
System Tests and the Development of Detector Control Systems for the ATLAS Inner Tracker and the FASER Experiment

Dissertation

for the award of the degree

“Doctor rerum naturalium” (Dr. rer. nat.)

Division of Mathematics and Natural Sciences

of the Georg-August-Universität Göttingen

within the Physics doctoral program

of the Georg-August University School of Science (GAUSS)

submitted by

Hans Ludwig Joos

Göttingen, 2025

Thesis Committee:

Prof. Dr. Stan Lai

II. Institute of Physics, Georg-August-Universität Göttingen

Dr. Benedikt Vormwald

CERN

apl. Prof. Dr. Jörn Große-Knetter

II. Institute of Physics, Georg-August-Universität Göttingen

Members of the Examination Board:

Reviewer: Prof. Dr. Stan Lai

II. Institute of Physics, Georg-August-Universität Göttingen

Reviewer: Dr. Benedikt Vormwald

CERN

Further members of the Examination Board:

Prof. Dr. Stefan Dreizler

Institute of Astrophysics, Georg-August-Universität Göttingen

apl. Prof. Dr. Jörn Große-Knetter

II. Institute of Physics, Georg-August-Universität Göttingen

Prof. Dr. Steffen Schumann

Institute of Theoretical Physics, Georg-August-Universität Göttingen

PD Dr. Benjamin Schwenker

II. Institute of Physics, Georg-August-Universität Göttingen

Date of the oral examination: 17 December 2025

Reference: II.Physik-UniGö-Diss-2025/04

System Tests and the Development of Detector Control Systems for the ATLAS Inner Tracker and the FASER Experiment

Abstract

To exploit the collider's full potential, the LHC will be upgraded to the High Luminosity-LHC. The instantaneous luminosity will increase five- to sevenfold from luminosities achieved during LHC Run 3. To cope with the resulting increase in occupancy, bandwidth and radiation damage, ATLAS will replace the Inner Detector with an all-silicon system, the Inner Tracker (ITk). The innermost part of the ITk will consist of a hybrid pixel detector with five layers that will consist of modules combined into serially powered (SP) chains and loaded on ring and stave shaped low mass carbon fiber local supports (LLS).

Validating this design requires rigorous system tests using prototypes. One such system test was conducted using the demonstrator for the Outer Barrel (OB) region. A prototype longeron LLS was intensively tested and its performance evaluated. This thesis presents the results of these tests, which validate the overall system design for the OB region.

Following successful validation, the project has entered the preproduction phase for LLSs. During preproduction, the Quality Control (QC) setups, that will be used during production, are being qualified. This thesis presents results from qualification steps of the CERN OB LLS Quality Control (QC) setup, as well as developments for all OB LLS QC setups to enable efficient testing during production.

For the safe operation of these QC setups, a sophisticated Detector Control System (DCS) was developed. The operational model of the detector was translated into the state model of a Finite State Machine (FSM) ensuring safe operation. This thesis presents the control system, including benchmarks showing its suitability for even larger systems.

FASER is a small forward experiment 480 m from the ATLAS interaction point and searches for new particles in regions not accessible by ATLAS. This thesis also presents contributions made to the FASER DCS to integrate a new Water Leak Detection system and to control the new Preshower upgrade of FASER with a completely newly developed control system.

Contents

List of Terms and Acronyms	xi
Glossary	xi
Acronyms	xiii
1 Introduction	1
2 Hadron Collider Physics at the LHC	5
2.1 The Standard Model of particle physics	6
2.1.1 Fundamental forces and particles	6
2.1.2 The Higgs mechanism	9
2.1.3 Limitations of the Standard Model	10
2.2 The Large Hadron Collider	10
2.3 The High Luminosity Large Hadron Collider upgrade	15
3 The ATLAS Experiment	17
3.1 Coordinate system and particle kinematics	18
3.2 Particle identification with the ATLAS detector	19
3.3 The ATLAS detector configuration for Run 3	20
3.3.1 The Inner Detector	21
3.3.2 The Calorimeters	25
3.3.3 The Muon Spectrometer	27
3.3.4 Triggering and Data Acquisition	29
3.3.5 Luminosity measurement	30
3.4 Detector and upgrades overview for the HL-LHC era	31
3.4.1 Trigger and Data Acquisition system upgrade	31
3.4.2 Upgrades in other areas	34
3.5 Requirements for the ITk upgrade and general overview	35

4	Particle Physics Detectors	43
4.1	Interactions of particles with matter	44
4.1.1	Interactions of charged particles	44
4.1.2	Multiple coulomb scattering	47
4.1.3	Interactions of photons	48
4.1.4	Strong interactions of hadrons	51
4.2	Principles of semiconductor pixel detectors	51
4.2.1	The pn-junction	52
4.2.2	General characteristics of pixel sensors	56
4.2.3	ITk Pixel sensors	64
4.2.4	Monolithic Active Pixel Sensors	65
4.3	ITk Pixel FE chips	66
4.3.1	The RD53A readout chip	68
4.3.2	The ITkPix readout chip	72
4.4	ITk Pixel modules	73
4.5	ITk Outer Barrel local supports	75
4.5.1	Serial powering scheme	78
4.5.2	MOPS monitoring	81
4.5.3	CO ₂ cooling	82
5	ITk Pixel System Tests and the Outer Barrel Demonstrator Program	85
5.1	The test structure	87
5.2	Setup	90
5.3	Powering	91
5.4	MOPSHUB4Beginners	92
5.5	Interlock and environment monitoring	93
5.6	CO ₂ cooling	95
5.7	Readout with the Opto System and FELIX	95
6	Results from the Outer Barrel Demonstrator	99
6.1	Setup commissioning	99
6.1.1	Detector Control System software and powering	100
6.1.2	Interlock	101
6.1.3	Test box and dry air flush test	103
6.1.4	CO ₂ cooling	103
6.1.5	Validation of the MOPS monitoring system	105
6.1.6	Fluctuations in the module temperature monitoring	106
6.1.7	Validation of the grounding and shielding scheme	108
6.1.8	Module readout with the Opto System	109

6.2	Performance studies	110
6.2.1	I(V) measurements	110
6.2.2	V(I) measurements	111
6.2.3	Module performance tests at different stages during assembly	112
6.2.4	Module performance comparison at different temperatures	138
6.2.5	Performance of the longeron under multi-SP-chain operation	140
6.2.6	Source scans	144
6.2.7	Longeron after thermal cycling	144
6.3	Summary and conclusions	149
7	Preparations for Preproduction of ITk Pixel OB Loaded Local Supports	151
7.1	Description of the CERN LLS QC setup	152
7.1.1	Detector Control System software and powering	153
7.1.2	CO ₂ cooling	153
7.1.3	Interlock	154
7.1.4	Online software framework	154
7.2	The preproduction longeron at CERN	155
7.2.1	Calibrations of the MOPS chips	157
7.2.2	Bit Error Rate Tests	161
7.3	Software developments toward uniform and automatic LLS testing	163
7.3.1	Retrieval of archived DCS data	163
7.3.2	Programmatic remote control of the DCS	164
8	Detector Control Systems	167
8.1	ATLAS DCS	167
8.2	The DCS Front-Ends	170
8.2.1	ELMB	170
8.2.2	Power supplies	170
8.2.3	Interlock	171
8.3	The DCS Back-Ends	171
8.3.1	SCADA software	172
8.3.2	JCOP framework	175
8.3.3	OPC-UA servers	175
8.3.4	Other communication protocols	176
8.4	Operation: Finite State Machines for Detector Control Systems	179
8.4.1	Theory of Finite State Machines	180
8.4.2	State Manager Interface framework	181
8.4.3	The Controls Hierarchy framework component	183
8.4.4	ATLAS-specific extensions	184

9	The Detector Control System for the ITk Pixel Detector	187
9.1	The architecture of the ITk Pixel DCS	188
9.1.1	Control and feedback path	188
9.1.2	Safety path	189
9.1.3	Diagnostics path	190
9.2	Requirements for the ITk Pixel DCS	190
9.3	Developments for the DCS for ITk Pixel OB LLS QC setups	192
9.3.1	System overview and project setup	192
9.3.2	Identifying information using aliases	198
9.3.3	Configuration of archive settings	199
9.3.4	Automatic procedures for HV and LV scans	201
9.3.5	Implementation of additional interlock strategies	201
9.3.6	Opto Box monitoring	202
9.3.7	MARTA monitoring	203
9.3.8	MHFB monitoring and control	205
9.3.9	Mass configuration and consistency checks of SP-chains	206
9.3.10	Metadata for detector hardware	207
9.3.11	Deployment and usage	208
9.3.12	WinCC OA HTTP REST API manager	208
9.4	The Finite State Machine for ITk Pixel setups	215
9.4.1	FSM control hierarchy	215
9.4.2	SP-chain Control Unit	217
9.4.3	Testing, operation and performance considerations	225
9.4.4	Discussion of the chosen implementation	227
9.5	Towards a final detector DCS	229
10	Controlling the FASER Experiment	233
10.1	The detector	234
10.1.1	The original detector layout	235
10.1.2	The Preshower upgrade	238
10.2	The Detector Control System of FASER	242
10.3	Commissioning and integration of a Water Leak Detection system	245
10.3.1	The RELIANCE box	245
10.3.2	Installation and calibration of the WALD system	247
10.3.3	Integration of the WALD system into the FASER DCS	251
10.4	The FASER Preshower DCS	251
10.4.1	The Preshower Interlock and Monitoring board	251
10.4.2	The FSM control hierarchy of the Preshower	253
10.4.3	The half-plane Control Unit	255

10.4.4 Testing, commissioning and initial operational experience	259
11 Summary and Conclusions	261
Acknowledgments	265
Bibliography	267
A Connection diagrams	285
B Additional material for the results from the Outer Barrel demonstrator	287
C Additional material for the ITk DCS developments	297
D Additional material for the FASER Preshower DCS	299

List of Terms and Acronyms

Glossary

BERT Bit Error Rate Test. Tests the signal transmission by sending a predefined known pattern and counting the number of bit errors received.

bPOL12V A radiation tolerant [DC-DC converter ASIC](#) for nominal 12 V input.

bPOL2V5 A radiation tolerant [DC-DC converter ASIC](#) for nominal 2.5 V input.

CANopen Communication protocol stack that implements the layers above and including the network layer.

Controller Area Network A vehicle bus standard designed to allow microcontrollers and devices to communicate with each other in applications without a host computer. It is a message-based protocol, designed originally for multiplex electrical wiring within automobiles, but is also used in many other contexts like in ATLAS slow control.

CROC Final [FE](#) chip for the CMS tracker upgrade, however different versions exists: A preproduction chip CROC-V1 (also known as [RD53B-CMS](#)) and a production chip CROC-V2 (also known as [RD53C-CMS](#)).

e-link Logical lower bandwidth electrical link. Several e-links make up a [GBT](#) or [lpGBT](#) link.

FELIX Front-End Link eXchange. FELIX is a network switch to convert [GBT](#)-like protocol data to Gigabit Ethernet or InfiniBand networking.

FwFsm Controls Hierarchy framework component for [WinCC OA](#).

GBCR Giga-Bit Cable Receiver. Radiation-hard equalizer to recover an electrical signal before inputting it into an [lpGBT](#).

GBT Gigabit Bidirectional Trigger and Data Link. Generic name of the “Gigabit transceiver”. Several radiation-hard chipsets are available.

IP Internet Protocol. A network layer communications protocol for relaying data from one host to another across network boundaries.

ITkPix Final [FE](#) chip for the ATLAS tracker upgrade, however different versions exists: A preproduction chip ITkPix-V1.1 (also known as [RD53B-ATLAS](#)) and a production chip ITkPix-V2 (also known as [RD53C-ATLAS](#)).

LDO Low Drop-Out voltage regulator.

lpGBT low power Gigabit Transceiver. Radiation tolerant [ASIC](#) for multipurpose high-speed bidirectional serial links.

MARTA Monoblock Approach for a Refrigeration Technical Application. A portable, evaporative CO₂ cooling system.

MODBUS An application layer communication protocol originally designed for use with [PLCs](#).

Monitoring Of Pixel System Radiation-hard [DCS ASIC](#) to monitor parameters of detector modules and [Opto Boxes](#).

MOPSHUB A communication interface between the [MOPS](#) and the [DCS](#) control station in the counting room also known as [PP3](#).

Opto Board Carries all opto-electrical components. Up to eight opto boards can be installed in one [Opto Box](#).

Opto Box The location of the Pixel opto-transceivers including [ASICs](#) for data handling.

Opto System The system of [Opto Boxes](#) with their [Opto Boards](#) for electrical-to-optical conversion of detector readout data to enable optical readout over long distances.

PLC Programmable Logic Controller. An industrial computer for activities that require high reliability.

PRBS Pseudorandom Binary Sequence. A binary sequence of pseudorandom bits generated by a deterministic algorithm. The PRBS7 algorithms uses a polynomial of degree 7.

Python An interpreted programming language with dynamic typing supporting object-oriented programming.

RAID Redundant Array of Independent Disks. A data storage technology for the purposes of data redundancy, performance improvement, or both.

RD53A Prototype [FE](#) chip.

RD53B RD53B-ATLAS ([ITkPix-V1.1](#)) or RD53B-CMS ([CROC-V1](#)). Preproduction [FE](#) chip in two variants for the experiments at the [LHC](#).

RD53C RD53C-ATLAS ([ITkPix-V2](#)) or RD53C-CMS ([CROC-V2](#)). Production [FE](#) chip in two variants for the experiments at the [LHC](#).

REST Representational State Transfer. A software architectural style for stateless interfaces ([HTTP](#)-based [APIs](#)) using [HTTP](#) methods.

ShuntLDO Shunt Low Drop-Out Voltage regulator. Central circuit for serial powering.

SP-chain Group of Pixel detector modules which are serially powered (SP) by a constant current source.

systemd System and Service Manager for Linux operating systems.

TCP Transmission Control Protocol. A transport layer communications protocol that provides reliable, ordered, and error-checked delivery of messages between applications running on hosts communicating via an [IP](#) network.

VCAN Voltage supply for the [MOPS](#) chip.

VPP3 Voltage supply for the [MOPSHUB](#).

VTRx⁺ Versatile Link⁺ Transceiver. Radiation tolerant electrical/optical transceiver to be placed on the detector.

YARR Yet Another Rapid Readout. [DAQ](#) software for readout based on [PCIe](#) or FELIX.

Acronyms

2PACL two-phase Accumulator Controlled Loop.

AC Alternating Current.

ADC Analog to Digital Converter.

AFE Analog Front-End.

ALICE A Large Ion Collider Experiment.

ALP Axion-Like Particle.

API Application Programming Interface.

APP Active Patch Panel.

ASIC Application Specific Integrated Circuit.

ATLAS A Toroidal LHC Apparatus.

BCM' Beam Conditions Monitor (prime).

BE Back-End.

CA Chasing Averages.

CAD Computer-Aided Design.

CAN [Controller Area Network](#).

CBA Central Buffer Architecture.
CDR Clock and Data Recovery.
CERN European Organization for Nuclear Research.
CI Continuous Integration.
CIC [CAN](#) Interface Card.
CMOS Complementary MOS.
CMS Compact Muon Solenoid.
CNT Carbon Nano-Tube.
COTS Commercially available Off-The-Shelf.
CP Charge conjugation and Parity.
CPU Central Processing Unit.
CSB Cable Saver Board.
CSS CERN Safety System.
CSV Comma-Separated Values.
CTP Central Trigger Processor.
CU Control Unit.

DAC Digital Analog Converter.
DAQ Data Acquisition System.
DBA Distributed Buffer Architecture.
DC Direct Current.
DCS Detector Control System.
DELPHI Detector with Lepton, Photon and Hadron Identification.
DEN Device Editor and Navigator.
DIM Distributed Information Manager.
DIP Data Interchange Protocol.
DP Data Point.
DPE Data Point Element.
DPT Data Point Type.
DSS Detector Safety System.
DU Device Unit.

EC (Outer) End Cap.
ECal Electromagnetic Calorimeter.
ELMB Embedded Local Monitor Board.
EMB Electromagnetic Barrel.
EMCI Embedded Monitoring and Control Interface.
EMEC Electromagnetic Endcap.
EMP Embedded Monitoring and control Processor.
EWSB Electroweak Symmetry Breaking.

EYETS Extended Year-End Technical Stop.

FASER ForwArd Search ExpeRiment.

FDR Final Design Review.

FE Front-End.

FFC Flat Flexible Cable.

FIFO First In, First Out.

FPGA Field Programmable Gate Array.

FSM Finite State Machine.

GCS Global Control Station.

GND Ground.

GUI Graphical User Interface.

HCal Hadronic Calorimeter.

HDD Hard Disk Drive.

HGTD High-Granularity Timing Detector.

HL-LHC High Luminosity Large Hadron Collider.

HLT High Level Trigger.

HTTP Hypertext Transfer Protocol.

HV High Voltage.

HVIF HV Interlock Fan-in.

IBL Insertable B-Layer.

ID Inner Detector.

IDE Integrated Development Environment.

IFT Interface Tracker.

IHR Inclined Half-Ring.

IMC Interlock Matrix Crate.

IP Interaction Point.

IS Inner System.

IST Inner Support Tube.

ITk Inner Tracker.

JCOP Joint Controls Project.

JSON JavaScript Object Notation.

L0 level 0 trigger (name of hardware trigger used in Run 4).

L0A level 0 accept signal.

L1 level 1 trigger (name of hardware trigger used in Run 3).

LAN Local Area Network.

LAr Liquid Argon.
LB Logic Board.
LED Light-Emitting Diode.
LEP Large Electron–Positron Collider.
LGAD Low Gain Avalanche Detector.
LHC Large Hadron Collider.
LHCb Large Hadron Collider beauty.
Linac4 Linear accelerator 4.
LISSY Local Interlock Safety System.
LLP Long-Lived Particle.
LLS Loaded Local Support.
LPE Low Power Enable (signal to enable low power mode in [RD53B](#) and [RD53C](#) readout chips).
LS Long Shutdown.
LU Logical Unit.
LV Low Voltage.

MAPS Monolithic Active Pixel Sensor.
MDT Monitored Drift Tube.
MHFB [MOPSHUB](#) for Beginners, or [MOPSHUB4Beginners](#).
MIC Main Interlock Crate.
MICROMEGAS micro-mesh gaseous structure.
MIP Minimum Ionizing Particle.
MOPS [Monitoring Of Pixel System](#).
MOS Metal-Oxide-Semiconductor.
MOSFET MOS Field-Effect Transistor.

NIC Network Interface Card.
NIEL Non-Ionizing Energy Loss.
NMOS n-type MOS Field-Effect Transistor.
NSW New Small Wheel.
NTC Negative Temperature Coefficient Thermistor.

OB Outer Barrel.
OPC-UA Open Platform Communications Unified Architecture.

PCB Printed Circuit Board.
PCIe PCI Express (Peripheral Component Interconnect Express).
PDR Preliminary Design Review.
PDU Power Distribution Unit.

PIM Preshower Interlock and Monitoring.
PMOS p-type MOS Field-Effect Transistor.
PMT Photomultiplier Tube.
PoP Proof of Principle.
PP Patch Panel (counted, starting at 0, along the connections from the module towards the service cavern).
PPB Patch Panel Board.
PRR Production Readiness Review.
PS Proton Synchrotron.
PSB Proton Synchrotron Booster.
PST Pixel Support Tube.
PSU Power Supply Unit.
PTC Positive Temperature Coefficient Thermistor.

QC Quality Control.
QCD Quantum Chromodynamics.
QED Quantum Electrodynamics.
QFT Quantum Field Theory.

RELIANCE Reliable Liquid Detection for Critical Environments.
RPC Resistive Plate Chambers.

SCADA Supervisory Control and Data Acquisition.
SCT Semiconductor Tracker.
SEE Single-Event Effect.
SM Standard Model.
SMD Surface-Mount Device.
SMI State Manager Interface.
SML State Manager Language.
SNMP Simple Network Management Protocol.
SPS Super Proton Synchrotron.
SQL Structured Query Language.
sTGC small-strip TGC.

TDAQ Trigger and Data Acquisition.
TGC Thin Gap Chambers.
TID Total Ionizing Dose.
Tilock Temperature Interlock.
TIM Tracker Interlock and Monitoring.
ToT Time over Threshold.

TRT Transition Radiation Tracker.

TSV Tab-Separated Values.

TTC Timing, Trigger and Control.

UDP User Datagram Protocol.

UI User Interface.

URL Uniform Resource Locator.

VME Versa Module Eurocard.

WALD Water Leak Detection.

WinCC OA SIMATIC WinCC Open Architecture.

XML Extensible Markup Language.

CHAPTER 1

Introduction

Questions such as “What is the world made of?” or “How does the universe work?” have driven people for centuries to invest time and resources into research in the field of physics. In early Greece, the philosopher Leucippus and his pupil Democritus coined the term “atom” (from the Greek word *atomos* meaning “uncuttable”) for the smallest building block of matter. The idea that matter is not a continuum that can be divided arbitrarily into smaller pieces was at first only of philosophical nature. But this idea caught on in science in later centuries, and many atomic models were developed that were able to explain chemical elements, for example. However, with the discovery of the electron, the first subatomic particle, the literal meaning of the word *atom* was proven wrong. The word is still used, but it is now known that the atom consists of, apart from the negatively charged electron, a positively charged nucleus. Scientists later discovered that the nucleus is also divisible and consists of protons and neutrons that also have substructure and consist of so-called quarks.

For the correct description of the physical quantities of the newly discovered subatomic particles and their behavior, a new theory was necessary. Quantum mechanics introduced new concepts such as quantization, the uncertainty principle and the wave-particle duality. It is now the fundamental basis to describe particles at microscopic scale. A theoretical framework that builds on quantum mechanics and also follows the principles of special relativity is Quantum Field Theory (QFT), in which particles are treated as excitations of their respective quantum fields.

One such QFT is the Standard Model (SM) of particle physics. It includes all known elementary particles, namely, particles without a substructure, and three of the four known

fundamental forces. Although it is successful with many precise predictions that were verified in experiments, it fails to answer questions like “Why is there more matter than antimatter?” or “How does dark matter fit into the framework?” There are theories beyond the SM that could explain these phenomena. Particle physicists are interested in probing the SM, exploring further the building blocks and fundamental laws of the universe and finding signs of these new theories.

This exploration is a continuous human effort that has already spanned several decades. To combine resources and bring together scientists from different European countries, the European Organization for Nuclear Research (CERN) was established in 1954. It houses several accelerators and collider experiments on its sites at the border of Switzerland and France. The largest collider at CERN is the Large Hadron Collider (LHC), one of the largest scientific instruments ever built. Since the start of its operation in 2010, a large community of scientists working in fundamental particle physics and high energy physics has been exploring the new energy domain made accessible by the LHC. Among the goals of the LHC are the validation of the SM, the study of known particles and the creation and study of new particles, postulated by theories beyond the Standard Model. One of the greatest successes so far was the experimental discovery of the Higgs boson in 2012 by research groups of the ATLAS and CMS experiments using datasets collected at collision energies of 7 TeV and 8 TeV [1, 2]. To extend the possibilities of new scientific discoveries and to account for radiation damage in current accelerator machinery, the LHC is foreseen to be upgraded to the High Luminosity Large Hadron Collider (HL-LHC) in the late 2020s. The aim is to achieve collision energies of 14 TeV corresponding to the LHC design value and instantaneous luminosities (rate of collisions) of a factor of five larger than the LHC nominal value [3]. [Chapter 2](#) gives a short description of the SM and an introduction to the LHC, including an overview of the upgrades planned for the LHC.

During the HL-LHC installation, also the ATLAS (“A Toroidal LHC Apparatus”) detector [4] will undergo a major upgrade. The Inner Detector (ID) of ATLAS will be replaced by a new all-silicon Inner Tracker (ITk) [5] that includes a hybrid pixel detector with more than four times the amount of detector modules and more than 50 times the amount of readout channels than the current pixel detector. Furthermore, upgrades to other detector subsystems and the readout system are planned to prepare for the HL-LHC era. [Chapter 3](#) introduces the ATLAS experiment, provides an overview of the detector’s technical design, and outlines planned upgrades for the HL-LHC, including a general overview of the requirements and design for the ITk.

Constructing a large detector requires meticulous planning and rigorous verification of every component destined for installation in the experimental cavern. In system tests, the detector’s design is validated by demonstrating that the smallest fully-functional detector

units consistently fulfill detector requirements. One such system test was the RD53A demonstrator program. Using RD53A prototype readout chips, prototype modules were built and loaded onto support structures to form detector units. One such demonstrator for the Outer Barrel (OB) region was a so-called longeron, set up in a lab at CERN and used for system tests. [Chapter 4](#) begins with the fundamentals of particle physics detectors, then describes semiconductor detectors in general and introduces the specific components needed to build the ITk Pixel Detector, providing a technical description of it. The system test setup with the RD53A OB demonstrator is described in [Chapter 5](#). The results of the system tests conducted within this thesis on the prototype longeron are presented in [Chapter 6](#) and validate the overall system design.

The project has since moved from the prototyping phase to preproduction. During preproduction, the Quality Control (QC) setups, which will be used during production to verify compliance of detector units with the requirements for the final detector, are qualified. [Chapter 7](#) presents preparations carried out within this thesis for the CERN QC setup specifically, as well as for other QC setups producing OB detector units.

For the safe operation of these QC setups, a Detector Control System (DCS) was developed with scalability in mind to support future development of a DCS for the final detector installed in ATLAS. [Chapter 8](#) introduces the DCS for ATLAS and the concept of a Finite State Machine (FSM) which models operational parts of a large detector in a hierarchy tree to provide operators with clear, summarized information while concealing overwhelming details. [Chapter 9](#) then focuses on the DCS that has been developed using a [SCADA](#) system as the controlling software for the ITk Pixel OB QC setups for detector units. This includes the Graphical User Interfaces (GUIs) and remote control capabilities that were developed, as well as an FSM hierarchy structure for easy operation. The chapter closes with an outlook on the final DCS of a large ITk detector system.

Additional contributions in the context of this thesis were made to the DCS of the FASER detector. FASER [\[6\]](#) is a small forward experiment 480 m from the ATLAS interaction point and searches for new particles in regions not accessible by ATLAS. With the work relevant to this thesis, a Water Leak Detection system has been installed, calibrated and commissioned for the FASER detector. Furthermore, a DCS has been developed for the Preshower upgrade of the detector. [Chapter 10](#) introduces the FASER detector and describes the Preshower upgrade and its scientific motivation. It then details the integration of the Water Leak Detection system into FASER and the development of the DCS for the Preshower upgrade, concluding with an operational assessment of the developed system during operation at the LHC.

Hadron Collider Physics at the LHC

The Standard Model (SM) [7] is the theoretical framework for nearly all modern particle physics. Even though it bears the name “model” in it, it is the *theory* that describes three of the four known fundamental forces in the universe and classifies all known elementary particles, the basic building blocks of which everything is made.

The Large Hadron Collider (LHC) [8] is the world’s largest particle accelerator and the nine experiments along the accelerator allow physicists to study the predictions of particle physics theories and perform measurements of elementary particles such as the Higgs boson.

This chapter provides an overview of the SM in [Section 2.1](#), with a focus on the role of the Higgs boson, to motivate the accelerators and experiments and their upgrades as presented later in this thesis. [Section 2.2](#) presents the LHC, the largest particle accelerator ever built and the machine that delivers particle collisions for the ATLAS and FASER experiments. Finally, [Section 2.3](#) gives an overview of the upgrades planned for the LHC, such that from 2030 on, the LHC can run as the High Luminosity Large Hadron Collider (HL-LHC). [Sections 2.2](#) and [2.3](#) will provide an explanation of what luminosity is, why physicists are interested in higher luminosity and how the higher luminosity for the HL-LHC will be achieved.

2.1 The Standard Model of particle physics

Particle physics is the study of fundamental particles, the basic building blocks of matter, and the forces that act between them. Four fundamental forces are currently known to exist: gravity, electromagnetism, weak interaction and strong interaction. Gravity is explained on a macroscopic level by Einstein’s general theory of relativity [9, 10] and is very weak at a microscopic level. The dominating theory that explains the other three fundamental forces on a microscopic level, namely electromagnetism, weak and strong interaction is called the Standard Model (SM) of Particle Physics. This theory also classifies all known elementary particles.

The SM was developed in the twentieth century through the work of many scientists worldwide. Through an interplay of theoretical advances on one side and experimental observations and discoveries on the other side, the theory has iteratively reached a state where it proves hugely successful in providing experimental predictions.

Mathematically, the SM is realized as a gauge QFT with global Poincaré symmetry (as in all relativistic QFTs) and local (i.e. spacetime-dependent) gauge symmetry with respect to the non-Abelian gauge group

$$\mathrm{SU}(3)_C \times \mathrm{SU}(2)_L \times \mathrm{U}(1)_Y \quad . \quad (2.1)$$

The meaning of this group structure will be briefly explained below, following to a large extent Refs. [11, 12, 13, 14, 15, 16].

2.1.1 Fundamental forces and particles

Starting from the requirement of local gauge symmetry as described by the symmetry group in Eq. (2.1), the theory predicts different types of interactions and generates the bosonic force carriers. Each subgroup corresponds to a fundamental force. The first major part is the (non-abelian) gauge group $\mathrm{SU}(3)_C$ that corresponds to the strong interaction. It has 8 generators which correspond to the 8 types of gluons, the force carriers of the strong interaction. The theory that describes this interaction is called Quantum Chromodynamics (QCD) [17–19] from the Greek word *chroma* meaning “color”. Color is used to describe the associated “charge” of the $\mathrm{SU}(3)_C$ gauge symmetry (which is emphasized by the index C) to which the gluons couple. Since the gluons themselves carry color charge, they can interact with each other. This is a typical property of a non-abelian gauge theory.

The second major part of the SM is the electroweak theory (sometimes called the Glashow-Weinberg-Salam model [20–22]) as a unification of electromagnetic and weak interaction.

The electroweak theory is a gauge theory with local $SU(2)_L \times U(1)_Y$ symmetry. Here again, the symmetry generates the interaction types and bosonic force carriers. Electroweak theory is, through the introduction of a scalar Higgs field, a spontaneously broken gauge theory with the Electroweak Symmetry Breaking (EWSB):

$$SU(2)_L \times U(1)_Y \xrightarrow{\text{EWSB}} U(1)_{\text{QED}}. \quad (2.2)$$

Before EWSB, the associated charges of $SU(2)_L$ and $U(1)_Y$ are given the name “weak isospin” (labeled T) and “weak hypercharge” (labeled Y) respectively. The index on the $SU(2)_L$ group indicates that the theory is chiral, and the coupling only occurs for left-handed particles which have a weak isospin of $T = \frac{1}{2}$. The generators correspond to the gauge bosons for the weak isospin W_1 , W_2 and W_3 and to the boson for the weak hypercharge B . The SM bosons are then obtained after EWSB due to the Higgs mechanism (see [Section 2.1.2](#)). The (massless) photon and (massive) Z boson emerge as linear combinations of W_3 and B , while the massive W^\pm bosons are combinations of W_1 and W_2 . The W^\pm and Z bosons therefore acquire their mass through coupling with the Higgs field and are the force carriers of the weak interaction. The remaining symmetry is that of Quantum Electrodynamics (QED) [23] with the photon as the force carrier for the electromagnetic interaction. The electric charge Q emerges as a conserved quantity with $Q = T_3 + \frac{1}{2}Y$, where T_3 is the third component of the weak isospin.

The groups of the gauge symmetry in [Eq. \(2.1\)](#) define the fundamental forces and the carrier of those forces as explained above. The carriers are elementary particles called bosons with integer spin quantum numbers. The matter particles, elementary particles with half-integer spin called fermions, are *representations* of those groups, where the representation defines the “charges” of the particles and how they interact with the fundamental forces. As particles with half-integer spin they obey Pauli’s exclusion principle. The particle content of the SM is summarized in [Fig. 2.1](#).

The SM includes three generations of fermions with an increase in their mass in each subsequent generation. Only the first generation is stable. This means that all particles from the second and third generation will eventually decay into particles from the first generation by radiating off a boson. Each fermion has an anti-fermion of the same mass but opposite charge associated with it. The fundamental fermions can be grouped into two types, quarks and leptons. Only the quarks (u , d , c , s , t , b) can take part in strong interactions. The leptons do not carry color charge and therefore do not experience strong interactions. From the six leptons, three carry an electric charge: the electron (e), the muon (μ) and the tau lepton (τ). Each of these leptons has an associated neutral neutrino.




Figure 2.1: Elementary particles of the Standard Model and their properties. The (electrical) charges are in units of the elementary charge. Data is taken from Ref. [24]. (© Carsten Burgard, licensed under CC BY 2.5, with adaptations)

The SM is a renormalizable theory [25, 26], meaning it handles the infinite terms that arise during calculations by relating them to experimentally determined, finite physical quantities. This results in a self-consistent theory with predictive power. As a side effect, the coupling “constants” become dependent on the energy of the interaction. This is particularly interesting, and important, for the phenomenology of proton-proton collisions as in the LHC, since QCD exhibits a behavior called *asymptotic freedom* for the coupling α_s of the strong interaction. This means, that the coupling, and thus the strength of the strong interaction, between particles becomes asymptotically weaker as the energy scale increases and the corresponding length scale decreases. Conversely, at low energies (or large distances), the interaction becomes strong, leading to *color confinement* of quarks and gluons within composite hadrons. As a result, quarks and gluons cannot be isolated, but always appear in color-neutral composites. If, for example, two quarks were to separate from each other far enough, a new quark-antiquark pair would be created to bind the free

quarks—a process known as *hadronization*. A hadron consisting of two quarks is called a meson, while one made up of three quarks is known as a baryon.

The colliding particles in the LHC are protons, which are baryons composed of three valence quarks (u, u, d) and other virtual quarks and gluons, that exist through vacuum fluctuations. As a consequence of QCD, when two protons collide at high energies, it is two partons (a quark or gluon from each proton) that undergo a hard scatter event involving a significant momentum transfer. After the collision, the outgoing partons can emit new partons in a process known as *parton shower*, continuing until the coupling constant α_s grows too large. At that point, hadronization occurs, converting the outgoing partons into collimated sprays of (color-neutral) particles called jets. The remnants of the proton also participate in QCD interactions. This process is usually referred to as the underlying event.

2.1.2 The Higgs mechanism

The electroweak theory is, as mentioned before, a gauge theory. As such, its gauge bosons are required to be massless to not break gauge symmetry. However, experimental evidence shows that the weak interaction is short ranged and thus involves massive bosons. A solution is provided by spontaneously breaking the electroweak symmetry (EWSB) via the associated Higgs mechanism [27–29]: A complex $SU(2)_L$ -doublet scalar field ϕ , called Higgs field, with weak hypercharge $Y = +1$ is added to the theory with potential

$$V(\phi, \phi^\dagger) = -\mu^2(\phi^\dagger\phi) + \lambda(\phi^\dagger\phi)^2. \quad (2.3)$$

The μ and λ are free parameters of the SM. The quartic term describes self-interactions of the Higgs field with the Higgs self-coupling λ . Provided that $\mu^2 > 0$ and $\lambda > 0$, the Higgs field has a vacuum expectation value $v = \sqrt{\mu^2/\lambda} \neq 0$. While the Lagrangian is gauge invariant, the vacuum state is not, and the symmetry is said to be spontaneously broken. In the broken phase, mass terms (dependent on v) for the W^\pm and Z bosons are generated, while the photon stays massless.

Also explicit fermion masses are not allowed in the theory, as they would break the chiral symmetry of the SM. The fermions of the SM acquire their mass through Yukawa coupling with the Higgs field. The nine Yukawa couplings of the Higgs field to the fermions are another set of free parameters of the SM. The Higgs mechanism also predicts a massive scalar boson, the Higgs boson H , whose mass depends on μ . This was the last missing particle of the SM and discovered at the ATLAS and CMS experiments in 2012 [1, 2].

2.1.3 Limitations of the Standard Model

While the SM has proven hugely successful in providing experimental predictions, there are, however, observations that the SM cannot explain. Some examples shall be given here to motivate the search for “new” physics, that could provide explanations to these observations and deepen our knowledge of the universe.

It is experimentally established, through the observation of neutrino oscillations, that neutrinos have mass [30]. Originally, the SM assumed that neutrinos are massless. There are two possible ways to introduce mass for the neutrinos into the SM depending on whether neutrinos are Dirac or Majorana particles, which is, however, up to now not known.

Another observation concerns gravitational effects that are only explained by general relativity when assuming dark matter. Dark matter, which does not interact with light and cannot be seen, makes up about 85 % of the matter content of the universe [31]. The SM does not contain any viable dark matter candidate.

Another example is the observed matter-antimatter asymmetry, that cannot be explained by the amount of CP violation present in the SM [32]. A related problem is that, while CP violation is observed in weak interactions, the experimentally not seen but theoretically possible CP violation in strong interactions poses a fine-tuning problem to the theory. One solution to this problem, the Peccei-Quinn theory [33], involves new scalar particles, called axions that have not been observed to this date.

There are more theoretical limitations to the SM: It was already mentioned before, that the SM does not include a quantum theory of gravity. The SM also does not explain the form of the Higgs potential. Currently, the simplest possibility for the potential in Eq. (2.3) is assumed, but every additional even power of the Higgs field would satisfy the symmetry requirements.

Several extensions of the SM and theories *beyond the SM* (BSM) exist to solve the shortcomings of the SM. To find significant deviations from the SM predictions, all the experiments at the LHC have BSM searches in their physics program.

2.2 The Large Hadron Collider

To perform measurements of elementary particles, such as measuring the properties of the Higgs boson and studying the shape of the Higgs potential in Eq. (2.3), particle physicists are using particle colliders that provide a controlled environment for the experiments. One such particle collider is the Large Hadron Collider (LHC) [8], the worlds largest particle

collider. It collides particles at record energies and is the only operating machine that reaches high enough energies to produce a sufficient amount of Higgs bosons for systematic studies.

The LHC was built by the European Organization for Nuclear Research (CERN). CERN's research facilities are located near Geneva in Switzerland. The site houses several accelerators, with the LHC being the largest and also the most powerful one in the world. The LHC is a ring accelerator and collider inside the 26.7 km long tunnel of the old Large Electron–Positron Collider (LEP) machine [34], between 45 and 170 m underground. It mainly collides proton beams, but can also accelerate lead or oxygen ions. One of the reasons for switching from light leptons at LEP to protons at the LHC is that it is easier to accelerate protons to higher energies E . This is because the energy loss ΔE due to synchrotron radiation is inversely proportional to m^4 [35]:

$$\Delta E \approx \frac{1}{3\varepsilon_0} \frac{q^2}{R} \left(\frac{E}{mc^2} \right)^4, \quad (2.4)$$

where m is the mass of the accelerated particle, q is the charge, ε_0 the vacuum permittivity and R the radius of the ring collider.

The position of the LHC in the larger CERN accelerator complex is shown in Fig. 2.2. The LHC is designed to create proton-proton collisions at center-of-mass energies of up to 13.6 TeV. The LHC is the only accelerator that can reach such high center-of-mass energies, therefore allowing the production of particles with masses up to the TeV-scale. In the acceleration chain of the protons, the LHC is only the last element.

Since 2020, the Linear accelerator 4 (Linac4) [37] is the proton source for the LHC, where negative hydrogen ions (H^-), consisting of a hydrogen atom with an additional electron, are accelerated to an energy of 160 MeV. The Linac4 (as well as the other accelerators) uses high frequency electromagnetic fields to accelerate the particles. These electric fields are stored in periodic structures along the beam line. As a consequence, the particles are not accelerated as a continuous beam but in well-defined packs called bunches. During injection into the next accelerator, the H^- -ions pass through a stripping foil, where the electrons are stripped from the ions leaving only protons. It follows a sequence of synchrotrons with increasing size: the Proton Synchrotron Booster (PSB), the Proton Synchrotron (PS) and the Super Proton Synchrotron (SPS).

The PSB consists of four pipes, one above the other, each forming a ring with a 25 m radius in which the protons are being accelerated. The protons leave the PSB at an energy of 2 GeV and enter the PS where they are accelerated to 26 GeV and bunches are grouped together to form a bunch train with a spacing of 25 ns between each bunch. The last stage

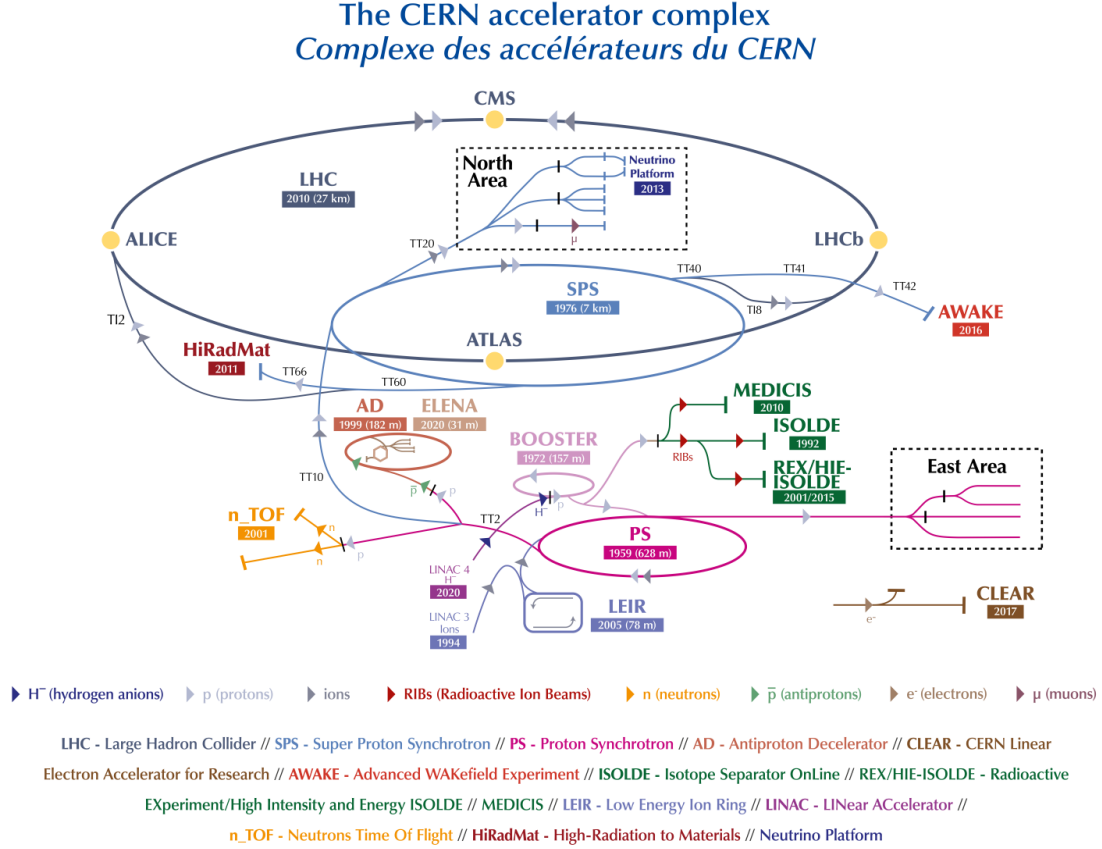


Figure 2.2: The CERN accelerator complex in 2022. (Reproduced from Ref. [36], licensed under CC BY 4.0)

before entering the LHC ring is the SPS, where the protons are accelerated to 450 GeV. From the SPS, the beam is injected into the two beam pipes of the LHC at two points. In one beam pipe, the beam circulates clockwise around the ring and in the other beam pipe, the beam circulates anti-clockwise. Two separate pipes are needed, because the magnetic field that keeps the protons on their circular orbit must have opposite orientation for the two different directions of circulation. The magnetic field is created by 1232 dipole magnets made out of niobium–titanium coils brought to superconductivity by being cooled with superfluid helium to a temperature of 1.9 K. The magnets produce a magnetic field with a strength of 8.3 T to bend the protons' path into an arc. This is done in 8 areas, such that the full LHC ring consists of 8 arcs and 8 straight sections. In one of the straight sections, there are RF cavities that compensate the energy loss due to synchrotron radiation and accelerate the protons to energies of up to 6.8 TeV. Additional beam optics are installed, e.g. quadrupole magnets to focus the beam before a collision point and separation dipoles to deflect it to collide with the other beam.

Depending on the filling scheme, the beam itself consists of up to 2808 bunches of protons with approximately 10^{11} protons per bunch. The nominal time between two bunches at a fixed point in the ring is 25 ns. At four points along the ring, the beams can cross and be brought to collision. The bunch-crossing frequency at these interaction points is then 40 MHz. Around these four interaction points, the four main detectors of the LHC have been built: ATLAS (A Toroidal LHC Apparatus) [4], CMS (Compact Muon Solenoid) [38], LHCb (Large Hadron Collider beauty) [39] and ALICE (A Large Ion Collider Experiment) [40]. ATLAS and CMS are general purpose detectors that were designed independently to allow for cross-confirmation of the measurements between the two experiments. The LHCb detector is a forward spectrometer and dedicated to heavy flavor physics. The ALICE detector is optimized to study heavy ion collisions and address the physics of strongly interacting matter and the quark-gluon plasma. In addition to the four main detectors, there are several much smaller detectors around the LHC that use the same collision points. One such experiment is FASER (ForwArd Search ExpeRiment) [6], an experiment dedicated to searching for light, extremely weakly-interacting particles in the very forward direction of LHC’s high-energy collisions.

The experiments take data of collision events when the beam is at a stable energy. A typical *fill* of the LHC with particle bunches is thus carried out as follows: First the beam has to be injected, that has come from the SPS. Then the energy ramps up and the particles are accelerated. This takes around 25 minutes. At the “flat top”, the particles have their final energy. After that, the bunches are squeezed to increase the so-called luminosity. For several hours following, there is a stable beam at constant energy during which the particle bunches are brought to collision and the detectors can take physics data. Once the beam intensity becomes too low due to collision losses, the beam is dumped and a new fill can be started.

A full *run* of the LHC consists of many fills, usually only interrupted by short technical stops. Run 1 of the LHC with protons at a beam energy of 3.5 TeV started in March 2011. The beam energy was later increased to 4 TeV, corresponding to a center-of-mass energy in the head-on collisions of 8 TeV. Run 1 ended in 2012. After the first Long Shutdown (LS), Run 2 started in 2015 with protons at an energy of 6.5 TeV and ended in 2018. Currently, the LHC is in Run 3, which started in July 2022 and is scheduled to end in June 2026. The LHC is now running at the record collision energy of 13.6 TeV corresponding to 6.8 TeV per proton beam.

During stable beams, when the proton bunches cross, particles can collide. The resulting collisions from one such bunch-crossing are called *events*. There can be many individual proton-proton collisions in a single bunch-crossing. *Pile-up* (μ) is then defined as the average number of particle interactions per bunch-crossing. The average pile-up measured

by ATLAS in Run 2 was $\langle\mu\rangle = 33.7$. The pile-up in Run 3 is expected to be between 52 and 57 collisions per bunch-crossing.

