no. 07, p. 1530022, 2015. DOI: 10.1142/S0217751X15300227.
- [11] R. Aaij *et al.*, “Measurement of the $B_s^0 \rightarrow \mu^+ \mu^-$ branching fraction and effective lifetime and search for $B^0 \rightarrow \mu^+ \mu^-$ decays”, *Phys. Rev. Lett.*, vol. 118, p. 191801, 19 May 2017. DOI: 10.1103/PhysRevLett.118.191801.

- [12] LHCb Collaboration, “Update of the LHCb combination of the CKM angle γ ”, CERN, Geneva, Tech. Rep. LHCb-CONF-2018-002, CERN-LHCb-CONF-2018-002, May 2018. [Online]. Available: <https://cds.cern.ch/record/2319289>.
- [13] R. Aaij *et al.*, “Test of lepton universality in beauty-quark decays”, CERN, Geneva, Tech. Rep. LHCb-PAPER-2021-004, CERN-EP-2021-042, Mar. 2021. [Online]. Available: <https://cds.cern.ch/record/2758740>.
- [14] L. Evans and P. Bryant, “LHC machine”, *Journal of Instrumentation*, vol. 3, no. 08, S08001–S08001, Aug. 2008. DOI: 10.1088/1748-0221/3/08/s08001.
- [15] ATLAS Collaboration, “The ATLAS Experiment at the CERN Large Hadron Collider”, *Journal of Instrumentation*, vol. 3, no. 08, S08003–S08003, Aug. 2008. DOI: 10.1088/1748-0221/3/08/s08003.
- [16] CMS Collaboration, “The CMS experiment at the CERN LHC”, *Journal of Instrumentation*, vol. 3, no. 08, S08004–S08004, Aug. 2008. DOI: 10.1088/1748-0221/3/08/s08004.
- [17] ATLAS Collaboration and CMS Collaboration, “Combined Measurement of the Higgs Boson Mass in pp Collisions at $\sqrt{s} = 7$ and 8 TeV with the ATLAS and CMS Experiments”, *Phys. Rev. Lett.*, vol. 114, p. 191 803, 19 May 2015. DOI: 10.1103/PhysRevLett.114.191803.
- [18] ALICE Collaboration, “The ALICE experiment at the CERN LHC”, *Journal of Instrumentation*, vol. 3, no. 08, S08002–S08002, Aug. 2008. DOI: 10.1088/1748-0221/3/08/s08002.
- [19] LHCb Collaboration, “The LHCb detector at the LHC”, *Journal of Instrumentation*, vol. 3, no. 08, S08005–S08005, Aug. 2008. DOI: 10.1088/1748-0221/3/08/s08005.
- [20] LHCb Collaboration, “LHCb Tracker Upgrade Technical Design Report”, Tech. Rep. CERN-LHCC-2014-001, LHCb-TDR-015, Feb. 2014. [Online]. Available: <https://cds.cern.ch/record/1647400>.
- [21] C. Elsässer. “ $\bar{b}b$ Production angle plots”, [Online]. Available: https://lhcb.web.cern.ch/speakersbureau/html/bb_ProductionAngles.html (visited on 25th Jan. 2022).
- [22] LHCb Collaboration. “LHCb Operations Plots Webpage”, [Online]. Available: <https://lbgroups.cern.ch/online/OperationsPlots/index.htm> (visited on 25th Jan. 2022).
- [23] LHCb Collaboration, “Letter of Intent for the LHCb Upgrade”, CERN, Geneva, Tech. Rep. CERN-LHCC-2011-001, Mar. 2011. [Online]. Available: <http://cds.cern.ch/record/1333091>.
- [24] LHCb Collaboration, “LHCb VELO Upgrade Technical Design Report”, Tech. Rep. CERN-LHCC-2013-021, LHCb-TDR-013, Nov. 2013. [Online]. Available: <https://cds.cern.ch/record/1624070>.