The number of events per unit time, the event rate R , depends on the number of particles in a bunch, on the bunch size and shape, on the bunch-crossing frequency and on the cross-section, a measure for the quantum mechanical probability of the event. This can be formulated as

$$R = \frac{dN_{\text{event}}}{dt} = \sigma_{\text{event}} \cdot \mathcal{L}, \quad (2.5)$$

where σ_{event} is the cross-section of the event of interest measured in barns ($1 \text{ b} = 10^{-24} \text{ cm}^2$). \mathcal{L} is called the (instantaneous) *luminosity*. It summarizes information about the beam and only depends on machine parameters [41]:

$$\mathcal{L} = \frac{f_{\text{rev}} n_b N_1 N_2}{2\pi \sqrt{\sigma_{x,1}^2 + \sigma_{x,2}^2} \sqrt{\sigma_{y,1}^2 + \sigma_{y,2}^2}} \cdot W \cdot F \cdot S, \quad (2.6)$$

where f_{rev} is the revolution frequency of the bunches, n_b the number of bunches, $N_{1,2}$ are the number of particles per bunch and $\sigma_{x,1,2}$ and $\sigma_{y,1,2}$ are the horizontal and vertical beam sizes (assuming a Gaussian profile). The reduction factors W , F and S contain losses in instantaneous luminosity due to beam offsets and crossing angles. In other words, luminosity quantifies the machine's ability to position particles in close proximity, thus determining the potential collision rate. The corresponding unit for luminosity is $\text{cm}^{-2} \text{ s}^{-1}$.

To obtain the total amount of events for a run, one has to integrate over time. Since the cross-section is time independent, this is given by

$$N_{\text{event}} = \sigma_{\text{event}} \cdot \int \mathcal{L} dt = \sigma_{\text{event}} \cdot L_{\text{int}}. \quad (2.7)$$

The integrated luminosity L_{int} can be used as a measure for the performance of a particle accelerator as the number of any event of interest is directly proportional to L_{int} . L_{int} is usually expressed in inverse femtobarns (fb^{-1}) for a complete run and measures the total amount of data that was delivered to the experiments. In Run 1, the LHC was able to deliver an integrated luminosity of around 30 fb^{-1} , while in Run 2 it delivered around 190 fb^{-1} with peak (instantaneous) luminosities of $2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. Until the winter break 2024/2025, the LHC delivered already more than 198 fb^{-1} of integrated luminosity in Run 3, surpassing the total integrated luminosity of the entire four-year duration of Run 2. Consequently, this gives physicists a chance to observe rare processes within an acceptable time and was one main design goal of the LHC. The choice to collide protons with protons and not protons with anti-protons was made because of the difficulties associated with accelerating a high-intensity beam of anti-protons.

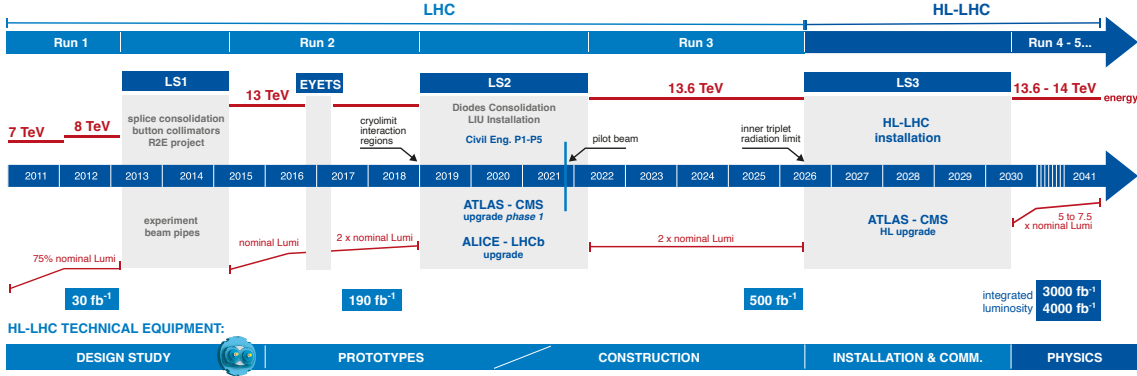


Figure 2.3: The HL-LHC upgrade schedule as of January 2025. The nominal luminosity of the LHC is $1 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. Note here, that the ATLAS and CMS upgrades during LS3 are also called phase 2-upgrades. EYETS stands for Extended Year-End Technical Stop. (Adapted from Ref. [43])

2.3 The High Luminosity Large Hadron Collider upgrade

After Run 3 it is projected that the LHC will have delivered an integrated luminosity of 500 fb^{-1} . Many critical components of the accelerator will have reached the end of their lifetime by then and will need to be replaced [42]. As an example, the inner triplet magnets close to the interaction point, which are responsible for generating small values of β^* , a value related to the transverse size of the particle beam at the interaction point, will reach their projected radiation limit by the year 2026. From a physics perspective there is also a problem: The statistical error for measurements scales with $1/\sqrt{N_{\text{event}}}$. The statistical gain in running the accelerator without an additional considerable luminosity increase beyond its design value will become marginal. To ensure scientific progress and to exploit the full capacity of the LHC, it was decided to upgrade the LHC in the years from 2026 to 2029 to the High Luminosity Large Hadron Collider (HL-LHC) [3]. The upgrade design goal is a five- to sevenfold increase of the instantaneous luminosity and a tenfold increase of the integrated luminosity with respect to the LHC nominal design values. This would mean an instantaneous luminosity of 5×10^{34} to $7.5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ and an integrated luminosity of 250 fb^{-1} per year, a significant increase to the LHC. The upgrade schedule for the HL-LHC project as of January 2025 is shown in Fig. 2.3. With the upgrade, the HL-LHC is to produce up to 4000 fb^{-1} of integrated luminosity by 2038, thereby enabling the experiments to enlarge their data sample by one order of magnitude compared to the LHC baseline program.

The increase in luminosity will be achieved by increasing the beam current and brightness, by reducing the beam spot size at the interaction point and by the introduction of crab cavities. The beam brightness is a beam characteristic that must be maximized at the

beginning of beam generation and then preserved throughout the entire injector chain and throughout the operation cycle in the LHC itself [3]. Therefore, independent of modifications on the LHC, the injector chain needs an upgrade as well [44]. The upgrade work then extends to the existing set of synchrotrons (PSB, PS and SPS) and a new proton source (Linac4) was put in operation in 2020. With this, it will be possible to increase the number of protons per bunch by a factor of two above the nominal design value. To reduce the beam spot size and with it β^* , new large-aperture inner triplet quadrupoles employing advanced superconducting technology based on niobium-tin will be installed. These new magnets generate more intense magnetic fields of 11.3 T, compared to 8.6 T, which allows them to better focus the beams. However, a reduction in β^* also requires larger crossing angles. An increased crossing angle causes a severe reduction of the instantaneous luminosity through the F and S factor in Eq. (2.6). This can be compensated by the use of special superconducting RF crab cavities, which can generate transverse electric fields that rotate each bunch longitudinally, such that the bunches effectively collide head on. It is projected that the effects through the upgrades are such that the number of events per crossing can be raised to around 130, while allowing a slight reduction of bunches from 2808 to 2760.

The higher radiation environment means that also the detectors will have to adjust in order to at least keep detector performance at the same level. The increased pile-up will lead to an increased occupancy and an increased readout bandwidth. But with the larger data sample, the experiments can study rare events and perform high-precision measurements. For example, the Higgs boson self coupling is an extremely important direct probe of the Higgs potential with implications on our understanding of the electroweak phase transition. The self coupling can be measured predominantly via Higgs pair production, a process with very small cross-section. And despite all the efforts so far, Higgs boson pair production has not been observed in the LHC [45]. With the HL-LHC, the estimated significance for the observation of the Standard Model Higgs boson pair production in ATLAS with (without) systematic uncertainties is 3.2σ (4.6σ) [46]. Combined with measurements from CMS, an observation of Higgs pair production is in reach. And although the energy will not increase by a large amount, it might also be possible to search for physics beyond the SM due to the significant increase in luminosity [47, 48].

The ATLAS Experiment

When the LHC accelerates particles to high energies and brings them into collision, experiments record the resulting debris and new particles flying out from the collision point in all directions. By analyzing this data, physicists of an experiment can investigate a wide range of phenomena in particle physics, from detailed studies of the Higgs boson to searches for particles that could constitute dark matter. One such experiment at the LHC is the ATLAS (“A Toroidal LHC Apparatus”) experiment [4]. The ATLAS detector is one of the two general-purpose detectors at the LHC, located at “Point 1” of the LHC near the CERN main site in Meyrin, Switzerland. The detector has a cylindrical shape and is housed approximately 100 m underground. With a length of 46 m, a diameter of 25 m, and weight of about 7000 t, ATLAS is the largest-volume detector ever constructed for a particle collider. A detector of this scale requires a large collaboration: The ATLAS collaboration comprises over 6000 members, including physicists, engineers, technicians, students and administrators, from more than 250 institutes worldwide, all contributing to detector operation, data analysis, and the development of new detector systems. The detector system tests, characterization studies, and DCS developments described in this thesis focus specifically on a new subsystem of the ATLAS detector, the ITk. This chapter introduces the ATLAS experiment, provides an overview of the detector’s technical design, and outlines planned upgrades to certain detector systems for the HL-LHC.

3.1 Coordinate system and particle kinematics

ATLAS uses a right-handed coordinate system with its origin at the nominal Interaction Point (IP) in the center of the detector and the z -axis along the beam pipe. The x -axis points from the IP to the centre of the LHC ring, and the y -axis points upwards. Given that the detector is cylindrical, with a nearly complete solid angle coverage, it is useful to introduce polar coordinates (r, ϕ) in the transverse plane, ϕ being the azimuthal angle around the z -axis. The z -axis is then the polar axis, and θ the polar angle. The coordinate system is shown in Fig. 3.1. The A-side of the detector is defined as the side with positive z and the C-side as the side with negative z .




Figure 3.1: Coordinate system of the ATLAS detector. The A-side (C-side) of the detector is defined as the side with positive (negative) z .

In any collision process, energy and momentum are conserved. For particle kinematics, it is therefore useful to consider the 4-momentum $p = (E/c, p_x, p_y, p_z)$, where c is the speed of light and $p_{x,y,z}$ are the components of the standard 3-momentum \mathbf{p} . It is important to note that in hadron colliders, where the colliding partons carry uncertain momentum fractions, the 4-momentum is generally defined in the lab frame rather than in the center-of-momentum frame of the collision event, which is typically unknown.

The direction of a final-state particle exhibits approximate rotational symmetry around the z -axis (beam axis). To describe this direction relative to the colliding beams, physicists use the polar angle θ , the longitudinal rapidity y or the pseudorapidity η .

Rapidity is defined as $y = \frac{1}{2} \ln \left(\frac{E/c + p_z}{E/c - p_z} \right)$ and has the key advantage that differences in rapidity are invariant under Lorentz boosts along the beam axis. Pseudorapidity is defined as $\eta = \frac{1}{2} \ln \left(\frac{|\mathbf{p}| + p_z}{|\mathbf{p}| - p_z} \right) = -\ln \tan(\theta/2)$ and, crucially, does not depend on the particle's energy. In the relativistic limit where the particle's energy is much greater than its mass, η and y coincide. Then, two special cases for values of pseudorapidity are:

- $\eta = 0$ corresponds to particles emitted perpendicular to the beam line ($\theta = 90^\circ$)
- $\eta \rightarrow \pm\infty$ corresponds to particles moving along the beam direction ($\theta \rightarrow 0^\circ$ or 180°)

Particles with small $|\eta|$ are classified as central, while those with large $|\eta|$ are termed forward.

3.2 Particle identification with the ATLAS detector

As a general-purpose detector, ATLAS is designed to comprehensively detect all possible reaction products. It achieves this through a layered structure surrounding the beam axis, with each layer serving a distinct function using specialized detector technologies.

The innermost layers are dedicated to tracking charged particles in a magnetic field. The magnetic field allows the determination of the charge sign based on the curvature of the particle's trajectory. Moving outward, the electromagnetic and hadronic calorimeters measure the energy of particle showers initiated by electrons, photons, and hadrons. The outermost layer is the muon spectrometer, which identifies muons. Since muons are (mostly) minimally ionizing particles, they traverse the calorimeters with minimal energy loss and without initiating showers. As a result, they are typically only detected in this outermost layer. [Figure 3.2](#) provides an overview of how different particles are detected within the ATLAS detector. Electrons and (relativistic) muons are considered stable particles as they traverse the detector, meaning they do not decay by themselves. In contrast, the tau lepton with a lifetime of 2.903×10^{-13} s travels a non-measurable distance before decaying. Its presence can be reconstructed by measuring the particles produced when it decays. Similarly, hadrons containing b -quarks (lifetime of about 1.5×10^{-12} s) also exhibit short lifetimes, but travel a noticeable path in the detector before decaying. The position of the decay can be reconstructed from the tracking of the decay products [\[49\]](#). This makes vertex detectors close to the interaction point necessary to observe secondary vertices and to enable the identification of such short-lived particles.

Electrons usually deposit their full energy within the electromagnetic calorimeter, where they are completely absorbed. Photons, while producing a similar energy signature in the calorimeter, are electrically neutral and therefore traverse the tracking detector without leaving a trace. Additionally, their trajectory remains unaffected by the magnetic field. Quarks, on the other hand, fragment into jets, collimated sprays of hadrons (e.g. protons or neutrons), due to the phenomenon of hadronization [\[51\]](#). While individual particles within a jet may not always be resolved, the total energy and momentum of the jet can be inferred from the combined energy deposits in the calorimeters.

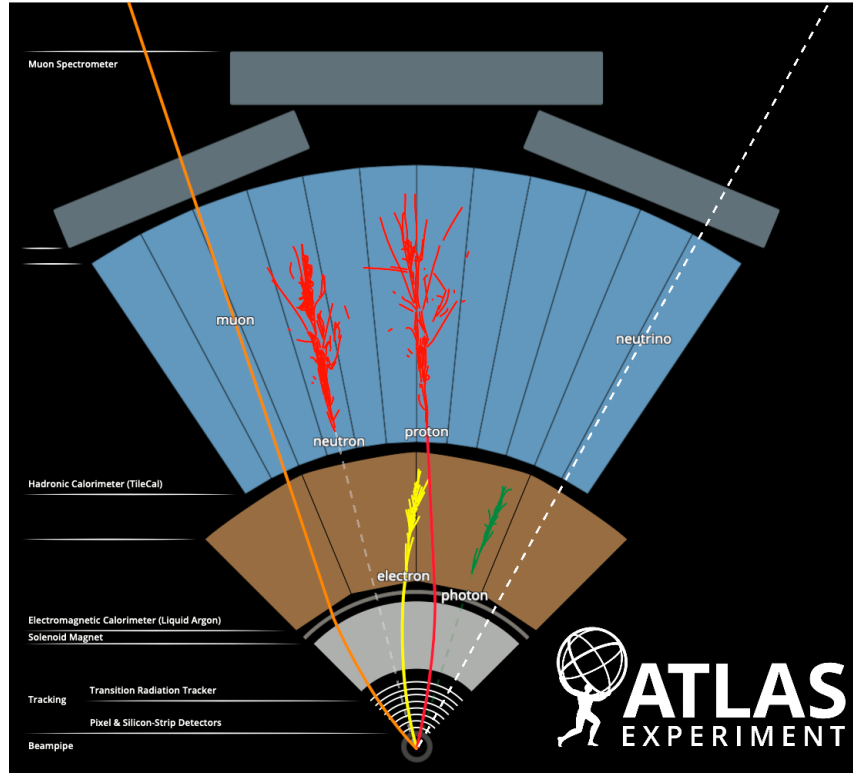


Figure 3.2: Detection of different particles by ATLAS and its detector systems in the transverse plane. The dashed tracks are invisible to the detector. The bent tracks of the charged particles are due to a magnetic field. (Reproduced from Ref. [50], © 2021 CERN)

Neutrinos are not directly detected. They have negligible mass and interact only via the weak force. Their presence must therefore be inferred through precise measurements of missing transverse energy (E_T^{miss}) in collision events, based on the principle of energy conservation.

3.3 The ATLAS detector configuration for Run 3

This section briefly introduces the different sub-detectors of ATLAS with a focus on the current ID. This is the subsystem that will be replaced for the HL-LHC era to meet the performance requirements in the new environment. The other subsystems currently in place for Run 3 (2022–2026) are also briefly introduced. A cut-away illustration of the full detector is shown in Fig. 3.3. For a more detailed description of the various subsystems, refer to Refs. [4, 52].




Figure 3.3: Schematic illustration of the ATLAS experiment in 2022. (Reproduced from Ref. [53], © 2022 CERN)

3.3.1 The Inner Detector — measuring momentum and charge

The purpose of the Inner Detector (ID) [54] is the tracking of charged particles from the interaction point until they reach the calorimeters. The ID is cylindrical around the beam pipe with an approximate length of 6.2 m and a diameter of 2.1 m. Tracking starts as close as practically possible to the beam. The data collected by the ID is then used for track reconstruction, momentum and (primary and secondary) vertex measurement and electron identification. The ID operates within a solenoid magnet that generates a magnetic field B of 2 T parallel to the beam axis and forces charged particles on a curved trajectory with radius R depending on their transverse momentum p_T and charge q :

$$R = \frac{p_T}{qB} \quad . \quad (3.1)$$

The acceptance in pseudorapidity is $|\eta| < 2.5$ for particles originating from the nominal interaction point, with full coverage in the azimuthal angle ϕ . At nominal luminosity, the ID must cope with approximately 1000 particles every 25 ns.

The ID consists of three independent, complementary sub-detectors arranged coaxially around the beam line:

1. the high-resolution silicon Pixel Detector [55], covering radial distances $r < 122.5$ mm,
2. the Semiconductor Tracker (SCT) [56], which uses stereo micro-strips at distances $299 \text{ mm} < r < 514$ mm, and

3. the Transition Radiation Tracker (TRT) [57] comprising several layers of gaseous straw tubes interleaved with transition radiation material.

The layout is shown in Fig. 3.4. The Pixel Detector features three concentric layers in its barrel region (layer B, 1 and 2) with their midpoints at the interaction point, and three end-cap disks on both the A- and C-side. It provides high granularity and high precision measurements close to the vertex region through the use of hybrid pixel sensors, where readout chips are bump-bonded to a silicon sensor. The FE-I3 [58] readout chips in the Pixel Detector have 2880 readout cells with an area of $50\,\mu\text{m} \times 400\,\mu\text{m}$ per cell and individual circuitry for each pixel element. Each chip covers an active area of $16.4\,\text{mm} \times 60.8\,\text{mm}$. The readout chips are oriented such that the $50\,\mu\text{m}$ pitch is used to measure the ϕ coordinate. Each chip is bump-bonded to its sensor substrate and incorporates radiation-hard electronics to withstand the LHC environment. To manage heat and mitigate radiation-induced leakage currents, the system employs evaporative cooling [59] using C_3F_8 (octafluoropropane) as the refrigerant. The basic building block of the active part of the Pixel Detector is the module, comprising 16 FE-I3 chips. There are in total 1744 Pixel modules with approximately 80.4 million readout channels. The total active area of the Pixel Detector amounts to $1.73\,\text{m}^2$.

The innermost pixel layer is called the IBL and was added to the detector between Run 1 and Run 2 of the LHC (2013–2015) at a radius of only 33.25 mm between Layer B of the Pixel Detector and a newly designed, smaller-radius beam pipe [61, 62]. This upgrade addressed performance degradation in the original Pixel Detector layers caused by radiation damage to silicon sensors and front-end electronics and to improve the resolution of track impact parameters¹, thereby enhancing the vertex reconstruction and flavour-tagging performance of the tracking system, which is important for b -hadron identification, as explained in Section 3.2. Due to its proximity to the interaction point, the IBL operates in an environment of extreme particle flux and radiation, necessitating sensors and readout components capable of withstanding these conditions. To meet these challenges, a new front-end readout chip, the FE-I4 [63], was developed. This chip satisfies the requirements of radiation tolerance and readout efficiency at higher luminosity. In comparison to the readout chip used in the other Pixel layers, the cell size was reduced to $250\,\mu\text{m} \times 50\,\mu\text{m}$ with the shorter side being in the transverse plane. In total there are 448 readout chips on 14 local support structures called staves. Each channel is bump-bonded to a single pixel of a sensor. The IBL is cooled by a two-phase CO_2 system which guarantees an operational

¹The *transverse* d_0 and *longitudinal* z_0 impact parameters are defined as the distances between the point of closest approach of the track to a primary vertex in the transverse and longitudinal plane, respectively. An impact parameter significantly larger than the experimental resolution of this quantity hints at the existence of a secondary vertex.




Figure 3.4: The structure of the ATLAS ID for Run 2 and Run 3 made of highly granular silicon pixels (Insertable B-Layer (IBL) and Pixel), silicon strips (SCT) and straw tubes (TRT). The red curved line represents a charged particle traversing the various layers and bending in the 2 T magnetic field. The innermost pixel layer is called the IBL and was added to the detector between Run 1 and Run 2 of the LHC. (Reproduced from Ref. [60], © 2020 CERN)

temperature in the range between -35 to 15 °C. This upgrade added 0.15 m^2 of active surface and approximately 12.04 million readout channels.

The IBL significantly improves the impact parameter resolution for low- p_T tracks in ϕ direction by providing an additional hit point close to the beam and in z direction due to the higher spatial resolution. An adequate number of measurements in the tracking system is essential for robust track reconstruction, with charged particles typically generating four hits in the Pixel Detector on average. The Pixel Detector status at the start of Run 3 was as follows: 10 modules (3.5 %) disabled in the B-layer, 14 modules (2.8 %) disabled in layer 1, 32 modules (4.7 %) disabled in layer 2, and 10 modules (3.5 %) disabled in the Discs. Additionally, three IBL front ends were disabled [52].

The SCT is a silicon microstrip detector positioned around the Pixel Detector comprising four barrel layers with radii from 299 mm to 514 mm and nine disks in each end-cap region.

The detector consists of 4088 modules in total. Each module is built up of 4 rectangular silicon-strip sensors [64] with a second pair of identical sensors glued back-to-back at a stereo angle of 40 mrad [56] enabling improved spatial resolution. A module features 768 strips, each approximately 12 cm in length. The SCT's total number of readout channels is approximately 6.3 million. In the barrel region, the strips are approximately parallel to the beam axis with a pitch of 80 μm . The small stereo angle of 40 mrad between the first and second pair of sensors enhances precision in the longitudinal (z) direction. The detector uses the same evaporative cooling system [59] as the Pixel Detector with C_3F_8 (octafluoropropane) as refrigerant fluid and a target operating temperature of 0 $^\circ\text{C}$ for the SCT modules [65]. Charged particles generate an average of eight hits in the SCT. The SCT was originally designed for an occupancy of 0.2% to 0.5% at LHC pile-up conditions of up to $\mu = 23$, with occupancy increasing with pile-up.

Running under higher pile-up conditions is desired. However, bandwidth limitations on data links from the modules constrain the maximum acceptable occupancy. Simulations during Run 1 indicated that bandwidth would be exceeded at pile-up levels of around $\mu = 87$. Through upgrades to the readout system and implementation of data compression techniques, the SCT now operates comfortably at pile-up levels of $\mu \sim 60$ in Run 3.

The TRT is the outermost part of the ID and used for continuous tracking and electron identification. Its sensitive volume in the barrel region spans radial distances from 563 mm to 1066 mm. Combined with the two end-cap regions it has an acceptance range of $|\eta| < 2.0$ [66]. The TRT's active components are proportional drift tubes, commonly referred to as “straws”, which have a diameter of 4 mm and are filled with a xenon- or argon-based gas mixture. Each straw's outer walls are at a potential of -1530 V with respect to a grounded center wire. The TRT straw layout ensures that charged particles with transverse momentum $p_{\text{T}} > 0.5\text{ GeV}$ and pseudorapidity $|\eta| < 2.0$ traverse an average of more than 30 straws [67]. As particles pass through the straws, they ionize the gas mixture, generating electrons that drift towards the central wire. The electrons cascade close to the wire and create a detectable signal. The TRT achieves a spatial resolution of approximately 110 μm .

Electron identification in the TRT relies on detecting transition radiation produced by high-energy particles. Radiation fibers and foils between the straws cause traversing particles to create transition radiation. The energy of the emitted photons is proportional to the Lorentz factor $\gamma = E/mc^2$ of the particle. They are soon absorbed by the gas in the straws, generating ionization signals. By measuring the energy deposited in this process, the TRT enables particle identification, effectively distinguishing electrons (which have a high γ) from pions.

TRT occupancy quantifies the probability of a straw signal exceeding a threshold in the 75 ns readout window, reflecting particle density and pile-up. Already at nominal luminosity, occupancy reaches 60 %. High occupancies degrade tracking via incorrect hit assignments, reduced efficiencies, and increased fake tracks. Firmware upgrades before Run 3 enable operation at 76–86 % occupancy while maintaining the 100 kHz level 1 trigger (L1) trigger rate.

Following detector alignment, the impact parameter resolutions for secondary vertex reconstruction of high-momentum tracks in the ID were determined to be $(22.1 \pm 0.9) \mu\text{m}$ and $(112 \pm 4) \mu\text{m}$ in the transverse and longitudinal directions, respectively. In this asymptotic limit of high p_{T} , the relative momentum resolution was measured to be $\sigma_p/p = (4.83 \pm 0.16) \times 10^{-4} \text{ GeV}^{-1} \times p_{\text{T}}$ [67].

3.3.2 The Calorimeters — measuring energy

The calorimeters in ATLAS measure the particles’ energy after they have passed through the ID. This process is destructive, as the particles interact inelastically with the calorimeter material, distributing their energy among many secondary particles and creating a particle “shower” until all particles are absorbed. During this process, their energy is converted into a detectable signal. The size of the calorimeters is such that almost all particles are completely stopped. Only muons, which interact minimally with the calorimeter material and are later detected in the muon chambers, and neutrinos, which remain undetected, pass through the calorimeters without being absorbed.

The ATLAS calorimeter system is illustrated in Fig. 3.5. It consists of Electromagnetic Calorimeters (ECals) and Hadronic Calorimeters (HCals) where the particles are slowed down and generate a particle shower primarily through the electromagnetic or the strong interaction, respectively. In both systems, ATLAS uses sampling calorimeters. These have layers of high density passive material to slow particles down, alternating with layers of active material for signal generation and energy measurement. The calorimeters are segmented to provide shape information of particle showers. In the η region where the ID provides tracking, highly-granular information is used for the precise identification and reconstruction of electrons, photons and τ leptons. The coarser granularity of the remaining calorimeter regions meets the physics requirements for jet reconstruction and missing transverse energy ($E_{\text{T}}^{\text{miss}}$) measurements. Additionally, some lower-granular information is fed into the trigger system (see Section 3.3.4).

The ECal [68] is a Liquid Argon (LAr) Calorimeter designed to measure the energy of electrons, positrons and photons. Lead serves as the passive material to initiate electromagnetic showers, while liquid argon at a temperature of -184°C acts as the




Figure 3.5: Cut-away illustration of the ATLAS calorimeter system. The EMB has an inner diameter of 2.8 m and an outer diameter of 4 m. The hadronic barrel has an inner diameter of 4.56 m and an outer diameter of 8.5 m. (Reproduced from Ref. [52], licensed under CC BY 4.0)

active material. The detector layers feature a characteristic accordion-like structure with alternating electrodes surrounded by liquid argon and lead absorber plates. This design ensures complete ϕ symmetry and a crack-free coverage. A high voltage is applied between the electrodes to collect ionization charges generated when particles interact with the liquid argon.

The ECal is divided into a barrel (EMB) and two end-caps (EMEC), covering a pseudorapidity range of $|\eta| < 3.2$. It achieves an energy resolution of $\sigma_E/E = 10\%/\sqrt{E(\text{GeV})} \oplus 0.7\%$ for electromagnetic showers. With a thickness of more than 22 radiation lengths² (X_0), electromagnetic showers are well contained within the ECal volume. The total material budget, expressed in units of nuclear interaction length³ (λ_I), is less than 2 for all pseudorapidities. Consequently, particles interacting predominantly via the strong force tend to initiate showers in the ECal that are not fully contained and instead extend into the HCal.

The HCals [69], which surround the ECal, consist of the Tile Calorimeter, the LAr hadronic end-cap calorimeter and the LAr forward calorimeter. All are sampling calorimeters employing distinct combinations of active and passive materials. Their function is to

²The radiation length X_0 is a characteristic of a material and defined as the distance after which the energy of an electron traversing this material is reduced by the factor $1/e$ due to bremsstrahlung.

³The nuclear interaction length λ_I of a material is the mean distance traveled by a hadronic particle before undergoing an inelastic nuclear interaction with the material.

measure the energy of hadronic showers, such as those initiated by protons, neutrons, charged pions or kaons. The Tile Calorimeter [70] is divided into a barrel and two end-cap regions. Together they cover a pseudorapidity range of $|\eta| < 1.7$. The calorimeter uses steel as the absorber material and scintillating plastic tiles as the active material. Hadrons interacting with the steel layers generate hadronic showers which subsequently produce photons in the scintillating tiles that are then transported by wavelength shifting fibers to Photomultiplier Tubes (PMTs), where they are detected. At $|\eta| = 0$, the Tile Calorimeter has a thickness of approximately $7.5 \lambda_I$, which is sufficient to reduce the number of punch-through particles entering the muon system well below the signals of prompt or decay muons.

The LAr end-cap calorimeter covers the pseudorapidity range $1.5 < |\eta| < 3.2$. It uses liquid argon as active material and copper as the passive material. The Tile Calorimeter and the LAr end-cap calorimeter are designed to achieve an energy resolution of $\sigma_E/E = 50\%/\sqrt{E(\text{GeV})} \oplus 3\%$ for hadronic showers.

The LAr forward calorimeter extends over a range of $3.1 < |\eta| < 4.9$ and was designed for an energy resolution of $\sigma_E/E = 100\%/\sqrt{E(\text{GeV})} \oplus 10\%$ for hadronic showers. Each end-cap comprises three modules: the first uses copper as passive material and is optimized for electromagnetic measurements, while the other two modules use tungsten as absorber material and measure predominantly hadronic interactions.

The calorimeter detectors require minimal modifications to operate at luminosities significantly higher than the original design, and are expected to remain operational throughout the LHC and HL-LHC lifetimes. For Run 3, the legacy analog trigger path was replaced with a new digital trigger path, offering finer granularity inputs to the upgraded trigger system.

3.3.3 The Muon Spectrometer — identifying muons

The outermost part of the ATLAS detector is the Muon Spectrometer [71]. It is designed to detect and track muons or any charged particles exiting the calorimeters and to measure their momentum within the pseudorapidity range $|\eta| < 2.7$ [4]. The structure of the muon spectrometer is illustrated in Fig. 3.6.

Muons behave as Minimum Ionizing Particles (MIPs) in energies ranging from $\mathcal{O}(1 \text{ GeV})$ to $\mathcal{O}(1 \text{ TeV})$ [49] and therefore lose minimal energy via ionization in the calorimeters. With a mean lifetime of approximately $2.2 \mu\text{s}$, these relativistic particles survive long enough to traverse the calorimeters and muon spectrometer, exiting the ATLAS detector before decaying. All other particles, except neutrinos (which are undetected), are absorbed by




Figure 3.6: Cut-away illustration of the ATLAS muon system. Both the barrel toroid magnets and the endcap toroid magnets create a magnetic field of 4 T. (Reproduced from Ref. [52], licensed under CC BY 4.0)

the calorimeters. This ensures that signals in the muon spectrometer are highly likely to originate from muons.

To distinguish muons by their momentum, the muon spectrometer contains large superconducting magnets. These generate magnetic fields (in the barrel region by the eight characteristic toroid magnets), that bend the paths of the traversing muons. Precision tracking along the bending coordinate is performed by Monitored Drift Tube (MDT) chambers with the goal to provide a good transverse momentum resolution for high- p_T muons. However, MDTs are unsuitable for fast measurements due to the relatively long drift time of ionization charges in each tube that can reach up to 700 ns, which is rather slow compared to the collision rate of 25 ns. Therefore, for (rapid) triggering on detected particles, the Muon Spectrometer employs Resistive Plate Chambers (RPC) and Thin Gap Chambers (TGC) as fast-responding trigger chambers. These chambers deliver track information within tens of nanoseconds after the passage of the particle.

In the forward region, where particle fluxes are particularly high, the original inner endcap wheels have been fully replaced for Run 3 by the New Small Wheels (NSWs) [72]. The NSWs provide rapid and precise track-segment matching to reject background in the middle wheel caused by relatively low-energy charged particles originating from hadronic showers in the forward shielding. The new detector employs two new chamber technologies: small-strip TGC (sTGC) detectors, and micro-mesh gaseous structure (MICROMEGAS) detectors. Both technologies can be used for triggering and have precision tracking capabilities with

spatial resolutions in the bending direction of the order of 100 μm . Such spatial precision is crucial to maintain the current muon momentum resolution in the high background environment of the upgraded LHC. The endcap muon spectrometer is designed to measure the transverse momentum (p_{T}) of traversing muons with a precision of $\sigma_{p_{\text{T}}}/p_{\text{T}} < 15\%$ for 1 TeV muons across the full pseudorapidity coverage $|\eta| < 2.7$.

3.3.4 Triggering and Data Acquisition — recording selected data

Physics data taking at ATLAS occurs at a bunch-crossing frequency of 40 MHz, corresponding to a 25 ns time interval between crossings, with an average of almost 60 proton-proton interactions per crossing. Storing all physics data continuously is infeasible, as this would require a data transfer rate of ca. 120 TB/s. To address this, ATLAS implements a two-level trigger system designed to identify interesting events for continued processing and to reduce the event rate in the first level to the maximum detector readout rate of 100 kHz and in the second level (high-level) to an event rate of 3 kHz that is saved to disk [73]. A schematic overview of the Trigger and Data Acquisition (TDAQ) system, including associated data rates, is provided in Fig. 3.7.

The first level trigger is called L1 and is mainly based on two independent systems which use custom electronics to trigger on reduced-granularity information from either the calorimeters (Level-1 Calo) or the muon detectors (Level-1 Muon). The final L1 trigger decision is formed by the Central Trigger Processor (CTP). The CTP reduces the peak event rate to about 100 kHz and defines crude regions of interest (RoIs) in the detector that are further passed to the High Level Trigger (HLT). The L1 trigger makes a decision within 2.5 μs , during which the event information is stored on dedicated Front-End (FE) buffers. If accepted by the L1 trigger, the FE buffers are read out and the event information is passed to the Data Acquisition System (DAQ) system. Additionally, information about RoIs is passed to the HLT.

In the HLT, online algorithms reconstruct events at progressively higher levels of detail compared to L1, either across the full detector volume or within specific RoIs. The HLT software is designed to replicate the offline selection as closely as possible, with the Run 3 implementation making use of the ATLAS offline software framework Athena [74]. The HLT operates on a dedicated server farm, further reducing the event rate to 3 kHz for permanent storage. The entire trigger chain is configured to cover all signatures critical to the ATLAS physics program including high- p_{T} electrons, photons, muons, τ -leptons, jets, b -jets and E_{T} which are essential for Standard Model precision measurements. Additionally, a set of low- p_{T} dimuon triggers is used to collect B -meson decays, which are important to the B -physics program of ATLAS [73].

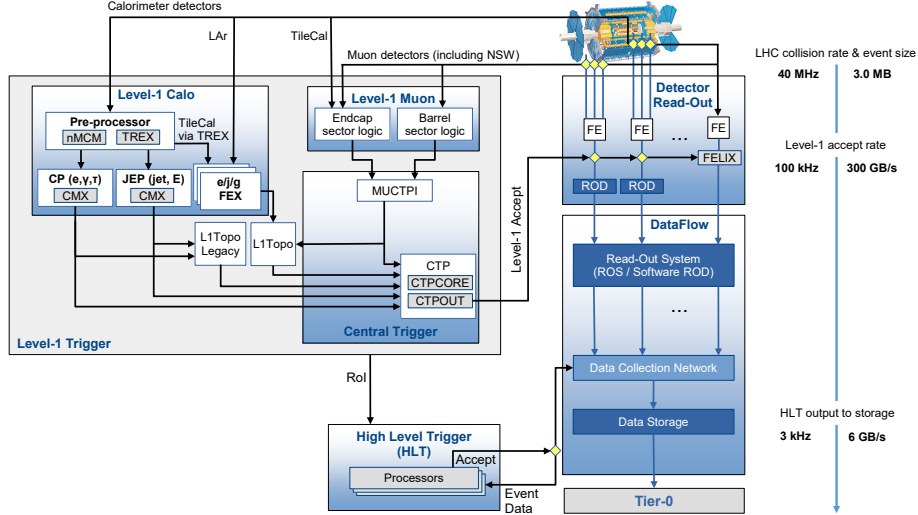


Figure 3.7: Schematic overview of the TDAQ system in Run 3. (Reproduced from Ref. [52], licensed under CC BY 4.0)

As previously described, data acquisition by the detector starts as soon as the FEs receive the Level-1 Accept signal from the L1 trigger. A key innovation in the Run 3 DAQ is the introduction of the FELIX⁴ [75] readout system and software, running on commodity servers (SW ROD). FELIX is integrated into the readout path for those detector systems with new or upgraded front-end electronics in Run 3 such as the NSW. It is foreseen to employ FELIX in all ATLAS detector subsystems for the HL-LHC era, as detailed in Section 3.4.1.

The final selection of events to be stored is made by the HLT. Following HLT processing, accepted events are routed to a dedicated server cluster for packing and compression before being transferred to the CERN Tier-0 computing facility for permanent (offline) storage. To accommodate the increased average rate of physics data, the Run 3 data transfer rate to permanent storage has been doubled, supporting up to 8 GB/s with an expected average rate of 6 GB/s.

3.3.5 Luminosity measurement—comparing experiment to theory

Precise luminosity determination is critical for estimating the expected number of collision events from a given process during data analysis or when performing cross-section measurements, as shown in Eq. (2.7). Multiple redundant luminosity detectors within ATLAS provide luminosity measurement results, with the LUCID-2 detector [76] serving as the primary instrument. LUCID-2 uses 20 PMTs on each side of the ATLAS detector,

⁴Front-End Link eXchange

positioned around the beam pipe at a distance of approximately 17 m from the IP. These PMTs detect Cherenkov light emitted by charged particles traversing thin quartz windows located in front of the PMTs, which act as the Cherenkov medium.

The LUCID-2 system is calibrated using van der Meer scans, a dedicated procedure where low-luminosity beams are swept transversely across each other. These scans establish a reference calibration under well-defined conditions. The results are subsequently extrapolated to the high-luminosity physics data-taking regime using complementary measurements from other luminosity sensitive detectors within ATLAS [77].

3.4 Detector and upgrades overview for the HL-LHC era

The previous chapter presented the detector configuration of the ATLAS experiment for Run 3. Prior to the start of Run 3 in 2022, the detector underwent several upgrades, including the NSW, which were performed during LS2 and are referred to as *Phase-I* upgrades, as published in 2011 [78]. Following the announcement of the HL-LHC project, the ATLAS collaboration initiated planning for further detector upgrades to prepare for the HL-LHC era. In 2012, the ATLAS collaboration first outlined the proposed *Phase-II* upgrades of the ATLAS detector that would enable operation at higher luminosity [79]. A “scoping document” [5] from 2015 describes in more detail an updated plan of the proposed upgrades, which includes also a description of the planned ITk upgrade.

The primary work conducted in context of this thesis focuses on the ITk upgrade. This section gives a brief overview of upgrades other than the ITk, that are planned for LS3, but where especially the TDAQ upgrade affects the design of ITk. Section 3.5 then provides a general overview of the requirements and design for the ITk. A detailed technical description of the ITk detector system will be presented in Chapter 4.

3.4.1 Trigger and Data Acquisition system upgrade

As described in Section 2.3, the high-luminosity phase of the LHC (HL-LHC) will provide an order of magnitude more data starting from 2030, allowing for improved sensitivity to a wide range of new-physics scenarios. To handle the demands of the HL-LHC era, upgraded detectors and trigger systems are being developed and constructed to meet the requirements imposed by an instantaneous luminosity of $\mathcal{L} = 7.5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ and pile-up of up to 200 inelastic collisions per bunch-crossing.

The design of the upgraded TDAQ system has evolved into a baseline architecture consisting of three components: An L0 hardware trigger, the DAQ system (comprising Readout and

Dataflow subsystems) and the software based Event Filter. The functional overview of the TDAQ system as planned for Phase-II is illustrated in [Fig. 3.8](#).

The single-level hardware trigger (L0) features a maximum rate of 1 MHz and 10 μ s latency [80, 81], compared to 100 kHz and 2.5 μ s in Run 3. The goal is to maintain low trigger thresholds and a high data volume in the high pile-up environment. The L0 trigger system comprises two independent subsystems, L0Calo and L0Muon, which are functionally similar to their Phase-I predecessors as described in the previous section. The L0Calo and L0Muon subsystems send their selected objects to the Global Trigger, a new subsystem that performs offline-like algorithms on full-granularity calorimeter data. The CTP forms the final L0 decision and transmits this decision as a L0A signal via FELIX to the detector systems.

From Run 4 onward, all ATLAS detector systems are expected to use FELIX (Front-End Link eXchange) [75] to interface with detector-specific electronics. The baseline design for FELIX uses a commodity server hosting two custom PCIe Interface Cards. Each FELIX card hosts optical transceivers for communication with the Timing, Trigger and Control (TTC) system (distributing clock and L0 trigger signals) and with detector electronics (up to 24 bidirectional optical links). These optical links are for example from ASICs like the lpGBT that aggregates several slow serial copper [e-links](#). The communication path from detector electronics to the FELIX readout system is called *uplink*, the opposite direction is called *downlink*. Additionally, the FELIX card includes a Field Programmable Gate Array (FPGA) to manage the routing of data. The FELIX transceivers and FPGA will support link speeds of up to 25 Gbit/s, while Network Interface Cards (NICs) in the FELIX server will be capable of supporting up to 400 Gbit/s of traffic. FELIX thus acts as a router between custom serial links (different protocols and encodings will be employed, depending on the detector system) and commodity multi-gigabit networks. It is largely detector-agnostic, encapsulating common functionality, though minimal detector-specific processing is required in FELIX to decode or route data beyond what is required to determine its destination. The architectural changes introduced by FELIX offer a key advantage: the use of industry standard Commercially available Off-The-Shelf (COTS) components early in the DAQ chain.

Upon receiving an L0A signal by the FEs, a full detector readout is triggered. Detector data is then sent from the FEs to the FELIX subsystem. Along the network, Data Handlers running on COTS servers, receive the data and perform detector-specific processing. Monitoring or control will be implemented before buffering the data in the Dataflow subsystem. The Readout subsystem, comprising FELIX and the Data Handlers, is designed to handle a 1 MHz event rate with an (uncompressed) event size of 5 MB, resulting in a total bandwidth of 5.2 TB/s. A large Event Filter processor farm of commodity servers

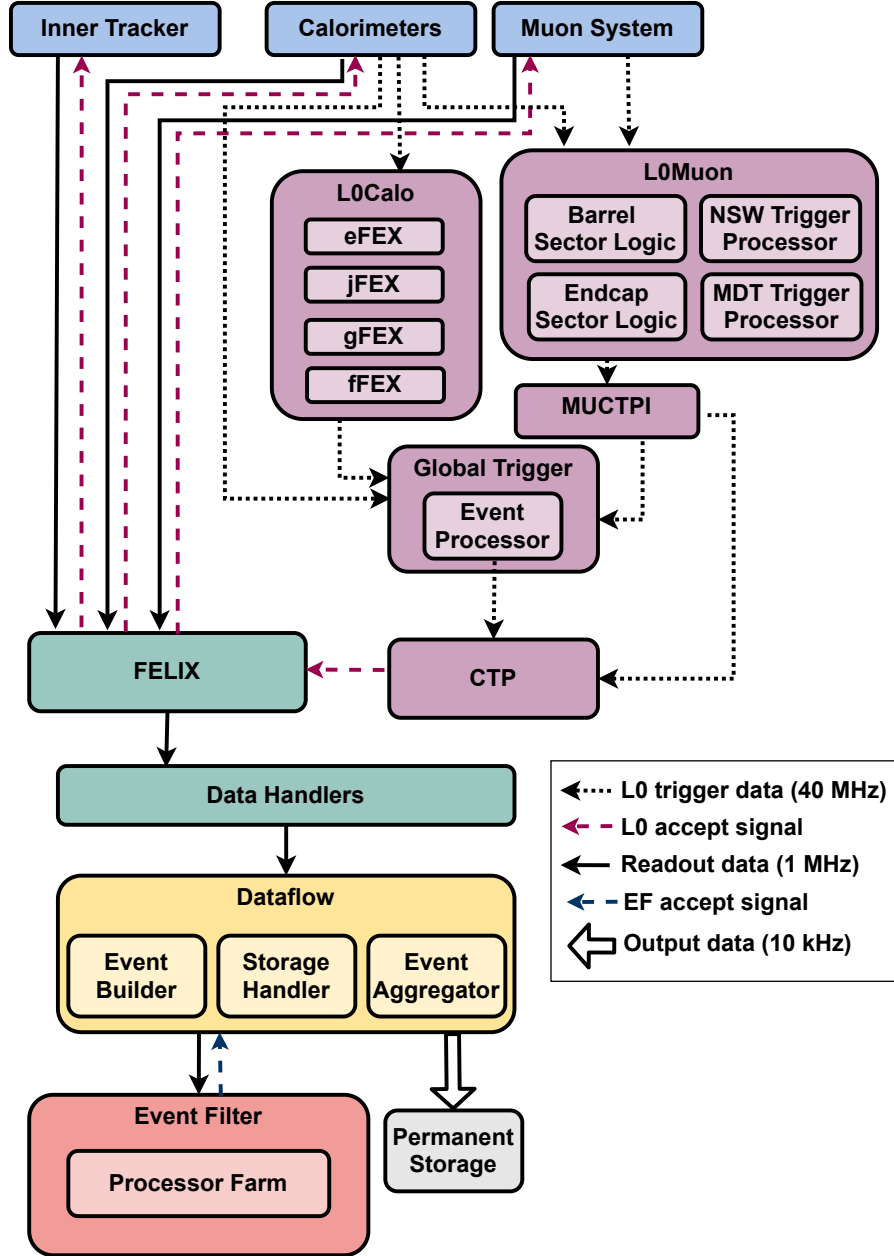


Figure 3.8: Schematic overview of the TDAQ system for Phase-II. In purple is the L0 trigger, in green the readout subsystem (together with the Dataflow subsystem part of the DAQ) and in orange the software Event Filter. The black dotted arrows indicate the L0 data flow from the detector systems to the L0 trigger system at 40 MHz, which must identify physics objects and calculate event-level physics quantities within 10 μ s. The result of the L0 trigger decision (L0A) is transmitted to the detectors as indicated by the red dashed arrows. The resulting trigger data and detector data are transmitted through the DAQ system at 1 MHz, as shown by the black solid arrows. Events are then selected by the Event Filter at a rate of 10 kHz and transferred from the buffers in the Dataflow system to permanent storage at an expected rate of 60 GB/s. (Reproduced from Ref. [80])

receives event data from the Dataflow system and provides the accept/reject decision and selected reconstruction data for offline storage based on a subset of offline-like reconstruction algorithms. The Event Filter reduces the event rate that is saved to offline storage to 10 kHz, corresponding to a bandwidth of 60 GB/s.

3.4.2 Upgrades in other areas

The upgrade of the trigger system is one reason why many readout electronics must also be upgraded. For the LAr Calorimeter, the existing readout electronics will be completely replaced due to both their incompatibility with the new trigger system and their limited radiation tolerance concerning especially on-detector components [82]. A new, more flexible readout architecture enables full calorimeter granularity and longitudinal shower information to be available to the lowest level trigger processors at the required 40 MHz rate, while data transmission and processing speed is kept below a latency of 1.7 μ s. The Tile Calorimeter will also lower the latency from 2.5 μ s to 1.7 μ s [83] using new readout electronics. Additional improvements include replacing current mechanical structures with shorter ones to allow easy access in all detector opening scenarios, making complicated opening procedures obsolete.

In addition to the NSW upgrade, the Muon spectrometer will also undergo Phase-II upgrades, including the replacement of a large fraction of the on- and off-detector readout and trigger electronics [84]. This is to ensure compatibility with the higher trigger rates and longer latencies required by the new L0 trigger. The Muon spectrometer will also receive new RPC chambers in the central barrel layer and will replace some of the MDT chambers and TGC chambers, both to increase the acceptance and robustness of the trigger and to reduce the muon fake rate.

The High-Granularity Timing Detector (HGTD) [85] adds a new subdetector to ATLAS to mitigate pile-up effects. Based on Low Gain Avalanche Detector (LGAD) technology [86] with 50 μ m thick sensors, the HGTD provides high-precision timing information to distinguish collisions close in space but well-separated in time, thereby helping to associate tracks to vertices. It enhances physics performance in the forward region by covering pseudorapidities between 2.4 and 4.0, complementing the capabilities of the ITk. The target average time resolution per track for a MIP is 30 ps initially, degrading to 50 ps at the end of HL-LHC operation. This can be compared to the beam-spot at the ATLAS interaction point (IP1), which has a spread in time of up to 200 ps.

Several upgrades aim to improve luminosity measurements. The HGTD contributes by measuring luminosity at the highest radial locations. A new radiation-hard beam monitor (BCM') based on polycrystalline Chemical Vapor Deposition diamonds, will monitor the

delivered luminosity and protect the inner detectors of ATLAS. LUCID-3 [87, 88] is a new high precision luminosity detector designed to achieve 1 % accuracy at the HL-LHC, which is critical for the ATLAS physics program, but could not be provided by the current LUCID-2 which would saturate under high pile-up. LUCID-3 will be installed about 30 cm from the beam pipe (compared to 12 cm for LUCID-2) and about 16 meters from the IP on the forward shielding pieces.

3.5 Requirements for the ITk upgrade and general overview

The fivefold increase in instantaneous luminosity at the HL-LHC will present numerous challenges, particularly for the ATLAS detectors close to the interaction region. With an average of 200 proton-proton interactions per bunch-crossing, both radiation levels and track densities are expected to be very high. To cope with these harsh conditions, detectors must be radiation-hard and need to have fine granularity. This is essential to keep channel occupancy, and consequently readout bandwidth, low, while also enabling precise reconstruction of primary vertices for effective pile-up rejection. The current ID subsystem (see [Section 3.3.1](#)) is already experiencing performance degradation due to radiation damage [89, 90]. This not only limits its current capabilities, but would make it unsuitable for operation at the HL-LHC. For these reasons, it was decided to replace the current ID with an all-silicon Inner Tracker (ITk) [5]. The ITk will consist of a Pixel Detector [91] utilizing hybrid pixel modules at a small radius around the beam pipe and a large-area Strip Detector [92] surrounding it. A key advantage of a large Pixel Detector close to the beam pipe is its low occupancy, because each pixel cell has only a small sensitive area. The TRT will be removed since it is not suited for a high luminosity environment.

The HL-LHC schedule (see [Section 2.3](#)) foresees the installation of all Phase-II upgrades, including the ITk, during LS3 from 2026 until 2030. Once installed, the ITk will operate in the 2 T magnetic field generated by the same solenoid used for the ID. The ITk design balances the tracking performance required for the Phase-II physics program and the associated construction costs.

The new Pixel Detector, which is the focus of this thesis, requires radiation-hard sensors and FE electronics. Due to the increased luminosity, the expected radiation levels in the ITk will be about ten times higher than those in the current ID. [Figure 3.9](#) shows radiation background in one quadrant of the Pixel detector as simulated for the ITk. The displayed neutron equivalent fluence⁵ can be used to estimate displacement damage. For more details on radiation effects in semiconductor detectors, refer to [Section 4.2.2](#). Given

⁵A fluence of 1 MeV neutrons that would cause the same amount of displacement damage in silicon as the actual mixed particle spectrum.




Figure 3.9: Fluence distribution for the Pixel Detector from simulations assuming an inelastic proton-proton cross-section of 79.3 mb. The 1 MeV neutron equivalent fluence is normalized to 4000 fb^{-1} . This graph is without safety factors taken into account. (Reproduced from Ref. [91], licensed under CC BY 4.0)

the high radiation levels in the innermost layers of the Pixel Detector, and the expected large radiation damage in this region, it is planned to replace the first two layers after an integrated luminosity of 2000 fb^{-1} . These simulations, combined with a safety factor of 1.5, establish a requirement on the radiation tolerance for the FEs and sensors in the innermost layer of $1.3 \times 10^{16} \text{ n}_{\text{eq}}/\text{cm}^2$ corresponding to a dose⁶ of 9.9 MGy.

Another requirement for the FEs electronics comes from the upgraded trigger system. With the Phase-II upgrade of the TDAQ system, the trigger rate will increase from previously 100 kHz to 1 MHz with a maximum latency of 10 μs . This performance has to be supported by the detector electronics. For the Pixel Detector, the requirement on the pixel size is to be small enough to keep occupancy below 1 %. By applying compression techniques for hit maps transmitted from the FEs to the FELIX system, and taking typical occupancies into account, this translates into an output bandwidth of up to 5.12 GB/s per Front-End chip.

To guarantee good tracking performance until the end-of-lifetime at 2000 fb^{-1} for the two innermost layers, and 4000 fb^{-1} for the other layers, sensors must achieve a high hit

⁶The *gray* ($1 \text{ Gy} = 1 \text{ J kg}^{-1}$) is the unit of ionizing radiation dose that measures the energy deposited by ionizing radiation in a unit mass of absorbing material

efficiency of greater than 97 %, where a random single-pixel inefficiency of 3 % is considered as part of aging. Regarding the track reconstruction efficiency, the new ITk should deliver a performance that is at least as good as the existing detector, but in the harsher environment of the HL-LHC. Specifically, for pseudorapidities up to $|\eta| = 2.7$ it is desired to have a track reconstruction efficiency greater than 99 % for muons with transverse momenta p_T above 3 GeV and an efficiency greater than 85 % for pions and electrons above 1 GeV [91]. These requirements translate directly into requirements for hit information needed for track reconstruction. To maintain a good track reconstruction efficiency, at least 9 (8, 7) hits in the ITk are required in regions $|\eta| \leq 2.0$ ($2.0 \leq |\eta| \leq 2.6$, $2.6 \leq |\eta| \leq 4.0$) with at least one hit in the Pixel subsystem and at most 2 holes [93]. This means the performance is robust against losses of up to 15 % of individual channels or modules that could conceivably occur over the time of HL-LHC operation.

The current layout and expected tracking performance is documented in Ref. [93]. A graphical representation of the layout is provided in Fig. 3.10. The ITk is an all-silicon detector, 6 m long with a diameter of approximately 2 m, occupying the entire available space within the solenoid magnet. Based on the previously mentioned requirements, the ITk was designed as a Pixel Detector with hybrid pixel modules in the inner layers and a Strip Detector with stereo strip modules in the outer layers. To maximize the physics potential, the tracking coverage in comparison to the ID is extended from $|\eta| < 2.5$ to $|\eta| < 4.0$. The Pixel subsystem covers a pseudorapidity range of $|\eta| < 4.0$ and consists of five flat barrel layers and five layers of inclined or vertical rings for coverage in the forward region. The Strip subsystem spans a pseudorapidity range of $|\eta| < 2.7$ and includes four strip layers in the barrel region and six disks in the endcaps, all utilizing double-sided modules.

A schematic display of the geometry of the active detector elements is given in Fig. 3.11, with Fig. 3.11(b) showing the geometry of one quadrant of the Pixel Detector and its subsystems. The 5 Pixel Detector layers are divided into three subsystems:

- The Inner System (IS), covering the inner-most two layers (layer 0 and 1) with rings and staves at radii 34 mm and 99 mm.
- The OB, covering the central part of layers 2 to 4 with inclined rings and staves at radii 160 mm, 228 mm and 291 mm.
- The (Outer) End Cap (EC) with rings in the forward region of layers 2 to 4.

The Strip and Pixel Detectors are separated by the Pixel Support Tube (PST). The inner two layers of the Pixel Detector are separated from the outer Pixel layers by the Inner Support Tube (IST). The mechanical design, and in particular the role of the IST in the support hierarchy, means that the outer Pixel subsystem is independent of the inner section, allowing for a replacement of the IS after 2000 fb^{-1} .




Figure 3.10: Longitudinal view of the most recent ITk Layout 03-00-00 presented in Ref. [93]. In green is shown the Pixel Detector and in blue the Strip detector. (Reproduced from Ref. [93], licensed under CC BY 4.0)

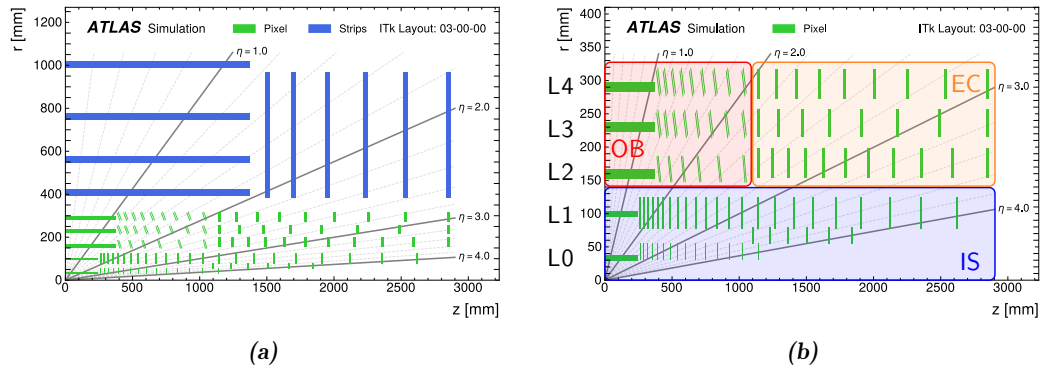


Figure 3.11: (a) A schematic depiction of the ITk layout as presented in Ref. [93] with the Strip detector in blue and the Pixel Detector in green. (b) A zoomed-in view of the Pixel Detector with its five layers L0 to L4. One quadrant and only active detector elements are shown. The horizontal axis is along the beam line with zero being the nominal interaction point. (Adapted from Ref. [93])




Figure 3.12: (a) Integrated radiation length (in units of X_0) and (b) integrated nuclear interaction length (in units of Λ_0 , in previous text λ_I) traversed by a straight track as a function of the absolute pseudorapidity $|\eta|$ at the exit of the ITk volume broken down by subsystem and material category. (Reproduced from Ref. [93], licensed under CC BY 4.0)

Simulations of the material budget indicate that the ITk design remains well below 1.5 radiation lengths (X_0) in the central region $|\eta| < 1.5$, comparable to the Run 3 ID. Significant reductions in the material were achieved in the more forward region [91, 93]. A low material budget, particularly in layers close to the interaction point, is crucial to reduce the negative effect of multiple scattering on the precision of vertex measurements. The low material budget for the current ITk design can be seen from Fig. 3.12 which shows the integrated radiation length and the integrated nuclear interaction length traversed by a straight track through the ITk volume as a function of pseudorapidity. The low material budget is achieved through the use of low-mass support material, the employment of high-performance and low-mass two-phase CO_2 cooling in the entire tracker and reduced cabling due to data link-sharing optimization and serial powering of modules. A more technical description of the Pixel detector is provided in Chapter 4.

With the described design, the performance of the ITk is expected to match or exceed the performance of the current ID [93]. The geometrical layout of the Pixel and Strip Detectors always guarantees at least 9 hits in the ITk per track, as can be seen from simulations. Figure 3.13 shows such a simulation. It displays on the x -axis the pseudorapidity and on the y -axis the number of hits in the ITk of single muons traversing the detector starting in the center.

The expected tracking efficiency, defined as the fraction of charged particles associated with a reconstructed track, is comparable to that of the ATLAS Run 3 detector. For muons with transverse momentum $p_T = 2 \text{ GeV}$ the efficiency is above 99.5 % and compatible with 100 % for larger p_T up to a pseudorapidity of $|\eta| = 3.6$. Outside this region (and until $|\eta| = 4.0$), the efficiency decreases slightly to 99 %, due to the smaller number of available




Figure 3.13: Number of combined potential Strip and Pixel measurements along a particle trajectory as a function of the truth particle pseudorapidity for the ITk. A sample of single-muon events with $p_T = 1$ GeV is used. The muons are produced with a uniform distribution between 0 and 2 mm in transverse distance to the beam line and at fixed values of $z = -15$ cm, 0 cm, and 15 cm, in equal amounts. (Reproduced from Ref. [93], licensed under CC BY 4.0)

measurements in that region (see Fig. 13 in Ref. [93]). The tracking efficiency for electrons and pions is lower due to their higher interaction rate with the detector material (see Fig. 14 in Ref. [93]). However, it is expected to remain above 85 % for all types of prompt and stable charged particles. In the forward region ($|\eta| > 3.0$), the expected tracking efficiency for pions is more significantly reduced than that for electrons, due to the increased nuclear interaction lengths that the tracks traverse.

The transverse momentum (p_T) resolution of the ITk is expected to surpass the resolution achieved with the ID. This is primarily due to the superior resolution in the bending plane provided by the Strip Detector at a large lever-arm for measurements compared to the TRT in the Run 3 detector. The expected resolution on the transverse (d_0) and longitudinal (z_0) impact parameters in simulations show an improvement ranging from 20 % to 300 % compared to the Run 3 detector depending on η and p_T . These improvements are mainly due to the similar radius of the innermost pixel layers and the reduced pixel pitch in the ITk Pixel Detector. High resolutions for these track parameters are critical for optimizing flavor-tagging and pile-up rejection performances.

Table 3.1: Comparison of selected parameters between the ITk and the ID showing the size differences.

Parameter	Pixels		Strips	
	ITk	ID	ITk	ID
Area (in m ²)	~13	1.9	~165	61
No. of modules	~8400	2000	~9700	4088
No. of channels (in M)	~5100	92	~60	6.3

[Table 3.1](#) summarizes key parameters comparing the size of the new ITk to the current tracking detector in Run 3.

Particle Physics Detectors

With the discovery of radioactivity by Henri Becquerel in 1896 using photographic plates, physicists became interested in systematically detecting particles. Advancements in nuclear and particle physics rely on the construction of detectors that can identify particles and radiation, enabling the measurement of their properties. Detectors can only sense particles if these particles somehow interact with the detector material. [Section 4.1](#) provides an overview of the different types of interactions of particles with matter. Since this thesis focuses on the ITk Pixel Detector, [Section 4.2](#) introduces the foundations of semiconductor detectors, explains the working principles of silicon sensors, and describes the sensors types employed by the ITk Pixel Detector. Constructing a fully functional large-scale detector requires integrating numerous components. In hybrid pixel detectors, sensors must be read out by FE electronics. [Section 4.3](#) introduces the RD53A and ITkPix FEs developed for the next pixel detector. [Section 4.4](#) presents the concept of a module, the smallest assembled unit comprising a sensor, FEs and connections. Modules are mounted on mechanical support structures that include a cooling system, support for aggregated cable routing and, in the case of the ITk Pixel detector, independent monitoring. [Section 4.5](#) details these support structures for the ITk Pixel detector, which represent the smallest, fully functional building blocks of the detector.

During a system test at CERN, two prototypes of support structures (loaded with modules) from the OB region were tested extensively to validate the detector design described in this chapter. [Chapter 5](#) describes these OB system tests using the RD53A demonstrator and details the experimental setup for tests that were conducted within this thesis, with the results of these tests presented in [Chapter 6](#).

4.1 Interactions of particles with matter

For particle detection, the particles under study must interact with the detector material. The task of a detector is to detect relatively long-lived or stable particles. These are elementary particles such as electrons, photons, neutrinos or muons, or composite particles like protons, neutrons, charged pions or charged kaons. Single quarks, for example, are not directly detected, because they are confined within hadrons as described in [Section 2.1.1](#).

Particles can be identified based on their mass, charge or the type and strength of their interaction with matter, as already mentioned in [Section 3.2](#). The main ways of particles to interact with matter are:

- Ionization of atoms by heavy charged particles.
- Bremsstrahlung emission by light charged particles in the fields of atomic nuclei.
- Photon scattering, photon absorption and pair production.
- Nuclear reactions of hadrons (charged and neutral) with nuclear matter.

Neutrinos do not interact via the strong or electromagnetic force and can only be detected through the weak interaction. Their presence is often inferred through energy imbalances in the detector.

4.1.1 Interactions of charged particles

Charged particles passing through matter lose kinetic energy by interacting with the outer electrons of an atom, thereby exciting bound electrons or ionizing the absorbing material. For heavy charged particles, which includes all charged particles except electrons and positrons, this behavior is approximately described by the *Bethe-Bloch formula* [[94–96](#)]. The Bethe-Bloch formula as given in Ref. [[49](#)] describes the average energy loss dE per length dX :

$$\left\langle -\frac{dE}{dX} \right\rangle = K z^2 \frac{Z}{A} \frac{1}{\beta^2} \left[\frac{1}{2} \ln \frac{2m_e c^2 \beta^2 \gamma^2 W_{\max}}{I^2} - \beta^2 - \frac{\delta(\beta\gamma)}{2} \right] \quad (4.1)$$

This is the same as saying a heavy, charged particle deposits the energy dE per length dX in the absorbing material when traversing it. In [Eq. \(4.1\)](#),

$\frac{dE}{dX}$ is the energy loss of the projectile particle usually given in $\text{eV}/(\text{g}/\text{cm}^2)$,

$X = \rho x$ is the density weighted length (in g/cm^2),

ρ is the density of the absorber material,

$K = 4\pi N_A r_e^2 m_e c^2 = 0.307075 \text{ MeV mol}^{-1} \text{ cm}^2$,

N_A is the Avogadro constant,

r_e is the classical electron radius,

m_e is the electron mass,

c is the speed of light,

z is the charge of the incident particle in units of the elementary charge,

$\beta = v/c$ the velocity of the incident particle,

Z, A are the atomic number and the mass number of the absorber,

$\gamma = 1/\sqrt{1 - \beta^2}$,

W_{\max} is the maximum possible energy transfer to an electron in a single collision,

I is the mean excitation energy, a characteristic of the absorber material (173.0 eV for silicon¹),

δ are density effect corrections.

Equation (4.1) is valid in the region $0.1 \lesssim \beta\gamma \lesssim 1000$ with an accuracy of a few percent [49]. An example for the average energy loss of muons passing through copper is shown in Fig. 4.1, where the regime for the Bethe-Bloch formula is indicated. At low energies the $1/\beta^2$ term is dominant, while at high energies the $\ln \gamma$ term is dominant. A general feature of this formula is the minimum at around $\beta\gamma \approx 3.5$ for the minimum deposited energy in the material. Particles with these velocities are called Minimum Ionizing Particles (MIPs). Every detector must be able to keep its noise well below this minimum energy to be able to detect MIPs.

The particles lose energy through ionization continuously, resulting in an energy dependent but fixed absorption range in the material. Usually, the energy lost by a charged particle traversing an absorber is deposited in the absorbing material. In rare cases, however, so-called δ -rays or δ -electrons are created. These are high-energy electrons that may escape the detector material. Consequently, the energy loss probability distribution for a projectile particle follows a Landau distribution [97, 98] with a long tail toward higher energies. While the Bethe-Bloch formula describes the *average* energy loss of a particle traversing matter, the *most probable* energy loss is lower.

In silicon detectors, this energy loss causes ionization. Since the energy deposited per collision typically exceeds the 3.7 eV required to produce an electron-hole pair in silicon [99], multiple electron-hole pairs are generated per interaction of an incident particle in silicon. The high density ($\rho = 2.3 \text{ g/cm}^3$) of silicon causes an average energy loss of about $\langle \frac{dE}{dX} \rangle = 390 \text{ eV}/\mu\text{m}$ for a MIP, yielding $106 \frac{\text{e-h}}{\mu\text{m}}$ (electron-hole pairs per μm) [99]. As the semiconductor detectors discussed later lack internal signal amplification in their sensors, these are the signals that the detector must be capable of processing.

¹https://pdg.lbl.gov/2024/AtomicNuclearProperties/HTML/silicon_Si.html




Figure 4.1: Average mass stopping power (dE/dX) for positive muons in copper as a function of $\beta\gamma = p/m$ over nine orders of magnitude in momentum. The region described by the Bethe-Bloch formula from Eq. (4.1) is indicated. At low energies the $1/\beta^2$ term is dominant, while at high energies the $\ln \gamma$ term is dominant. A minimum is at around $\beta\gamma \approx 3.5$. Particles with these velocities are called MIPs. (Reproduced from Ref. [49], licensed under CC BY 4.0)

The previous discussion focused on *heavy* charged particles, like the muon. For electrons and positrons, ionization energy loss at lower energies arises from interactions with atomic shells via Møller scattering (for electrons) or Bhabha scattering (for positrons). Additionally, the possibility for positrons to undergo complete annihilation with the electron of the absorber material producing photons must be accounted for.

Due to their significantly smaller mass, electrons and positrons also experience substantial energy loss when traversing matter through photon radiation (*bremsstrahlung*), even at relatively low energies. This energy loss does not originate from interactions with atomic shells (as previously described), but from the emission of photons in the Coulomb field of the nuclei [100, 101]. For electrons in lead, bremsstrahlung losses exceed ionization losses already above 7 MeV, in air this threshold changes to 100 MeV [102]. Notably, at higher energies, the average energy loss $\langle dE/dX \rangle$ due to bremsstrahlung scales linearly with the electron/positron energy. This behavior defines the *radiation length* X_0 , the characteristic

length for energy loss through bremsstrahlung by:

$$\left\langle \frac{dE}{dX} \right\rangle = -\frac{E}{X_0}. \quad (4.2)$$

Integration of Eq. (4.2) yields:

$$\langle E(X) \rangle = E_0 e^{-\frac{X}{X_0}}, \quad (4.3)$$

where E_0 is the initial energy. This means that after a path length $X = X_0$, an electron possesses only $1/e$ of its initial energy. X_0 is a critical parameter for quantifying detector stopping material and was previously used in describing the calorimeters of ATLAS and the requirements for the ITk (see Fig. 3.12(a) for a simulation of ITk). In general, for silicon, the value² is $X_0 = 21.82 \text{ g/cm}^{-2}$. With a silicon density of 2.329 g/cm^{-3} , this corresponds to 9.370 cm. Consequently, electrons easily pass through thin silicon sensors with thicknesses much smaller than X_0 .

4.1.2 Multiple coulomb scattering

When charged particles traverse material, they undergo Coulomb scattering in the electric fields of the nuclei according to the Rutherford cross-section. In thick materials, multiple scattering events produce a statistical distribution of scattering angles relative to the original trajectory. If the number of scatters exceeds about 20, this is termed multiple or Molière scattering and the distribution of the scattering angle projected onto a plane perpendicular to the direction of motion of the incoming particle is well approximated by a Gaussian distribution [103].

The root-mean-square (RMS) of the projected scattering angle, equivalent to the standard deviation of the Gaussian distribution, is given by [49, 104]:

$$\theta_{\text{plane}}^{\text{RMS}} = \frac{13.6 \text{ MeV}/c}{\beta p} z \sqrt{\frac{X}{X_0}} \left(1 + 0.038 \ln \left(\frac{X}{X_0} \right) \right). \quad (4.4)$$

Here, $\beta = v/c$, p and z are the speed, momentum, and charge number of the scattering particle, respectively, and X/X_0 is the fractional radiation length (with X_0 the radiation length). The fractional radiation length can also be replaced by the material budget defined as $\sum_i X_i/X_{0,i}$ for layered materials.

Multiple Coulomb scattering degrades secondary vertex reconstruction precision, especially scattering in the innermost detector layer since it dominates the extrapolation error to the

²https://pdg.lbl.gov/2024/AtomicNuclearProperties/HTML/silicon_Si.html

secondary vertex. Since pixel detectors are the closest to the interaction point, minimizing their material budget is critical.

Knowledge of the material budget is also a key ingredient to simulations that will later be compared to measurements in physics analyses. Equation (4.4) can be rearranged to measure material budget via multiple scattering. The RMS scattering angle $\theta_{\text{plane}}^{\text{RMS}}$ can be extracted from scattering angle distribution fits. Such a measurement was performed for an ITk Pixel module at the CERN PS, where the average X/X_0 across a module was found to be approximately 0.89 %, which agrees within uncertainties with an estimate of 0.88 % that was derived from material component expectations for a 150 μm thick planar silicon sensor and a module PCB [105].

4.1.3 Interactions of photons

For photon detection, a photon needs to interact with the detector material to create an electron (or positron) and generate a detectable signal. High-energy photons primarily interact with matter via three processes:

1. Photoelectric effect,
2. Compton scattering,
3. Pair production,

with pair production dominating at collider experiments.

All three processes are discussed below. Figure 4.2 illustrates the cross-sections for these interactions in carbon and lead. At low energies, photons may also undergo coherent scattering off a whole atom without excitation or ionization. This process is called Rayleigh scattering. But since almost no energy transfer to the atom occurs, this process is irrelevant for photon detection.

Photoelectric effect

At energies below ~ 100 keV, the photoelectric effect is the dominant interaction process of photons with matter. In the photoelectric effect the photon transfers all its energy to a bound atomic electron. If this energy exceeds the electron's binding energy, the electron is ejected from the atomic shell. The cross-section for the photoelectric effect increases with increasing atomic number Z (at fixed photon energy), and decreases with increasing photon energy. Silicon also has the ability to absorb photons by the photoelectric effect, but only at photon energies below 20 keV for a sensor of 300 μm thickness [102]. The CMS




Figure 4.2: Photon total cross-sections as a function of energy in carbon and lead. The photoelectric effect is dominant in the low keV energy regime, Compton scattering is dominant for medium energies up to 1 MeV and pair production (in the Coulomb field of the nucleus, κ_{nuc} , or an electron, κ_e) for higher energies above 1 MeV. Depending on the atomic number Z of the material, the number of steps in the cross-section for the photoelectric effect differs. (Reproduced from Ref. [49], licensed under CC BY 4.0)

experiment, for example, uses infrared laser beams with a wavelength of $\lambda = 1075 \text{ nm}$ (corresponding to 1.2 eV) for alignment of their tracker [38].

Compton scattering

The Compton effect is dominant in an extended energy region around 1 MeV . It describes the process of a photon scattering on a free or quasi-free electron from an outer valence shell of an atom. Through the scattering process, the photon loses energy and changes direction. The energy lost depends on the scattering angle and is transferred to the electron that is kicked out of the atom, carrying away the recoil momentum. The maximum energy transfer can be achieved when the photon scatters completely backward. This case leads to the so-called Compton edge in the energy spectrum of the recoil electron. For detectors, it is important that the kinetic energy of the recoil electron can be detected.

Pair production

If the energy of the incoming photon is at least 1.022 MeV (twice the mass of electrons³), pair production in the Coulomb field of a charge (usually the Coulomb field of a nucleus) can occur. The photon converts into an electron-positron pair. Momentum conservation requires the nucleus or another electron to take the recoil. This process dominates at high energies. The process is closely related to the process of bremsstrahlung as described earlier and is therefore treated together [100]. The outbound positrons and electrons then radiate off photons again due to bremsstrahlung as described in Section 4.1.1. The photon absorption length (or mean free path) due to bremsstrahlung is [49]

$$\lambda_\gamma = \frac{9}{7} X_0 , \quad (4.5)$$

such that the intensity of a photon beam in material follows an exponential law

$$N(X) = N_0 e^{-\frac{X}{\lambda_\gamma}} , \quad (4.6)$$

where N_0 is the initial photon count and $1/\lambda_\gamma$ is the mass attenuation coefficient.

Implications for detectors

When either a photon or an electron enters an electromagnetic calorimeter, it starts an electron-photon shower, where pair-production and bremsstrahlung are the two alternating

³See Fig. 2.1.

processes in the shower development. In each step, the particles lose energy until a critical energy below which electrons and photons only lose their energy through ionization. For tracking detectors, pair production is problematic: If the photon gets absorbed, it cannot be detected in the calorimeter. Additionally, secondary electron-positron pairs degrade track reconstruction. Minimizing tracker material budget to reduce pair production is thus essential as can be seen in Eq. (4.5).

4.1.4 Strong interactions of hadrons

In interactions of hadrons (p , n , π , K , ...) with matter, in addition to ionization processes for charged particles as described in Section 4.1.1, nuclear interactions play an important role. A multitude of possible interactions can take place. Inelastic processes, where secondary, strongly interacting particles are produced, can be described with a cross-section σ_{inel} . This cross-section is responsible for inducing a hadron shower in a material. The nuclear interaction length λ_{I} (mean path length for inelastic interactions) is defined as:

$$\lambda_{\text{I}} = \frac{\mathcal{A}}{N_{\text{A}} \rho \sigma_{\text{inel}}} , \quad (4.7)$$

where \mathcal{A} is the atomic mass per mole of the material (in g/mol), N_{A} is the Avogadro constant (in mol^{-1}) and ρ the density of the material (in g/cm^3). The nuclear interaction length λ_{I} has units of cm, though the literature often reports $\rho\lambda_{\text{I}}$ in g/cm^2 .

Material with small λ_{I} (high interaction probability) in the tracking system can disrupt low-energy charged pion measurements, as they may undergo inelastic scattering before reaching the calorimeter. Simulations of the ITk indicate a detector depth of 0.2 to $0.6 \lambda_{\text{I}}$ (see Fig. 3.12(b)). In the hadronic calorimeter itself, the material budget (measured in units of λ_{I}) should be large enough to contain the whole particle shower. The ATLAS hadronic calorimeters are approximately 7.5 to $15 \lambda_{\text{I}}$ deep [4]. Crucially, for high- Z materials the interaction length $\rho\lambda_{\text{I}}$ is generally much larger than the radiation length X_0 (for example $\rho\lambda_{\text{I}} \approx 10 X_0$ for iron⁴), making hadronic calorimeters significantly larger than electromagnetic calorimeters.

4.2 Principles of semiconductor pixel detectors

As detectors with high spatial resolution, semiconductor detectors have been used in experiments in high-energy physics since the 1980s [106]. The first semiconductor detectors were strip detectors, while pixel detectors with high achievable granularity were

⁴https://pdg.lbl.gov/2024/AtomicNuclearProperties/HTML/iron_Fe.html

first implemented in experiments during the 1990s. All experiments at the LHC now employ large-scale silicon-based tracking detectors with pixel and strip geometries. When constructed from 2005 onward, those represented the state of the art at that time [107]. In comparison to gaseous detectors, semiconductors have a higher density, resulting in larger energy deposition in the material through ionization and thus more charge is available for signal generation. Additionally, the energy required to create free charges is significantly lower in semiconductor detectors compared to gaseous ones, enabling a much higher energy resolution in semiconductor detectors. Through micro-structuring, semiconductor detectors achieve high granularity and excellent spatial resolution, which are critical properties in a detector to mitigate challenges in a high pile-up scenario as previously described.

With the ITk upgrade, the ATLAS experiment will deploy a silicon detector that is nearly three times larger in sensitive area than the current pixel and strip system. This section first gives a general introduction to semiconductor detectors, then details the sensor choices for ITk Pixel modules and briefly presents the solution for the FASER Preshower upgrade, a focus of this thesis discussed further in [Chapter 10](#).

4.2.1 The pn-junction

Charged particles traversing semiconductor material ionize atoms, creating weakly bound electron-hole pairs. An electric field separates these pairs, inducing a signal current as charges drift toward electrodes. In silicon sensors, this occurs within the depletion region of a reverse-biased pn-junction, which is the fundamental building block of semiconductor detectors. The pn-junction builds up an electric field that collects signal charge while suppressing leakage current, an important source of noise. A pixel sensor is a reversely biased pn-diode with a highly segmented cathode or anode.

The behavior of the pn-junction can be explained using band theory. Electrons bound in the electric field of a nucleus can only take on certain discrete values of energy, called energy levels, where the Pauli principle limits each energy level to two electrons of opposite spin. Materials like silicon crystallize into a periodic lattice structure, causing atomic orbitals to overlap and energy levels to split into many more, where some levels lie so close together they form quasi-continuous energy bands separated from each other by a band gap with energy E_G . The closest band above the band gap is called the conduction band, and the closest band beneath the band gap is called the valence band. The band structure of a material determines its electrical properties:

- Insulators have a fully filled valence band with tightly bound electrons and a large band gap to the next higher, empty band, such that no conduction is possible.

- Conductors have either a partially filled conduction band or an overlapping valence and conduction band leading to some quasi-free electrons with high mobility.
- Semiconductors have a smaller band gap (e.g. $E_G = 1.12$ eV for silicon), such that it can be overcome by thermal excitations or by external electric fields to enable conduction.

Charged particles traversing a silicon sensor lose energy through ionization which can be understood as the excitation of valence-band electrons into the conduction band, creating electron-hole pairs. The energy needed to produce an electron-hole pair was empirically found to be roughly three times the band gap energy⁵ [108], which is only 3.7 eV for silicon, as mentioned earlier. This is small enough to also create electron-hole pairs through thermal excitation. The charge carriers can recombine easily, unless an electric field acts on them, forcing them to drift apart.

Silicon (Si) dominates tracking detectors due to its small band gap, short radiation length and industrial-scale availability of high-purity crystals which significantly reduces production costs. Alternatives like silicon carbide (SiC, $E_G = 3.26$ eV), gallium nitride (GaN) and amorphous SiC (a-SiC) were under review already in 2005 as possible radiation-hard alternatives for an application at an upgraded LHC by the RD50 collaboration [109]. Silicon carbide's larger band gap suppresses thermal noise compared to silicon. Its production has matured in recent years, allowing to produce high-quality SiC substrates, making it interesting for application in particle detectors again [110]. For other materials used in semiconductor detectors with application for example in radiation detection, see Ref. [111, Chapter 2.7] or Ref. [102, Chapter 8.13].

Undoped silicon has an intrinsic charge carrier concentration of $n_i = 9.65 \times 10^9 \text{ cm}^{-3}$ at 300 K due to thermal excitations [112]. The charge carriers are electrons in the conduction band and holes in the valence band. In the undoped silicon case, the concentration of electrons and holes are the same ($n_i = n_e = n_h$). For a typical sensor of size $40 \times 40 \times 0.15 \text{ mm}^3$ this means there are approximately 2×10^9 free charge carriers. A MIP passing perpendicular through such a sensor, will create on average a signal of $106 \frac{\text{e-h}}{\mu\text{m}} \cdot 150 \mu\text{m} = 15\,900$ free electron-hole pairs, where the most probable value is even lower since the deposited energy follows a Landau distribution. To resolve signals, free charge carriers must be reduced via cooling or depletion using doped silicon in reverse-biased pn-junctions.

Doping silicon (Group IV in the periodic table) with Group III elements (for example boron) creates p-type semiconductors: acceptor levels just (in the order of 0.05 eV) above the valence band allow electrons from the valence band being lifted into the acceptor level leaving holes behind and generating hole majority carriers. Doping with Group V elements

⁵The rule that the electron-hole pair creation energy is roughly three times the band gap energy is explained in Ref. [108] as coming from energy and momentum conservation, assuming a free-particle approximation for the thermalized final particles.

(for example phosphorus or arsenic) yields n-type semiconductors: donor levels just below the conduction band provide electrons that are usually lifted into the conduction band and act as electron majority carriers. Both types of extrinsic semiconductor remain electrically neutral.

When p- and n-type semiconductors are brought into contact, this is called a pn-junction. A schematic representation of the pn-junction can be seen in Fig. 4.3. Electrons from the n-type semiconductor diffuse into the p-type area, and holes from the p-type semiconductor diffuse into the n-type area. At the boundary, recombination of both carrier types occurs leading to a *depletion zone* which is free of charge carriers but rich in locally fixed ionized donors and acceptors. The depletion zone is no longer neutral but features a space charge. This space charge establishes an electric field, corresponding to a varying electrostatic potential across the depletion zone with a potential difference between the two ends of the depletion zone leading to a built-in voltage V_{bi} which is typically around 0.5 V. Electron-hole pairs that would be created in the depletion zone are forced apart by the electric field and electrons and holes drift in opposite directions out of the region of the electric field. The depletion width is [111]

$$d = \sqrt{\frac{2\varepsilon_{Si}\varepsilon_0}{e} V_{bi} \left(\frac{1}{N_A} + \frac{1}{N_D} \right)}, \quad (4.8)$$

where $\varepsilon_{Si} = 11.9$ is the dielectric constant of silicon [49], ε_0 is the vacuum permittivity, e the charge of the electron and N_A and N_D the doping concentrations with acceptor atoms and donor atoms, respectively. Typical doping concentrations for silicon sensors are $N_A = 10^{19} \text{ cm}^{-3}$ (highly-doped p-type) and $N_D = 10^{12} \text{ cm}^{-3}$ (low-doped n-type) leading to a depletion zone width of $d \approx 20 \mu\text{m}$.

The pn-junction as described above has the properties of a diode. An external voltage V_{ext} (also called *bias* voltage) applied to the pn-junction, where the positive voltage is applied to the p-type area (forward bias, $V_{ext} > 0$), causes the built-in voltage V_{bi} to reduce by V_{ext} and the depletion zone to become thinner. Once the depletion area shortens, a large current I_{ext} can flow over the pn-junction. At *reverse bias*, when the positive external voltage is applied to the n-type region ($V_{ext} < 0$), such that the external voltage adds to the built-in voltage, the depletion zone gets wider. The width of the depletion zone is then

$$d = \sqrt{\frac{2\varepsilon_{Si}\varepsilon_0}{e} (V_{bi} - V_{ext}) \left(\frac{1}{N_A} + \frac{1}{N_D} \right)}, \quad (4.9)$$

where V_{ext} has to be taken negative for reverse bias. The various cases of external voltages applied to the pn-junction are graphically represented in Fig. 4.4. Since the depletion zone




Figure 4.3: Doping and space charge densities, electric field strength and electric potential at an abrupt pn transition. (Reproduced from Ref. [102], © 2015 Kolanoski, Wermes)

does not have free charge carriers, it can be regarded as a capacitor with capacitance

$$C = \epsilon_{\text{Si}} \epsilon_0 \frac{A}{d}, \quad (4.10)$$

where A is the area and d the width of the depletion zone. The depletion width d can therefore be measured by measuring the capacitance. When particles pass through the depletion zone, they ionize the material and electron-hole pairs are created and drift in the electric field of the depletion zone inducing a detectable signal. The depleted region

must reach the electrode where the readout electronics are connected. Applied reverse bias is the operation condition for semiconductor detectors with the goal to reach complete depletion. A large depletion zone is essential for high detection efficiency.

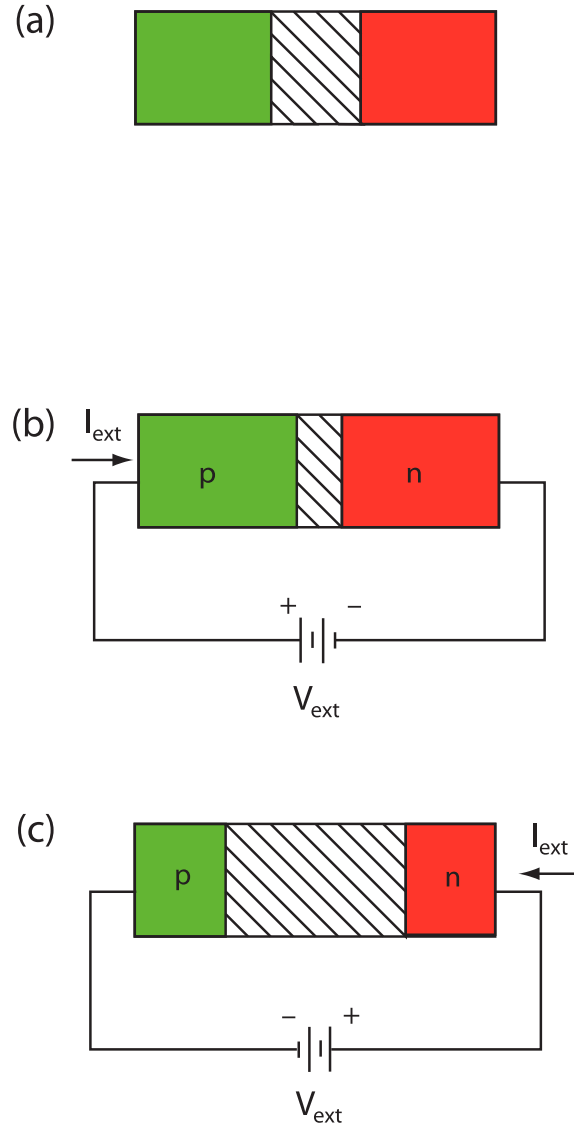
Other junctions

Another important interface is the Metal-Oxide-Semiconductor (MOS) structure. The MOS structure comprises three layers (metal, insulating oxide and semiconductor) with two distinct interfaces. MOS structures can be used to build MOS Field-Effect Transistors (MOSFETs), the most widely used transistors in digital circuits. Two basic configurations exist, determined by substrate doping and terminal types: **NMOS** and **PMOS**. When both types are integrated into the same substrate, where one of them is embedded in a doped “well”, the design is termed Complementary MOS (CMOS) technology. CMOS transistors offer significantly lower power consumption than single-type implementations, as they dissipate power only during switching. These are the dominant MOSFET variants and are commonly used in ASICs such as the ITk Pixel readout chip (see [Section 4.3](#)). The FASER Preshower upgrade [113] employs a pixel detector where the sensing elements are implemented directly into CMOS technology on the same wafer as the readout circuitry (see [Section 4.2.4](#)). For further details on operation principles and characteristics of semiconductor devices, refer to the literature (e.g. Ref. [114]).

4.2.2 General characteristics of pixel sensors

Silicon semiconductor detectors are commonly employed as tracking detectors in high-energy physics applications, such as at the LHC. By using highly segmented anodes or cathodes, it is possible to construct *pixel* detectors with high granularity. Alternative geometries, such as strip detectors, typically offer lower position resolution and granularity but are often deployed in outer tracking layers surrounding a pixel detector, where the dead time of the strips has minimal impact on tracking efficiency. For pixel detectors in the innermost layers, fast readout of hit information is essential, as collisions can occur every 25 ns. Sufficient time resolution is also critical for pile-up rejection and correct association of hits to the originating bunch crossing. Since pixel detectors are positioned closest to the IP, they are exposed to the highest radiation levels and must be radiation-hard. They are then used to detect relativistic charged particles in the keV to MeV range.

This section briefly outlines different possible pixel sensor geometries and then discusses characteristics of pixel sensors in operation in high radiation environments like at the LHC.



Kolanoski, Wermes 2015

Figure 4.4: The pn-junction under different voltage conditions: (a) without external voltage, (b) with forward voltage applied, (c) with reverse voltage applied (operation condition for semiconductor detectors). The shaded area is the depletion zone. The total current I_{ext} exhibits an exponential dependency on the externally applied bias voltage V_{ext} as shown later in Eq. (4.11). (Reproduced from Ref. [102], © 2015 Kolanoski, Wermes)




Figure 4.5: Schematic cross-sections of (left) planar sensor, and (right) 3D sensor, emphasizing the decoupling of active thickness (Δ) and collection distance (L) in 3D sensors. (Reproduced from Ref. [116], © 2012 CERN)

Both pixel and strip detectors build on the characteristics of the pn-junction as described earlier. Pixel sensor geometries can be planar or 3D [115], with the latter featuring doped columns reaching into the substrate. Figure 4.5 shows a comparison of planar and 3D sensors. The fabrication of 3D sensors is more complex and results in higher costs and lower production yields compared to planar sensors. The IBL marked the first adoption of 3D sensors in a collider tracking detector [62]. Their design is illustrated in Fig. 4.6. The 3D sensors are employed in regions with high pseudorapidity η , as for normal incidence tracks partial column inefficiency was observed. A key advantage of the 3D geometry is that sensor thickness becomes decoupled from drift distance. This reduces charge trapping, a negative effect caused by radiation damage, and lowers the depletion voltage required. In summary, 3D sensors are more radiation-hard, while being less demanding in terms of bias voltage and cooling.

Planar sensors can be categorized by thickness. Thin planar sensors (100 to 150 μm) also have the advantage of shorter drift paths, which reduces bias voltage requirements, thus relaxing the demands on the cooling system to limit leakage currents. However, these advantages are less pronounced than those of 3D sensors. Thin sensors also minimize detector material, reducing the probability of multiple scattering. Their primary drawback relative to thick planar sensors (300 μm) is a smaller signal amplitude, as fewer electron-hole pairs are generated, potentially degrading energy resolution. Planar sensors are generally much cheaper in production than 3D sensors.

The choice of sensors for the ITk Pixel detector is detailed in Section 4.2.3.