- [25] P. Kopciewicz, T. Szumlak, M. Majewski *et al.*, “The upgrade I of LHCb VELO—towards an intelligent monitoring platform”, *Journal of Instrumentation*, vol. 15, no. 06, pp. C06009–C06009, Jun. 2020. DOI: 10.1088/1748-0221/15/06/c06009.
- [26] K. Hennessy, “LHCb VELO upgrade”, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 845, pp. 97–100, 2017, Proceedings of the Vienna Conference on Instrumentation 2016, ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2016.04.077>.
- [27] I. Carli, “Upstream Tracker - The silicon strip tracking detector for the LHCb Upgrade”, *PoS*, vol. ICHEP2020, 724. 5 p, 2021. DOI: 10.22323/1.390.0724.
- [28] LHCb Collaboration, “LHCb PID Upgrade Technical Design Report”, Tech. Rep. CERN-LHCC-2013-022, LHCb-TDR-014, Nov. 2013. [Online]. Available: <https://cds.cern.ch/record/1624074>.
- [29] M. Fiorini, “The upgrade of the LHCb RICH detectors”, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 952, p. 161 688, 2020, 10th International Workshop on Ring Imaging Cherenkov Detectors (RICH 2018), ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2018.12.003>.
- [30] M. Baszczyk, P. Carniti, L. Cassina *et al.*, “CLARO: An ASIC for high rate single photon counting with multi-anode photomultipliers”, *Journal of Instrumentation*, vol. 12, no. 08, P08019–P08019, Aug. 2017. DOI: 10.1088/1748-0221/12/08/p08019.
- [31] LHCb Collaboration. “Trigger schemes”, [Online]. Available: <http://lhcb.web.cern.ch/lhcb/speakersbureau/html/TriggerScheme.html> (visited on 28th Jan. 2022).
- [32] LHCb Collaboration, “LHCb Trigger and Online Upgrade Technical Design Report”, Tech. Rep. CERN-LHCC-2014-016, LHCb-TDR-016, May 2014. [Online]. Available: <https://cds.cern.ch/record/1701361>.
- [33] P. Moreira, A. Marchioro and Kloukinas, “The GBT: A proposed architecture for multi-Gb/s data transmission in high energy physics”, 2007. DOI: 10.5170/CERN-2007-007.332.
- [34] R. Schwemmer, personal communication, 22nd Oct. 2020.
- [35] J. Cachemiche, P. Duval, F. Hachon, R. L. Gac and F. Réthoré, “The PCIe-based readout system for the LHCb experiment”, vol. 11, no. 02, P02013–P02013, Feb. 2016. DOI: 10.1088/1748-0221/11/02/p02013.
- [36] T. Colombo, A. Amihalachioaei, K. Arnaud *et al.*, “The LHCb online system in 2020: Trigger-free read-out with (almost exclusively) off-the-shelf hardware”, *Journal of Physics: Conference Series*, vol. 1085, p. 032 041, Sep. 2018. DOI: 10.1088/1742-6596/1085/3/032041.

- [37] F. Alessio and R. Jacobsson, “System-level Specifications of the Timing and Fast Control system for the LHCb Upgrade”, CERN, Geneva, Tech. Rep. LHCb-PUB-2012-001, CERN-LHCb-PUB-2012-001, Feb. 2014. [Online]. Available: <https://cds.cern.ch/record/1424363>.
- [38] R. Arink *et al.*, “Performance of the LHCb Outer Tracker”, *JINST*, vol. 9, P01002, 2014. DOI: 10.1088/1748-0221/9/01/P01002.
- [39] P. d’Argent *et al.*, “Improved performance of the LHCb Outer Tracker in LHC Run 2”, *JINST*, vol. 12, P11016, 2017. DOI: 10.1088/1748-0221/12/11/P11016.
- [40] J. Luisier, “Performance of LHCb Silicon Tracker Detector in the LHC”, *Physics Procedia*, vol. 37, pp. 851–858, 2012, Proceedings of the 2nd International Conference on Technology and Instrumentation in Particle Physics (TIPP 2011), ISSN: 1875-3892. DOI: <https://doi.org/10.1016/j.phpro.2012.04.097>.
- [41] LHCb Collaboration, “LHCb outer tracker: Technical Design Report”, Tech. Rep. CERN-LHCC-2001-024, 2001. [Online]. Available: <https://cds.cern.ch/record/519146>.
- [42] LHCb Collaboration, “LHCb inner tracker: Technical Design Report”, Tech. Rep. CERN-LHCC-2002-029, 2002. [Online]. Available: <https://cds.cern.ch/record/582793>.
- [43] LHCb Scintillating Fibre Tracker Collaboration. “Material for conferences”, [Online]. Available: https://twiki.cern.ch/twiki/bin/view/LHCb/SciFiConference#Material_for_conferences (visited on 6th May 2021).
- [44] J.B. Birks, “Chapter 3 - the scintillation process in organic materials”, in *The Theory and Practice of Scintillation Counting*, ser. International Series of Monographs in Electronics and Instrumentation, Pergamon, 1964, pp. 39–67, ISBN: 978-0-08-010472-0. DOI: <https://doi.org/10.1016/B978-0-08-010472-0.50008-2>.
- [45] L. Gavardi, “Studies for the LHCb SciFi Tracker”, Dissertation, Technische Universität Dortmund, Oct. 2017. DOI: 10.17877/DE290R-18775.
- [46] A. Cavalcante, B. Dey, L. Gavardi *et al.*, “Refining and testing 12,000 km of scintillating plastic fibre for the LHCb SciFi tracker”, *Journal of Instrumentation*, vol. 13, no. 10, P10025–P10025, Oct. 2018. DOI: 10.1088/1748-0221/13/10/p10025.
- [47] Kuraray Co., Ltd. “Plastic Scintillating Fibers (Materials and Structures)”, [Online]. Available: <http://kuraraypsf.jp/psf/index.html> (visited on 10th Jun. 2021).
- [48] T. Förster, “Transfer mechanisms of electronic excitation energy”, *Radiation Research Supplement*, vol. 2, pp. 326–339, 1960, ISSN: 04858611. [Online]. Available: <http://www.jstor.org/stable/3583604>.