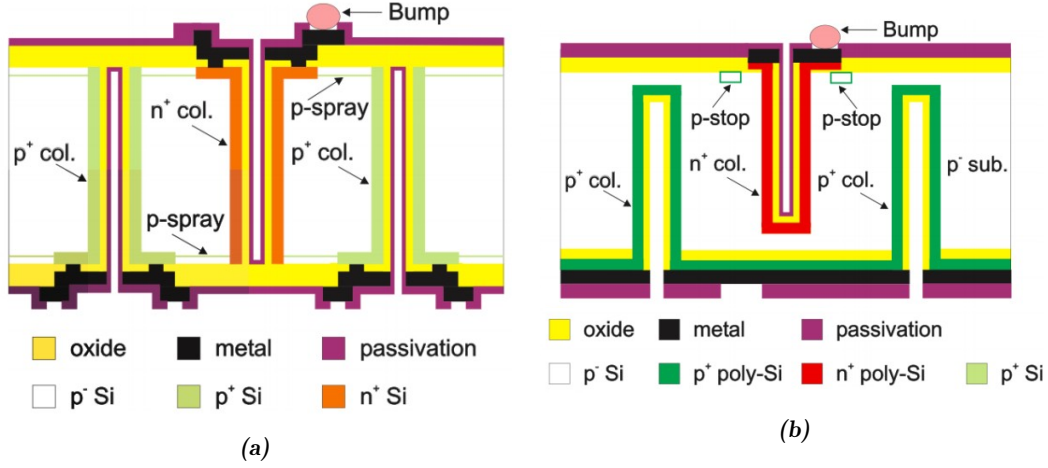


Figure 4.6: Design of the columns of 3D sensors for the ATLAS IBL detector manufactured by (a) the FBK (Fondazione Bruno Kessler) and (b) the CNM (Centro Nacional de Microelectrónica) research institutions. The FBK sensor is a full 3D sensor (the columns reach the opposite side of the sensor), while the CNM sensor is only partially-through 3D. This sketch is for illustration only and is not to scale. (Reproduced from Ref. [62], licensed under CC BY 3.0)

Depletion voltage, breakdown voltage and leakage currents

The working principle of silicon sensors is based on the principle of the pn-junction as described in Section 4.2.1. Under external voltage, the pn-junction exhibits the characteristics of a diode. For an ideal diode, the current I follows the Shockley equation under external voltage V_{ext} [117]:

$$I(V_{\text{ext}}) = I_S \left(e^{eV_{\text{ext}}/k_B T} - 1 \right), \quad (4.11)$$

where I_S is the reverse saturation current, k_B the Boltzmann constant, e the electron charge and T the temperature. In forward direction ($V_{\text{ext}} > 0$), the current quickly increases. In reverse bias ($V_{\text{ext}} < 0$), the current saturates at I_S . The saturation current I_S depends on the intrinsic charge carrier concentration n_i and is highly temperature dependent.

To create a large depletion region for charge collection, an external reverse bias voltage is applied, defined as $V_{\text{bias}} = -V_{\text{ext}} > 0$. The voltage required to fully deplete the sensor in reverse bias is the *depletion voltage* V_{depl} . Any reverse bias voltage on a pn-junction will generate a small leakage current as can be seen from Eq. (4.11). For the ideal diode, leakage current arises solely from the reverse saturation current I_S . The reverse saturation current I_S comprises two components [111]. The first component is a *diffusion current* from a diffusion movement of free charge carriers from the undepleted region into the sensitive space charge region. The second component is a *volume generation current*

caused by thermal generation of electron-hole pairs within the depleted volume. In reality, additional contributions arise from impurities, contaminants, defects in the sensor bulk, and interface defects at silicon-oxide boundaries. Some originate from radiation damage, which affects leakage current through two components [118]: First, an additional contribution to the volume generation current due to electron-hole pair generation at radiation induced defects in the space-charge region. The defects are caused by neutron radiation and lead to additional energy levels in the band gap which can be used for the electron-hole pair generation. Second, ionizing radiation (photons or charged particles) causes defects at silicon-oxide interfaces and leads to *surface generation currents*. Consequently, leakage current increases with radiation damage. Radiation effects on sensors are discussed further in [Section 4.2.2](#).

In reverse bias, leakage current behavior is characterized as follows:

- For $V_{\text{bias}} < V_{\text{depl}}$, bulk contributions dominate. Leakage current scales with the depleted volume, proportional to the depletion width d from [Eq. \(4.9\)](#). Neglecting the built-in voltage V_{bi} ,

$$I(V_{\text{bias}}) \propto d \propto \sqrt{V_{\text{bias}}} , \quad \text{for } V_{\text{bias}} < V_{\text{depl}} . \quad (4.12)$$

- For $V_{\text{bias}} > V_{\text{depl}}$, surface generation currents contribute, leading to a leakage current plateau. Beyond this, the electric field, first very localized then in the whole sensor, becomes so large that electrical breakdown occurs: at a certain voltage accelerated electrons and holes themselves start ionizing atoms, creating avalanches of electron-hole pairs and causing abrupt current surges. This voltage is called *breakdown voltage*.

The progression of the IV-curve (or IV characteristics) is shown in [Fig. 4.7](#). The normal operating mode for a detector is in the plateau region of the leakage current to avoid breakdown.

The leakage current from bulk effects depends on the carrier concentration n_i and thus exhibits a strong temperature dependence [119, 120]:

$$I(T) \propto T^2 \exp\left(-\frac{E_{\text{eff}}}{2k_{\text{B}}T}\right) , \quad (4.13)$$

where k_{B} is the Boltzmann constant and E_{eff} is called the effective energy with a value of 1.21 eV. This relationship enables the temperature normalization of IV characteristics. Only normalized IV-curves can be meaningfully compared. The typical leakage current in a silicon detector at room temperature is a few nA/cm² [121].

A large current, potentially caused by radiation damage, poses two risks: Firstly, thermal runaway can happen when the excess power is not dissipated well enough by the cooling




Figure 4.7: Typical shape of an IV-curve of a silicon sensor. The positive voltage axis corresponds to the reverse bias voltage $V_{\text{bias}} = -V_{\text{ext}}$ from the previous discussion. The origin of different leakage current contributions are indicated. The voltage V_{max} indicates the hard breakdown voltage of the sensor. (Adapted from Ref. [111])

system, heating up the sensor and causing a larger leakage current in return entering a positive feedback loop. Secondly, a localized large current causes localized but extreme noise in the readout FEs, making it necessary to mask affected pixels.

Given their sensitivity to sensor defects, IV-curves are powerful diagnostic tools. During production, IV-curves are regularly taken and any deviation from the expected shape is indicative of a problem in the production process.

Signal formation by moving charges

Charge carriers in semiconductors move from regions of higher concentration to regions of lower concentration via random thermal motion, a process known as *diffusion*. Under the influence of an external, time-independent electric field, charges also move in a directed way called *drift*. This drift occurs at a constant velocity because the acceleration from the electric field is balanced by frictional forces from collisions within the semiconductor lattice. A simple model for the drift motion of electrons in a semiconductor is given by the Drude model [122]. In the case of a constant drift velocity (i.e. no acceleration or deceleration), the drift velocity \vec{v}_D in the Drude model is proportional to the electric field

$$\vec{v}_D = \mu \vec{E} , \quad (4.14)$$

where μ is the *mobility* of the charge carrier, which depends on the temperature, doping concentration and imperfections of the semiconductor lattice. In practice, μ becomes also field dependent for large electric field strengths [123].

Mobility μ is different for electrons and holes in a given material. In silicon at room temperature, electron mobility is about three times larger than hole mobility [124]. In an external magnetic field, the Lorentz force causes the charge carriers to drift on a curved trajectory. For the situation of a magnetic field perpendicular to the electric field, the deviation from the straight line is characterized by the *Lorentz angle* α_L between two collisions and is dependent on the mobility of the charge carrier. Consequently, the Lorentz angle is also dependent on the electric field and therefore the external bias voltage.

Free charge carriers are generated when particles traverse the sensor. They drift toward electrodes and induce a measurable current signal $i_{S,i}$ on electrode i according to the Shockley-Ramo theorem [125, 126]:

$$i_{S,i} = q \vec{E}_{w,i} \vec{v}_D , \quad (4.15)$$

where $\vec{E}_{w,i}$ is the *weighting field* that determines the signal induced on electrode i . For a highly-segmented pixel detector with small pixels, the weighting field is confined to a small region around the collecting electrode, meaning that most of the signal is induced in the last part of the drift path. Integrating the induced current over time yields the collected charge.

When particles cross the pixel sensor at an angle or a specific position, the created electron-hole pairs might drift to different electrodes. This phenomenon is called *charge sharing*. Another cause for charge sharing can be the drift in magnetic fields, as explained earlier. Charge sharing can improve the spatial resolution of the detector by incorporating the charge distribution within clusters of hit pixels into fits for position measurements. Charge sharing enhances the detector's position resolution by incorporating the charge distribution within pixel clusters into position measurements. However, it is important that at least one pixel in a cluster collects sufficient charge to overcome the threshold for detection.

Readout systems may operate in binary mode, where only information about a signal surpassing a certain threshold is processed, or analog mode, where the collected charge is measured. Analog processing can be circumvented by measuring the Time over Threshold (ToT), which typically exhibits an almost linear relationship with charge [111]. Figure 4.8 shows the dependence of the ToT on the collected charge and threshold. High bias voltage and thick sensors can provide a large space charge region in which a signal can be formed. A perpendicular track through 300 μm of silicon will produce about 25 000 electron-hole pairs, equivalent to 4 fC [107].




Figure 4.8: Measurement of the Time over Threshold (ToT) for two different particles hitting the sensor at the same time. A higher energetic particle creates a larger charge to be collected. The larger the charge, the larger is the signal and the measured ToT. A conversion of ToT to collected charge will depend on the threshold. Note, that even though the two particles hit the sensor at the same time, their signal starts to exceed the threshold at different times. This time difference is called *time walk*.

Radiation-hardness and radiation damage

Since highly granular silicon pixel detectors are usually closest to the interaction point, they experience the highest radiation levels and radiation damage can be substantial. The radiation environment is characterized by three quantities:

1. Non-Ionizing Energy Loss (NIEL) of hadrons measured in 1 MeV equivalent neutron fluence (n_{eq}), where the NIEL of 1 MeV neutrons in silicon is defined as 95 MeV mb.
2. Total Ionizing Dose (TID) measured in Gy, defined as the energy deposited by ionization per unit mass of the material.
3. Fluence of hadrons with energy $E > 20$ MeV serving as a measure for the general particle flux.

For the sensor, radiation damage primarily occurs in the silicon bulk [127] and is caused by NIEL-induced displacement of lattice nuclei. Large energy transfers displace lattice atoms from their original position creating vacancies or interstitials. These defects introduce additional energy bands in the band gap between the valence band and conduction band of the semiconductor, thereby allowing for easier excitation of electrons into the conduction band. This increases the leakage current of the sensor. The increased current raises power dissipation and heats the sensor, which can increase the leakage current further, potentially triggering thermal runaway of the sensor. Proper cooling mitigates this positive feedback loop.

Furthermore, additional energy bands can cause charge trapping, where charge carriers become trapped in defects and are released only after the system's charge collection time. This reduces the charge collection efficiency, defined as the collected charge on the electrode divided by the deposited charge. The irradiated 3D sensors of the IBL feature a smaller decrease in charge collection efficiency than their planar sensors, due to the shorter charge collection path [128].

Radiation-induced point defects in the silicon bulk can be charged, leading to an increase in space charge in the depleted region, thereby altering the effective doping concentration. Consequently, higher bias voltages are required to fully deplete the sensor compared to nominal doping concentrations. However, bias voltage cannot be raised indefinitely due to limits imposed by electric breakdown. Extreme doping concentration changes may even lead to type-inversion: When n-type silicon is exposed to high fluences, charged acceptor-like defects are generated, gradually changing the depletion behavior to that of p-type material [111]. These effects were taken into account for the design of the current ATLAS Pixel detector [4, 55].

Bulk defects can be repaired at high temperatures through *beneficial annealing* [118]. At high temperature, the lattice mobility increases and vacancies and interstitials can recombine. Nevertheless, a prolonged time at high temperatures activates other, more severe defects, a process called *reverse annealing* [129], and must be avoided during operation since it degrades detector performance.

Electronics positioned within the detector volume must also be radiation-hard. TID effects can alter the internal electric field of MOSFETs, degrading transistor characteristics [127]. A high particle flux increases the likelihood of Single-Event Effects (SEEs), where ionizing particles deposit sufficient charge in transistors or memory cells to cause for example a simple bit flip from a logical 1 to 0 or vice versa. Mitigation strategies, such as redundant circuit elements, can prevent issues from these effects. Numerous techniques to make electronics radiation-hard exist, see for example Ref. [130] for details.

4.2.3 ITk Pixel sensors

The sensors for the ITk Pixel detector must meet the requirements of the environment, while still balancing the cost for the project. The ITk collaboration chose 3D sensors [131] for the innermost layer, due to their radiation-hardness and low power dissipation, while planar sensors were selected for the outer layers owing to their high fabrication yield and lower costs [91].

The double-sided 3D sensor process was already used in the IBL [62] (see also Fig. 4.6). For the ITk Pixel detector, however, a single-sided n-in-p technology is used to reduce cost and increase yield. The 3D sensor is reduced in size compared to the IBL design, with a reduced pixel pitch of $25 \times 100 \mu\text{m}^2$ (for the barrel) and $50 \times 50 \mu\text{m}^2$ (for the rings), and an active thickness of $150 \mu\text{m}$ (with additional $100 \mu\text{m}$ support) [132]. The sensor measures approximately $19.5 \times 20.3 \text{ mm}^2$ in size [133]. For 3D sensors, one FE chip is bump-bonded to each sensor (see Section 4.3 for more details on hybrid detectors). The High Voltage (HV) bias is applied from the backside which is metallized with aluminum. Requirements for the 3D sensors specify that the depletion voltage V_{depl} must not exceed 10 V before irradiation and that the leakage current at -25°C must not exceed $40 \mu\text{A}/\text{cm}^2$ ($160 \mu\text{A}$ for a 4 cm^2 sensor) after irradiation to $1.3 \times 10^{16} \text{ n}_{\text{eq}}/\text{cm}^2$ [132]. The leakage current *per pixel* must never exceed 10 nA during the lifetime of the detector, as this is a FE limitation [132]. The detector will be operated at around -25°C . The power dissipation performance of the 3D sensors at operating temperature and after irradiation must be less than $10 \text{ mW}/\text{cm}^2$ and the maximum operational bias voltage shall be less than 250 V. Extensive test beam campaigns are studying the 3D sensor technology, see for example Ref. [134].

In layers 1 to 4, the ITk Pixel detector uses planar sensors produced in single-sided n-in-p technology with a thickness of $100 \mu\text{m}$ in the layer 1 barrel region and a thickness of $150 \mu\text{m}$ in layer 1 rings and in other layers [135]. The p-type silicon bulk is doped with boron, while n-type implants using phosphorus doping form the electrodes. These are DC coupled to aluminum read-out pads to form a pixel with pixel isolation achieved via p-stop or p-spray implants. All planar sensors have a pixel pitch of $50 \times 50 \mu\text{m}^2$. The sensor size is approximately $41.1 \times 39.5 \text{ mm}^2$ [133], requiring four FEs per sensor. Requirements for the planar sensors specify that the depletion voltage V_{depl} must not exceed 50 V (90 V) before irradiation for sensors with thickness of $100 \mu\text{m}$ ($150 \mu\text{m}$) [135]. The power dissipation performance of the planar sensors at operating temperature and after irradiation shall be less than $10 \text{ mW}/\text{cm}^2$ at the maximum planned operational bias voltage of 400 V for sensors with a thickness of $100 \mu\text{m}$, and less than $22 \text{ mW}/\text{cm}^2$ at the maximum planned operational bias voltage of 600 V for sensors with a thickness of $150 \mu\text{m}$.

3D sensors and thin planar sensors are used in the innermost layers, due to their suitability for high radiation environments. However, replacement of the two innermost layers after 2000 fb^{-1} is foreseen.

4.2.4 Monolithic Active Pixel Sensors

The pixel detector solution chosen for the ITk is based on a hybrid concept of a separate sensor (without signal amplification) and readout chip (for more information on the ITk FE

and hybridization, see [Section 4.3](#)). However, alternative developments aim to integrate the sensing element directly into CMOS technology on the same wafer as the readout circuitry, eliminating the need for dedicated sensor wafers [136]. This approach is known as CMOS Monolithic Active Pixel Sensor (MAPS) technology. The advantages of monolithic sensors include a lower material budget, smaller pixel pitches in the order of $10\mu\text{m}$, simplified assembly (since hybridization is unnecessary) and significantly lower costs compared to hybrid detectors.

The ALICE collaboration adopted the CMOS MAPS technology for their tracker upgrade during LS2 [137, 138], operating in a radiation environment significantly milder than ATLAS or CMS tracking detectors. Similarly, early studies for the ITk explored using CMOS MAPS in the outer layers of the pixel detector [139, 140]. For Run 4 of the LHC, the ALICE collaboration will deploy bent, wafer-scale stitched CMOS MAPS to realize a truly cylindrical vertex detector [141, 142].

Traditionally, MAPS faced limitations in radiation-hardness under NIEL, as early designs collected charges primarily via diffusion, which is rather slow and highly dependent on the health of the silicon bulk [136]. Recent developments therefore focus on increasing the electric field within the sensor to accelerate signal charge collection via drift motion and minimize charge trapping. The FASER Preshower upgrade [113] employs monolithic pixel detectors based on SiGe BiCMOS technology [143]. This thesis details the work on the FASER Preshower DCS in [Section 10.4](#).

4.3 ITk Pixel FE chips

As mentioned previously, the ITk Pixel system employs hybrid pixel detectors across all layers, meaning that sensor and readout electronics are on separate wafers and later connected through a process called hybridization [144]. After hybridization, each pixel in the sensor is electrically connected to the readout circuitry in the FE by a bump bond. A schematic cross-section of such a connection is shown in [Fig. 4.9](#). The primary advantage of hybrid pixel detectors is that both the sensor and the readout chip can be independently engineered for radiation-hardness. However, these detectors are comparatively expensive, and the pixel pitch cannot be made arbitrarily small due to limitations imposed by the bump bonds. For the ITk, the connection between the sensor and the readout chip will be made with SnAg/SnPb or Indium bumps for a Direct Current (DC)-coupling between the sensor and the chip circuitry [133].




Figure 4.9: Cross-section of a hybrid pixel detector with an n-in-p planar sensor using single-sided processing. Only one end with two connections between sensor and the Front-End electronics of the readout chip is shown. To maintain the cost reduction through single-sided processing, HV guard rings have to be located on the same side as the n-type pixel implants, bringing the sensor edges to high voltages.

For the Phase-II upgrades of their pixel detectors, ATLAS and CMS joined efforts and established the RD53 collaboration⁶ to develop the hybrid pixel readout ASIC in the selected 65 nm CMOS technology [145, 146]. Key requirements for the final chip include radiation-hardness up to 10 MGy of TID, a high hit rate (3.5 GHz/cm^2), high-speed data readout (5.12 Gbit/s) and low power consumption ($<1 \text{ W/cm}^2$) [147]. For the chips used in the ITk, the readout chips must support a trigger rate of 1 MHz and a trigger latency of $10 \mu\text{s}$ as dictated by the TDAQ upgrade for Phase-II as described in Section 3.4.1. The RD53 collaboration has led to the production of three chip iterations:

1. A prototype chip called RD53A [148, 149],
2. a preproduction chip called RD53B [150–152] in two variants: ITkPix-V1.1 for ATLAS and CROC-V1 for CMS, and
3. a production chip called RD53C [153], again in two variants: ITkPix-V2 for ATLAS and CROC-V2 for CMS.

The RD53A prototype, submitted in 2017 for production, was used to finalize the design for the production chips. It underwent extensive single-chip and module testing, and was also integrated into the demonstrator for system test measurements. Results from these tests are presented in Chapter 6. A comprehensive overview of the RD53 chip can be found in Ref. [154]. This section describes the RD53A chip, as well as gives a brief introduction to the preproduction ITkPix-V1.1 and production ITkPix-V2 variants.

⁶<https://rd53.web.cern.ch/>

4.3.1 The RD53A readout chip

The RD53A [149] is a half-size pixel chip demonstrator featuring three different Analog Front-End (AFE) designs and two different pixel readout architectures. It was used to study various implementations and determine the most appropriate design for the final production chip. The ASIC comprises a matrix of 400×192 pixels with $50 \times 50 \mu\text{m}^2$ pitch and a total size of $20 \times 11.6 \text{ mm}^2$. The challenge was to design a readout chip with low noise and good threshold dispersion performance to operate at low thresholds around 1000 electrons [155], which guarantees good detection efficiency with the chosen sensors even after irradiation.

The three variants of integrated AFEs are called synchronous, linear and differential. Their arrangement on the chip is shown in Fig. 4.10. The pixel matrix consists of pixel cores with 8×8 pixels each. Within a pixel core, 2×2 pixels form a so-called pixel quad that contains an “island” for analog pixel circuitry and a “sea” for the digital pixel logic. Each pixel core is identical, and pixel cores are grouped into core columns. At the bottom of the chip is the chip periphery consisting of the digital chip bottom, analog chip bottom and the pad frame with pads for wire-bond connections and ShuntLDOs for serial powering.

The digital chip bottom contains the blocks that implement all control and processing functionality. It contains the command decoder block (to receive commands from the control link and relay them to the pixel matrix), a data builder block (to build event data from the pixel matrix), an Aurora 64b/66b [156] encoder block (to encode data on 1 to 4 parallel lanes at 1.28 Gbit/s per lane), and the channel synchronizer block (to generate and distribute the 40 MHz bunch crossing clock to the pixel matrix). The command decoder block, the data builder block and the channel synchronizer block all use a 160 MHz clock generated by the Clock and Data Recovery (CDR) block in the analog chip bottom. Furthermore, the analog chip bottom contains the output serializer that uses a high speed clock of 1.28 GHz from the CDR block, and sends the Aurora encoded data. Additional blocks in the analog chip bottom are for signal building for calibration, temperature and radiation sensors, and for monitoring of analog quantities using voltage and current multiplexers and a 12 bit Analog to Digital Converter (ADC).

For readout, trigger pulses are generated by the command decoder block in response to a trigger command. An internal self-trigger was not implemented in RD53A. Similarly absent in RD53A (but present in RD53B and RD53C) is the block for data merging of data links from two or four FEs into one.

For commands to work, the chip’s input path must first initialize and lock the phase, which happens automatically on startup and when an input command serial stream is supplied.




Figure 4.10: Arrangement of AFE flavors in RD53A. The pixel column number range of each flavor is shown along the bottom. The digital readout architecture in the synchronous FE is Central Buffer Architecture (CBA) and in the linear and differential FE Distributed Buffer Architecture (DBA) is used. (Reproduced from Ref. [149], licensed under CC BY 4.0)

On-chip power distribution

A novel powering scheme with constant current, in which modules are powered in series, has been adopted for the new pixel chips [157]. The analog and digital power domains of the chip are powered in parallel by on-chip ShuntLDO [158] power regulators. In a quad module with four FE chips, there are 8 ShuntLDOs in parallel. The ShuntLDO regulator combines a specialized linear Low Drop-Out (LDO) voltage regulator with a shunt element. The power is provided by a constant current source. The chips VI-characteristic under constant current is shown in Fig. 4.11(a). Above a given minimum operation current, the ShuntLDO actively regulates a constant output voltage V_{out} tuned to be 1.2 V, independent of load current, while maintaining a well-defined input impedance to allow for current sharing among other consumers in parallel [154]. The impedance is achieved by dissipating excess current through a regulated shunt. However, this represents a key drawback of serial powering: to allow for headroom in the operation, the total power overhead on average is 20 to 40 % [154], which must be dissipated by the cooling system. Notably, from the power supply's perspective, neither the internal voltage V_{out} nor the load size can be inferred as the voltage drop V_{in} remains constant for a given input current I_{in} . A simplified schematic of a ShuntLDO is shown in Fig. 4.11(b).

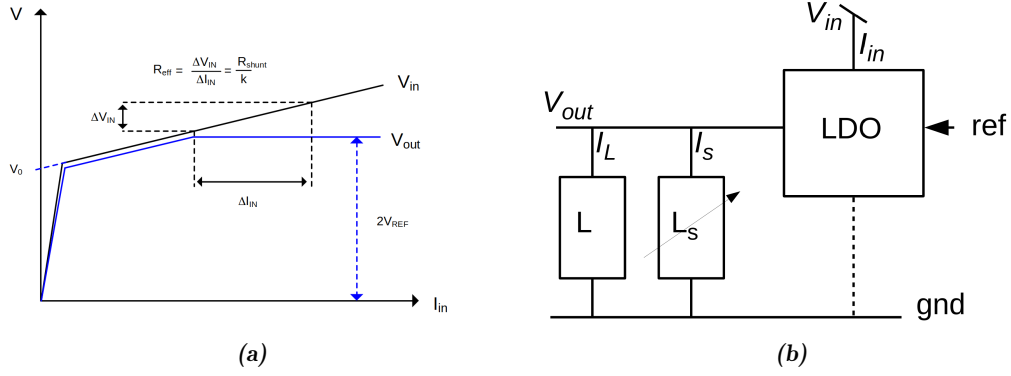


Figure 4.11: (a) Ideal ShuntLDO VI-characteristics without load. The operational range starts from the point where the output voltage V_{out} is stabilized. (b) Simplified schematic of the ShuntLDO made from a special LDO with regulated shunt current (I_S) in parallel to the load current (I_L) such that $I_{in} = I_S + I_L$. (Reproduced from Ref. [154], licensed under CC BY 4.0)

The ShuntLDO of the RD53A experienced difficulties starting up correctly at cold temperatures. Ensuring reliable and safe startup has been a major challenge. Several incremental improvements to the ShuntLDO were implemented into subsequent chip generations, including under-shunt current protection (to avoid a faulty module load drawing too much current, $I_L > I_{in}$) and over-voltage protection [150, 159].

Analog front ends

Three different AFE variants, i.e. the circuitry containing the signal amplifier (which converts an input charge to a voltage), the signal shaper and the discriminator (which compares voltage pulses to a threshold voltage), are implemented in one RD53A chip (see Fig. 4.10) to enable detailed performance comparisons.

The synchronous AFE uses a baseline “auto-zeroing” scheme that periodically acquires a baseline reference, eliminating the need for pixel-by-pixel threshold trimming [149]. The linear AFE includes a linear pulse amplification prior to the discriminator, while the differential AFE uses a differential gain stage in front of the discriminator and implements a threshold by unbalancing the two branches [149]. The pulse width at the discriminator, also called Time over Threshold (ToT), can be converted into a representative charge measurement of the charge deposited in the sensor. For the digital readout, RD53A implements two different architectures for storing ToT information. The linear and differential AFEs use an architecture called Distributed Buffering Architecture which stores ToT data per pixel, while the synchronous AFE employs an architecture called Centralized Buffering Architecture, where ToT information is stored in a shared memory alongside a

hit map to suppress zero ToT values. ToT information is then later used to enhance the precision of cluster position measurement [93].

To equalize the pixel response across the pixel matrix in terms of threshold and ToT, the discriminator threshold voltage has to be tuned individually for each pixel in each chip. This is usually done during module production. A *tuning procedure* uses electrical tests as explained in the next section and alternates them with a changing configuration to find the best setting. For details on the tuning procedure, see Ref. [160].

Injection circuitry and electrical testing

All AFEs share the same calibration injection circuitry, which is essential for direct performance comparisons [149]. Every pixel contains the same circuit. By applying a controlled voltage to an injection capacitor, a known charge is injected into the amplifier. The RD53A injection circuitry uses two 12 bit voltage output Digital Analog Converters (DACs) to charge the injection capacitor with a differential voltage and enable charge injections up to approximately 45 000 electrons. This functionality enables verification of the operability of the analog circuitry of the chip even before a sensor is connected or before particles are detected.

The injection circuit receives input from the two adjustable global DC voltages and two global control signals. However, each pixel includes an individual “enable” bit to activate only a selected subset of pixels for injection. This is a necessity, as most readout concepts cannot handle too many simultaneously generated hits.

For the readout in the digital domain of the pixel circuitry, it is also possible to send a control enable signal to be able to only test the digital part of the circuitry. Additionally, any pixel may be individually masked (disabled) during readout, which is valuable for mitigating faulty AFEs or noisy sensor pixels that would otherwise consume the available bandwidth. Timing information, i.e. when a hit occurred, and ToT information is stored by pixel region. Timing information can go back up to 511 bunch crossings, while ToT is quantized in bunch crossings with a maximum value of 15 using the 40 MHz clock. Upon trigger reception, a trigger tag is stored to label the hits, and readout proceeds one trigger tag at a time [149].

Each pixel features up to 8 bits of local configuration for masking and threshold tuning. Global configurations are stored in global registers. For example for tuning of ToTs, global registers are used. Prior to electrical testing of demonstrator chips, as described later, all chips were configured using settings obtained from tuning during module production. The target thresholds during tuning were 1400 electrons for the linear and differential AFEs

and 2500 electrons for the synchronous AFE. The ToT was tuned for 7 bunch crossings to correspond to a charge of 10 000 electrons. For the final chips and the final detector, it is expected to use a threshold of 900 electrons for modules in the innermost layer of the barrel, with all other modules using a threshold of 600 electrons [93].

Typical testing procedures using readout software include checking the response of the digital FE in a digital scan, checking the general response of the AFEs in an analog scan, and testing the response of the pixels in the matrix in a threshold and ToT scan. Additionally, disconnected bump bond scans validate the chip-to-sensor connections, while cross-talk scans evaluate interference between neighboring pixels. These tests were performed with the modules in the demonstrator. Their scan routines are detailed and results presented in [Section 6.2.3](#).

4.3.2 The ITkPix readout chip

Using the knowledge gained from tests with the RD53A chip, preproduction (RD53B) and later production (RD53C) chips were developed by the RD53 collaboration. The preproduction and production chips were separated into two variants to meet the specific requirements of the ATLAS and CMS experiments: ATLAS adopted the differential AFE, while CMS selected the linear AFE [154]. The pixel matrix uses the full size, which means for the ATLAS chips a size of 400×384 pixels. The RD53B-ATLAS (ITkPix) was submitted in 2020 to a foundry for production. A first version (ITkPix-V1) exhibited non-functional ToT measurement and a high current bug, prompting resubmission as ITkPix-V1.1 with changes in the metal layer to mitigate the high current bug. The revised chip is also used for preproduction Loaded Local Supports (LLSs) to qualify the LLS QC test setups. The final production chip (ITkPix-V2) was submitted in 2023.

Compared to the prototype, the preproduction and production chip, have only one AFE variant across the whole matrix and incorporate improvements to the ShuntLDO, including a low-power mode (controlled via the [LPE](#) signal) for integration and data-merging functionality. This allows a primary chip to aggregate serial data from one or more secondary chips, then time-multiplex it with its own output for transmission via the uplink data stream [151]. This data-merging capability is extensively used in the outer layers of the ITk Pixel detector to reduce cabling requirements within the detector volume. For additional differences between the RD53A and RD53B chips, see Ref. [150, Chap. 3].

The final RD53 chip implementations contain 660 M transistors, 56 M standard cells and 12 M memory elements [154].

4.4 ITk Pixel modules

The sensors and readout chips, as described previously, are bump bonded together in a process called hybridization [144]. The resulting hybrid pixel module is termed a *bare module*. Two types of bare hybrid pixel modules exist for the ITk: a single-chip bare module consisting of one 3D sensor of size $19.5 \times 20.3 \text{ mm}^2$ and one FE, and a quad-chip bare module consisting of a single, but larger, planar sensor of size $41.1 \times 39.5 \text{ mm}^2$ and four FE chips bump bonded to it in a 2×2 array. The bare module is then attached to a module Printed Circuit Board (PCB) to form the hybrid pixel module. The module PCB is electrically connected to the FE chips via aluminum wire-bonds using ultrasonic wedge bonding [161].

Two pixel module configurations are deployed:

1. Triplets: Three single-chip bare modules are attached to a PCB in three distinct geometries. For the barrel region in layer 0, the three single-chip bare modules are arranged in a straight line, while for the rings in layer 0 and 0.5 they are arranged circularly with varying radii.
2. Quads: One quad-chip bare module is attached to a PCB. Two variants exist with different sensor thicknesses of $100 \mu\text{m}$ for layer 1 barrel and $150 \mu\text{m}$ for layer 1 rings and layers 2, 3, and 4.

Within a module, FEs are powered in parallel, while modules are powered in series in a serially powered (SP) chain. For the quad modules, commands are sent on a single downlink and broadcast to all FEs. For the uplink, data-merging functionality can be used in the preproduction and production versions of the FE chip. The PCB provides the Low Voltage (LV) power domains for the FE chips, routing and filtering of the sensor HV bias, application of the HV bias to the sensor through a hole in the PCB, connection and distribution of the data signals to and from the FEs, housing of Negative Temperature Coefficient Thermistors (NTCs) with one for external, independent monitoring and four additional ones, each read out by a FE. Also implemented on the module PCB are interfaces to external electrical services for powering and monitoring of the voltage levels in front of the module. Connections are grouped into two PCB connectors: one for data signals and another that integrates LV power, HV bias, NTC connection and connection for voltage measurement. The Flat Flexible Cables (FFCs) that connect to the module PCB are termed “pigtailed”.

Since the sensor edges are at high voltages during operation, all modules require protection against HV discharge, that can occur at sensor bias voltages above 500 V. This is achieved via a uniform parylene coating over the module, providing electrical insulation and corrosion




Figure 4.12: Schematic cross-section of a quad module in the OB region. In the detector, cooling will be provided from the bottom, as this is where the module will attach to the local support.




Figure 4.13: Top-view photograph of an ITkPix preproduction quad module after parylene coating. FE1 is on the top left, FE2 below. FE3 is on the bottom right and FE4 on the top right.

resistance. For modules in the OB region an additional mechanical wire-bond protection is implemented to protect the wire-bonds against contact with the module pigtails. A cross-section of a quad module is shown in [Fig. 4.12](#). [Figure 4.13](#) shows a photograph of a quad module with the position of the FE chips indicated by labels on the module PCB.

The final distribution of pixel modules within the ITk Pixel volume is summarized in [Table 4.1](#) for the barrel region and in [Table 4.2](#) for the ring or endcap region. In total, 8372

Table 4.1: Distribution of modules in the ITk barrel region. Note, that triplets contain three sensors with a size of approximately $2 \times 2 \text{ cm}^2$ each, while quad modules contain one sensor with a size of approximately $4 \times 4 \text{ cm}^2$. (Adapted from Ref. [93])

Barrel layer	Radius [mm]	Flat section		Inclined section				Module type	Pixel dim. [μm^2]	Sensor type	Sensor thick. [μm]
		Rows of sensors	Flat barrel $ z $ [mm]	Sensors per row	Incl. rings $ z $ [mm]	Incl. rings	Sensors per incl. ring				
0	34	12	0–245	24	–	–		triplets	25×100	3D	150
1	99	20	0–245	12	–	–		quads	50×50	planar	100
2	160	32	0–372	18	380–1035	2×6	32	quads	50×50	planar	150
3	228	44	0–372	18	380–1035	2×8	44	quads	50×50	planar	150
4	291	56	0–372	18	380–1035	2×9	56	quads	50×50	planar	150

Table 4.2: Distribution of modules in the ITk ring or endcap region. Note, that triplets contain three sensors with a size of approximately $2 \times 2 \text{ cm}^2$ each, while quad modules contain one sensor with a size of approximately $4 \times 4 \text{ cm}^2$. (Adapted from Ref. [93])

Ring layer	Radius [mm]	$ z $ [mm]	Rings	Sensors per ring	Module type	Pixel dim. [μm^2]	Sensor type	Sensor thick. [μm]
0	33.20	263–1142	2×15	18	triplets	50×50	3D	150
0.5	58.70	1103–1846	2×6	30	triplets	50×50	3D	150
1	80.00	263–2621	2×23	20	quads	50×50	planar	150
2	154.50	1145.5–2850	2×11	32	quads	50×50	planar	150
3	214.55	1145.5–2850	2×8	44	quads	50×50	planar	150
4	274.60	1145.5–2850	2×9	52	quads	50×50	planar	150

pixel modules will be installed in the new detector: 7976 quad modules and 396 triplet modules. The total sensitive area in the Pixel detector will be approximately 13.2 m^2 .

4.5 ITk Outer Barrel local supports

The tests conducted for this thesis focused on the flat prototype Loaded Local Support (LLS) of the OB region, also called a longeron. LLSs incorporate all necessary on-detector services, including structures for CO_2 cooling to meet the thermal requirements of the detector. By precisely positioning modules within the detector volume, LLSs ensure complete coverage in the azimuthal angle ϕ . This section describes the design of the OB LLSs. Further information can be found in Ref. [162].

The OB region encompasses the central area of the Pixel detector at layers 2, 3 and 4 (see Fig. 3.11(b)). Its local support design relies on two primary elements: the module cells




Figure 4.14: Module loading in OB as a two-step process: **(a)** Cell loading creates a loaded cell. **(b)** Cell integration of a loaded cell onto a functional local support with a cooling pipe and a base block using screws creates a LLS (not shown is the connected pre-bent module pigtail).

and the functional local supports. The functional local support comprises a carbon fiber reinforced plastic (CFRP) truss structure with an integrated titanium cooling pipe for CO₂ cooling and aluminum-graphite base blocks for mounting module cells. Two different types of local supports are employed in the OB:

1. *Longerons*: Stave-like structures in the central flat section, featuring two rows of modules;
2. *Inclined Half-Rings (IHRs)*: structures at higher $|z|$ where modules are mounted on both sides with inclinations of 67°, 58° and 55° in layers 2, 3 and 4, respectively.

Before loading onto the functional local support to form an LLS, modules are equipped with a *bare* cell (comprising a pyrolytic graphite tile and a cooling block with internal threads for secure mounting) to form a *loaded* cell. The process of gluing a bare cell to a module is called *cell loading* and is depicted in Fig. 4.14(a). In a subsequent step, pre-bent pigtails are attached to the loaded cell, which is then mounted onto the functional local support using screws to allow reworkability. This process is called *cell integration* and is shown in Fig. 4.14(b). PP0s⁷ are added to the local support and the module pigtails are connected to them. These pigtails carry both the data and powering/monitoring signals. Module pigtails are also called Type-0 services, as cable bundles between Patch Panels are sequentially labeled Type-n bundle (counting from the module toward the service cavern). The resulting local support with loaded modules is called an LLS and is shown in Fig. 4.15. All IHRs on one side and in one layer form an inclined unit; two inclined units per layer plus the connecting longerons between them constitute a half-layer shell as shown in Fig. 4.16.

Titanium pipes carry boiling CO₂ to cool the FEs and the sensor of the quad module to suppress leakage current. The PP0s shown in Fig. 4.15 implement the on-detector services. They are realized in flex-rigid technology (rigid PCBs interconnected by FFCs).

⁷Patch Panel 0, patch panels are numbered from the module toward the service cavern




Figure 4.15: The cell integration step brings the PP0, the loaded cells and the functional local supports together to form a Loaded Local Support (LLS). Shown are both types of local supports in OB: An Inclined Half-Ring (IHR) and a longeron. In the IHR, modules are mounted on both sides of the functional local support to allow for gap-less coverage. In the longeron, the modules are mounted at different angles on a top and a bottom row to be able to create a small overlap.




Figure 4.16: CAD representation of the layer 4 *half layer shell* in the OB. On both ends are the inclined units, which will hold the IHRs, and in the middle are the longerons connecting them both. The modules are shown in yellow. For simplicity, the services are not shown.

Each IHR hosts two PP0s, and each longeron hosts four. Modules connect to the PP0 via “pigtailed” (Type-0 cable), as noted earlier. Power and DCS services use Type-1 cable bundles to link a PP0 to a PP1, carrying Low Voltage (LV), High Voltage (HV), Low Power Enable (LPE), Temperature Interlock (Tilock), VCAN and CAN signals for the Monitoring Of Pixel System (MOPS) chip. Both module pigtailed and the Type-1 power bundles connect to the PP0 via ultra-flat Samtec Z-Ray[®] interposers. The Type-1 data services, that route the modules electrical data, clock and command signals between the PP0 and Opto Boards, use shielded twinax cables and connect to the PP0 via a Samtec FireFly[®] connector. To reduce the material in the outer layers, where hit rates are smaller compared to the Inner System, modules make use of the data-merging functionality of the ITkPix chips. The layer 2 PP0 of the OB differs from the layer 3/4 PP0 in that it allows electrical connections in the uplink from two out of four FEs (versus one FE for layer 3/4). Full quad modules are read out in layer 2 in 2-to-1 data-merging mode and in layers 3 and 4 in 4-to-1 data-merging mode which reduces the amount of necessary data cables [162].

The PP0 distributes one LV and two HV lines to the modules. All modules connected to a PP0 are powered in series via the LV line using a constant current source, forming a so-called *SP-chain*. The two HV lines partition the modules into two HV groups, where modules in each group are biased in parallel by one HV channel from a HV Power Supply Unit (PSU) capable of handling high leakage currents. MOPS monitoring provides independent monitoring information of the temperatures and voltages of the modules to the DCS. A complete connection diagram of an SP-chain (or, equivalently, of a PP0) is given in the appendix in Fig. A.1. A simplified version of a connection diagram is given here in Fig. 4.17. It illustrates the connections for the LV FE power, the HV bias lines and the MOPS chip. The FE readout is handled through the Opto System with Opto Boards in Opto Boxes, where the electrical signals from the FEs are converted into optical signals. All Opto Boards associated to an SP-chain are powered in parallel by one powering channel. The serial powering scheme of the modules and the MOPS monitoring on the PP0 are detailed in subsequent sections.

Table 4.3 lists the total LLS count in the final detector. The OB subsystem will contain in total 4472 pixel quad modules, which corresponds to approximately 7 m^2 of active silicon. Therefore, the OB subsystem represents 54% of the ITk Pixel’s active area.

4.5.1 Serial powering scheme

The PP0 groups 6 to 14 modules into an *SP-chain*. All quad modules within an *SP-chain* are powered in series via a constant current source through the LV line. Within each quad module, the four FE chips (and even the analog and digital domains within each FE) are

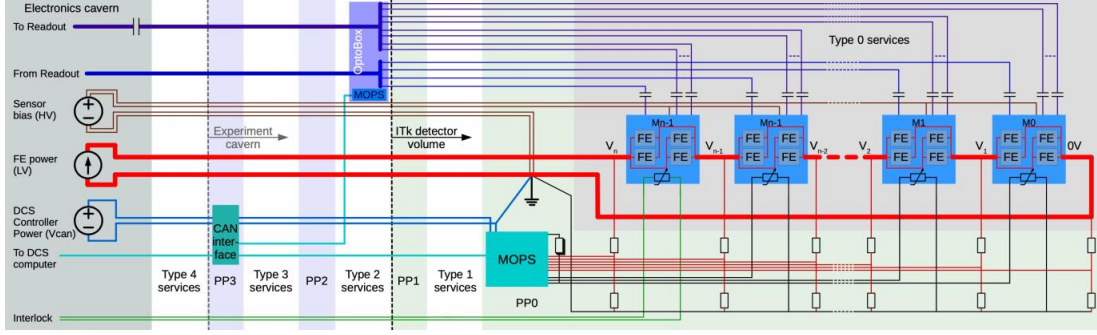


Figure 4.17: Connection diagram of an SP-chain with quad modules including the MOPS chip. The CAN interface at PP3 is called MOPSHUB and connects to several MOPS'. (Reproduced from Ref. [163])

Table 4.3: LLSs and SP-chains in the OB subsystem of the ITk Pixel detector. The numbers in round brackets after the SP-chain specification give the number of modules per SP-chain. (Adapted from Ref. [162])

Barrel layer	Radius [mm]	Flat section			Inclined section		
		Longerons	Modules per longeron	SP-chains per longeron	IHRs	Modules per IHR	SP-chains per IHR
2	160	16	36	4 (12 + 6 + 6 + 12)	$2 \times 2 \times 6$	16	2 (8 + 8)
3	228	22	36	4 (12 + 6 + 6 + 12)	$2 \times 2 \times 8$	22	2 (11 + 11)
4	291	28	36	4 (12 + 6 + 6 + 12)	$2 \times 2 \times 9$	28	2 (14 + 14)

powered in parallel through ShuntLDOs in the chip, as detailed in Section 4.3.1. A generic SP-chain with n quad modules is illustrated in Fig. 4.18(a). The SP-chain improves power efficiency compared to parallel powering as sketched in Fig. 4.18(b). For the simplified case shown in Fig. 4.18, the efficiency gain is [158]:

$$\frac{\varepsilon_s}{\varepsilon_p} = \frac{1 + \frac{nIR}{V}}{1 + \frac{nV}{IR}}, \quad (4.16)$$

where ε_s is the power efficiency (i.e. the useful power output over the total power input) for the serial powering scheme, and ε_p denotes the efficiency in the parallel powering scheme.

Because of their high granularity, the FE chips demand a relatively high supply current at their operating voltage of 1.2 V. For an RD53A chip, the required supply current reaches up to 1 A, while for the final chip design it can reach even up to 2 A depending on digital activity. Powering modules in parallel would result in excessively high total current, leading to either significant power losses or increased cable mass within the detector volume [165]. The use of switched mode DC-DC power conversion on the pixel chip, or on the pixel




Figure 4.18: Schematic comparison of (a) serial powering versus (b) parallel powering of a chain of n quad modules. The voltage V and the current I are the voltage and current needed for one quad module. The load voltage (i.e. the voltage over all modules) is V in the parallel configuration and nV in the serial configuration. The lost voltage drop over the cable is shown above the resistor (with a resistance R in both cases). (Adapted from Ref. [164])

module, for power efficiency gains in a parallel powering scheme was investigated. However, such converters must be radiation-hard, compact and should minimally impact the material budget. No sufficiently mature solution exists for the use in the HL-LHC [164]. The serial powering scheme with a constant current source achieves efficiency gains without added mass using a constant current source, which also easily tolerates large voltage drops across power cables ensuring consistent power delivery to the modules.

The serial powering scheme has been extensively validated [157, 166–168] and adopted for the pixel detector upgrades in ATLAS and CMS. A potential hazard is the breaking of an SP-chain due to a faulty module. The initial proposal to implement additional chips on the PP0 that would allow bypassing a faulty module was abandoned, as redundancy exists within the quad module’s parallel powering path to dissipate excess current from broken paths. Functional ShuntLDOs can handle up to 200% of nominal chip current [154]. If there is only an intermittent fault with an FE in an SP-chain that requires a reset, the reset to return the chip to its default configuration and communication state has to be implemented in a way that does not require power cycling the whole chain, as this would be a major disruption of operation.

Serial powering also affects module biasing. The modules in an SP-chain do not share a common Ground (GND); each module has a local module GND. This creates a characteristic ladder of monitored module voltages along the SP-chain with single module voltage drops

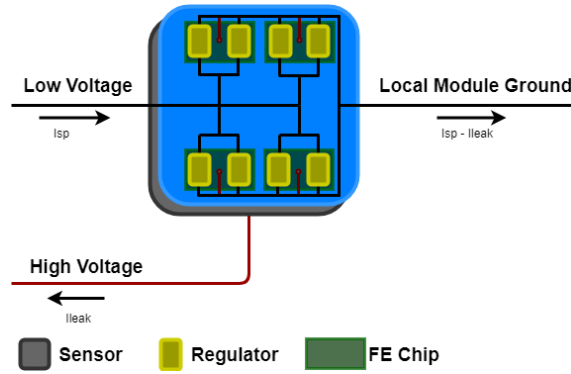


Figure 4.19: Power distribution and sensor biasing in an SP-chain with quad modules. Shown is one module and the current flow through the module. (Reproduced from Ref. [163])

of the order of 1.7 V. Consequently, the HV reverse bias differs for each module in an HV group, such that each module will have its own effective HV reverse bias equal to the sum of the voltage difference coming from the HV PSU and the local module GND offset. For SP-chains with many modules, the GND potential difference between the first and last module can reach up to 24 V. To limit reverse bias variation, modules are split into two HV groups along the center of the LV line.

Critically, the sensor leakage current caused by the HV bias acts against the LV current, as can be seen in Fig. 4.19, which shows the power distribution and sensor biasing in an SP-chain with LV and HV turned on. All LV and HV power supplies are floating and their return lines are tied together and referenced at PP0 (see Fig. 4.17). To have well-defined leakage current paths, LV must be switched on when HV is on. This is especially important when taking sensor IV-curves by scanning HV. Even with HV off, LV power alone creates a small, module-specific reverse bias. A HV PSU with low ohmic off mode, as employed in the ITk, prevents the return of the leakage current in forward direction through the last module of the SP-chain [168].

4.5.2 MOPS monitoring

The Monitoring Of Pixel System (MOPS) chip is a radiation-hard Application Specific Integrated Circuit (ASIC) designed for independent monitoring of module voltages and temperatures. It aggregates and digitizes monitoring data before transmitting it to the central monitoring system, that can be up to 60 m away from the location of the MOPS [169]. The chip is mounted directly on the PP0 to reduce the number of services in the detector. Consequently, the MOPS chip must be radiation-hard up to 5 MGy of TID [169].


The MOPS chip integrates a 12 bit ADC to measure module temperatures (via an NTC on the module PCB) and voltages (LV input voltage relative to PP0 GND) of all modules in an SP-chain. For power, the chip uses the LDO part of a modified version of the RD53A ShuntLDO. The MOPS is powered through a dedicated powering line called VCAN. For communicating the measurement results, the chip makes use of two standardized communication protocols: Controller Area Network (CAN) [170] and CANopen [171]. A differential CANbus (with two wires for CAN-H and CAN-L) connects to the MOPS readout system, where the recessive state uses 0 V difference and the dominant state uses 1.2 V, deviating from the industry standard of up to 5 V difference in the dominant state. This choice was made to stay below the maximum applicable voltage of 1.32 V for the transistors of the 65 nm CMOS technology used for the chip without the use of cascaded transistors [172]. Consequently, the communication with standard CAN modules requires a voltage level shifter. The transfer rate of the bus is 150 kbit/s [169]. The MOPS chip has seen three versions during its development. Since version 2, the chip supports automatic trimming of its internal (autonomous) oscillator, that provides the clock for the CAN interface and the remaining digital circuitry, to a reference signal to compensate for process variation and temperature [173].

As part of the on-detector DCS, the MOPS chip provides continuous, but not safety critical, monitoring of the temperatures and voltages of the SP-chain modules to detect over-temperature or over-voltage situations. The NTC of the last module in an SP-chain along the cooling line always bypasses the MOPS chip and connects directly to the interlock system, as can be seen in the connection diagram in Fig. 4.17. Further details on the MOPS readout system and the interlock are provided in Chapter 9, as both are integral to the ITk Pixel DCS.


4.5.3 CO₂ cooling

Evaporative CO₂ cooling was selected for the silicon-based detector upgrades of CMS and ATLAS for HL-LHC [174–176] since CO₂ evaporates under high pressure, enabling small-diameter evaporator tubes that minimize the material in the detector. The titanium cooling pipes are integrated into the functional local support and carry boiling CO₂ to cool the FEs and the sensor of a quad module.

Cooling relies on the two-phase Accumulator Controlled Loop (2PACL) [177] thermodynamic cycle. Figure 4.20(a) shows a schematic of a cooling plant using the 2PACL cycle, with its pressure-enthalpy diagram shown in Fig. 4.20(b). Detector cooling is achieved through evaporation of subcooled liquid CO₂. When heat is absorbed, CO₂ can change from its liquid to gaseous phase. In the two-phase state, CO₂ temperature is strictly



(a)



(b)

Figure 4.20: (a) Schematic of a cooling plant using the 2PACL principle. Not shown here is a chiller, that would be attached to the condenser to provide cooling power. The evaporator goes through the detector and cools the experiment. (b) The 2PACL cycle in the pressure-enthalpy diagram of CO_2 . The nodes in the diagram correspond to the locations in the schematic. (Reproduced from Ref. [177], licensed under CC BY 3.0)

dependent on pressure, as indicated by isothermal lines in Fig. 4.20(b), though the diagram neglects the small pressure drop along the cooling tube (from node 4 to 5) causing a slight CO_2 temperature decrease in the flow direction. Critically, CO_2 and tube temperatures align only when evaporation starts. Beyond a critical fraction of vapor to liquid, liquid detaches from the tube wall, which degrades heat transfer and triggers a rapid increase in tube temperature. Tube temperature then starts to rise drastically the further to the right one moves in the pressure-enthalpy diagram. This phenomenon is called dry-out.

To prevent dry-out and protect the detector against sudden cooling failure, each SP-chain routes the NTC of the last module in the cooling loop sequence to the interlock instead of the MOPS chip. The interlock triggers if the temperature of the module exceeds 40°C .

ITk Pixel System Tests and the Outer Barrel Demonstrator Program

Constructing a large detector requires meticulous planning and rigorous verification of every component destined for installation in the cavern. The ITk upgrade project is organized into functional work packages, with each component undergoing a defined verification process prior to production [178]. The Preliminary Design Review (PDR), Final Design Review (FDR) and Production Readiness Review (PRR) milestones of this verification process are illustrated in Fig. 5.1.

The OB system test with the RD53A demonstrator program supported the FDR of the OB LLSs, the smallest fully functional detector units used to construct the final detector. The goal of the system test was to evaluate the performance of the modules after loading and integration to verify the design of the LLSs. Many tests were performed within the scope of this thesis, and the test results are presented in Chapter 6. This chapter provides an overview of the RD53A demonstrator and the setup for the system tests.




Figure 5.1: Milestones of a verification process. Shown are the PDR, FDR and PRR. Only after passing a formal review, the activities within a work package can continue to the next phase.

The goal of the system tests (for all three ITk Pixel subsystems OB, EC and IS) was design validation and verification of the correct interplay between different detector components. Within the RD53A program, *all* types of LLSs of the OB and EC subsystems were built and tested. These setups provided the first opportunity to test a complete detector system and demonstrate that detector units consistently fulfill the requirements. Since different subsystems employ different approaches for loading modules onto local supports, the system tests also validated these procedures and established their equivalence. The system tests included real services and cables, production-type power supplies and a DCS, with setups following the grounding and shielding scheme planned for the final detector. The system tests verified the following:

- Verification of the readout path
- Verification of the the powering scheme (LV and HV)
- Testing the serial powering scheme
- Testing modules on LLSs
- Validation of the services
- DCS functionality checks
- Operation under realistic conditions
- Cross-talk checks between modules
- Source tests of modules with radioactive sources for signal generation

Specifically, the OB system test was conducted using the *OB demonstrator* in the [radiation lab](#) in the SR1 building at CERN, where testing occurred from summer 2022 until summer 2023. A Computer-Aided Design (CAD) model of the originally planned demonstrator is shown in [Fig. 5.2](#). However, due to the limited availability of prototype RD53A quad modules, only one half-filled longeron and one half-filled IHR were built. As the tests in [Chapter 6](#) focus on the longeron, the following section describes its configuration and its modules.




Figure 5.2: CAD model of the RD53A demonstrator as it was originally planned. The plan foresaw a half-filled longeron (layer 3 functional local support and layer 2 electrical services) and two half-filled layer 3 IHRs. However, due to a shortage of modules, only one half-filled IHR in addition to the half-filled longeron was built.

5.1 The test structure

In total, 18 non-irradiated RD53A quad modules were loaded and integrated on the A-side of the longeron in two SP-chains: a central, smaller SP-chain (A-SP-1, also referred to as the $M6$) with 6 modules, and a larger SP-chain (A-SP-2, also referred to as the $M12$) with 12 modules. The module positions on the longeron are shown in Fig. 5.3, which also illustrates the SP-chain current flow and HV group separation as defined by the PP0 for all OB subsystem longerons. Figure 5.4 shows the longeron in the test box. The modules used in A-SP-1 are listed in Table 5.1, and those in A-SP-2 are listed in Table 5.2.




Figure 5.3: Serially powered (SP) chains of the OB longeron. The central SP-chains (C-SP-1 and A-SP-1) feature six modules each, while the longer chains at higher $|z|$ (C-SP-2 and A-SP-2) comprise twelve modules each. Indicated is also the current flow of LV for each SP-chain, the HV groups and the names of the modules as used in the DCS during testing. (Adapted from Ref. [162])




Figure 5.4: (a) Photograph of the demonstrator longeron half-filled with RD53A quad modules in two SP-chains (A-SP-1 and A-SP-2). (b) Photograph from below of the longeron while only the short SP-chain (A-SP-1) was mounted. The MOPS chip is mounted on the PP0 facing inwards and visible just to the right of the middle pigtail. (c) Photograph of the top of the PP0 showing the power and DCS connection with the golden cable, as well as the two data connectors with the blue twinax cables.

Table 5.1: Modules of the M6 SP-chain (A-SP-1) on the demonstrator longeron.

Position in Longeron	Name in DCS	Module	HV channel	Position in SP-chain	Notes
A-M-B-1	M01	ThinQ10	HV2	5	
A-M-T-1	M02	Paris10	HV2	4	FE4 non-functional
A-M-B-2	M03	Paris3	HV2	6	not compatible with FELIX readout (replaced by ThinQ9 later)
A-M-B-2	M03	ThinQ9	HV2	6	FE4 non-functional
A-M-T-2	M04	SiegenQ1	HV1	3	
A-M-B-3	M05	Paris12	HV1	1	FE4 non-functional
A-M-T-3	M06	SiegenQ2	HV1	2	

Table 5.2: Modules of the M12 SP-chain (A-SP-2) on the demonstrator longeron.

Position in Longeron	Name in DCS	Module	HV channel	Position in SP-chain	Notes
A-M-B-4	M01	SiegenQ3	HV2	8	
A-M-T-4	M02	ThinQ11	HV2	7	
A-M-B-5	M03	Paris9	HV2	9	
A-M-T-5	M04	SiegenQ4	HV1	6	
A-M-B-6	M05	Paris7	HV2	10	
A-M-T-6	M06	Paris13	HV1	5	
A-M-B-7	M07	Goe10	HV2	11	
A-M-T-7	M08	ThinQ8	HV1	4	FE4 non-functional
A-M-B-8	M09	Goe4	HV2	12	
A-M-T-8	M10	ThinQ4	HV1	3	
A-M-B-9	M11	KEKQ20	HV1	1	
A-M-T-9	M12	Liv5	HV1	2	

The longeron in [Fig. 5.4](#) uses a layer 3 functional local support, but with layer 2 electrical services, meaning that only FE2 and FE4 are electrically connected to the readout system. [Figure 5.4\(a\)](#) also shows that only two modules (bottom right and top second-to-right) had mechanical wire-bond protection installed to protect the wire-bonds from the pigtail. The other modules had their wire-bonds potted with SYLGARD™ for protection.

The use of prototype RD53A modules limits the demonstrator's informative value compared to a final detector with production modules. Specifically, the RD53A quad modules use half-size RD53A chips:

- with connections only at the ends of the sensor, leaving the middle of the sensor connected to additional, non-powered chips, which affects power consumption and power dissipation;
- without data-merging support, restricting the readout to FE2 and FE4 due to layer 2 electrical services;
- with three AFEs variants (synchronous, linear, differential). But multi-module readout for the synchronous FE (dropped by CMS and ATLAS) lacked robust software support;
- with ShuntLDOs exhibiting a known cold start-up issue;
- without low-power mode support.

Additionally, schedule constraints and limited module availability required selecting modules for integration onto a functional local support that did not meet all quality criteria. Some

modules were already known to have non-functional FEs or bad HV behavior. This information is also noted in [Table 5.1](#) and [Table 5.2](#).

5.2 Setup

Operating the OB RD53A demonstrator required a complex and sophisticated setup. [Figure 5.5](#) shows a schematic diagram of the test infrastructure in the radiation lab in SR1. The demonstrator was installed in a metallic, lightproof test box, which was continuously flushed with dry air to maintain a dew point of -60°C . The metal pieces of the box were connected to the building’s electrical GND acting as a Faraday cage for the detector. Feedthroughs accommodated cooling pipes, cables for environmental sensors (light, temperature and humidity), and the Type-1 power bundle (LV, HV, MOPS connection and Temperature Interlock (Tilock)). Temperature and pressure sensors monitored the inlet and outlet of the CO_2 cooling; a flowmeter was installed in the CO_2 inlet.

The Type-1 power bundle terminated at a custom-designed Cable Saver Board (CSB), a PCB that replaced PP1 and PP2 while adhering to the ITk grounding scheme. From the CSB, 60 m long cables connected to the PSUs in the rack area of SR1, replicating closely the real off-detector service chain’s functionality as specified in Ref. [\[163\]](#) and previously shown in [Fig. 4.17](#).

For the control of the power supplies, the monitoring of the system, including all infrastructure components, and the storage of the monitoring data into the CERN-wide ORACLE¹ database, a DCS based on SIMATIC WinCC Open Architecture (WinCC OA) was developed [\[179\]](#). This DCS system also forms the foundation of the developments presented in [Chapter 9](#).

¹A relational database management system, <https://www.oracle.com/database/>

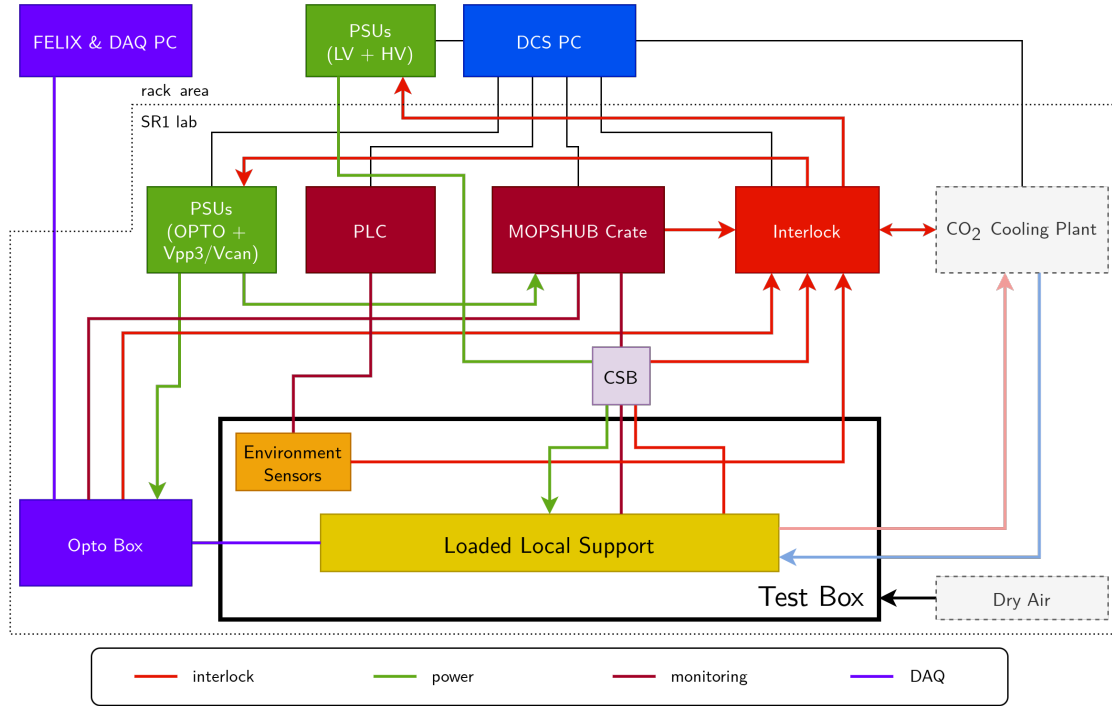


Figure 5.5: Overview of the demonstrator setup in SR1. Shown are the various types of connections between the setup infrastructure and the LLS under test.

5.3 Powering

The demonstrator was powered by large, production-type power supplies installed in the SR1 rack area. While these were not the final ITk power supplies, there were either pre-series models or very similar devices. The power supplies connected via 65 m long Type-3 cables to the CSB in the radiation lab:

- LV: Wiener PL512² prototype current source with 4 channels (up to 10 A and 50 V) for FE power (nominal: 4.8 A)
- HV: iseg HV modules³ (type EHS F607n-F; up to 700 V) for sensor bias, mounted in a Wiener MPOD Mini Crate⁴

Within the radiation lab:

- OPTO: Wiener PL512 voltage source for powering the Opto System (nominal: 9 V)

²<https://www.wiener-d.com/product/pl512-power-supply-system/>

³<https://iseg-hv.com/en/products/detail/EHS>

⁴<https://www.wiener-d.com/product/mpod-mini-crate/>

- MOPSSYS: Wiener LV module⁵ in a Wiener MPOD Mini Crate to power the MOPSHUB4Beginners (nominal: 15 V)

5.4 MOPSHUB4Beginners

To monitor modules on the demonstrator via the MOPS chip (as described in [Section 4.5.2](#)) and to monitor the Opto System (with one MOPS per Opto Box), an interim solution, the MOPSHUB4Beginners [\[180\]](#) (also abbreviated as MHFB), was used, which replaced the MOPSHUB foreseen in the final design. The MOPSHUB4Beginners is shown in [Fig. 5.6](#). This system comprises a Raspberry Pi that connects via a CANbus to multiple MOPS chips to receive data. For DCS readout, the Raspberry Pi is connected to the LAN via Ethernet and runs an Open Platform Communications Unified Architecture (OPC-UA) client. The client updates MOPS data in an OPC-UA server that was developed using the quasar framework⁶. From there, the information can be read by WinCC OA for monitoring and archiving.

⁵<https://www.wiener-d.com/product/mpod-lv-module/>

⁶<https://github.com/quasar-team/quasar>




Figure 5.6: The MOPSHUB4Beginners system installed in a rack next to the test box in the radiation lab in SR1.

5.5 Interlock and environment monitoring

The interlock serves as the final safeguard against overheating of hardware or other hazardous operating conditions. This can, for example, be an open test box, light exposure, high humidity, or a cooling system failure. It triggers protective measures when monitored values exceed predefined thresholds and operates independently of other systems. The main element of the interlock system was the IBL Interlock Matrix Crate (IMC) [181], populated with modules as shown in Fig. 5.7. An FPGA implemented the interlock logic that defines the systems responses for given inputs, with input cards adapted for various kinds of sensors and output cards transmitting an interlock signal to power supplies or the cooling system. The interlock logic mapped input signals to corresponding outputs, typically implementing logical OR connections between multiple inputs per output. The entire logic, i.e. which inputs act on which outputs, could therefore be represented by a matrix with matrix elements representing a connection.

Input signals included one module NTC per SP-chain, NTCs in the Opto Box, switches on the test box (to check for open doors), light sensors, two Vaisala DMT143L dew point sensors⁷, a signal from the cooling system and an emergency stop button.

Although hardware-based and completely independent from other systems, the interlock system communicated its actions to the DCS via ELMBs, which monitored the incoming and outgoing signals and reported them to the control system. All monitoring signals

⁷<https://www.vaisala.com/en/products/instruments-sensors-and-other-measurement-devices/instruments-industrial-measurements/dmt143>




Figure 5.7: The Interlock Matrix Crate (IMC) in the SR1 setup for use with the OB demonstrator. Behind the plate labeled “IMC C” is the FPGA. Different interlock cards are used for different input and output signals. The signals are monitored by Embedded Local Monitor Boards (ELMBs), which are all connected via a CANbus.

are latched, meaning the information about a value exceeding a threshold is kept until manually reset, to allow the user to reconstruct the behavior of the system even in case of very short interlock signals. Critically, all interlock signals followed negative safety logic where 0 V indicates an active interlock signal, meaning that in case of a power loss or a broken cable, the interlock becomes active, preventing the powering of a channel under unsafe conditions.

While the IMC was capable of reading many different sensor types, that not necessarily needed to feed into an interlock decision, it was decided to extend the monitoring capabilities by using a Siemens S7-1500 PLC to take over non-critical environment monitoring functionality [182]. In the system test, the PLC read the temperature and pressure sensors at the CO₂ cooling inlet and outlet as well as the flowmeter at the inlet. Additionally, there was a flowmeter for the dry air flow into the test box, which was read by the PLC. While the system test demonstrated the monitoring capabilities of the PLC, a similar system would be installed in the LLS QC setup (see [Section 7.1](#)), where it would also take over the interlock functionality.

5.6 CO₂ cooling

The demonstrator utilized the ATLAS SR1 CO₂ cooling station for detector upgrades. The cooling plant has an operating range from 18 °C down to approximately −28 °C and was previously used for the commissioning of the IBL detector. Like the CO₂ cooling system for the IBL detector [52], its primary refrigeration system responsible for cooling the CO₂ is based on chillers using a hydrofluorocarbon mixture (R-404A) as the refrigerant. However, due to insufficient cooling power, it cannot achieve the operational temperature of −40 °C required for the final ITk CO₂ cooling system [176].

5.7 Readout with the Opto System and FELIX

The demonstrator validated the optical readout path as planned for the final detector. As shown in Fig. 5.8, a realistic module readout chain was realized in the demonstrator comprising 3 m long commercial twinax cables connecting the PP0 to Opto Boards within an Opto Box. The Opto Box was housed in a small metal enclosure that connected to the test box via a flexible metal tube to extend the Faraday cage of the test box.

The Opto System converts the electrical signals from the detector to optical signals, which are less susceptible to degradation, before transmitting the data, in the final detector, to the ATLAS electronic rooms. To minimize signal loss in the electrical path, this conversion occurs as close as possible to the pixel modules. In the final detector, the Opto System will reside just outside the ITk detector volume with 6 m long electrical cables [183].

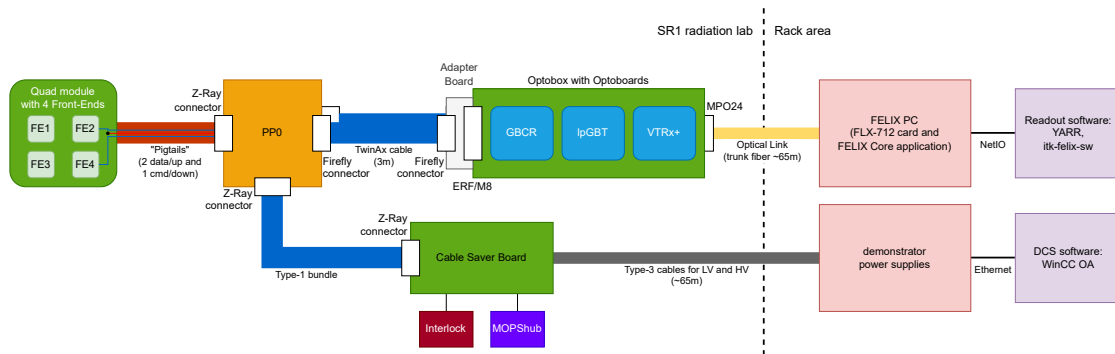


Figure 5.8: Module readout and powering path in the OB RD53A demonstrator. The readout path is shown on top with a commercial twinax cable. The powering path in the bottom shows the CSB which replaces PP1 and PP2 boxes.

The hardware of the Opto Box for the demonstrator closely followed the final design supporting up to eight Opto Boards. In the demonstrator configuration, the box housed seven Opto Boards. Each board contains the following set of data transmission ASICs:

- GBCRs [184]: up to 4 GBCRs receive electrical signals from the modules via so-called e-links, restoring signal integrity after attenuation in the twinax cable.
- lpGBTs⁸ [185]: serializes data from multiple e-links after passing through a GBCR with an output data rate up to 10.24 Gbit/s. One *master* lpGBT handles sending commands on the downlink and receiving data on the uplink, while up to 3 additional *secondary* lpGBTs are available for uplinks only.
- VTRx⁺ [186, 187]: one VTRx⁺ finally converts the electrical signals to optical signals.

Critically, e-links are AC-coupled to the GBCR, due to varying module voltage levels in an SP-chain. If Opto Boards are not powered while the SP-chain is on, coupling capacitors can charge the Opto Board inputs through their finite leakage resistance. Thus, Opto Boards must be powered before SP-chains.

The power delivery concept for the Opto Box with its Opto Boards uses a two-stage DC/DC converter⁹ system. Five 9 V supply lines enter the Opto Box, stepped down to 2.5 V by bPOL12Vs [188] on a power board. A connector board distributes the five 2.5 V power lines to up to eight Opto Boards in various possible configurations powering the VTRx⁺ drivers and the next stage of DC/DC converters on the Opto Boards. A bPOL2V5 [189] on each Opto Board converts the 2.5 V to 1.2 V for the ASICs on the Opto Board. The connector board enables granularity on the SP-chain level: one powering channel into the Opto Box controls all Opto Boards belonging to one SP-chain.

A 65 m long optical trunk fiber carried the signal from the VTRx⁺ to a Phase-I FELIX card (as introduced in Section 3.4.1) in a FELIX computer in the SR1 rack area. The card was a FLX-712 card¹⁰ running specific firmware for RD53A data.

The readout software ran directly on the FELIX computer. Due to the early stage of FELIX-based readout development, no single software package for data readout supported all electrical scans across all different AFE variants. Consequently, multiple packages were used:

- YARR [190] (support for linear and differential FE)
- VakYARR (support for synchronous FE, but only readout of one chip at a time)
- itk-felix-sw (support for linear and differential FE)

⁸<https://ep-esf.web.cern.ch/project/lpgbt-and-versatile-link>

⁹<https://power-distribution.web.cern.ch/ASICS/>

¹⁰https://atlas-project-felix.web.cern.ch/public/docs/hardware/manuals/FLX-712_manual.pdf

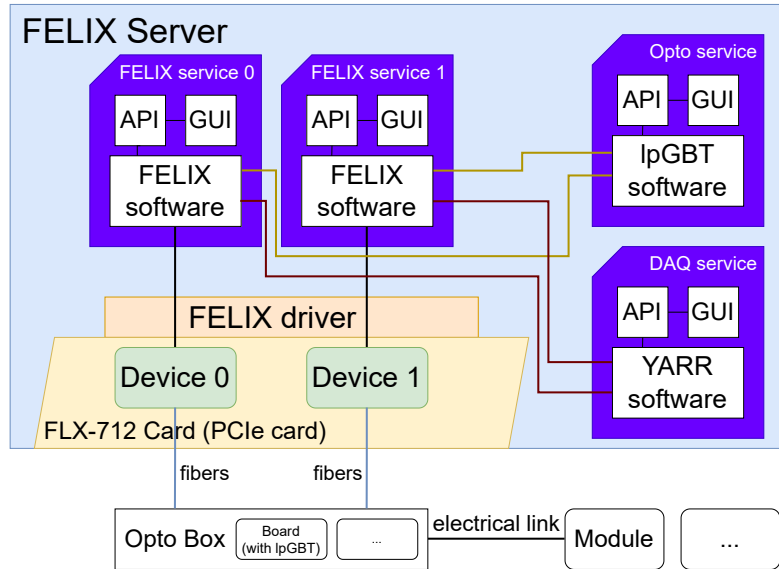


Figure 5.9: The online software framework based on a microservice architecture as used for the demonstrator tests.

For orchestrating the online software components, development of a framework based on the microservice architecture started with the demonstrator program. Each service (for FELIX configuration, Opto Board configuration and readout software) is packaged and deployed independently using immutable Docker¹¹ containers that include all necessary dependencies¹². Figure 5.9 shows the microservice architecture required for reading out modules in the demonstrator. The FELIX software configures the FLX-712 card, prepares it for readout of RD53A modules, and enables communication with downstream hardware. Additionally, software is required to configure the IpGBTs on the Opto Boards (for example setting correct e-link polarity). All services provide a REST API to enable programmatic access to the software components and support automation. The development targets widespread deployment during production LLS testing at multiple sites.

¹¹<https://www.docker.com/>

¹²Containers are sandboxed processes on a host machine, isolated from the rest of the machine.

Results from the Outer Barrel Demonstrator

This chapter presents the results from the tests with the RD53A demonstrator, as described in [Chapter 5](#). This includes initial commissioning tests of the setup, presented in [Section 6.1](#), which validate the interlock and MOPSHUB system to verify the detector infrastructure, which forms together with the detector the detector system. Subsequently, the performance studies with the demonstrator longeron are presented in [Section 6.2](#). A demonstrator IHR was also studied, with the results presented in Ref. [191]. These performance studies include sensor and FE tests. Comparisons with results from previous assembly stages successfully validate the detector design as presented in [Section 4.5](#) by demonstrating no degradation of detector components between stages. The results presented in this chapter are also included in the supporting documentation [162] for the Final Design Review (FDR) of the ITk Pixel Loaded Local Supports.

6.1 Setup commissioning

The system test setup used to study the demonstrator was already described in [Section 5.2](#). It encompasses significant infrastructure to test the detector within a complete detector system close to the final design. This includes a DCS, production-grade power supplies, long off-detector cables (as used in the final system), an interlock system, an environmental monitoring system, CO₂ cooling, a monitoring system for the on-detector components using MOPS chips (see [Section 5.4](#)), a data transmission system approaching the final design by employing an Opto Box, and a DAQ system with FELIX and readout software. The system test also follows the detector grounding scheme to prove its proper functioning.

In order to test the interplay between different components of the detector system, a good understanding of each component and its behavior, also in isolation, is essential. Significant effort was put into setup commissioning and characterizing the infrastructure before installing the demonstrator. The following sections present results from infrastructure tests and characterization.

6.1.1 Detector Control System software and powering

The system test provided a test bed for DCS developments even already before the installation of the demonstrator, which are part of the developments presented in [Section 9.3](#). During the demonstrator tests, the DCS was in a stable state and was used to control the power supplies as described in [Section 5.3](#) and monitor the setup. For environment monitoring, the Interlock Matrix Crate (IMC) and the PLC as described in [Section 5.5](#) continuously sent their sensor readout data to the DCS software. Similarly, the Opto Box monitoring data and the on-detector monitoring data from MOPS chip readout with a MOPSHUB4Beginners was received by the DCS software. The DCS was configured to archive all relevant DCS parameters, that is, all monitoring data and control data for the power supplies into the CERN ORACLE database.

To reduce the amount of data stored permanently, smoothing procedures were configured for each parameter. Value smoothing was applied, meaning that new data for a parameter was archived only if it changed by a certain minimum amount from the last stored value. Additionally, time smoothing was applied, meaning that even if the data did not change significantly to trigger archival, it was archived only if it remained unchanged already for a long period, for example for 1 h. The data stored in the ORACLE database can be retrieved by WinCC OA and displayed in its trend panels. To test the archival procedure, the DCS was configured to archive all relevant parameters and ran for a short time. After a restart, it was checked that the previously archived data could be retrieved by the WinCC OA trend panels.

It was also possible to retrieve archived data directly from the ORACLE database using a normal, authenticated database connection and Structured Query Language (SQL) queries. This could be done with Python using the official Python extension module from ORACLE¹. For some diagrams, the data was retrieved from the database using this approach and therefore contains smoothed data. This is indicated in the respective plots. Some tests conducted from the DCS (for example I(V) scans to obtain IV-curves of sensors for performance evaluation) use the live data directly available in the DCS software.

¹<https://pypi.org/project/oracledb/>

For powering the RD53A quad modules in the demonstrator, a prototype PL512 current source was used (also called LV power). This prototype received a firmware update before the installation of the demonstrator to fix an issue with incorrect output voltage and current monitoring when a channel was switched off. The prototype had 4 output channels, but the first channel showed large voltage fluctuation of ± 100 mV around its central value during testing with a high-power resistor. Since at maximum 2 channels were needed at a time, it was decided to always use the last two powering channels of this PL512 prototype.

For sensor biasing, two relatively old iseg HV modules were initially available (one for operation and one as a spare), but one module exhibited non-working voltage monitoring and was therefore not used in the demonstrator tests. The other iseg HV module, used later in the demonstrator tests, was connected via a bare PP0 to a test module (Thick9) to perform an $I(V)$ scan and record the IV characteristic of the sensor (see Section 4.2.2). The result was compared to the measurement obtained using a Keithley 2410 source meter as the HV power supply. The Keithley 2410 states a measurement accuracy of $0.029\% \cdot I_{\text{Out}} + 300$ pA in its manual², while for the iseg HV module the accuracy³ is at least $\pm(0.05\% \cdot I_{\text{Out}} + 8$ μ A) guaranteed for one year, which is significantly worse. Using an adapter on the CSB (see the “powering path” in Fig. 5.8), the Keithley was connected to supply the HV. A series of measurements was manually taken at output voltages from -5 to -165 V in steps of 10 V and the measured output current was recorded. The comparison of the two $I(V)$ curves is depicted in Fig. 6.1. The graph shows that the iseg HV module lacks good current monitoring capabilities at small output voltages. Additionally, the figure suggests a systematic offset in the iseg HV module’s current monitoring, reporting values that are shifted to higher levels by about 0.075 μ A.

6.1.2 Interlock

The IMC, as described in Section 5.5, was configured to trigger an interlock signal for high module temperatures greater than 40 °C and high Opto Box temperatures greater than 30 °C to protect the hardware. The modules and Opto Box temperatures were monitored with “interlock” NTC thermistors. Both NTC types, for the modules (NTCG103JF103FTDS) and Opto Box (NTCG163JF103FT1S), share the same B constant from the manufacturers data to convert from their resistance R to temperature T :

$$\frac{1}{T} = \frac{1}{T_0} + \frac{1}{B} \ln \left(\frac{R}{R_0} \right) , \quad (6.1)$$

²https://download.tek.com/datasheet/1KW-2798-4_2400_SourceMeter_SMU_Datasheet_100324.pdf

³<https://iseg-hv.com/files/iseg-high-voltage-power-supplies.pdf>




Figure 6.1: Comparison of the $I(V)$ curves of the Thick9 module obtained with the iseg HV module and the Keithley 2410 for power supply validation. For reference, an additional curve of the Keithley 2410 result but shifted by $0.075 \mu\text{A}$ is also displayed, showing strong agreement with the iseg HV result.

where $T_0 = 25^\circ\text{C}$ and $R_0 = 10 \text{ k}\Omega$. To achieve more accurate temperature measurements, especially at lower temperatures, Steinhart-Hart coefficients for these NTCs were fitted in Ref. [192]. The Steinhart-Hart equation [193] as given below provides higher accuracy at lower temperatures:

$$\frac{1}{T} = A + B \ln R + C(\ln R)^3, \quad (6.2)$$

where T is the temperature (in K), R is the NTC resistance at T (in Ω), and A , B and C are the Steinhart-Hart coefficients characteristic for a semiconductor material.

The fitted parameters A , B , and C were implemented in the IMC for its temperature monitoring. The thresholds for the interlock logic were tested with a potentiometer. The found thresholds are given in Table 6.1, where each resistance value triggered the interlock, but slightly lower values (within the accuracy of the potentiometer) did not. This validated the interlock function of the IMC.

It should be noted that the interlock module on the longeron PP0s can be selected via switches, as the CO_2 flow direction and side of the cooling connection determine which of the four corner modules in an SP-chain serves as the interlock module (see also Section 4.5.3). The longeron demonstrator was tested in two stages: first with only the M6 installed and later with both the M6 and M12. For the second stage, the CO_2 flow changed direction and consequently a different interlock module was selected in the M6 than before.

Table 6.1: Approximate thresholds for the interlock logic of the IMC. Interlock inputs of type NTC were tested using a potentiometer and the found thresholds are given here.

interlock input	threshold for interlock	corresponding NTC temperature
interlock module 1	5.63 k Ω	41.0 °C
interlock module 2	5.55 k Ω	41.4 °C
interlock module 3	5.16 k Ω	43.5 °C
interlock module 4	5.28 k Ω	42.8 °C
interlock Opto Box 1	8.23 k Ω	30.3 °C
interlock Opto Box 2	8.16 k Ω	30.5 °C

6.1.3 Test box and dry air flush test

The environment was continuously monitored by the IMC and PLC. To prevent ice formation on sensitive detector hardware should the air become too humid, the dew point measurement serves as an interlock input and the CO₂ cooling plant is stopped if the dew point exceeds -40°C . The large-volume test box is always flushed with dry air to keep the dew point low. In a dry air flush test, the time required to reach dew points below -40°C and -59°C was measured. Figure 6.2 shows the dry air flow rate and box dew point as a function of time during this test. From the data it can be seen that with a constant dry air flow of around 540 L/min, it takes approximately 50 min to reach a dew point of -40°C and approximately 265 min to reach a dew point of -59°C . Based on these results, to avoid wait times for drying out the test box, the test box was kept closed to with a constant dry air flow at all times and during long breaks only the CO₂ plant was set to warm.

6.1.4 CO₂ cooling

After installing an LLS in the test box and connecting the CO₂ cooling pipes, a pressure test was always performed first to check for leaks. During this test, the CO₂ cooling pipe was isolated from the cooling plant and pressurized to approximately 50 bar. The pressure was then monitored over an extended period of time to verify stability and confirm the absence of leaks.

Cooling down to the “warm” operating temperature of 10°C was fast. Figure 6.3 shows the CO₂ temperature in the cooling pipe during cooldown to 10°C as a function of time. The data shows that cooling down required less than 10 min. During stable operation, the pressure at the CO₂ inlet inside the cooling pipe remained at around 45 bar.




Figure 6.2: Dry out test with the test box. The flow of dry air was turned on at 446 s. The small bumps in dew point measurement come from the autocalibration procedure of the dew point sensor. The data shows that with a constant dry air flow of around 540 L/min, it takes approximately 50 min until a dew point of -40°C and approximately 265 min until a dew point of -59°C is reached. These measurements were retrieved directly from the ORACLE database and represent smoothed data.




Figure 6.3: CO_2 cooling temperature over time when setting the target temperature to 10°C for normal operation. The cool down takes less than 10 min. These measurements were retrieved directly from the ORACLE database and represent smoothed data with different smoothing parameters for CO_2 and module temperatures.

6.1.5 Validation of the MOPS monitoring system

After the installation of the demonstrator and a pressure test, the electrical services were connected and a simple connection test was performed. It was then verified whether the MOPS chips on the PP0s could be successfully read by the MHFB (see [Section 5.4](#)). The MOPS chips monitor module temperatures and voltage drops across a module, thereby providing useful information about the operational status of the demonstrator. The MOPS monitoring system is completely independent of the rest of the LLS: It is powered via dedicated power lines and read out by a separate readout system.

The correct functioning of the MHFB could be confirmed already before the installation of the demonstrator. On the MHFB, jumpers were used to select a CAN polarity. Using a bare PP0 with bridging PCBs (which include mounted NTCs) installed on the module connectors, the communication of the MHFB with the MOPS on the PP0 (see also [Section 4.5.2](#)) could be verified. For functioning and accurate MOPSV1 readout (installed on both the M6 and M12), knowledge of the MOPS CAN transceivers' bit rate (this chip version did not support automatic trimming), internal ADC reference voltage, external regulated 1V2 output for resistance measurements, voltage divider resistor values and, for resistance-to-temperature conversion, Steinhart-Hart coefficients of NTCs is essential. While voltage divider resistor values can be taken from PP0 schematics, the ADC reference voltage required calibration due to variations from the manufacturing process. The ADC (together with the NTC) was calibrated as reported in Ref. [\[192\]](#).

When connecting the demonstrator, the MOPS chips communicated successfully with the MHFB and data could be read. [Figure 6.4](#) illustrates the monitoring data for the M12 longeron demonstrator operated at a CO₂ set point of 17°C. As shown in the upper plot ([Fig. 6.4\(a\)](#)), module temperatures rose when the SP-chain was turned on (time: 18:11:50). The slightly higher temperature of the orange line corresponds to module Paris9, which had a known issue with its thermal interface. For the M6, it was observed that the connection to the NTC of the SiegenQ1 module (M04) was broken and it was not possible to obtain temperature measurements from this module. The lower plot ([Fig. 6.4\(b\)](#)) shows the monitored input voltage to the SP-chain modules. The expected ladder of voltages with equal voltage steps between the different modules formed as soon as the chain was turned on. The voltage fluctuations observed before time 18:11:50 are an effect of small leakage currents during the off state of the prototype current source power supply. This property builds up a small voltage on the ShuntLDOs of the modules in the SP-chain even when the power supply is off.

Additionally, a quick “temperature scan” was performed with the M6 to validate the MOPS readings at different temperatures. [Figure 6.5](#) shows the evolution of the module and

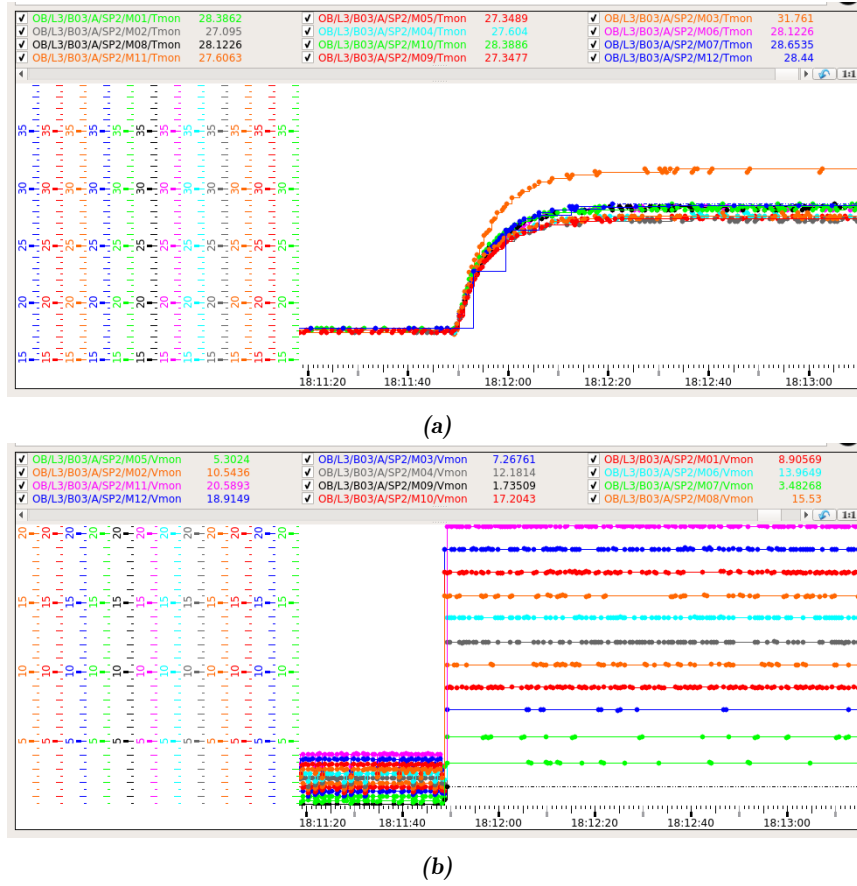


Figure 6.4: Validation of the MOPS monitoring system on the M12 SP-chain. The graph in (a) shows the temperature monitoring when turning on LV. The orange line with higher temperature belongs to the Paris9 module (M03), which had a known issue with its thermal interface. The graph in (b) shows the voltage monitoring with the characteristic ladder of monitored module voltages when turning on LV in the SP-chain. When LV is off, a small leakage current from the prototype current source is visible as voltages do not return to zero.

CO₂ temperature during the temperature scan. The minimum temperature of -28°C was reached after approximately 80 min. The MHFB successfully monitored the SP-chain at all temperatures and regardless of the module status. Notably, temperature monitoring exhibited large fluctuations especially at cold temperatures when modules were first off and then turned on. This prompted more detailed investigations, presented in [Section 6.1.6](#).

6.1.6 Fluctuations in the module temperature monitoring

During the qualification of the M6 and M12, fluctuations in the module temperature measurements up to 10 K on a time scale of a second have been observed. An example for




Figure 6.5: Module and CO₂ temperature during the temperature scan with the M6. The target CO₂ temperature was changed in steps of 10 K from 10 °C to −30 °C, but only −28 °C were reached. During cool down, the SP-chain was kept switched on. While warming up again, the modules were in each step first turned off and then turned on again. The module temperatures were measured by the MOPS chip, except for the M05, where the temperature was measured by the interlock system. These measurements were retrieved directly from the ORACLE database and represent smoothed data.

this behavior is shown in Fig. 6.6, which shows fluctuations in the temperature monitoring. The time constant of the fluctuation is clearly shorter than that of the thermalization of the module, as can also be observed in the same plot. For that reason, it is assumed that the fluctuations are not real temperature variations but are instead caused by an unknown mechanism.

Several observations suggest that the effect is correlated with serial powering:

- Rapid power-cycling of the SP-chain can, in some cases, eliminate the fluctuations (see Fig. 6.6).
- The fluctuations disappear when the SP-chain is turned off.
- No fluctuations are visible on the LV supply current when connecting a current probe to the LV cable.
- The temperature fluctuations can be easily reproduced when operating the SP-chain at low supply currents (for example 3 A instead of 4.8 A).
- At low temperatures, the fluctuations cannot be eliminated, even with rapid power-cycling.




Figure 6.6: Observed fluctuations in the monitored temperature of the modules in the M12. Shown are the temperatures of modules M03 and M09 at a CO₂ set point of 10 °C as seen live in the DCS. Additionally, the LV power channel voltage output is plotted to see the times when LV was on or off.

- The fluctuations are visible in both the MOPS monitoring data and the interlock system.
- When performing a threshold scan (see also [Section 6.2.3](#)) and comparing noise distributions between the fluctuation state and the no-fluctuation state, no differences are visible.

It can be excluded that the effect arises from a problem in the MHFB in the system test setup, as this component was exchanged for test reasons and the fluctuations are also visible in the interlock system.

It should be noted again, that the ShuntLDO in the RD53A FEs has known start-up problems at low temperatures. No satisfactory explanation for the observed fluctuations with the RD53A modules was found. However, a test setup using a MOPSV2 and ITkPix-V1.1 modules did not exhibit this behavior, and it was also not observed in the preproduction longeron at CERN (see [Section 7.2](#)).

6.1.7 Validation of the grounding and shielding scheme

As part of the system-level validation, the distribution of the reference ground to the LLS has also been tested in the demonstrator setup. In the referencing scheme of the OB detector, the shields of the data cable DC connect the global GND to the 0V-plane of the PP0 (see also [Fig. A.1](#)). The shields of the CAN cable for the MOPS chip provide Alternating Current (AC) coupling of the global GND to the 0V-plane. All power supply channels are floating and the return lines are referenced only on the 0V-plane of the PP0.

Table 6.2: Validation of the DC referencing scheme of the M6 SP-chain (A-SP-1) on the demonstrator longeron. The power supply return lines were probed on the CSB with a digital multimeter.

Measurement (to setup GND)	M6 (data cables connected)	M6 (no data cables connected)
HV1 return	9.6 Ω	no connection
HV2 return	9.6 Ω	no connection
LV return	0.7 Ω	no connection
VCAN return	4.2 Ω	no connection

Even though the final off-detector service chain was not available (the CSB was used to replace the PP1 and PP2 of the OB design), the DC referencing scheme could be validated for the demonstrator setup. For this, probe points on the CSB were used to measure the resistance between the power supply return lines and the setup GND using a digital multimeter. The results of this test are shown in [Table 6.2](#).

The resistance measurements with the data cables connected are consistent with the different wire gauges of the return lines in the Type-1 power bundle. As expected, no DC connection exists when the data cable is disconnected. These results validate the DC referencing scheme of the OB design. Impedance measurements remained pending at the time of the tests until a more realistic off-detector chain became available, including a final version of the data cable.