- [49] R. C. Ruchti, “The use of scintillating fibers for charged-particle tracking”, *Annual Review of Nuclear and Particle Science*, vol. 46, no. 1, pp. 281–319, 1996. DOI: 10.1146/annurev.nucl.46.1.281.
- [50] J. Menne, “The LHCb SciFi Tracker: studies on scintillating fibres and development of quality assurance procedures for the SciFi serial production”, Dissertation, Technische Universität Dortmund, Mar. 2018. DOI: 10.17877/DE290R-18868.
- [51] A. B. Rodrigues Cavalcante, F. Garnier, L. Gavardi *et al.*, “Shrinking of bumps by drawing scintillating fibres through a hot conical tool”, CERN, Geneva, Tech. Rep. LHCb-PUB-2016-010, CERN-LHCb-PUB-2016-010, Apr. 2016. [Online]. Available: <http://cds.cern.ch/record/2144875>.
- [52] C. Joram, U. Uwer, B. D. Leverington *et al.*, “LHCb Scintillating Fibre Tracker Engineering Design Review Report: Fibres, Mats and Modules”, CERN, Geneva, Tech. Rep. LHCb-PUB-2015-008, CERN-LHCb-PUB-2015-008, Mar. 2015. [Online]. Available: <https://cds.cern.ch/record/2004811>.
- [53] Y. LeCun and Y. Bengio, “Convolutional networks for images, speech, and time series”, in *The Handbook of Brain Theory and Neural Networks*. Cambridge, MA, USA: MIT Press, 1998, pp. 255–258, ISBN: 0262511029. DOI: 10.5555/303568.303704.
- [54] C. Joram and T. Schneider, “Mirroring of fibre ends for the LHCb SciFi project”, CERN, Geneva, Tech. Rep. LHCb-PUB-2014-020, CERN-LHCb-PUB-2014-020, LHCb-INT-2013-060, Feb. 2014. [Online]. Available: <https://cds.cern.ch/record/1662543>.
- [55] U. Uwer, “Revised SciFi stereo-layer geometry”, CERN, Geneva, Tech. Rep. LHCb-INT-2017-010, CERN-LHCb-INT-2017-010, May 2017. [Online]. Available: <https://cds.cern.ch/record/2264922>.
- [56] Particle Data Group, “Review of Particle Physics”, *PTEP*, vol. 2020, no. 8, p. 083C01, 2020. DOI: 10.1093/ptep/ptaa104.
- [57] O. Girard, G. Haefeli, A. Kuonen, L. Pescatore, O. Schneider and M. E. Stramaglia, “Characterisation of silicon photomultipliers based on statistical analysis of pulse-shape and time distributions”, 2018. arXiv: 1808.05775 [physics.ins-det].
- [58] A. Ghassemi, K. Kobayashi and K. Sato, “A technical guide to silicon photomultipliers (MPPC)”, Hamamatsu Photonics K.K., Tech. Rep., Jan. 2018. [Online]. Available: <https://hub.hamamatsu.com/us/en/technical-note/sipm.html>.
- [59] A. Kuonen, “Development and Characterisation of Silicon Photomultiplier Multichannel Arrays for the Readout of a Large Scale Scintillating Fibre Tracker”, Ph.D. dissertation, IPHYS, Lausanne, 2018, p. 163. DOI: 10.5075/epfl-thesis-8842.