6.1.8 Module readout with the Opto System

When powering up the modules of an LLS, the FELIX server displays an indication that “decoding alignment” is working for each connected FE chip during good communication. Here, decoding alignment for e-links means that the FE chip receives and phase-locks to the clock from the readout system on the downlink and sends an Aurora-encoded data stream back to the readout system on the uplink, where FELIX successfully aligns the Aurora header in its decoding block.

It was found that the decoding for e-links connected to secondary lpGBTs (see [Section 5.7](#)) was sometimes unstable, with decoding alignment disappearing and reappearing suddenly again. This issue was reported to the responsible firmware developers. Through power-cycling it was possible to reach a stable state, where decoding alignment did not disappear.

If the polarity setting in the lpGBT for the downlink to the module was correct but incorrect for the uplink, decoding alignment would fail. In such cases, Bit Error Rate Tests

(BERTs) could identify the incorrect polarity. BERTs check for errors in the logical path of the data transmission by sending a known bit sequence (pattern) from the FE and checking it at the receiving lpGBT. For this test, FEs are configured to output a Pseudorandom Binary Sequence 7 (PRBS7) pattern [149], and the lpGBT performs the BERT as described in Ref. [194]. Due to finite testing time, results provide an upper limit on the bit error rate at 95 % confidence level. If BERT results show 100 % errors with PRBS7, the lpGBT has incorrect uplink polarity configured. If the results show 50 % errors with PRBS7, data transmission to the lpGBT has completely failed.

BERTs are also used to check for possible cross-talk on the data lines, as detailed in [Section 6.2.5](#).

6.2 Performance studies

This section presents the results of performance tests conducted on modules installed in the longeron (as described in [Section 5.1](#)). Where applicable, results were compared with earlier measurements in previous stages of the assembly to assess potential performance degradation during the cell loading and integration steps. Both the sensor and FE chips, including their connection, were tested to validate the OB detector design.

6.2.1 I(V) measurements

To evaluate sensor performance after integration, I(V) scans (HV scans) were performed. Because multiple modules are connected in parallel to a single HV channel, individual sensor I(V) characteristics cannot be obtained. I(V) curves for all HV groups in the longeron were obtained by varying the HV voltage set point and recording the monitored voltage and current values from the HV power supply. [Figure 6.7](#) shows the I(V) curves for the HV channels of the M6, while [Fig. 6.8](#) displays those for the M12. These curves are compared against the expected leakage current derived from results obtained during module QC. The expected leakage current in each HV group is estimated by summing individual contributions. To account for temperature variations, individual module curves were normalized to the testing temperature in the demonstrator using [Eq. \(4.13\)](#), which works well in the plateau region. Individual I(V) curves were additionally shifted along the x-axis to compensate for module-specific GND offsets caused by serial powering (see also [Fig. 6.4\(b\)](#)). The effective HV reverse bias on each module under serial powering is equal to the sum of the voltage difference coming from the HV PSU and the local module GND offset (see [Section 4.5.1](#)). In both SP-chains, HV channel 2 biases the modules closer to the SP-chain GND. The summed contributions are presented as a stacked histogram.

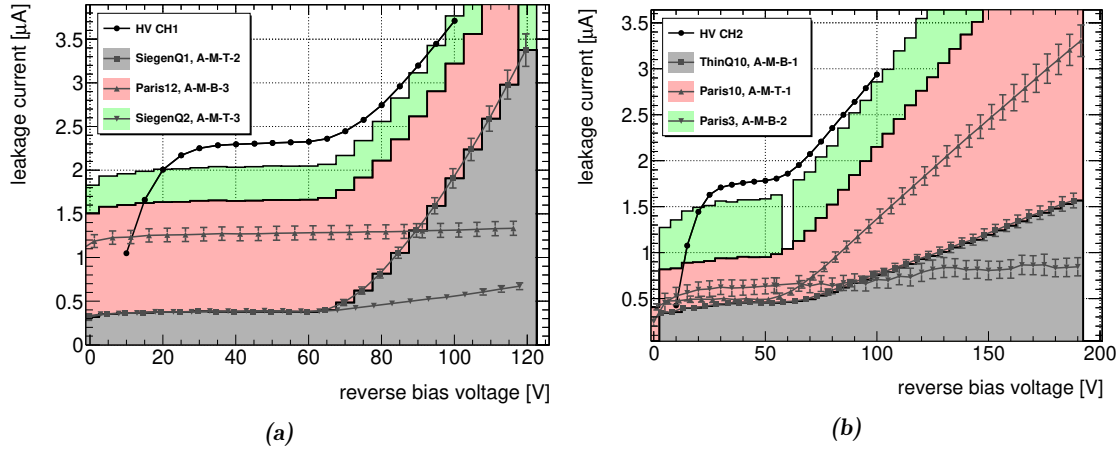


Figure 6.7: $I(V)$ curves obtained for the two HV groups in the M6. Channel 1 is shown in (a), while channel 2 is shown in (b). Shown is the leakage current behavior dependent on the reverse bias voltage for the integrated modules in comparison to the single module behavior before integration. After integration, three modules are biased in parallel by one HV channel. The leakage current of this group of modules is shown with points and compared against the individual $I(V)$ curves of the modules that were taken before integration. The comparison should be made against the (stacked) histogram which shows the expected behavior of the group of modules as a sum of the individual contributions.

In general, both the plateau region and start of breakdown are well reproduced in the demonstrator setup. A distinct plateau is unidentifiable in Fig. 6.8(a) due to the early breakdown voltage of the ThinQ8 module. The current level of the plateau appears to be measured too high by the HV power supply in the demonstrator setup, an effect previously observed (albeit less pronounced) in Section 6.1.1. Nevertheless, no sensor performance degradation is evident in the $I(V)$ curves.

6.2.2 $V(I)$ measurements

Although during normal operation SP-chains are not operated at low currents, $V(I)$ scans (LV scans) were performed with the M6 at different temperatures to obtain the $V(I)$ characteristic of the SP-chain and verify monitoring capabilities through the MOPS chip. The CO_2 cooling temperatures for these tests were 10°C , -10°C and -20°C . In each scanning step, the input current setting to the SP-chain was adjusted, the LV channel turned on and the corresponding input voltages of the modules were recorded by the MOPS chip and the DCS. Figure 6.9 shows the $V(I)$ curves of the M6 at different temperatures. The external ohmic behavior of the ShuntLDOs is visible as a straight line in the voltage monitoring data. Additionally, the percentage of operational FEs is plotted, where an operational FE is defined as one exhibiting successful decoding alignment in the FELIX

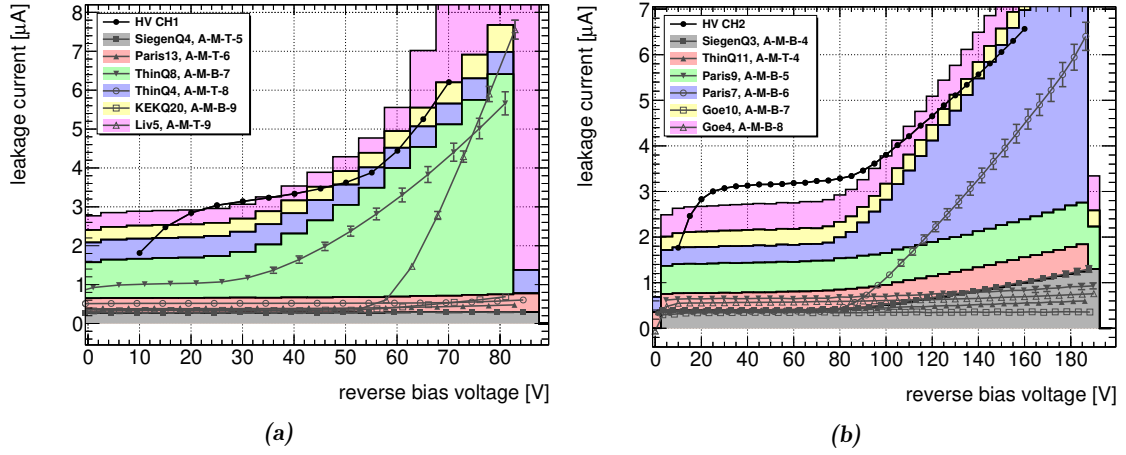


Figure 6.8: $I(V)$ curves obtained for the two HV groups in the M12. Channel 1 is shown in (a), while channel 2 is shown in (b). Shown is the leakage current behavior dependent on the reverse bias voltage for the integrated modules in comparison to the single module behavior before integration. After integration, three modules are biased in parallel by one HV channel. The leakage current of this group of modules is shown with points and compared against the individual $I(V)$ curves of the modules that were taken before integration. The comparison should be made against the (stacked) histogram which shows the expected behavior of the group of modules as a sum of the individual contributions.

system. At 3, 4 and 5 A, the switch-on was performed four times to record a possible random component in the successful startup of a FE.

At lower temperatures, an increasing number of FEs have problems in starting up, which is a known limitation of the ShuntLDO in the RD53A chips. As expected, reducing the input current prevents the ShuntLDOs from generating the necessary operating voltage, causing FEs to stop working. Nevertheless, monitoring of the status of the SP-chain through the MOPS chip remains functional even when communication with the modules is lost.

6.2.3 Module performance tests at different stages during assembly

FE chip functionality and module performance were evaluated using electrical *scans*. During module production, Quality Control (QC) tests continuously verify the quality of modules against requirements specified in Ref. [161]. These electrical testing procedures follow regulations [195] to ensure consistency across all ITk sites. By performing identical FE scans (using *tuned* FE configurations from module QC), FE functionality can be tracked through many assembly stages to catch potential performance degradation in a particular assembly step. A tuned FE configuration here means the pixel matrix is tuned for uniform pixel threshold and ToT and the ShuntLDOs are trimmed to provide the required 1.2 V operating voltage.

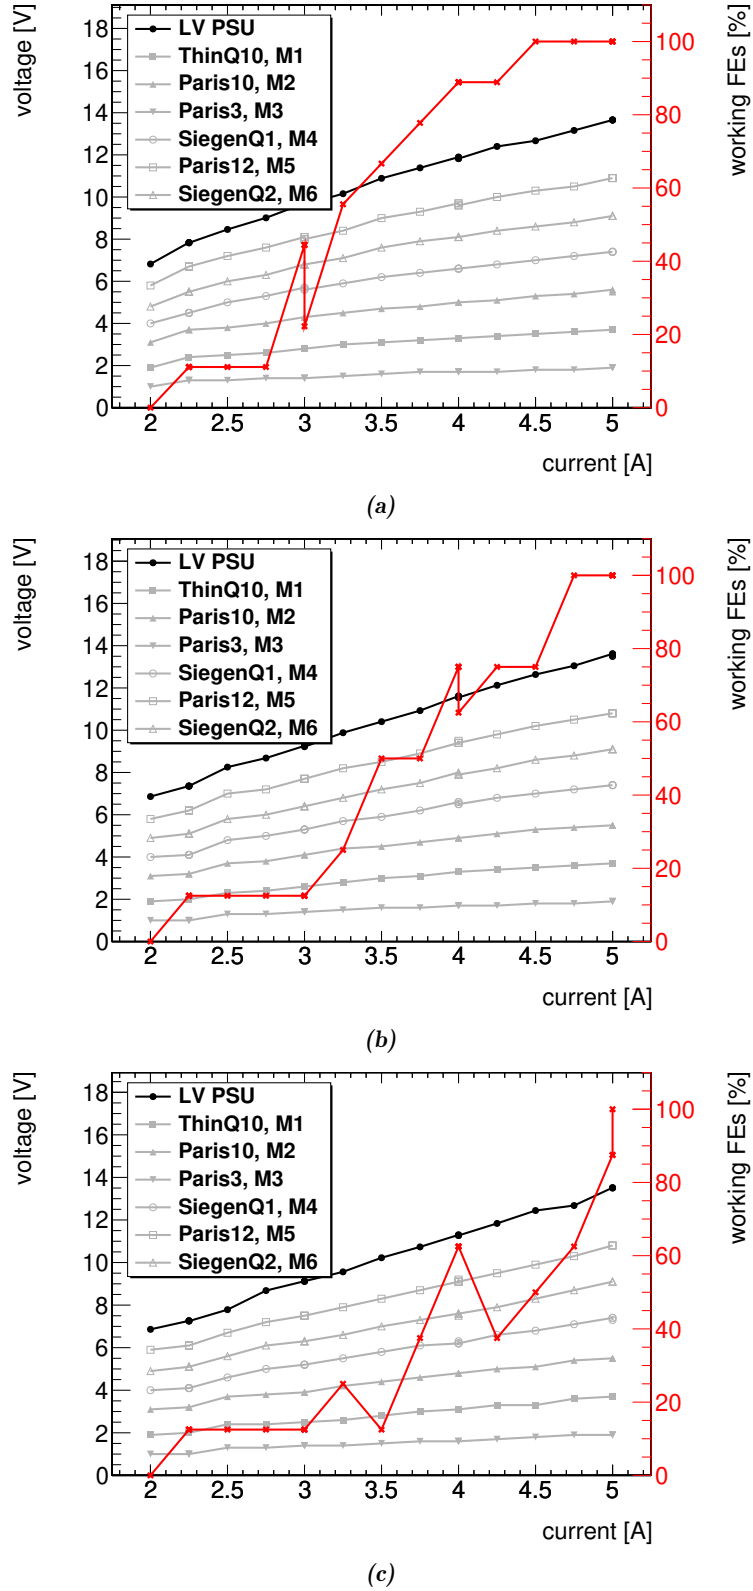


Figure 6.9: $V(I)$ curves for the M6 at temperatures (a) 10 °C, (b) -10 °C and (c) -20 °C. The left y -axis (black) shows the input voltages to the modules measured by the MOPS chip to the PP0 0V-plane. The right y -axis (red) shows the percentage of working FEs in the SP-chain.

The electrical services connected two of the four FEs (FE2 and FE4) of each module listed in [Table 5.1](#) and [Table 5.2](#) to the readout system. Electrical scans were performed by the DAQ readout software and executed with FEs, that could communicate with the DAQ. Due to a shortage of modules, the M6 was populated with the Paris10 and Paris12 modules, which contained non-functional FE4. For those modules, only the results obtained with FE2 have been included in the analysis. Furthermore, after cell integration and installation of the demonstrator in the test setup, it was found that the Paris3 module requires special readout settings, which are compatible with a standard module QC readout system but not with a FELIX based readout. Consequently, Paris3 was excluded from the subsequent analysis and later replaced by the ThinQ9 module in a successful re-work campaign. Thus, 8 out of the available 12 FEs in the M6 demonstrator were analyzed.

On the M12, ThinQ8 (with known non-functional FE4 from production) and Paris13 (where FE2 exhibited significant bit errors in its data transmission to FELIX which rendered communication impossible) were installed. Attempts to recover Paris13 FE2 with increased driver settings on the module side failed. The Liv5 module did not communicate with FELIX at all: It showed no decoding alignment but it could also not be configured to send PRBS7 data streams for BERTs. It was therefore also not possible to configure the module for electrical scans. Thus, 20 out of the available 24 FEs in the M12 demonstrator were analyzed.

The modules were tested at a CO₂ set point of 10 °C, which translates into module temperatures around 21 °C. The precise, per-module temperatures during performance tests are given in [Appendix B](#) in [Table B.1](#) for the M6 and in [Table B.2](#) for the M12. The nominal LV current for powering the RD53A quad modules was 4.8 A. For the M6, the reverse bias voltage was set to 42 V for channel 1 (HV1) and 56 V for channel 2 (HV2). For the M12, HV1 was set to 36 V and HV2 was set to 41 V. These values had been chosen taken the maximum HV setting of each module into account that was determined during module QC and the additional local module GND shift due to serial powering. I(V) curves obtained earlier in [Section 6.2.1](#) confirm these HV settings depleted all sensors.

[Table 6.3](#) describes the assembly stages used for module performance comparisons to identify performance degradation. The following electrical tests characterized module performance:

- Digital response
- Analog response
- Threshold/noise scan
- Time over Threshold measurement
- Disconnected bump scan

Table 6.3: Stage description for the module performance comparison.

Stage	Description
1	after module building (“module QC results”)
2	after cell loading
3	after connection of final pigtail
4	after cell integration (“measurement on the demonstrator”)

These scans are detailed below.

Electrical scans with FEs

The readout system uses internal injection circuitry of the FE chips to verify installed detector modules. Testing procedures include simple checks of the digital and analog circuitry, along with tests of the response of the pixel matrix under varying artificially injected charges. Scans refer to procedures executed by the readout system. All scans were performed using the tuned FE configurations obtained during module production. This section details the FE scans performed on the modules of the longeron, with results compared against identical single module tests conducted at previous stages in the assembly of the modules and LLS.

DIGITAL SCAN The digital scan is a function check of the digital circuitry in a pixel and can be used to test the readout handling of the readout system. This scan injects 100 digital test pulses to the digital circuitry and counts the output signals for all pixels individually. A perfect digital scan shows a uniform occupancy with 100 hits in each pixel. If the occupancy in a pixel is 0, the pixel is termed “digital dead”. If the occupancy in a pixel is smaller than 98 or larger than 102, the pixel is called “digital not-good”.

ANALOG SCAN In analog scans, the analog circuitry of each pixel (amplifier and discriminator) is tested. A known charge is injected into the amplifier using the injection circuitry (see [Section 4.3.1](#)). The injected charge is well above the threshold setting, such that the registering of a hit is expected, but low enough to not cause cross-talk between neighboring pixels. The scan injects the charge 100 times in each pixel. A perfect analog scan then yields a uniform occupancy with 100 hits in each pixel. If the occupancy in a pixel is 0, the pixel is termed “analog dead”. If the occupancy in a pixel is between 1 and 97, the pixel is called inefficient and if the occupancy is larger than 102, the pixel is called noisy. An inefficient or noisy pixel is also called “analog not-good”.




Figure 6.10: Turn-on behavior of a single pixel described by an S-curve as stated in Eq. (6.3). The x -axis is the injected charge Q and the y -axis the probability to register a hit. A sharp S-curve for each pixel (low σ_{noise}) is a design goal to enable setting low thresholds without getting noise hits. Tuning tries to equalize the mean threshold $\langle Q_{\text{th}} \rangle$ across the whole pixel matrix.

THRESHOLD SCAN The discriminator in the AFEs is used to suppress noise in the sensor and electronics like amplifier below a defined threshold that can arise from thermal fluctuations, leakage currents and semiconductor devices. An ideal discriminator would never fire for signals below the threshold and always fire for signals above, a behavior described by a step function. A real discriminator, however, is influenced by electrical noise in the system that can arise from cross-talk within the chip, noise on the supply lines and on the substrate, voltage drops, etc. [111]. The threshold Q_{th} in a pixel is therefore a random variable following a Gaussian distribution $Q_{\text{th}} \sim \mathcal{N}(\langle Q_{\text{th}} \rangle, \sigma_{\text{noise}}^2)$ with $\langle Q_{\text{th}} \rangle$ being the mean threshold and σ_{noise} being the standard deviation from the mean arising from the electrical noise. In a series of measurements with injected charges, the probability that the threshold Q_{th} is below the injected charge Q , i.e. the injected charge was larger than the threshold and therefore the discriminator fired to register a hit, is

$$P_{\text{hit}}(Q) = P(Q_{\text{th}} < Q) = \frac{1}{2} + \frac{1}{2} \operatorname{erf} \left(\frac{Q - \langle Q_{\text{th}} \rangle}{\sqrt{2}\sigma_{\text{noise}}} \right), \quad (6.3)$$

where erf is the error function defined as $\operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt$. Equation (6.3) therefore describes the turn-on behavior of a pixel by a so-called S-curve. The mean threshold $\langle Q_{\text{th}} \rangle$ is the charge where 50% of the injections fire the discriminator, $P_{\text{hit}}(Q = \langle Q_{\text{th}} \rangle) = 0.5$. The smaller the standard deviation σ_{noise} (the discriminator noise), the sharper is the S-curve. Figure 6.10 displays the shape of the S-curve given in Eq. (6.3).

Charges are injected into the AFE to measure threshold and discriminator noise using the injection circuitry described in [Section 4.3.1](#). The injection capacitor (with capacitance C_{inj}) is charged with a differential voltage V_{inj} generated by two voltage output DACs that can be controlled through FE registers. The difference of the DAC inputs is called V_{cal} and can be turned into the voltage V_{inj} by a linear conversion when knowing the DAC reference voltage. The injected charge is then

$$Q_{\text{inj}} = C_{\text{inj}} V_{\text{inj}} , \quad (6.4)$$

and with a nominal injection capacitance of $C_{\text{inj}} = 8.2 \text{ fF}$ the relationship is approximately $Q_{\text{inj}} \approx 10 V_{\text{cal}}$, where Q_{inj} is given in electrons. The conversion from DAC input to injected charge is performed by the readout system.

In a threshold scan, multiple charges (usually 50 times per step) are injected into the amplifier in the analog circuitry of a pixel with increasing amplitude at a fixed threshold setting. In each step, the fraction of injections that fired the discriminator is used as an estimate for $P_{\text{hit}}(Q)$. An S-curve as described in [Eq. \(6.3\)](#) is fitted to the resulting distribution, where $\langle Q_{\text{th}} \rangle$ and σ_{noise} are fit parameters. With this procedure, one obtains the mean and standard deviation for the threshold distribution of each pixel. During a tuning procedure, the mean thresholds across the whole pixel matrix are equalized (they can vary due to variations in the production process, etc.) such that the distribution of mean thresholds across the pixel matrix is narrow in order to guarantee a constant hit efficiency⁴. The target mean thresholds during tuning were 1400 electrons for the linear and differential AFEs and 2500 electrons for the synchronous AFE. A narrow threshold distribution *in each pixel* is a design goal to enable setting low thresholds without getting noise hits.

TIME OVER THRESHOLD SCAN The ToT scan is used to check the uniformity of the tuned mean ToT for a given charge across the pixel matrix. A specified charge is injected 100 times into a pixel and the ToT is measured. Similar to the threshold scan, the ToT scan gives the distribution of mean ToT values. The ToT was tuned to 7 bunch crossings when a charge of 10 000 electrons is injected.

NOISE OCCUPANCY SCAN The noise occupancy scan is used to find noisy pixels (pixels firing without a real signal) to be able to mask them. In a noise scan, no charges are explicitly injected. Instead, pixels are configured with their tuned threshold setting and readout triggers are sent to the chip at a fixed frequency. Some pixels may register hits,

⁴Hit efficiency is the number of registered hits over the number of expected hits.

arising from noise in the sensor or amplifier, and will send these hits to the readout system. The lower the threshold, the more likely it is to register noise hits. A noisy pixel is defined as a pixel with an average occupancy larger than 10^{-6} per bunch crossing. These pixels should be masked. This type of scan was not supported by any readout software used with the demonstrator during the early tests and therefore not performed. A variant of this scan was later used during the source scan, explained hereafter.

DISCONNECTED BUMP SCAN A disconnected bump for a pixel is identified by injecting in neighboring pixels and checking the occupancy in the central pixel. The disconnected bump scan is essentially an analog scan with a very large injected charge. Due to the sensor connection, a cross-talk between neighboring pixels is expected if bump bonds properly connect sensor and FE. A pixel is defined as having a disconnected bump when the pixel occupancy is less than 50 % of the number of injections. Note that disconnected bump scans are not reliable for single pixels, but only for large delaminated areas.

SOURCE SCAN In a source scan, modules are exposed to a radioactive source to detect real hits. If an external trigger, for example from a scintillator, is available, readout of the module can be triggered through the external trigger. Without an external trigger, the ITkPix chips can self-trigger on their own hits and transmit data to the readout system. Since the RD53A chip does not support self-triggers, the source scan is run with random readout triggers. Digital bad, analog bad and noisy pixels are masked for this scan beforehand. Running over a longer time, real particle hits are collected in each pixel. This scan can be used to find disconnected bumps, i.e. pixels that are not connected to the sensor and therefore will not register hits.

Results of electrical tests and comparison with previous stages

This section presents electrical scan results from the demonstrator longeron, analyzed and benchmarked against earlier assembly stages. For reference, the comparison stages were listed in [Table 6.3](#) and the module temperatures during the demonstrator stage were around 21 °C with details given in [Table B.1](#) and [Table B.2](#).

The analysis consisted of systematic comparisons of all obtained distributions. For brevity, only representative subsets are shown in the following. Additional plots are deferred to the appendix.

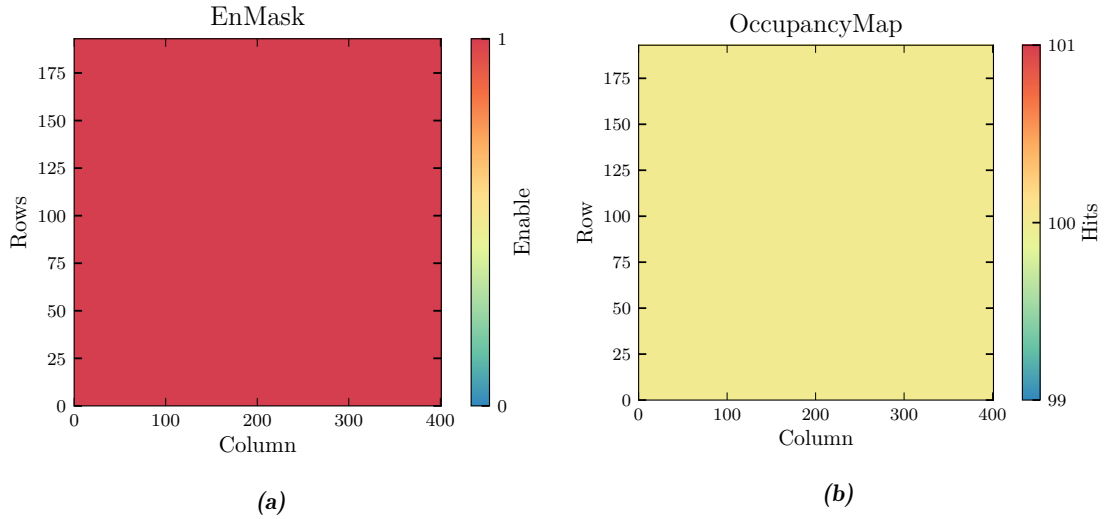


Figure 6.11: Example result of a digital scan of FE2 of the Paris10 module obtained using VakYARR. The full 400×192 pixel matrix of the FE is shown with the FE column on the x -axis and the FE row on the y -axis. (a) shows the enable mask, i.e. which pixels were enabled (1) or disabled (0) for the scan, and (b) the occupancy map after 100 triggers, i.e. the number of registered hits in each pixel.

DIGITAL SCAN The digital scans were performed using VakYARR (see Section 5.7) across the entire pixel matrix of the FE, simultaneously testing synchronous, linear and differential AFE. Figure 6.11 shows a representative hit map of a digital scan for the Paris10 module, FE2. Also shown is the *enable mask* defining enabled pixels for this scan. The hit map shows a perfect scan result with an occupancy of 100 in each pixel after 100 digital “injections”. No problematic pixels in the digital scans of any module in the M6 and at any stage were observed.

Figure 6.12 tracks the evolution of not-good pixels seen in the digital scan for the differential AFE of the modules in the M12. Only one FE, namely FE4 of the SiegenQ4 module exhibited an area of not-good pixels, a feature that was already observed before loading. The comparisons for all AFEs of the M6 and M12 are provided in the appendix in Fig. B.1.

ANALOG SCAN For the analog scan, it was necessary to separate the scan routine and perform one scan for all linear AFEs and one scan for all differential AFEs using YARR. The tests of the synchronous AFEs was of lesser interest, as ATLAS had already chosen the differential AFE for the final design chip. However, the synchronous AFEs of the M6 were tested using VakYARR. Only in the M12, the synchronous AFEs were not scanned due to time constraints. A representative example result for an analog scan with its enable mask and occupancy map is shown for a linear AFE in Fig. 6.13. The enable mask shows

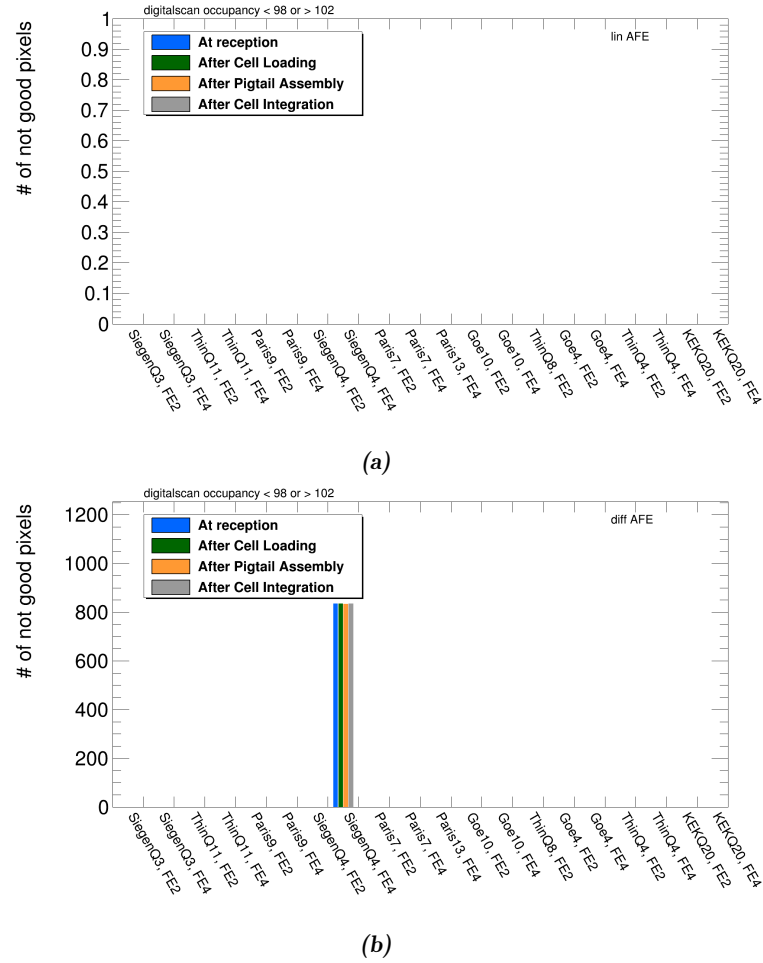


Figure 6.12: Comparison of the number of not-good pixels in a digital scan over the four assembly stages for various modules. In (a) the evolution in the linear AFEs of the modules of the M12 is shown. In (b) the evolution in the differential AFEs of the modules of the M12 is shown. A not-good pixel when 100 triggers are sent is defined as having less than 98 or more than 102 registered hits.

the disabled noisy pixels found during module QC tests. The occupancy map shows a near-perfect occupancy with 100 hits after 100 injections in the enabled pixels.

Also for the analog scan, the evolution of the amount of not-good pixels was analyzed. Figure 6.14 shows the evolution of not-good pixels in the M6 (Fig. 6.14(a)) and M12 (Fig. 6.14(b)). The results for the linear AFEs are displayed at the top, the results for the differential AFEs at the bottom. In most of the cases the number of pixels with a problematic analog response remains constant throughout all stages. The differential AFEs of the SiegenQ2 module behave not normally, but this behavior is constant throughout all stages. The bad performance is already seen in the first stage, and therefore the bad performance in the demonstrator cannot be attributed to the loading process. For the

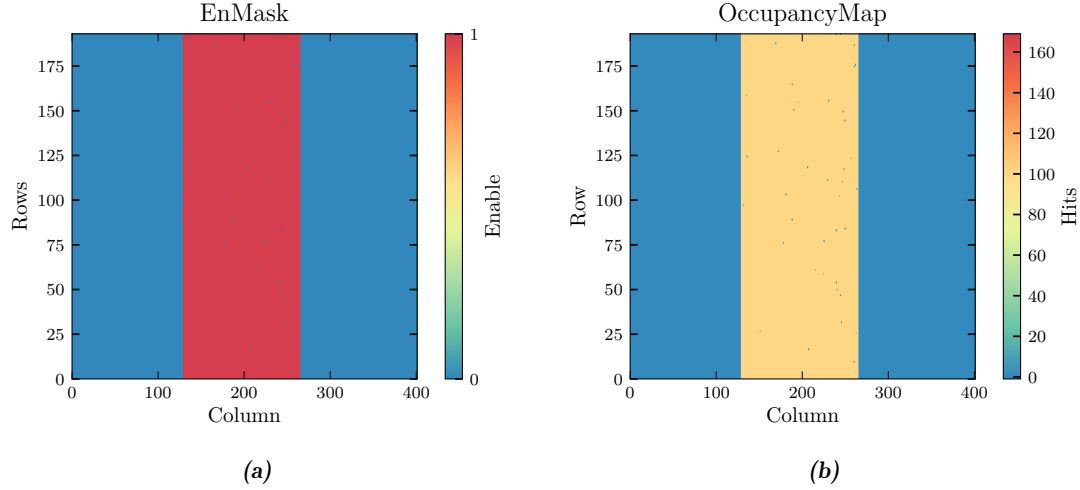


Figure 6.13: Example result of an analog scan of the linear AFE (columns 128 to 263) of FE2 of the SiegenQ1 module obtained using YARR. The full 400×192 pixel matrix of the FE is shown with the FE column on the x -axis and the FE row on the y -axis. (a) shows the enable mask, i.e. which pixels were enabled (1) or disabled (0) for the scan, and (b) the occupancy map after 100 injections, i.e. the number of registered hits in each pixel.

ThinQ4 module the number of dead pixels in the differential AFE decreased after the connection of the pigtail while it increased in the linear AFE. This effect can be attributed to an inconsistency in the FE configurations that were used for the scans before and after pigtail installation. The measurements after pigtail assembly and after cell integration in the demonstrator setup are very consistent once again. Similar results, i.e. no strong increase in not-good pixels, are also seen in the other AFEs. Plots for evolution of not-good pixels broken down to dead, inefficient and noisy are shown in the appendix in Fig. B.3 and Fig. B.4 for the M6 and in Fig. B.6 for the M12.

THRESHOLD SCAN For all modules mounted on the longeron demonstrator, configurations had been obtained during module QC tests with a threshold tuned to $1400 e^-$. Performing a threshold scan with a module in the demonstrator would produce a threshold ($\langle Q_{th} \rangle$) distribution for an AFE, a distribution of the discriminator noise (σ_{noise}) and an S-curve for all pixels as shown in Fig. 6.15 for the differential AFE of FE2 of the Paris10 module in the M6.

To compare the performance, the threshold and noise distributions as shown in Fig. 6.15(b) and Fig. 6.15(c) are obtained for the different stages and overlaid. The thresholds are temperature dependent, and they might shift in a different setup. If the threshold distribution widens too much, a threshold tuning would be performed again. It was decided to not tune the thresholds again in the demonstrator setup, but to reuse the tuned configurations

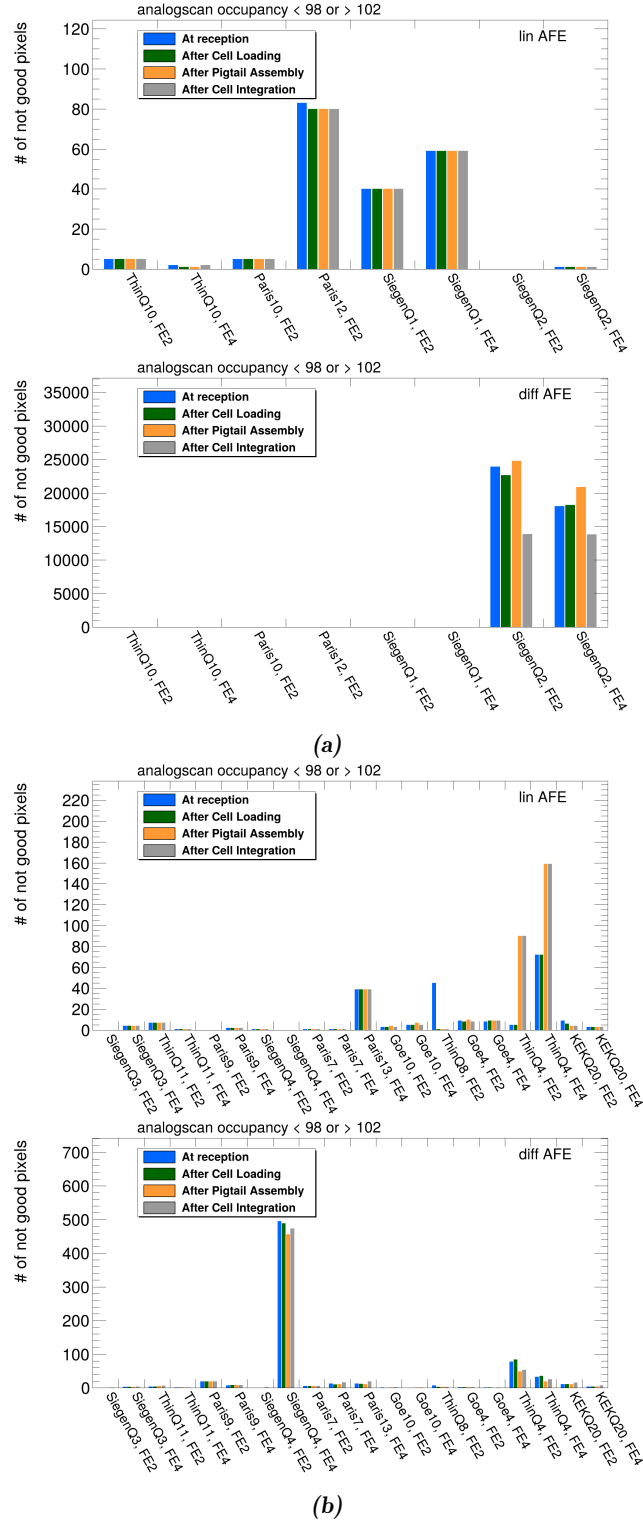


Figure 6.14: Comparison of the number of not-good pixels in an analog scan over the four assembly stages. In (a) the evolution for the FEs of the M6 is shown; in (b) the evolution for the FEs of the M12. In both cases, the results for the linear AFE are shown at the top and the results for the differential AFE at the bottom. A not-good pixel when 100 injections are made is defined as having less than 98 or more than 102 registered hits.

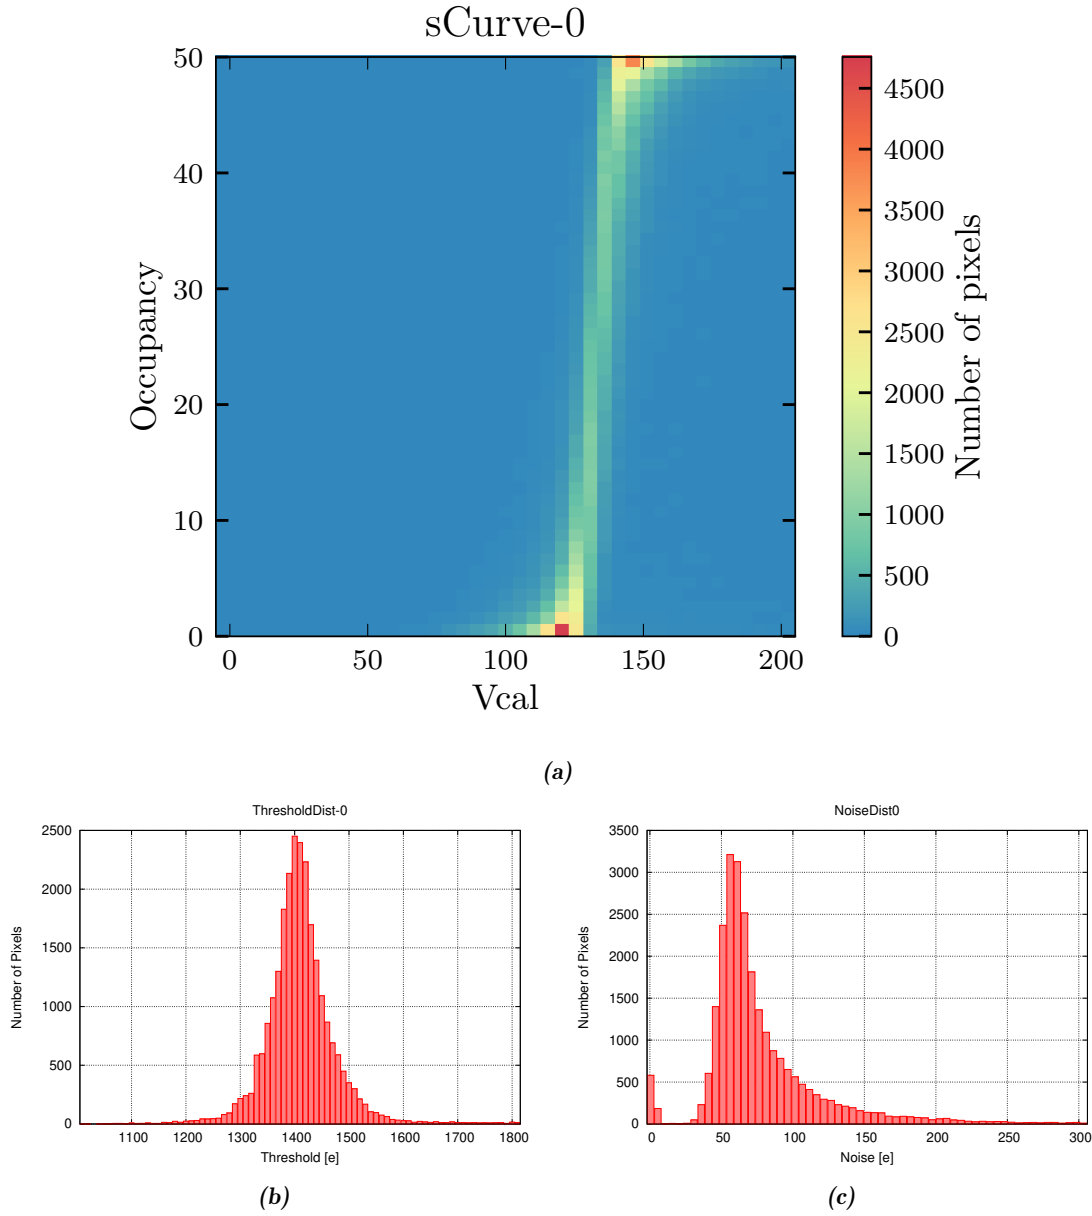


Figure 6.15: Example result of a threshold scan of the differential AFE of FE2 of the Paris10 module in the M6. The scan was performed with YARR using a fixed threshold setting obtained from tuning after module production. (a) shows the number of pixels on the z -axis that register a given number of hits at 50 charge injections and at a given charge size (measured by Vcal). Vcal is plotted on the x -axis and the occupancy, i.e. the number of registered hits, is plotted on the y -axis. The S-curve is fitted for each pixel to obtain a mean threshold ($\langle Q_{th} \rangle$) and a noise (σ_{noise}) value (measured in electrons). (b) shows the distribution of the mean thresholds across the whole pixel matrix. (c) shows the distribution of the electrical noise (σ_{noise}) across the whole pixel matrix.

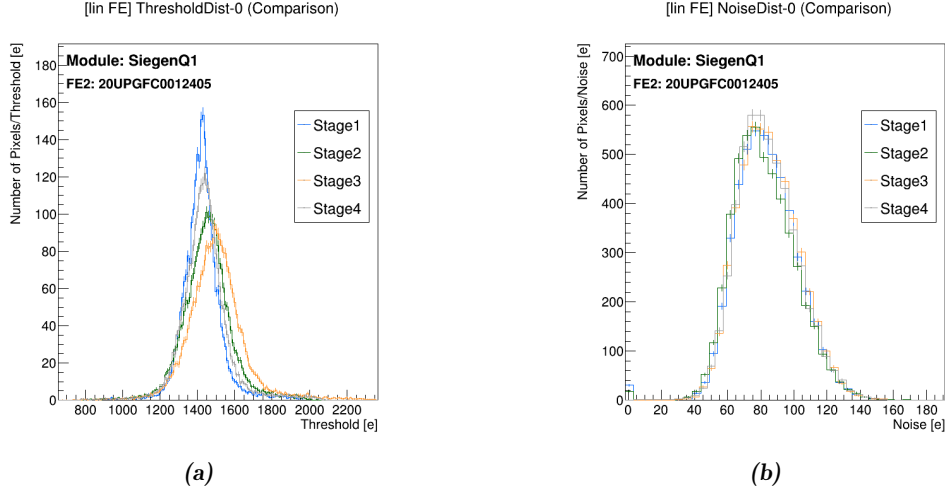


Figure 6.16: Threshold and noise stage comparison of the linear AFE of FE2 of the SiegenQ1 module in the M6. (a) shows the mean threshold ($\langle Q_{th} \rangle$) distributions for all four stages as defined in Table 6.3. (b) shows the noise (σ_{noise}) distributions for all four stages. The distributions obtained at different stages are superimposed for comparison.

obtained during module QC tests to allow for more meaningful comparisons and not tune away features developed during assembly. The discriminator noise (σ_{noise}) is also temperature dependent but not a tunable property. It is an intrinsic property of the FE and important for the performance of the chip. Only a low noise value allows setting small thresholds to achieve high hit detection efficiency with small sensor noise. The danger comes from handling the module during integration, potentially increasing the noise of the FE chips, and from integration the module in an SP-chain on a local support. To validate the loading procedure, it is important to show that it does not cause a systematic, strong increase in noise.

An example of the threshold and noise distributions in the four different stages is given in Fig. 6.16 for the linear AFE of FE2 of the SiegenQ1 module in the M6. The results are overlaid and the distributions match very well. The threshold distribution became a little wider and shifted slightly to higher values. In the noise distribution no significant changes between the stages can be observed.

The noise distributions shown in Fig. 6.15(c) and Fig. 6.16(b), as well as all other noise distributions, exhibit a feature where bins are filled with pixels having a noise less than $10 e^-$. The readout software does not expose the full per-pixel scan result, but rather performs the S-curve fitting, and therefore threshold and noise analysis, per pixel, returning only the fit results. The YARR software uses the Levenberg-Marquardt least-squares curve fitting algorithm⁵ for its S-curve fitting, which, in addition to the fitted parameters, also

⁵<https://jugit.fz-juelich.de/mlz/lmfit>

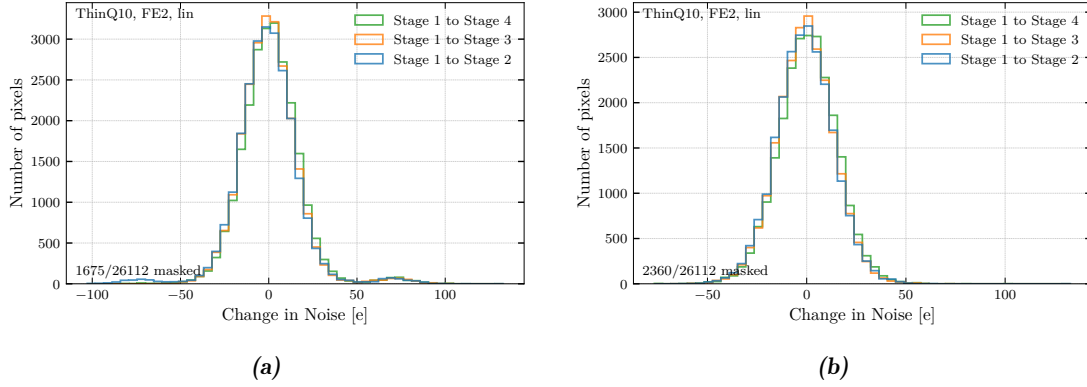


Figure 6.17: Per pixel noise difference distribution of the linear AFE of FE2 of ThinQ10 in the M6. The noise difference per pixel is calculated as $\sigma_{\text{noise, Stage } i} - \sigma_{\text{noise, Stage 1}}$ for stages $i = 2, 3, 4$ as defined in Table 6.3. The graph in (a) shows the per pixel noise difference distribution using the raw data from the readout software. Visible are two smaller peaks on the sides of the main peak. The graph in (b) shows the per pixel noise difference distribution when masking all pixels with a noise less than $10 e^-$. Note, that to be able to compare different stages, if a pixel was masked in at least one stage due to the above criteria, it was set to “masked” in all stages.

returns if the fit was successful or failed. In the noise distributions shown earlier, pixels where fits failed are already excluded.