Bibliography

- [60] Hamamatsu Photonics K.K., “SiPM S13552-H2017 Delivery Specification Sheet”, CERN Engineering Data Management Service (EDMS), Internal access, Oct. 2016. [Online]. Available: <https://edms.cern.ch/document/1742017/1>.
- [61] O. G. Girard, “Development of the scintillating fibre tracker technology for the LHCb upgrade and the LHC beam profile monitoring system”, Ph.D. dissertation, IPHYS, Lausanne, 2018, p. 140. DOI: 10.5075/epfl-thesis-8851.
- [62] W-IE-NE-R Power Electronics GmbH, “MARATON Power Supply System - Technical Manual”, Tech. Rep., Jan. 2019. [Online]. Available: https://file.wiener-d.com/documentation/MARATON/WIENER_MARATON_HighRad_Manual_A2.pdf.
- [63] CAEN S.p.A., “A1539, AG539, A154X, AG54X Power Supply Boards Manual (Rev. 11)”, Tech. Rep., Feb. 2021. [Online]. Available: <https://www.caen.it/products/a1539b>.
- [64] CAEN S.p.A., “SY4527 - SY4527LC Power Supply Systems Manual (Rev. 21)”, Tech. Rep., May 2021. [Online]. Available: <https://www.caen.it/products/sy4527>.
- [65] C. Joram, I. C. and S. A., “Requirements for the LHCb SciFi water cooling system”, CERN Engineering Data Management Service (EDMS), Internal access, Mar. 2018. [Online]. Available: <https://edms.cern.ch/document/1919838/2>.
- [66] B. D. Leverington, “C-Frames: Mechnics and Services”, SciFi General Meeting, Sep. 2017. [Online]. Available: <https://indico.cern.ch/event/664962/>.
- [67] D. Banerjee, J. Bernhard, M. Brugger *et al.*, “The North Experimental Area at the Cern Super Proton Synchrotron”, Jul. 2021. [Online]. Available: <https://cds.cern.ch/record/2774716>.
- [68] P. Durante, J.-P. Cachemiche, G. Vouters *et al.*, “MiniDAQ1: A Compact Data Acquisition System for GBT Readout over 10G Ethernet at LHCb”, in *Proceedings of International Conference on Technology and Instrumentation in Particle Physics 2017*, Z.-A. Liu, Ed., Singapore: Springer Singapore, 2018, pp. 332–336, ISBN: 978-981-13-1313-4.
- [69] K. Akiba, P. Ronning, M. van Beuzekom *et al.*, “The Timepix Telescope for high performance particle tracking”, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 723, pp. 47–54, 2013, ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2013.04.060>.
- [70] W. Vink, A. Pellegrino, G. Ietswaard, J. Verkooijen, U. Carneiro and A. Massefferi, “LHCb Scintillating Fiber detector front end electronics design and quality assurance”, *Journal of Instrumentation*, vol. 12, no. 03, pp. C03053–C03053, Mar. 2017. DOI: 10.1088/1748-0221/12/03/c03053.

- [71] LHCb Scintillating Fibre Tracker Collaboration, “PACIFICr5 Carrier Board (v4)”, CERN Engineering Data Management Service (EDMS), Internal access, 2018. [Online]. Available: <https://edms.cern.ch/document/2001668/1>.
- [72] J. M. de Cos, H. Chanal, A. Comerma Montells *et al.*, “PACIFIC: SiPM readout ASIC for LHCb upgrade”, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 912, pp. 354–358, 2018, New Developments In Photodetection 2017, ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2017.12.044>.
- [73] A. Comerma, “PACIFIC: The Readout ASIC for the LHCb Scintillating Fibre Tracker”, in *2017 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*, 2017, pp. 1–3. DOI: [10.1109/NSSMIC.2017.8532693](https://doi.org/10.1109/NSSMIC.2017.8532693).
- [74] A. Comerma Montells and J. Mazorra De Cos, “PACIFICr5 ASIC Data Sheet”, CERN, Geneva, Tech. Rep. LHCb-PUB-2021-006, CERN-LHCb-PUB-2021-006, May 2021. [Online]. Available: <https://cds.cern.ch/record/2766810>.
- [75] A. Comerma Montells, “Development of a multichannel integrated circuit for Silicon Photo-Multiplier arrays readout”, Presented on Jan 2014, Ph.D. dissertation, Universitat de Barcelona, 2013. [Online]. Available: <http://hdl.handle.net/10803/134876>.
- [76] H. Chanal, A. Comerma, D. Gascón *et al.*, “PACIFIC: SiPM Readout ASIC for the LHCb Upgrade”, 11th International Meeting on Front-End Electronics, May 2018. [Online]. Available: <https://indico.cern.ch/event/688153/>.
- [77] NXP Semiconductors, “I²C-bus specification and user manual”, Oct. 2021. [Online]. Available: <https://www.nxp.com/docs/en/user-guide/UM10204.pdf>.
- [78] R. W. Hamming, “Error detecting and error correcting codes”, *The Bell System Technical Journal*, vol. 29, no. 2, pp. 147–160, 1950. DOI: [10.1002/j.1538-7305.1950.tb00463.x](https://doi.org/10.1002/j.1538-7305.1950.tb00463.x).
- [79] L. Witola, “Calibration and performance studies of the readout ASIC for the LHCb SciFi Tracker”, M.S. thesis, University of Heidelberg, 2019. [Online]. Available: https://www.physi.uni-heidelberg.de/Publications/MA_Witola.pdf.
- [80] S. Gentile, E. Kuznetsova and F. Meddi, *Photon detection efficiency of Geiger-mode avalanche photodiodes*, 2010. arXiv: 1006.3263 [physics.ins-det].
- [81] P. C. Consul and G. C. Jain, “A Generalization of the Poisson Distribution”, *Technometrics*, vol. 15, no. 4, pp. 791–799, 1973. DOI: [10.1080/00401706.1973.10489112](https://doi.org/10.1080/00401706.1973.10489112).
- [82] L. Witola, “Readout Electronics for the LHCb Scintillating Fibre Tracker”, Poster, Jun. 2019. [Online]. Available: <https://indico.cern.ch/event/827947/>.

- [83] A. Caratelli, S. Bonacini, K. Kloukinas *et al.*, “The GBT-SCA, a radiation tolerant ASIC for detector control and monitoring applications in HEP experiments”, *JINST*, vol. 10, no. 03, p. C03034, 2015. DOI: 10.1088/1748-0221/10/03/C03034.
- [84] P. Moreira, R. Ballabriga, S. Baron *et al.*, “The GBT Project”, 2009. DOI: 10.5170/CERN-2009-006.342.
- [85] A. Caratelli, D. H. Montesinos and K. Kloukinas, “GBT-SCA User Manual”, Tech. Rep., 2019, Internal access. [Online]. Available: <https://espace.cern.ch/GBT-Project/GBT-SCA>.
- [86] “IEEE Standard for Test Access Port and Boundary-Scan Architecture”, *IEEE Std 1149.1-2013 (Revision of IEEE Std 1149.1-2001)*, pp. 1–444, 2013. DOI: 10.1109/IEEESTD.2013.6515989.
- [87] W. Vink, “LHCb SciFi Clusterization Board specifications and manual”, Nikhef Redmine, Internal access, Jun. 2021. [Online]. Available: https://redmine.nikhef.nl/et/projects/clusterization_board.
- [88] G. Haefeli, “Contribution to the development of the acquisition electronics for the LHCb experiment”, Presented on Aug 2004, Ph.D. dissertation, 2004. [Online]. Available: <http://cds.cern.ch/record/800810>.
- [89] S. Esen, J. Van Tilburg, L. Del Buono, P. Billoir and L. Henry, “Clustering and rawbank decoding for the SciFi detector”, CERN, Geneva, Tech. Rep. LHCb-INT-2018-024, CERN-LHCb-INT-2018-024, Jul. 2018. [Online]. Available: <https://cds.cern.ch/record/2630154>.
- [90] W. Vink, “Master Board, technical manual”, Nikhef Redmine, Internal access, Feb. 2021. [Online]. Available: https://redmine.nikhef.nl/et/projects/master_board_v2.
- [91] W. Vink, “Master Board (v4)”, CERN Engineering Data Management Service (EDMS), Internal access, May 2019. [Online]. Available: <https://edms.cern.ch/document/2150024/1>.
- [92] G. Ripamonti, M. Ursino, S. Saggini, S. Michelis and F. Faccio, “Regulated resonant switched-capacitor point-of-load converter architecture and modeling”, *IEEE Transactions on Power Electronics*, vol. 36, no. 4, pp. 4815–4827, 2021. DOI: 10.1109/TPEL.2020.3019841.
- [93] M. Menouni, P. Gui and P. Moreira, “The GBTIA, a 5 Gbit/s Radiation-Hard Optical Receiver for the SLHC Upgrades”, 2009. DOI: 10.5170/CERN-2009-006.326.
- [94] G. Mazza, F. Tavernier, P. Moreira *et al.*, “High-Speed, Radiation-Tolerant Laser Drivers in 0.13 μm CMOS Technology for HEP Applications”, *IEEE Transactions on Nuclear Science*, vol. 61, no. 6, pp. 3653–3659, 2014. DOI: 10.1109/TNS.2014.2361679.