When examining *per-pixel* noise differences across stages, we see that pixels with small noise do not necessarily retain its small noise. In subsequent stages, they might show more expected noise (around $70 e^-$). Conversely, a pixel exhibiting noise around $70 e^-$ might suddenly feature a noise of a few electrons in the next stage. Thus, noise differences for these pixels appear as small peaks left and right of the main peak in the per-pixel noise difference distribution shown in Fig. 6.17(a). The origin of these small noise values reported by the readout software is unknown. Since small noise can appear and disappear over time, this suggests it could stem from a successful fit with poor quality, perhaps finding a local minimum instead of the global one. When masking pixels with noise less than $10 e^-$ (assuming poor S-curve fitting), the per-pixel noise difference distribution loses its two smaller peaks, as shown in Fig. 6.17(b). In general, the distribution in Fig. 6.17(b) is rather wide. This is even the case for the results of subsequent scans in the same setup, suggesting that the noise measurements do not perfectly reproduce results for single pixels.

Since masking pixels reduces the number of pixels contributing to the stage comparison analysis, it is important to ensure, that this does not hide systematic performance degradation. Figure 6.18 shows the evolution of masked pixels in the noise analysis of the M6 due to failed fits or noise being less than $10 e^-$. This evolution shows no systematic increase in masked pixels across stages but suggests that the mask has a random component. Notably,

in linear AFEs, the number of masked pixels is consistently small in stage 4 compared to earlier stages. Differences may arise from the readout software version or the machine used for S-curve fits. It is also noticeable, that the number of masked pixels is overall larger in the differential FE.

The per-pixel noise difference distribution for all stages for a differential AFE is shown in Fig. 6.19. Compared to the distribution of a linear AFE (shown in Fig. 6.17(b)) it has much longer tails. This is possible due to the large tail in the noise distribution of the differential AFEs shown for example in Fig. 6.15(c). The linear AFEs have a much narrower noise distribution as can be seen from Fig. 6.16(b). The longer tails in Fig. 6.19 also suggest that noise per pixel in the differential AFEs varies more from test to test than in linear AFEs. This could be a feature of the readout system and how it performs the threshold scan or of the AFE itself. The origin of these features is not known and requires deeper investigation. For this thesis, data from unmasked pixels was analyzed to compare stages and identify potential performance degradation.

For a systematic comparison of the noise evolution across different assembly stages, both absolute noise values and relative changes are presented. The mean (in electrons) of the noise distribution in stage 3 (after pigtail assembly) and in stage 4 (demonstrator tests) is shown for the M6 in Table 6.4 and for the M12 in Table 6.5. For *per-pixel* noise differences, the relative differences, i.e. $(\sigma_{\text{noise, Stage } i} - \sigma_{\text{noise, Stage 1}}) / \sigma_{\text{noise, Stage 1}}$ for stages $i = 2, 3, 4$ as defined in Table 6.3, were calculated (yielding histograms similar to Fig. 6.19) and the evolution of the mean of these distributions is shown for the M6 in Fig. 6.20 and for the M12 in Fig. 6.21. Relative changes are always referenced to stage 1. Error bars in subsequent stages indicate the standard deviation of the relative per-pixel noise difference distributions. The modules in the figures are sorted by their position in the SP-chain, i.e. module close to the PP0 0V-plane on the left and modules with large local GND offsets on the right.

Comparing linear and differential AFE, a consistent slight increase in the noise levels is visible in the differential FE. The strongest noise increase typically occurs from stage 1 to stage 2, which is the cell loading step. For the M6, the largest noise increases occur (in both linear and differential AFE) in the Paris10 and Paris12 modules. The noise increase in stage 4 (compared to stage 1) is less than 10 % for all modules except Paris12 and the differential AFE of Paris10. In the linear AFE of both the SiegenQ2, FE2 and Paris12, FE2 the relative change of the per-pixel noise decreases in the later stages after an initial large increase. In the M12, the pattern is similar, with the largest increases in the noise seen in stage 2. Also here the noise increases are more consistent in the differential AFE. In addition, the Paris7 and Paris9 modules show the largest changes with almost 15 % noise increase for Paris9, while the SiegenQ3 and SiegenQ4 modules are the most stable.




Figure 6.18: Evolution of masked pixels (due to failed fits as reported and masking of pixels with noise less than $10 e^-$) over the four assembly stages (see Table 6.3) in the noise analysis for (a) the linear AFEs and (b) the linear AFEs of the M6. Masking these pixels removes the secondary peaks as shown in Fig. 6.17(b).




Figure 6.19: Per pixel noise difference distribution of the differential AFE of FE2 of ThinQ10 in the M6. The noise difference per pixel is simply calculated as $\sigma_{\text{noise, Stage } i} - \sigma_{\text{noise, Stage 1}}$ for stages $i = 2, 3, 4$ as defined in Table 6.3. The graph shows the per pixel noise difference distribution when masking all pixels with a noise less than $10 e^-$. Note, that to be able to compare different stages, if a pixel was masked in at least one stage due to the above criteria, it was set to “masked” in all stages.

Table 6.4: Means of noise distributions of FEs in M6 in stage 3 (after pigtail assembly) and stage 4 (demonstrator tests).

FE	Linear AFE			Differential AFE		
	$\bar{\sigma}_{\text{noise}}$ in Stage 3 [e^-]	$\bar{\sigma}_{\text{noise}}$ in Stage 4 [e^-]	Change [%]	$\bar{\sigma}_{\text{noise}}$ in Stage 3 [e^-]	$\bar{\sigma}_{\text{noise}}$ in Stage 4 [e^-]	Change [%]
ThinQ10, FE2	76.8	77.9	1.4	74.9	75.7	1.0
ThinQ10, FE4	73.9	75.2	1.8	72.9	74.7	2.5
Paris10, FE2	77.0	79.1	2.7	75.6	80.9	7.0
SiegenQ1, FE2	89.7	89.0	-0.7	65.5	68.2	4.1
SiegenQ1, FE4	76.3	78.3	2.6	77.3	78.7	1.8
SiegenQ2, FE2	86.6	85.9	-0.8	81.4	80.3	-1.3
SiegenQ2, FE4	80.1	81.0	1.1	86.8	84.2	-3.0
Paris12, FE2	88.9	80.9	-9.0	66.2	68.9	4.1




Figure 6.20: Relative per pixel noise difference evolution for the **(a)** linear AFEs of the M6 and **(b)** differential AFEs. Plotted are the means of the distributions of the per pixel relative changes (always with respect to stage 1) of the noise. The error bars represent the standard deviation of this distribution. The stages are defined in [Table 6.3](#).

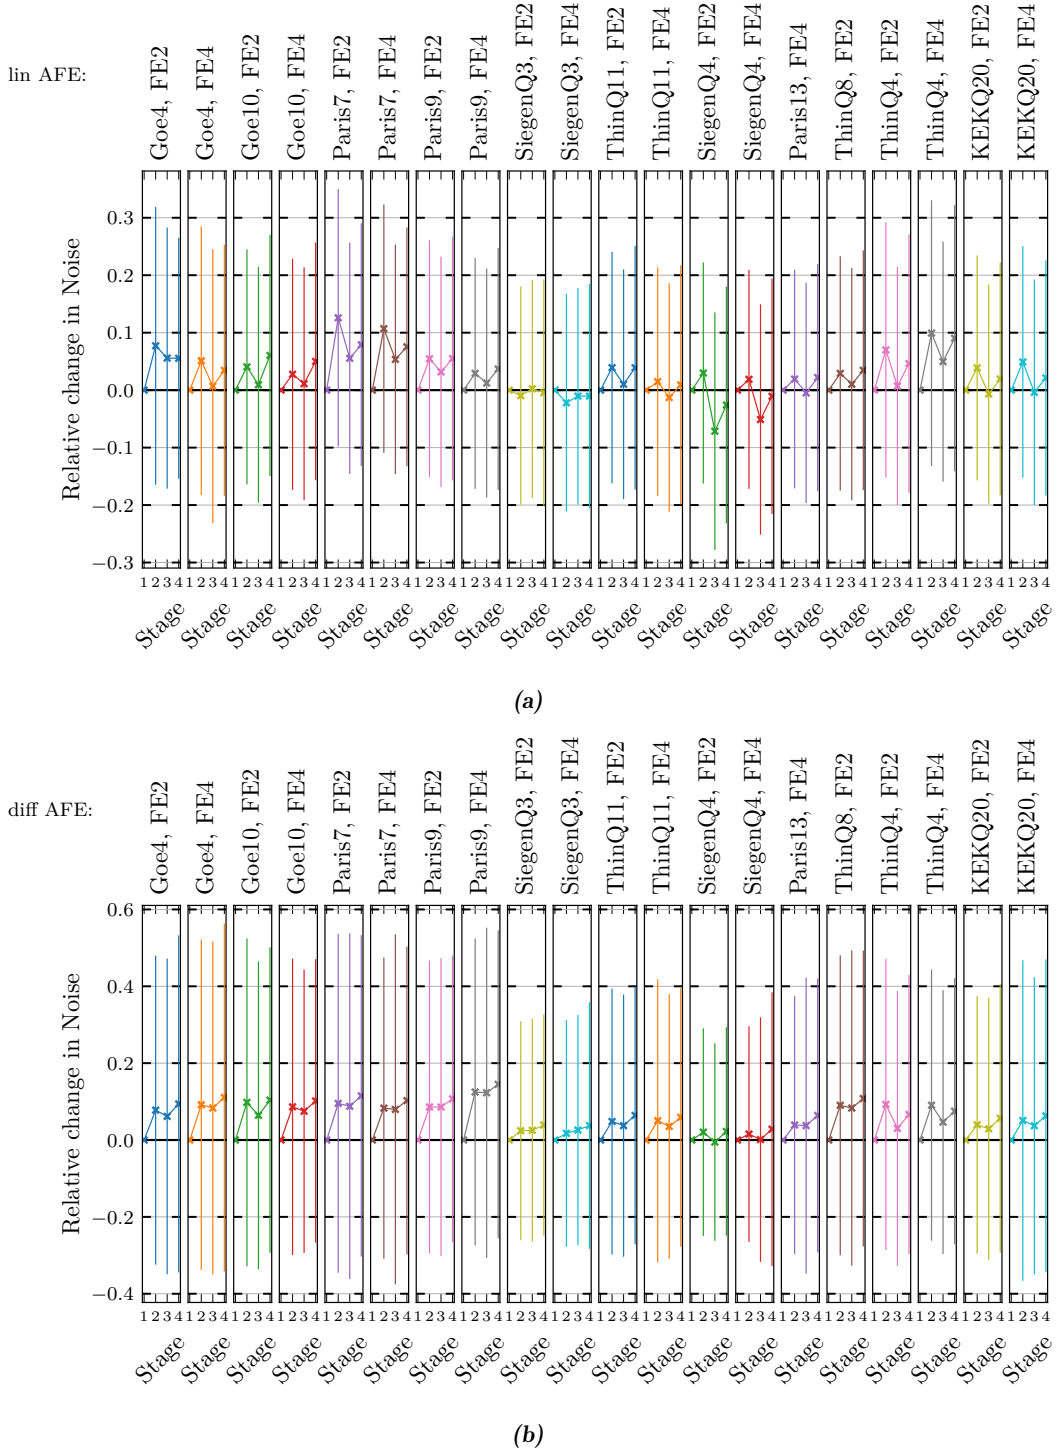


Figure 6.21: Relative per pixel noise difference evolution for the (a) linear AFEs of the M12 and (b) differential AFEs. Plotted are the means of the distributions of the per pixel relative changes (always with respect to stage 1) of the noise. The error bars represent the standard deviation of this distribution. The stages are defined in Table 6.3.

Table 6.5: Means of noise distributions of FEs in M12 in stage 3 (after pigtail assembly) and stage 4 (demonstrator tests).

FE	Linear AFE			Differential AFE		
	$\bar{\sigma}_{\text{noise}}$ in Stage 3 [e ⁻]	$\bar{\sigma}_{\text{noise}}$ in Stage 4 [e ⁻]	Change [%]	$\bar{\sigma}_{\text{noise}}$ in Stage 3 [e ⁻]	$\bar{\sigma}_{\text{noise}}$ in Stage 4 [e ⁻]	Change [%]
Goe4, FE2	88.2	88.6	0.4	73.3	76.0	3.7
Goe4, FE4	85.8	88.6	3.3	77.2	78.6	1.8
Goe10, FE2	77.5	81.4	5.1	68.9	73.4	6.5
Goe10, FE4	76.0	79.0	3.9	71.4	75.9	6.3
Paris7, FE2	76.0	77.5	2.0	72.4	74.3	2.6
Paris7, FE4	76.3	77.7	1.9	74.2	76.1	2.5
Paris9, FE2	74.0	75.6	2.3	70.9	71.6	1.0
Paris9, FE4	74.9	76.6	2.3	73.1	73.6	0.6
SiegenQ3, FE2	77.8	77.0	-1.1	85.7	80.8	-5.7
SiegenQ3, FE4	77.1	76.7	-0.4	87.8	83.6	-4.7
ThinQ11, FE2	73.9	75.8	2.7	70.3	75.4	7.2
ThinQ11, FE4	75.8	77.4	2.1	67.0	70.5	5.1
SiegenQ4, FE2	76.7	80.4	4.8	86.4	82.4	-4.6
SiegenQ4, FE4	76.5	79.6	4.1	84.9	80.3	-5.4
Paris13, FE4	76.3	78.4	2.7	81.3	83.0	2.0
ThinQ8, FE2	74.9	76.6	2.3	77.8	78.6	1.0
ThinQ4, FE2	73.9	76.7	3.8	74.9	77.5	3.5
ThinQ4, FE4	76.5	79.3	3.7	74.7	76.2	2.0
KEKQ20, FE2	75.0	76.8	2.4	75.4	76.7	1.8
KEKQ20, FE4	73.4	75.1	2.3	70.0	71.6	2.3

The standard deviations of the noise difference distributions (represented as error bars in the figures) are larger for the differential than for the linear AFE. In both cases they are large in comparison to the mean change, allowing no conclusions on single pixel level. No trend correlates noise change with the position in the SP-chain.

Noise is temperature-dependent, meaning that an increase in noise at warmer temperatures is expected. The tests in the demonstrator stage were performed at module temperatures around 21 °C, as were prior tests. A more detailed comparison would require better control and knowledge of the temperature of the modules during tests and better understanding of the readout software versions used. Software coherency across many module production sites is very difficult and QC procedures were only being defined.

However, when particles traversing the (non-irradiated) sensor create more than 10 000 free electrons, noise increases of less than 10 % at a threshold of $1400 e^-$ and typical noise levels of $75 e^-$ do not impact efficiency. Even after irradiation with fewer charges collected, the impact on efficiency is negligible as noise levels remain small relative to thresholds. While measurable, the noise increase is insufficient to affect module efficiency. Observed changes are thus compatible with a working detector system.

The other result of a threshold scan is the threshold itself, a tunable value. No re-tuning was performed during the assembly stages. The threshold changes (in electrons) can be seen for the M6 in [Fig. 6.22](#) with the M12 results given in the appendix in [Fig. B.7](#).

Changes are more consistent in the differential AFE, with the largest shift again in stage 2. For the linear AFEs, the thresholds in stage 4 often changed relative to stage 3 but reverted toward stage 1 values. The changes in threshold are below $50 e^-$ for all differential AFEs. Larger changes up to $100 e^-$ occurred in the linear AFEs of SiegenQ1 (FE2 and FE4), SiegenQ2(FE4) and Paris12 (FE2). All changes are compatible with a functional loading procedure and no significant deviation or degradation of the module performance was found.

In contrast to the *per-pixel* threshold and noise differences, the absolute mean thresholds and noises in the various stages are shown in [Appendix B](#) in [Fig. B.8](#) and [Fig. B.9](#) with very consistent results for all AFE types and across all stages.

TIME OVER THRESHOLD SCAN The mean ToT was tuned at module QC to 7 bunch-crossings when injecting a charge of $10\,000 e^-$. A ToT scan in the demonstrator returns a distribution of mean ToT values in a FE as shown in [Fig. 6.23](#). Similar to the threshold scan, the obtained distributions are compared between the four stages for all AFEs. An example of this ToT distribution comparison is given in [Fig. 6.24](#) for the differential AFE of FE2 of the SiegenQ1 module in the M6.

For a more systematic comparison, the per pixel ToT changes are again calculated and presented in [Fig. 6.25](#) for the M6 and in [Fig. 6.26](#) for the M12. In the M6, the largest differences are observed in the SiegenQ2 module, with a change of the mean ToT by more than 1 bc in the differential AFE. Overall the ToT values are more consistent in the differential AFE. Especially the threshold has an impact on the ToT if everything else is kept constant. In general, for FEs in which the threshold increases, the ToT decreases, which is visible when compared to the threshold evolutions before. No strong deviations are visible and no significant degradation of the module performance was found.

The evolution of the *most common* ToT over the different stages is shown in [Appendix B](#) in [Fig. B.10](#).

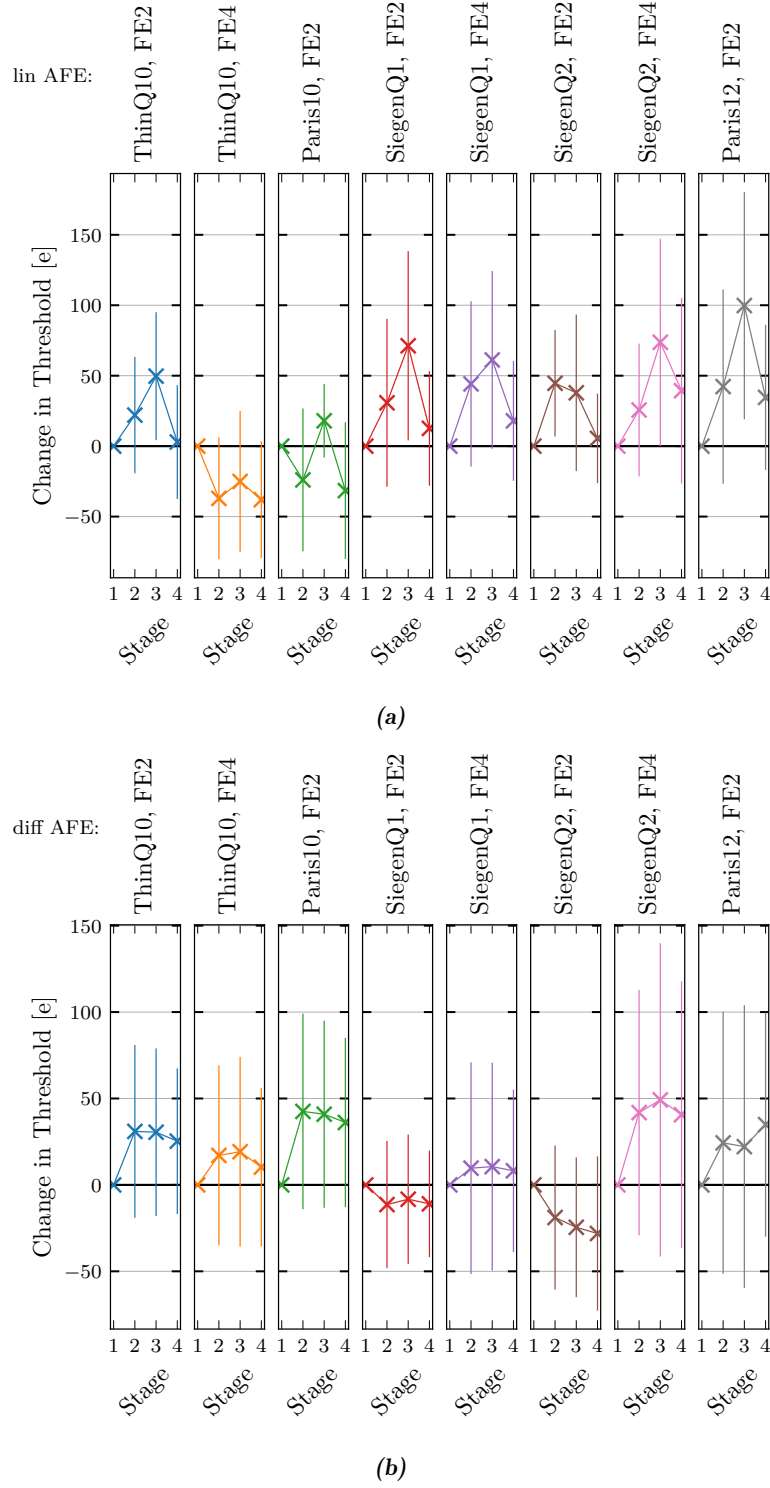


Figure 6.22: Per pixel threshold difference evolution for the (a) linear AFEs of the M6 and (b) differential AFEs. Plotted are the means of the distributions of the per pixel changes (always with respect to stage 1) of the threshold. The error bars represent the standard deviation of this distribution. The stages are defined in Table 6.3.




Figure 6.23: Example result of a ToT scan of the differential AFE of FE2 of the SiegenQ1 module. Shown is the distribution of mean ToT values (measured in bunch-crossings) per pixel across the differential pixel matrix when a charge of $10\,000\,e^-$ is injected.




Figure 6.24: ToT stage comparison of the differential AFE of FE2 of the SiegenQ1 module in the M6. The distributions of the mean ToT obtained at different stages are superimposed for comparison. The stages are defined in [Table 6.3](#).

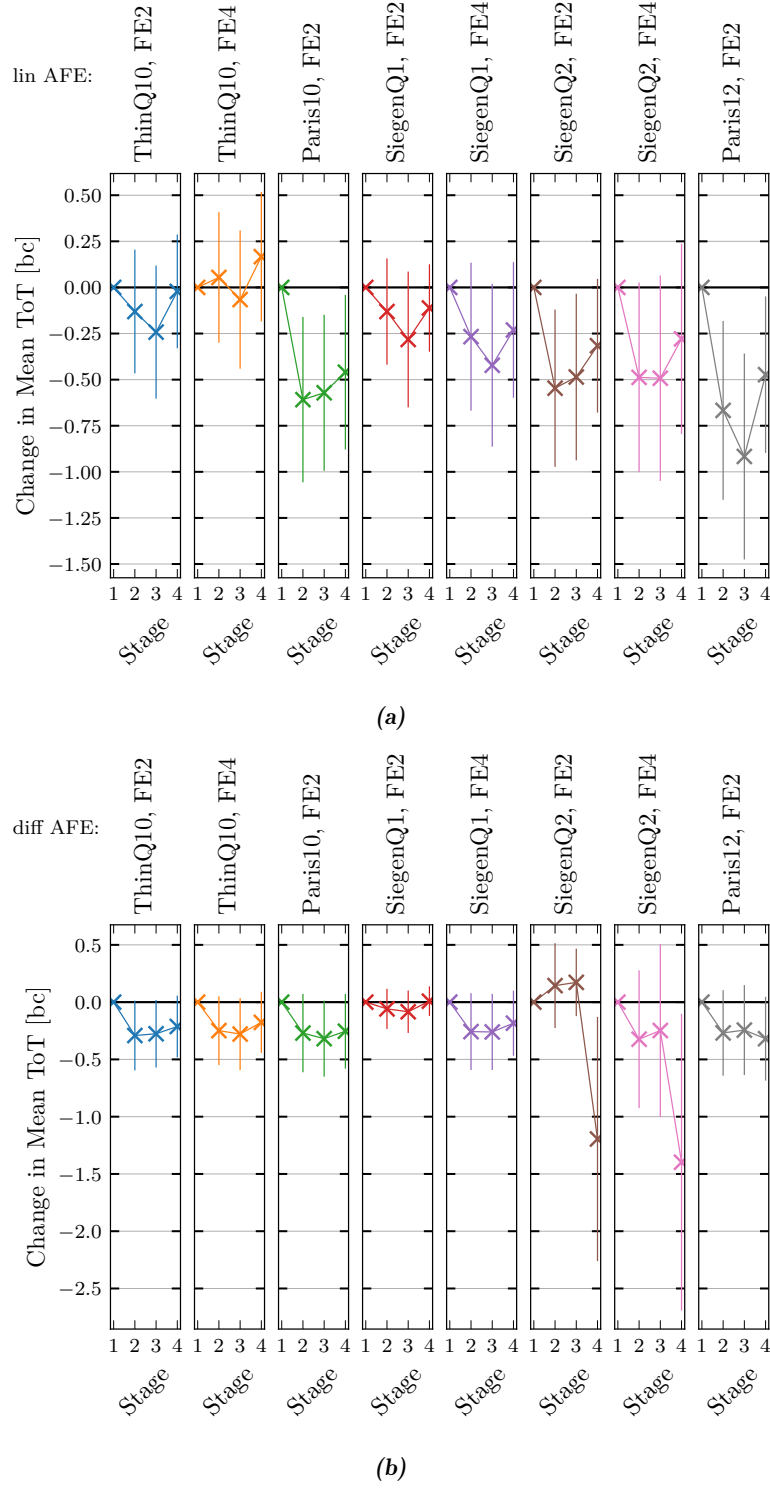


Figure 6.25: Per pixel mean ToT difference evolution for the **(a)** linear AFEs of the M6 and **(b)** differential AFEs. Plotted are the means of the distributions of the per pixel mean ToT changes (always with respect to stage 1) of the mean ToT. The error bars represent the standard deviation of this distribution. Stages are defined in [Table 6.3](#).

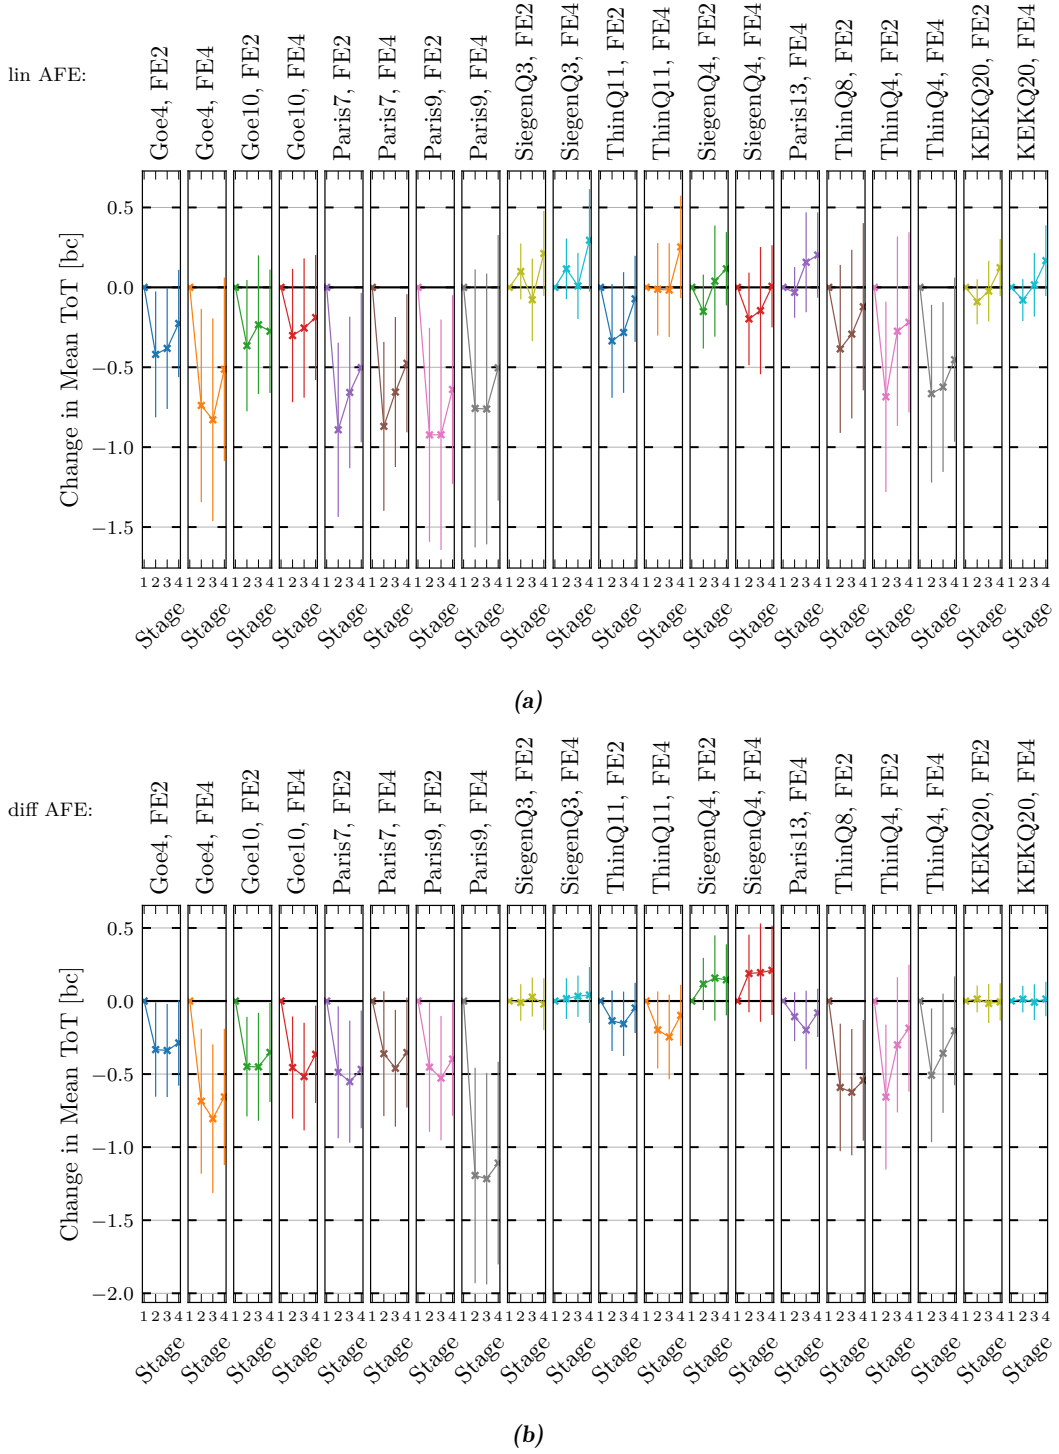


Figure 6.26: Per pixel mean ToT difference evolution for the (a) linear AFEs of the M12 and (b) differential AFEs. Plotted are the means of the distributions of the per pixel mean ToT changes (always with respect to stage 1) of the mean ToT. The error bars represent the standard deviation of this distribution. Stages are defined in Table 6.3.




Figure 6.27: Example result of a disconnected bump scan of the differential AFE (columns 264 to 399) of FE2 of the Paris12 module in the M6 in stage 4 (demonstrator tests). The full 400×192 pixel matrix of the FE is shown with the FE column on the x -axis, the FE row on the y -axis and the registered hits on the z -axis. The scan shows on the right a large area in blue with 0 registered hits in the disconnected bump scan indicating an area of delamination of the sensor from the FE chip. This delamination was known before integration into the demonstrator. Columns 0 to 263 belong to synchronous and linear AFEs and were not included in the scan.

DISCONNECTED BUMP SCAN The analysis of the disconnected bump bonds is a crucial test for the qualification of the loading process. It should be stressed that the results of the disconnected bump scan are not meaningful in the individual pixel-by-pixel comparison, but they are very powerful to identify large, disconnected areas, which is the main concern. Disconnected bumps then should show up in a disconnected bump scan as pixels with zero occupancy. A representative example result for a hit map obtained by a disconnected bump scan is given in Fig. 6.27 for the differential AFE of FE2 of the Paris12 module in the M6. This module exhibits a large area where the sensor disconnected from the FE. This delamination was visible from the first stage, before integration into the demonstrator.

The disconnected bump scan was performed for the linear and differential AFEs of the M6 and the M12 (with YARR), as well as for the synchronous AFEs of the M6 (with VakYARR, FE by FE). Due to time constraints, the synchronous AFEs of the M12 were

not scanned with VakYARR. [Figure 6.28](#) shows the evolution of the number of disconnected bump bonds for the M6 as derived from the disconnected bump bond scan. The test is very sensitive to the environmental conditions and therefore some fluctuation in between the different stages and measurements setups are expected. No significant degradation in the different stages can be seen. The large delamination in the differential AFE of FE2 of the Paris12 module that is visible in the demonstrator tests (see [Fig. 6.27](#)) was already present before loading. The slight increase of disconnected bump bonds observed in the synchronous AFE in stage 4 in [Fig. 6.28\(d\)](#) is a consequence of the specific readout software that had to be used. This readout software treats triggering and the timing of the scans differently than the other software, leading to an apparently higher number of disconnected bumps. However, no delamination areas were visible in the actual 2D maps. [Figure 6.29](#) shows the evolution of the number of disconnected bump bonds for the M12. Overall, the results show that no new delaminated areas were introduced during the cell integration step.

In order to validate the results of the disconnected bump scan, a source scan was also carried out on the M6 as presented in [Section 6.2.6](#).

Conclusions

Electrical testing of the M6 and M12 demonstrator has shown no indication for any degradation while moving from individual module cells to a complete detector system on the local support. No systematic issues were found. On the M12, communication problems with the Liv5 module remain unexplained. The readout performance of all other FE chips was excellent and consistent across all testing stages.

6.2.4 Module performance comparison at different temperatures

The M6 demonstrator has also been tested electrically at a lower CO₂ temperature of -10°C corresponding to module temperatures of around 5°C . At this temperature, none of the FEs that were tested previously at a CO₂ temperature of 10°C exhibited fatal start-up problems in their ShuntLDOs, that were known to occur in RD53A chips. The results of the performance scans at the lower temperature are in excellent agreement with the results presented in [Section 6.2.3](#). [Figure 6.30](#) shows the evolution of the mean of the measured per pixel noise difference at 10°C (stage 1) and -10°C (stage 2). As expected, the noise is reduced slightly at the lower operation temperature. This effect is larger for the linear AFE than for the differential AFE. No other differences are visible.

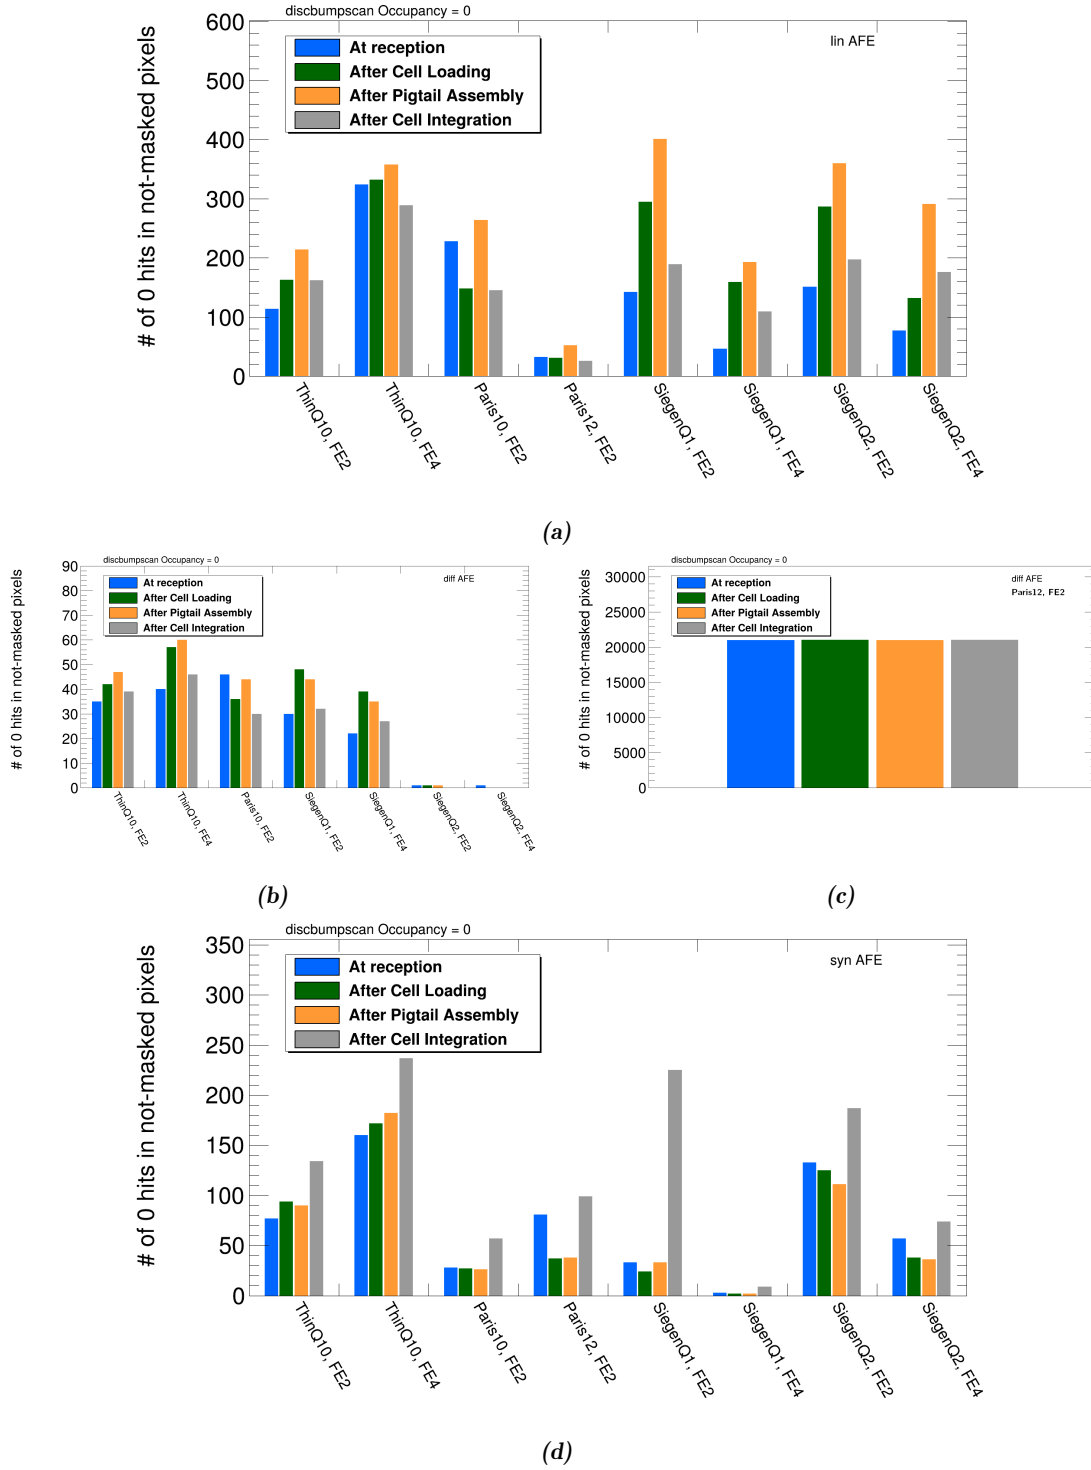


Figure 6.28: Shown is the evolution of non-masked pixels with 0 hits in the disconnected bump scan for the (a) linear AFEs of the M6, (b) and (c) differential AFEs and (d) synchronous AFEs. For the differential AFE, the result of Paris12 FE2 with its known large delamination is shown in (c) in a separate plot. No new delaminations appeared. There is a large delamination in the Paris12 module (see also Fig. 6.27) which was already there before. For the synchronous AFE, VakYARR was used in stage 4 (after cell integration).

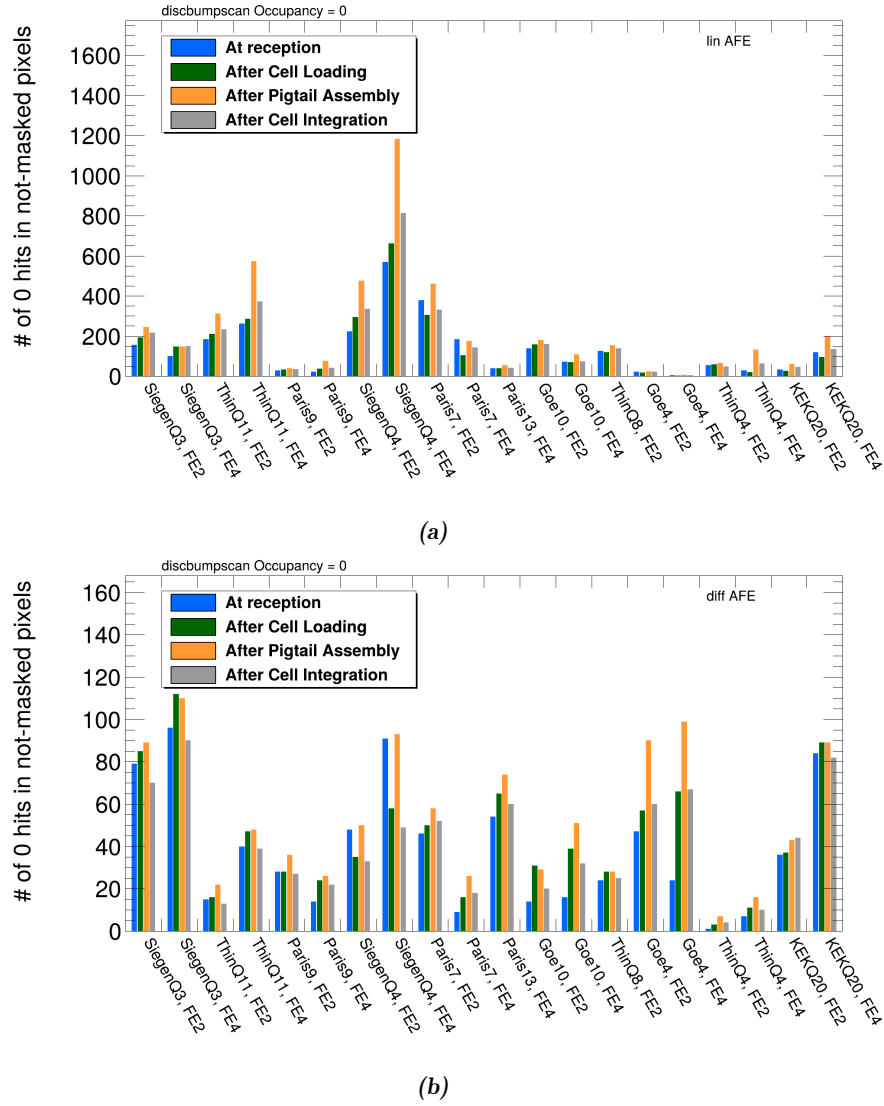


Figure 6.29: Shown is the evolution of non-masked pixels with 0 hits in the disconnected bump scan for the (a) linear AFEs of the M12 and (b) differential AFEs. No new delaminations appeared.

6.2.5 Performance of the longeron under multi-SP-chain operation

After the separate validation of the M6 and M12 SP-chains of the longeron demonstrator, studies have been carried out to check for any cross-talk during the parallel operation of two SP-chains on the same local support.

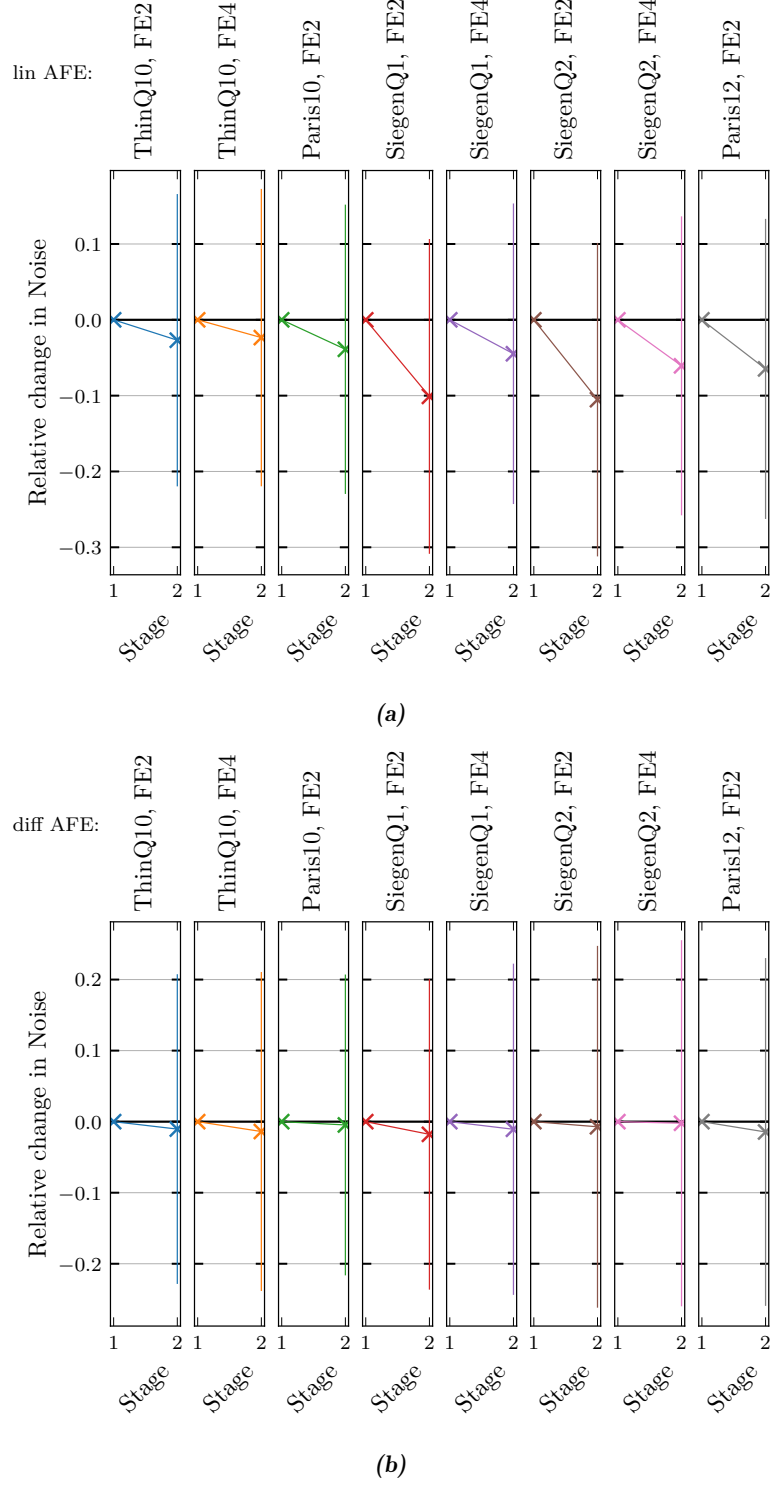


Figure 6.30: Relative per pixel noise difference evolution at cooling temperatures of 10°C (stage 1) and -10°C (stage 2) for the (a) linear AFEs of the M6 and (b) differential AFEs. Plotted are the means of the distributions of the per pixel relative changes (with respect to stage 1) of the noise. The error bars represent the standard deviation of this distribution.