- [95] L. Amaral, S. Dris, A. Gerardin *et al.*, “The versatile link, a common project for super-LHC”, *Journal of Instrumentation*, vol. 4, no. 12, P12003–P12003, Dec. 2009. DOI: 10.1088/1748-0221/4/12/p12003.
- [96] P. Moreira, J. Christiansen and K. Wyllie, “GBTX Manual”, Tech. Rep., 2021, Internal access. [Online]. Available: <https://espace.cern.ch/GBT-Project/GBTX>.
- [97] S. Lin and D. Costello, *Error Control Coding*. Jan. 2004, ISBN: 978-1-4613-6787-1. DOI: 10.1007/978-1-4615-3998-8_3.
- [98] G. Papotti, “An Error-Correcting Line Encoding ASIC for a HEP Rad-Hard Multi-GigaBit Optical Link”, in *2006 Ph.D. Research in Microelectronics and Electronics*, 2006, pp. 225–228. DOI: 10.1109/RME.2006.1689937.
- [99] F. Vasey *et al.*, “The Versatile Link Application Note”, Tech. Rep. Version 2.7, Jan. 2017, CERN Engineering Data Management Service (EDMS). [Online]. Available: <https://edms.cern.ch/document/1309737/1>.
- [100] W. Karpinski, “Light injection system performance and production”, SciFi Front-End Electronics Production Readiness Review, 2018. [Online]. Available: <https://indico.cern.ch/event/690897/>.
- [101] W. Vink, P. Perret and X. Han, “Master Board, Cluster Board and PACIFIC Board Production”, SciFi Front-End Electronics Production Readiness Review, 2018. [Online]. Available: <https://indico.cern.ch/event/690897/>.
- [102] “LHCb SciFi Production Database”, Internal access. [Online]. Available: <https://scifi.physi.uni-heidelberg.de>.
- [103] MapCERN, Building 3852, Oct. 2021. [Online]. Available: [https://maps.cern.ch/?n=\[%273852%27\]](https://maps.cern.ch/?n=[%273852%27]).
- [104] A. Massafferri, A. Pellegrino, D. Berninghoff *et al.*, “FEE Test System - Read-Out Box Test System Protocol for SciFi Detector”, Jul. 2020.
- [105] ETM professional control GmbH, “SIMATIC WinCC Open Architecture”, Eisenstadt, Austria. [Online]. Available: <http://etm.at>.
- [106] C. Gaspar. “WinCC-OA Introduction for Newcomers”, [Online]. Available: <https://lhcb-online.web.cern.ch/ecs/pvssintro.htm> (visited on 25th Feb. 2022).
- [107] L. G. Cardoso, C. Gaspar, J. V. V. Barbosa and F. Alessio, “Controlling DAQ electronics using a SCADA framework”, in *2016 IEEE-NPSS Real Time Conference (RT)*, 2016, pp. 1–4. DOI: 10.1109/RTC.2016.7543123.
- [108] O. Holme, M. González-Berges, P. Golonka and S. Schmeling, “The JCOP Framework”, CERN, Geneva, Tech. Rep. CERN-OPEN-2005-027, Sep. 2005. [Online]. Available: <https://cds.cern.ch/record/907906>.

Bibliography

- [109] D. Berninghoff, T. Damen, M. Feo Rivello *et al.*, “SciFi Tracker Controls and Monitoring - Design of the SciFi ECS”, CERN Engineering Data Management Service (EDMS), Internal access, Dec. 2018. [Online]. Available: <https://edms.cern.ch/document/2060160>.
- [110] W. Vink, “LV 10mF Capacitor Board production files”, Nikhef Redmine, Internal access, May 2020. [Online]. Available: <https://redmine.nikhef.nl/et/documents/192>.
- [111] L. Del Buono, H. Chanal, O. Le Dortz, A. Pellegrino and W. Vink, “LHCb Upgrade SciFi Tracker Front-End Data Format”, Tech. Rep. 1898940 v.4, May 2021, CERN Engineering Data Management Service (EDMS). [Online]. Available: <https://edms.cern.ch/document/1898940/4>.
- [112] R. Schwemmer, J. P. Cachemiche, N. Neufeld, C. Soos, J. Troska and K. Wyllie, “Evaluation of 400 m, 4.8 Gbit/s Versatile Link lengths over OM3 and OM4 fibres for the LHCb upgrade”, *Journal of Instrumentation*, vol. 9, no. 03, pp. C03030–C03030, Mar. 2014. DOI: 10.1088/1748-0221/9/03/c03030.
- [113] Avago Technologies, “MiniPOD™ AFBR-811VxyZ, AFBR-821VxyZ Datasheet”, Tech. Rep. AV02-2842EN, Mar. 2013. [Online]. Available: <https://docs.broadcom.com/doc/AV02-2842EN>.
- [114] Fluke Networks, “MultiFiber™ Pro Datasheet”, Tech. Rep., Sep. 2020. [Online]. Available: <https://www.flukenetworks.com/content/datasheet-multifiber-pro>.
- [115] K. Wyllie, “Status of VTRX issue”, LHCb Technical Board Meeting, May 2021. [Online]. Available: <https://indico.cern.ch/event/1040629/>.
- [116] Semiconductor Components Industries, “Understanding Data Eye Diagram Methodology for Analyzing High Speed Digital Signals”, Tech. Rep., Jun. 2015. [Online]. Available: <https://www.onsemi.com/pub/collateral/and9075-d.pdf>.
- [117] C. Insa, “Front-End Electronics Mechanics”, CERN Engineering Data Management Service (EDMS), Mar. 2019. [Online]. Available: <https://edms.cern.ch/document/1759287/1>.

Acknowledgements

Nowadays, scientific research is hardly possible without a strong team effort. And so also the thesis presented here would not have been possible without the great support of many people, to whom I would like to express my gratitude on the last page.

First and foremost, this applies to my supervisor Prof. Dr. Ulrich Uwer who gave me the opportunity to join his group and work on an exciting research topic. During the many years of this thesis, he always gave useful advice and provided constant support and encouragement, which really meant a lot to me.

Furthermore, I would like to thank Prof. Dr. Klaus Reygers for the interest in my work and for reviewing this dissertation, as well as Prof. Dr. Tilman Plehn and Prof. Dr. Ulrich A. Glasmacher for completing the examination committee.

This thesis has been carried out in the context of the LHCb SciFi Tracker, which involves people from more than a dozen institutes actively working together on the project. For this reason I would like to express my gratitude to all members of the SciFi collaboration who directly or indirectly contributed to this work.

A special word of thanks goes to Xiaoxue Han from whom I have learned a lot over the last years of tightly working together. I could always count on her support and would hereby also like to offer my appreciation for proofreading parts of this thesis.

In addition, I would like to thank my colleagues from the SciFi electronics team – in particular Magali Magne, Wilco Vink, Maurício Féo, José Mazorra de Cos, Olivier Le Dortz, André Massafferri and Diogo Ayres Rocha for a close collaboration on various topics throughout this work. Special thanks goes to Lukas Witola with whom I worked closely together during the commissioning and who has gradually taken over many of my responsibilities. In the context of the commissioning, I would also like to express my gratitude to the LHCb Online group for providing extensive support in the areas of the DAQ system and the network infrastructure.

I would like to thank the entire Heidelberg LHCb group for a great working environment and the nice get-togethers for coffee and cake. Thanks to my (former) Heidelberg SciFi colleagues Blake Leverington, Sebastian Bachmann, Albert Comerma, Michał Dziewiecki and David Gerick for fruitful discussions in our hardware meetings in earlier phases of the project.

And last but not least, I would like to thank my family and friends for their continuous support and for always believing in me. Without having them at my back, this thesis would not have been possible.