Bit Error Rate Tests (BERTs)

BERTs were performed to find the correct e-link polarities on the uplink as reported previously in [Section 6.1.8](#). During combined testing of the M6 and the M12, BERTs were also used to check for cross-talk in the data transmission lines between the two different SP-chains. In the first step, the SP-chains have been tested individually. That meant, that all modules of the M6 have been configured to transmit a PRBS7 pattern instead of the standard Aurora encoded data while the M12 was turned off and vice versa. In the second step, both the M6 and the M12 were turned on and another set of BERTs was performed. For the tests, in total 68 719 476 736 bits were sent at a data rate of 1.28 Gbit/s, corresponding to a test time of approximately 54 s per FE. [Figure 6.31](#) shows the obtained upper limits on the bit error rate to 95 % confidence level comparing the two steps mentioned before.

The results clearly show the known non-functional FE chips in the M6 as mentioned in [Tables 5.1](#) and [5.2](#) (Paris10 FE4, ThinQ9 FE4 and Paris12 FE4 in the M6, as well as ThinQ8 FE4 in the M12) and the FE-chips Paris13 FE2 and Liv5 that have communication problems as already mentioned in [Section 6.2.3](#). The communication with ThinQ9 FE2 could later be successfully improved to receive valid data. Beyond that, there is no indication of any impact on the quality of the communication when operating the SP-chains alone or together.

After the experience with BERTs gained in the system test, it was decided to perform longer tests in the LLS QC setups during LLS production to verify the transmission path for each module.


Threshold scan to investigate potential influence on noise

In addition to the BERTs, the discriminator noise from a threshold scan has been measured in the different powering combinations for the M6 and M12. The noise levels are compared on the pixel level and the summary results are shown in [Fig. 6.32](#) for the M6 and [Fig. 6.33](#) for the M12. The stages for the M6 are defined as follows:


- Stage 1: M6 on, M12 off
- Stage 2: M6 on, M12 on

The stages for M12 are defined like:

- Stage 1: M12 on, M6 off
- Stage 2: M12 on, M6 on



(a)



(b)

Figure 6.31: The upper limit on the bit error rate (BER) to 95 % confidence level as recorded from the FE chips in the (a) M6 and the (b) SP-chain while the other SP-chain was turned off and on.

No relevant increase in noise is visible in all FEs. For the differential AFE of ThinQ9, FE2 a reduction in noise by approximately $4 e^-$ was measured. Though on average, the noise does not increase, the standard deviation of the noise difference distribution is approximately $14 e^-$ as shown as error bars on the respective plots. These large errors indicate, that on a per-pixel level, noise can change significantly in both directions. At noise levels of around $80 e^-$, this corresponds to changes of about 18 % on a per-pixel level. The averaged results

do not indicate cross-talk or an effect on the efficiency of the detector during the parallel operation of several SP-chains on one LLS.

6.2.6 Source scans

In order to validate the results of the disconnected bump scan, a source scan was carried out on the M6. Strontium-90 (^{90}Sr) was used as the radioactive source. ^{90}Sr undergoes β -decay, producing a beta particle (high-energetic electron) with a mean energy⁶ of 195.7 keV. The source scan was performed with the VakYARR readout software. It is a definite procedure to find disconnected bumps.

The source was placed on a movable stage very close to the modules due to geometrical constraints. A photo of the setup is shown in Fig. 6.34. The stage could be controlled from outside the box, enabling continuous testing without opening the test box. Due to the small distance and the strong collimation of the source, a homogeneous illumination of the entire module was not possible. Instead, the source was focused on one AFE after the other and the results were merged together. Measurements at the three positions per FE were necessary in order to reach an acceptable coverage. During the tests, the CO_2 cooling temperature was set to -10°C .

The measurements were taken over several hours in one position with random triggers due to the missing capabilities of self-triggering of the RD53A chip, which renders the source scans very time consuming and inefficient. Due to additional geometric constraints with the stage, FE4 of the ThinQ10 module could not be reached with the source. The resulting hit maps of the source scan are shown in Fig. 6.35. Areas with reduced hits can be attributed to Surface-Mount Device (SMD) components on the module PCB. No large areas of delamination (areas with zero hits) are visible, except for the delamination in the Paris12 FE2 differential AFE. This was observed already before in the disconnected bump scan (see Fig. 6.27) and was present from before the cell loading step (see Fig. 6.28). No new delaminated areas were therefore introduced during the cell loading and cell integration steps and the results of the disconnected bump scan could be confirmed.

6.2.7 Longer on after thermal cycling

The results presented so far were shown at the FDR for the OB LLSs. After the FDR, the longer on was installed in a climate chamber and thermally cycled 25 times between -55°C and 60°C . It was then retested in the test box in SR1. This test was conducted after all

⁶<https://www.nndc.bnl.gov/nudat3/decaysearchdirect.jsp?nuc=90Sr&unc=NDS>

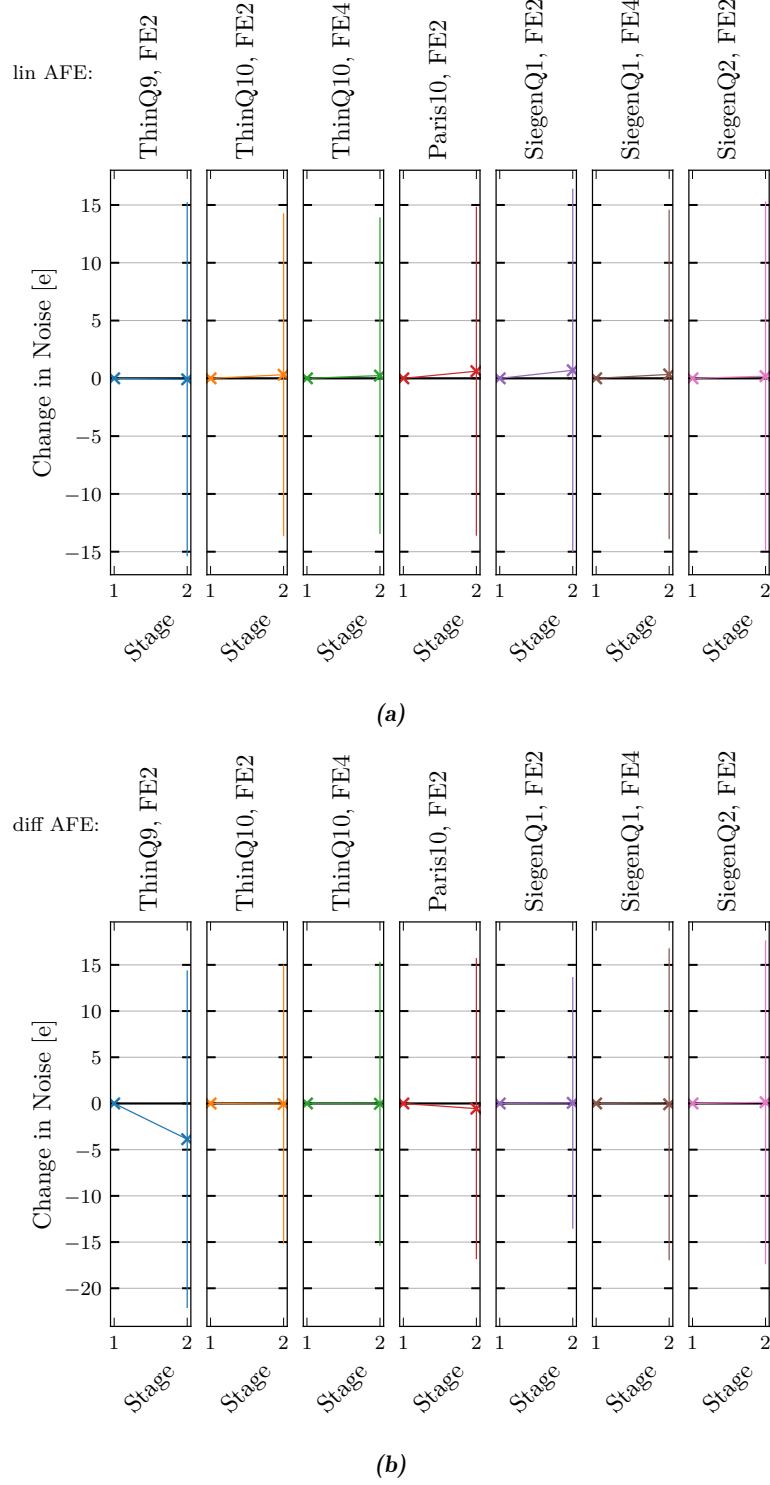


Figure 6.32: Per pixel noise difference evolution of the M6 without (stage 1) and with (stage 2) the M12 powered on for the (a) linear AFEs and (b) differential AFEs. Plotted are the means of the distributions of the per pixel noise changes (with respect to stage 1). The error bars represent the standard deviation of this distribution.

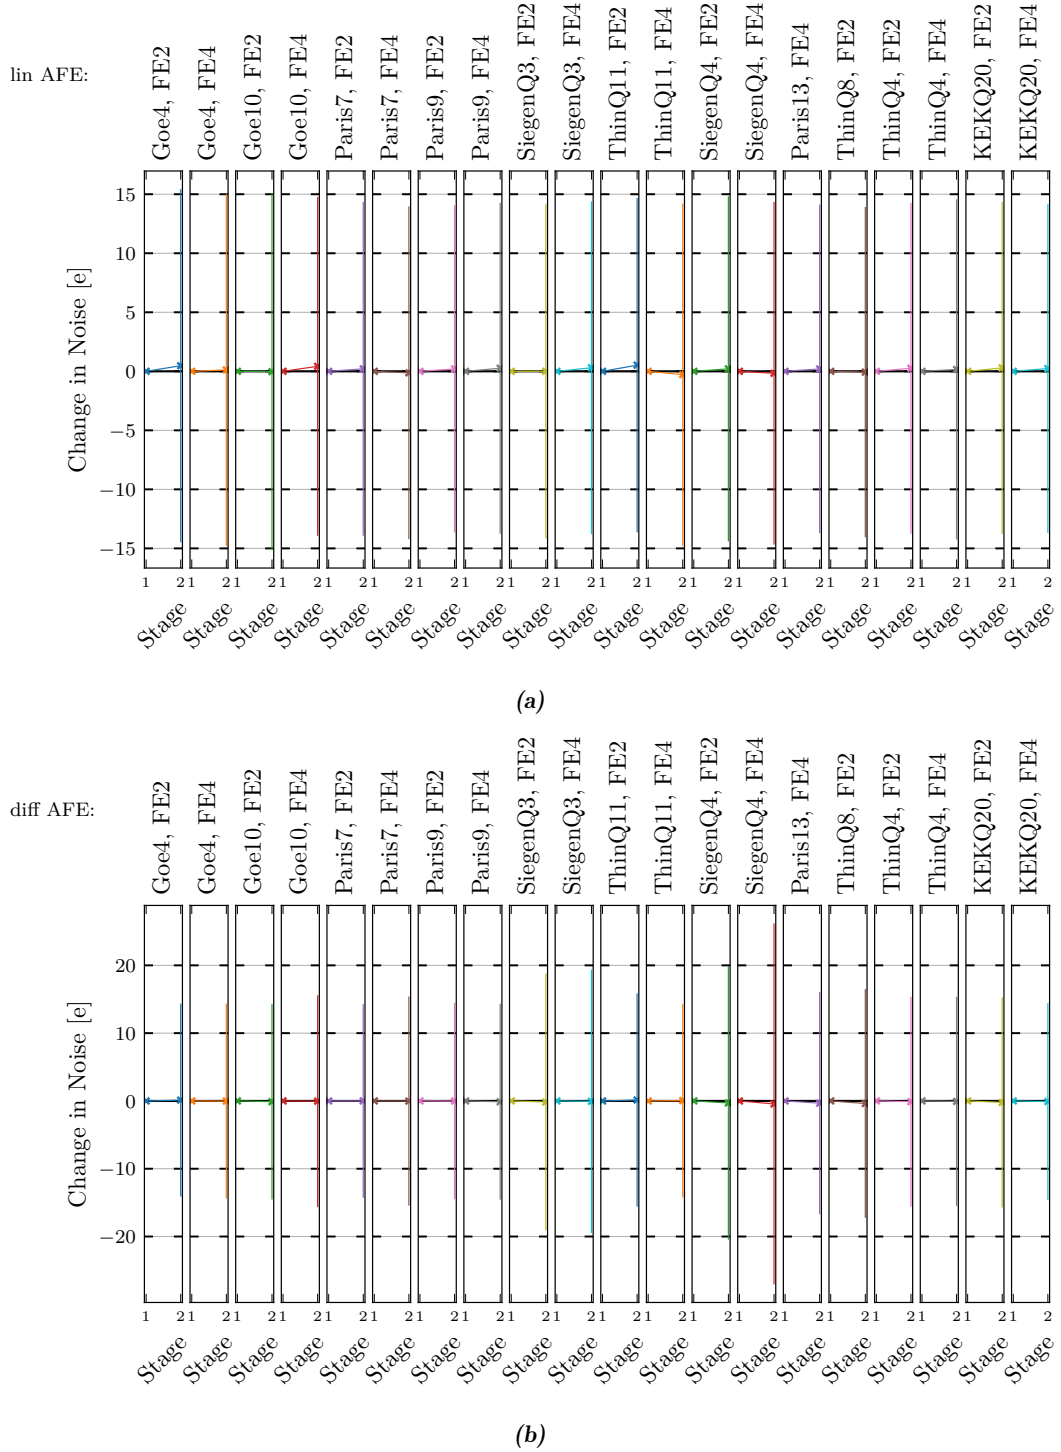


Figure 6.33: Per pixel noise difference evolution of the M12 without (stage 1) and with (stage 2) the M6 powered on for the (a) linear AFEs and (b) differential AFEs. Plotted are the means of the distributions of the per pixel noise changes (with respect to stage 1). The error bars represent the standard deviation of this distribution.




Figure 6.34: Setup for the source scan. The radioactive source shown in the bottom was placed very close to the modules of the M6. The source was attached to a movable stage that could be controlled from the outside while keeping the test box closed during the whole time of the test.

other tests because it was known to be potentially destructive. The findings after thermal cycling are reported in Ref. [191]. Starting with the test of M6, the LV power supply channel showed abnormal high voltage values compared to the results before the thermal cycling and the modules could not be powered. As the SP-chain was not operational anymore, investigations into the cause started. Probing with a digital multimeter directly on the module PCB, measured voltages indicated a problem with the ShuntLDOs. It was attempted to verify the wire-bond connections under a microscope. However, the SYLGARD™ potting of the wire-bonds made it impossible to see any disconnections. It was noted, that the potting material turned soft after thermal cycling. For further investigations, 8 modules were removed from the longeron and tested individually. The tests included an HV scan of the modules' sensor and a digital scan to verify the operational capability of the FE. Results from these tests are given in Table 6.6.

In the given test cases, no thermal stress damage was observed in modules with mechanical wire-bond protection, supporting the hypothesis, that the potting was responsible for the failure after thermal cycling. For production modules, only mechanical wire-bond protection will be used and no potting.

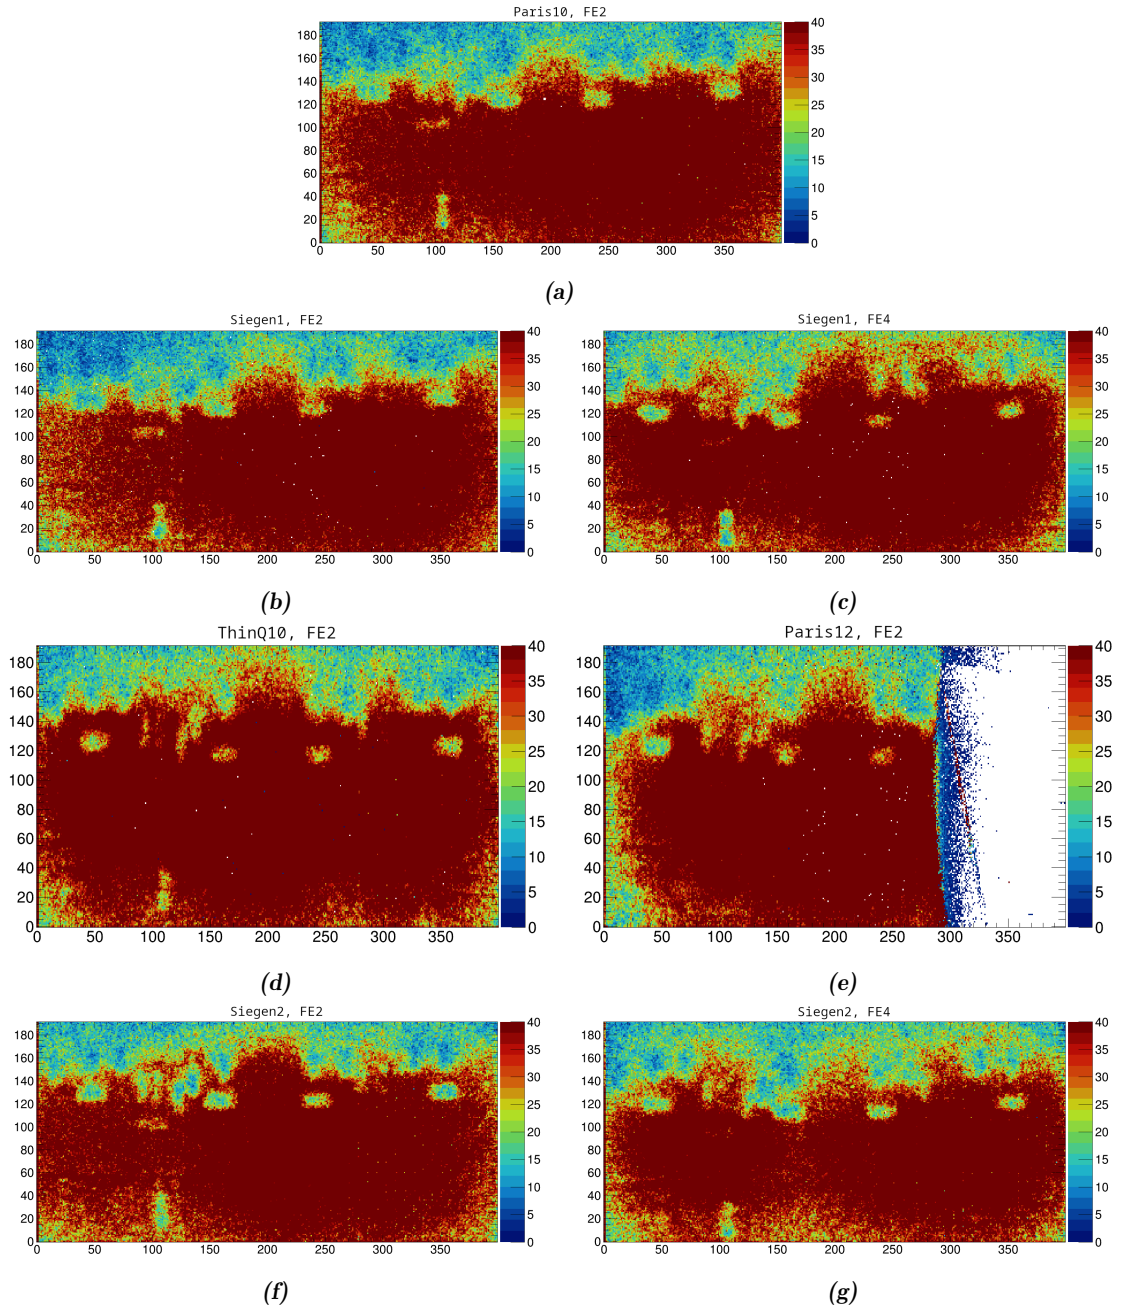


Figure 6.35: Hit map of a source scan at -10°C taken with VakYARR. On the x - and y -axes are the columns and rows of the FE chip. The z -axes display the occupancy that was measured during the source scan. For each FE, the source was positioned in front of each AFE and the results were merged together.

Table 6.6: Summary of the individually tested modules of the longeron after 25 thermal cycles. (From [191])

Module	SP-chain	Wire-bond protection	V(I) curve	Operational?
ThinQ10	M6	Potting	OK	Yes
Siegen1	M6	Potting	Abnormal	No
Siegen2	M6	Potting	Abnormal	No
ThinQ4	M12	Mechanical protection	OK	Yes
KEKQ20	M12	Mechanical protection	OK	Yes
ThinQ8	M12	Potting	Abnormal	No
Goe4	M12	Potting	Abnormal	No
Liv5	M12	Potting	Abnormal	No

6.3 Summary and conclusions

The tasks of the system test, to demonstrate that detector units consistently fulfill the detector requirements and to validate the loading procedure, were successfully completed. In addition to the demonstrator longeron, a demonstrator IHR was also tested [191]. The tests of the OB RD53A demonstrators produced valuable results of a detector prototype. The experimental results validated both the cell integration process and the overall system design. No fundamental or systematic problems were observed in electrical module performance. Individual issues with single modules were not caused by the integration process, but all evidence suggests that they can be attributed to known shortcomings of the RD53A readout chip and the sub-optimal quality of the modules used in the demonstrators.

Beyond validating the OB LLS design, preparing and operating the demonstrator test setup represented a critical milestone, as it helped finalizing the concept for the test stands used for the QC of the LLS. Many hardware and software components were developed and commissioned while testing the demonstrators, and significant knowledge in the operation of such a complex system was built. This knowledge and all associated developments will be transferred to the QC setups and sites.

The test results underwent expert review, and the OB successfully “passed” the FDR. With this, testing sites are preparing their LLS QC test setups for use during preproduction and production to ensure the quality of detector components installed in the final detector. Clear testing procedures and systemic analyses are being defined with rigorous documentation to enable easy comparison of module performances across production steps. Testing sites are qualifying their setups, with parts of this thesis contributing to the qualification of the CERN setup as detailed in [Chapter 7](#).

Preparations for Preproduction of ITk Pixel OB Loaded Local Supports

Having passed the FDR, the ITk OB community proceeded to prepare for production of OB LLSs. Cell loading and cell integration will be undertaken in different clusters. Five clusters will perform the cell integration step and build OB LLSs. Each LLS will undergo quality control in the LLS QC test setup of its cluster to verify compliance with the requirements for the final detector. The five clusters are CPPM¹, LAPP² & LPSC³, University of Bonn, University of Geneva and CERN. Accordingly, a LLS QC test stand has been established at CERN. Before the production can start, all five clusters will have to build a LLS and demonstrate a full QC procedure. A preproduction longeron assembled with ITkPix-V1.1 modules has been constructed at CERN and is being used to qualify the setup prior to the Production Readiness Review (PRR).

This chapter describes the CERN LLS QC setup, details MOPS calibrations conducted before the assembly of the preproduction longeron and presents initial electrical tests performed with the preproduction longeron. Additionally, it presents software developments that were carried out to enable automated testing procedures on the LLSs across all five LLS clusters.

At the time of writing, the preproduction longeron at CERN has already undergone intensive testing and partial results were presented at the first half of the PRR, while more results are being collected and prepared for the second half of the PRR.

¹Centre de Physique des Particules de Marseille

²Laboratoire d'Annecy de Physique des Particules

³Laboratoire de Physique Subatomique et de Cosmologie de Grenoble




Figure 7.1: The ITk OB LLS QC setup at CERN.

7.1 Description of the CERN LLS QC setup

The experience gained from the RD53A demonstrator program as described in [Chapter 5](#) informed the design of the ITk OB LLS QC setups. The system test served as the foundation for LLS QC setup development, with significant synergies between the demonstrator and QC setups. This influenced developments such as the DCS, the microservice-based online software framework, the CSB and other hardware solutions. The community also standardized infrastructure hardware across all five clusters to enable shared expertise, common problem-solving approaches if issues arise, and efficient test stand preparation. [Figure 7.1](#) shows the LLS QC setup at CERN with a test box for the LLS and infrastructure needed for testing.

The test box is substantially smaller than the one used for the RD53A demonstrator, accommodating at most one OB longeron or IHR for QC tests. The test box can be flushed with dry air to reduce the humidity inside. For the detector readout, an Opto Box similar to the one in the RD53A demonstrator is used. It resides within a case called Opto Panel. Optical fibers connect to a FELIX computer in the same structure. Power supplies for Opto, LV, HV and the MHFB are similar to the hardware in the demonstrator. Also a similar PLC that was used for the environment monitoring in the demonstrator is used in the LLS QC setup, though taking over additional interlock functionality. The Opto Box is actively cooled to around 15 °C by an external mono-phase chiller with a cooling capacity of about 30 W. The two-phase CO₂ cooling for the detector modules is provided by a portable CO₂ cooling plant called MARTA.

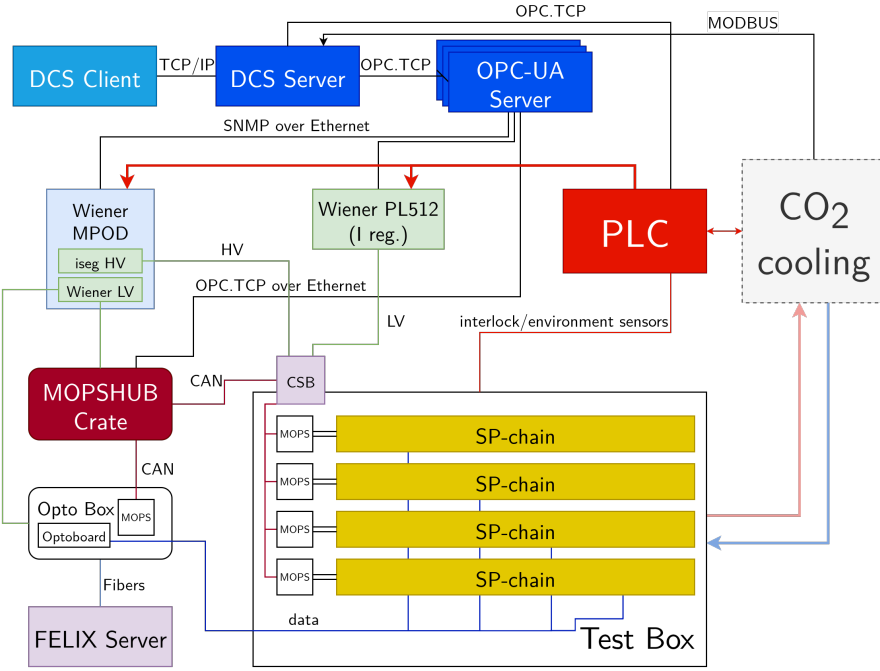


Figure 7.2: Schematic overview of the ITk OB LLS QC setup at CERN with focus on the DCS. Shown are the various types of connections between the setup infrastructure and the LLS under test as they are important for the DCS.

7.1.1 Detector Control System software and powering

A schematic overview of the infrastructure hardware of the setup is given in Fig. 7.2. When compared to the demonstrator setup in Fig. 5.5, significant similarities are evident. The DCS software controls power supplies, receives environmental and interlock data from the PLC, monitors the MARTA cooling unit and processes MOPS monitoring data from the MHFB.

The deployed DCS solution evolved from the DCS used in the demonstrator. These DCS-specific developments constitute a major component of this thesis and are detailed in Section 9.3.

7.1.2 CO₂ cooling

A portable, small CO₂ cooling plant called MARTA [196] serves as the cooling solution for detector modules. The cooling plant has a temperature range from 17°C to −22°C. It can be controlled via a touchscreen interface and monitored remotely, for example by the DCS, via a MODBUS connection.

7.1.3 Interlock

A similar PLC as that already used for the demonstrator is employed in the LLS QC setups for environmental monitoring and, additionally, as an interlock system [182]. The PLC implements interlock logic and can send interlock signals to power supplies and the MARTA cooling plant. This logic is represented by an action matrix, defining which input signals trigger which interlock outputs. The interlock matrix for the LLS QC setups is provided in Table 7.1. On-detector interlock NTCs should trigger when temperatures exceed 40 °C. Additional Positive Temperature Coefficient Thermistors (PTCs) in the test box and cooling pipes monitor the temperature and disable cooling if the dew point inside the box exceeds safe levels.

The temperature thresholds for the interlock NTCs were successfully tested by simulating temperature changes using a potentiometer. The determined thresholds are:

- SP1 NTC: between 36 °C and 43 °C
- SP2 NTC: between 36 °C and 40 °C
- SP3 NTC: between 39 °C and 41 °C
- SP4 NTC: between 38 °C and 41 °C
- OPTO NTC1: between 37 °C and 40 °C
- OPTO NTC2: between 39 °C and 42 °C

where the range reflects the potentiometer's sensitivity. The actions on the light sensors, the door switches, and emergency buttons were also successfully tested.

Additionally, tests confirmed that power supplies turn off when receiving an interlock signal. For voltage sources, this was verified with open lines. For the LV current source, the tests used high power resistors connected to the power supplies outputs to be able to start with a switched-on current source.

7.1.4 Online software framework

The microservice-based online software framework that was used in the demonstrator tests underwent continued development and is now deployed at all QC test setups. Enhancements include updated GUIs and adaptations to support (pre-)production ITkPix modules. New components provide standardized user interfaces for analyses of test results. The provision of simple Application Programming Interfaces (APIs) for all software components has proven valuable, enabling automated tests by programs using the APIs.

Table 7.1: Interlock matrix for the LLS QC setups in OB. An interlock on the output (rows) is triggered, if any connected input (columns) goes bad. The connection is indicated by a “x”. The condition for the MARTA interlock is $\max(d1, d2) + 10K > \min(c1, c2, c3, c4, c5, c6, c7, c8, c9, c10)$. (From Ref. [197])

	SPC1 NTC	SPC2 NTC	SPC3 NTC	SPC4 NTC	OPTO NTC1	OPTO NTC2	Dewpoint Inlet	Dewpoint Envbox	Light Sens. Top	Light Sens. Bottom	Door Switch 1	Door Switch 2	Door Switch 3	Door Switch 4	Emergency Stop 1	Emergency Stop 2	MARTA status	PTC CO2 In	PTC CO2 Out	PTC CO2 Load In	PTC CO2 Load Out	PTC Box Top	PTC Box Bottom
LV1	x				x	x					x	x	x	x	x	x	x						
LV2		x			x	x					x	x	x	x	x	x	x						
LV3			x		x	x					x	x	x	x	x	x	x						
LV4				x	x	x					x	x	x	x	x	x	x						
HV1.1	x				x	x			x	x	x	x	x	x	x	x	x						
HV1.2	x				x	x			x	x	x	x	x	x	x	x	x						
HV2.1		x			x	x			x	x	x	x	x	x	x	x	x						
HV2.2		x			x	x			x	x	x	x	x	x	x	x	x						
HV3.1			x		x	x			x	x	x	x	x	x	x	x	x						
HV3.2			x		x	x			x	x	x	x	x	x	x	x	x						
HV4.1				x	x	x			x	x	x	x	x	x	x	x	x						
HV4.2				x	x	x			x	x	x	x	x	x	x	x	x						
OPTO					x	x										x	x						
MHFB																x	x						
MARTA	c1	c2	c3	c4			d1	d2							x	x		c5	c6	c7	c8	c9	c10

7.2 The preproduction longeron at CERN

The first preproduction LLSs were a longeron built at CERN and an IHR built at the University of Bonn. These preproduction LLSs use modules with ITkPix-V1.1 FE chips and preproduction versions of the PP0s. The longeron built at CERN is, like the demonstrator longeron, half-filled due to the shortage of preproduction modules. A picture of this preproduction longeron is shown in Fig. 7.3. It uses electrical services for layer 3 and layer 4, with only FE4 of every module electrically connected to an Opto Board. The modules used in the preproduction longeron are listed in Table 7.2 for the modules of the short SP-chain (M6) and in Table 7.3 for the modules of the long SP-chain (M12). The PP0s in this longeron use MOPSV2 chips for module temperature and voltage monitoring. These chips were calibrated before module integration onto the local support, as described in Section 7.2.1.




Figure 7.3: The ITk Pixel OB preproduction longeron in the LLS QC setup at CERN. Just barely visible from the side are the mechanical wire-bond protections on all modules.

Table 7.2: Modules of the M6 SP-chain (A-SP-1) on the preproduction longeron.

Position in Longeron	Name in DCS	Module serial number	HV channel	Position in SP-chain
A-M-B-1	M01	20UPGM22601160	HV2	5
A-M-T-1	M02	20UPGM22601182	HV2	4
A-M-B-2	M03	20UPGM22601181	HV2	6
A-M-T-2	M04	20UPGM22601146	HV1	3
A-M-B-3	M05	20UPGM22601158	HV1	1
A-M-T-3	M06	20UPGM22601171	HV1	2

Table 7.3: Modules of the M12 SP-chain (A-SP-2) on the preproduction longeron.

Position in Longeron	Name in DCS	Module serial number	HV channel	Position in SP-chain
A-M-B-4	M01	20UPGM22601136	HV2	8
A-M-T-4	M02	20UPGM22601161	HV2	7
A-M-B-5	M03	20UPGM22601149	HV2	9
A-M-T-5	M04	20UPGM22601155	HV1	6
A-M-B-6	M05	20UPGM22601167	HV2	10
A-M-T-6	M06	20UPGM22601168	HV1	5
A-M-B-7	M07	20UPGM22601159	HV2	11
A-M-T-7	M08	20UPGM22601143	HV1	4
A-M-B-8	M09	20UPGM22601162	HV2	12
A-M-T-8	M10	20UPGM22601147	HV1	3
A-M-B-9	M11	20UPGM22601176	HV1	1
A-M-T-9	M12	20UPGM22601153	HV1	2




Figure 7.4: PP0s used for the preproduction longeron and for which a MOPS chip calibration was performed.

7.2.1 Calibrations of the MOPS chips

Due to process variations in the MOPS chips, their temperature and voltage measurements can vary between chips. To guarantee accurate temperature and voltage monitoring of the preproduction longeron, the MOPSv2 chips on the PP0s were calibrated before module integration. The PP0s for which a MOPS chip calibration was performed are shown in [Fig. 7.4](#). Note, that the MOPS chip is surface-mounted to the backside of the PCB and not visible in the picture.

For temperature readings, the MOPS chip performs a resistance measurement by measuring the output voltage in a voltage divider with a known input voltage (V_{dd1V2}) and a known reference resistance:

$$R_{NTC} = R \cdot \frac{1}{\left(\frac{V_{dd1V2}}{V_{ADC}} - 1\right)} \quad (7.1)$$

[Figure 7.5](#) shows a simplified circuit diagram for this measurement. Knowledge of the internal ADC reference voltage V_{ref} (assuming an ideal, linear ADC) and the external regulated V_{dd1V2} MOPS output is necessary for accurate resistance measurement and temperature readings. The uncalibrated values for V_{ref} are 0.9 V and for V_{dd1V2} 1.2 V.




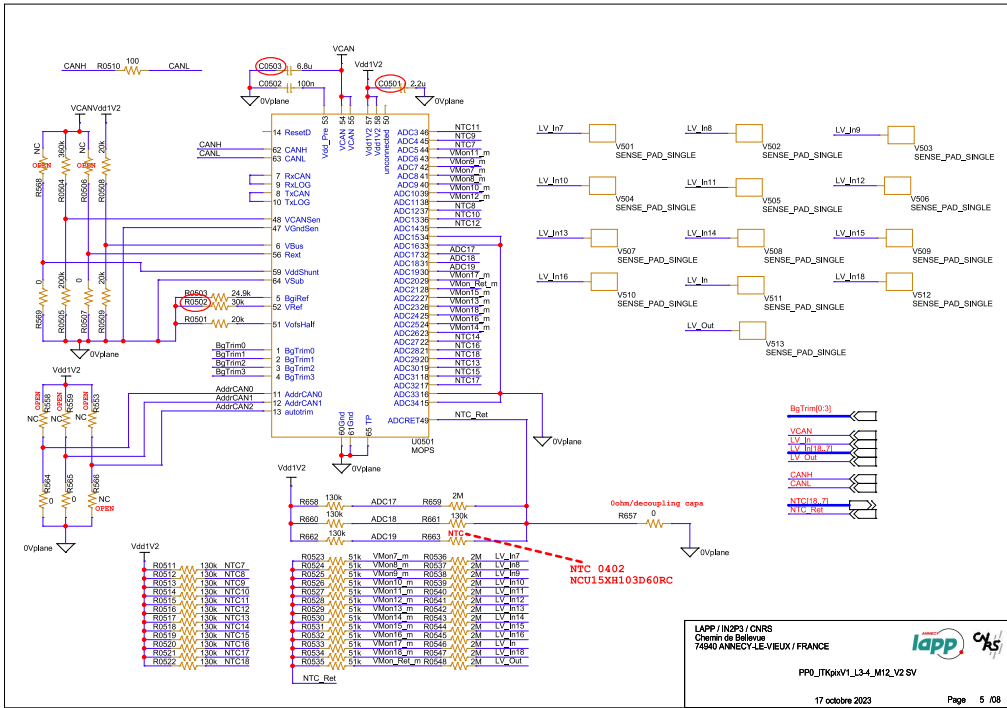
Figure 7.5: Simplified circuit diagram for temperature measurements of a MOPS chip via an NTC. The resistance R is known. The V_{dd1V2} voltage is provided by the MOPS chip and known. The MOPS measures the voltage with one of its ADC channels. The resistance R_{NTC} can then be calculated using [Eq. \(7.1\)](#).

In a first step, the bare PP0 was connected to a CSB and the MHFB for readout of monitoring information. The MHFB provided power to the MOPS chip at approximately $V_{CAN} \approx 1.8$ V. Schematics of the MOPS and pictures of probe points on the PP0 for the MOPS calibration are given in [Fig. 7.6](#). Using a digital multimeter, V_{Ref} (on MOPS pad 52) across R0502 and V_{dd1V2} (on MOPS pad 57) across C0501 were measured. Note, that V_{Ref} is the voltage reference for the definition of the regulator output and the ADC reference voltage V_{ref} [198]. However, the exact relationship between V_{Ref} and V_{ref} was unclear as multimeter measurements yielded values of around 0.6 V for V_{Ref} . It was therefore decided to estimate V_{ref} by comparing the MOPS measurement of V_{CAN} (via V_{CANSen} , an ADC channel for monitoring of the voltage supply) to the V_{CAN} measurement with a multimeter. The MOPS measurement of V_{CANSen} is given in ADC counts and provided by the MHFB. Using $V_{V_{CAN}}/V_{ref} = V_{CANSen}/4096$, V_{ref} was calculated taking into account that $V_{V_{CAN}}$ is the V_{CAN} voltage at the V_{CANSen} ADC channel after a voltage divider. This test could even be performed in an already loaded local support, where the backside of the PP0 is inaccessible, but where V_{CAN} can be taken from the measurement at the MHFB.



Consequently, V_{ref} and V_{dd1V2} were measured as follows:

- M6 (MOPS 9004): $V_{ref} \approx 0.82$ V, $V_{dd1V2} \approx 1.1753$ V
- M12 (MOPS 9506): $V_{ref} \approx 0.83$ V, $V_{dd1V2} \approx 1.184$ V

These results deviate significantly from the default values. The need for calibration is demonstrated, when comparing uncalibrated temperature measurements from an NTC (with default values for V_{ref} and V_{dd1V2}) to calibrated measurements, as shown in [Fig. 7.7](#). The NTC temperature was calculated using the Steinhart-Hart equation in [Eq. \(6.2\)](#) with Steinhart-Hart coefficients obtained in a separate calibration. [Figure 7.7](#) displays the uncalibrated and calibrated temperatures of an NTC as a function of its resistance.



(a)



(b)

Figure 7.6: (a) MOPS schematics of an M12 PP0 for layer 3 or 4. Highlighted are the interesting probe points. The MOPS chip is powered via VCAN. (b) Probe points on the PP0 to perform manual MOPS calibration. On the left is the backside of the PCB with the MOPS visible, and the right is the top of the PCB.




Figure 7.7: Comparison of the uncalibrated temperature measurements of the two MOPS chips used in the preproduction longeron to the calibrated measurement using an NTC. (a) shows the calibrated and uncalibrated temperature as measured from an NTC with given resistance. (b) shows the difference of the uncalibrated temperatures T^* to the calibrated temperature T .

Without calibration, monitored temperatures would be between 1.3 and 1.6 K higher than the actual temperature, as evident in Fig. 7.7(b). Accurate temperature readings are important however, as they are necessary for correctly estimating the thermal performance of the modules.

The calibration described in the paragraph before used only one measurement point (VCAN on the VCANSen channel) for the calibration of all ADC channels. In a second step, multiple ADC channels were calibrated using several calibration points. For this purpose, small bridging PCBs were placed on all module connections on the PP0 except the last one, and a voltage source was connected to the LV line of the PP0 through the CSB. By varying the voltage source output, the ADC input voltage could be adjusted while recording the ADC responses for voltage-monitoring channels. The resulting data set was fitted using a linear model:

$$V = \text{offset} + \text{slope} \cdot \text{ADC counts} \quad (7.2)$$

This was performed for 6 ADC channels in the M6 MOPS and 7 ADC channels in the M12 MOPS. The ADC channel responses to varying input voltages are shown in Fig. 7.8. No significant differences in behavior between channels within a single MOPS chip are visible. An average fit for all channels in one MOPS was created and is also shown in Fig. 7.8. The fit residuals to this average fit are shown as well, further supporting the use of a simple linear model. Since differences are always less than 7 ADC counts and no channel behaves much differently than others, it is justified to treat all channels identically when converting ADC counts to voltages. The one-point calibration using VCANSen is also shown in Fig. 7.8. Its close alignment with the average fit indicates that the single-point calibration

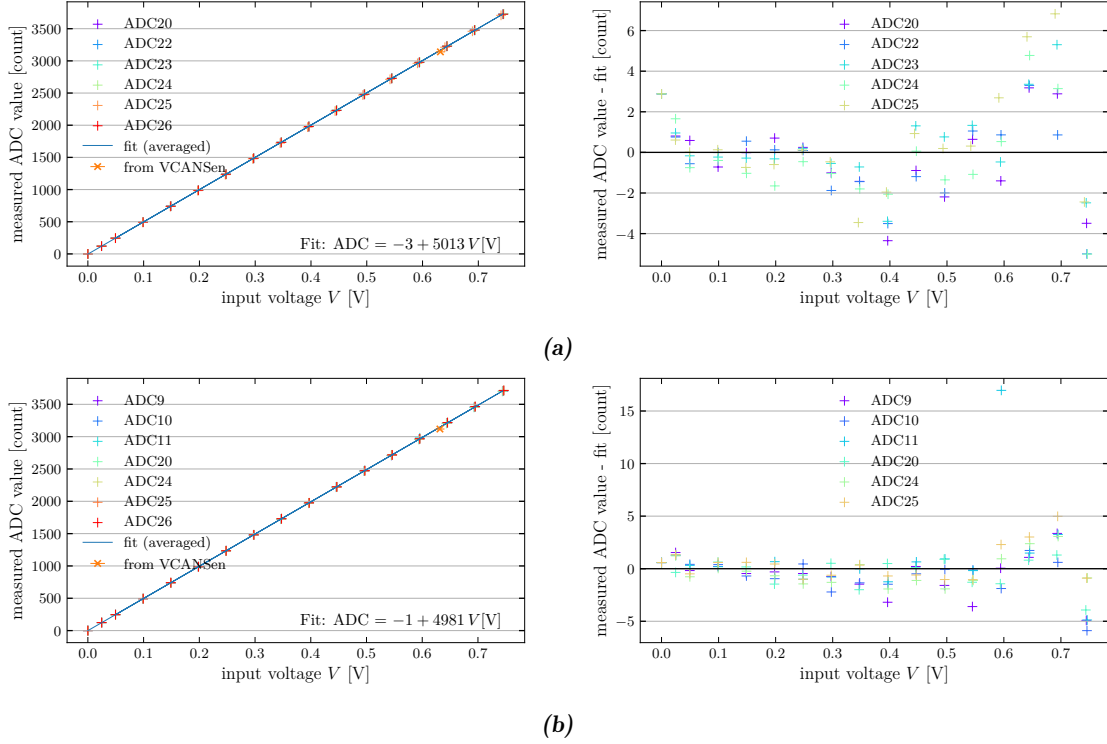


Figure 7.8: On the left is shown the single ADC channel response to varying voltage input. For each ADC channel, a linear fit was performed. For simplicity, an average of the fits is also shown here. The one-point measurement using VCANSen is also indicated. On the right are the fit residuals. **(a)** shows the results for the MOPS chip in the M6, **(b)** shows the results for the MOPS chip in the M12. (Note: The single outlier in the residuals for the M12 fit was likely an error in the manual data recording.)

is already highly accurate. Therefore, the single-point calibration is the recommended method for MOPS calibration during preproduction.

The resulting calibration constants for these PP0s were implemented in the MHFB software and used in subsequent tests with the loaded preproduction longeron to ensure accurate temperature and voltage monitoring of the modules.

7.2.2 Bit Error Rate Tests

After loading the longeron with 18 ITkPix-V1.1 quad modules and installing it in the LLS QC setup, it is now being thoroughly tested. Most tests mirror those performed on the RD53A demonstrator, and the experience gained from those procedures is used to define the QC tests.

FELIX API - itk-felix-sw

Home

Monitoring

FelixRunning

Links :

0

1

2

3

4

5

6

7

8

9

10

11

12

13

14

15

16

17

18

19

20

21

22

23

Optical Power

TX:

✓

✓

✓

✓

✓

✓

✓

✓

✓

✓

✓

✓

✓

✓

✓

✓

✓

✓

✓

✓

✓

✓

✓

✓

Read

Advanced

RX:

✓

✓

✓

✓

✓

✓

✓

✓

✓

✓

Optical link alignment

Read

✓

✓

✓

✓

✓

✓

Electrical link alignment

Read

0 :

✓

✓

✓

4 :

✓

✓

✓

8 :

✓

✓

✓

12 :

✓

✓

✓

16 :

✓

✓

✓

20 :

✓

✓

✓

24 :

✓

✓

✓

25 :

✓

✓

✓

Figure 7.9: Web interface of the FELIX online software service. All 18 electrically connected FE of the preproduction longeron show decoding alignment as check marks in the electrical link alignment section in channels (rows) 0 to 20. Channels 24 and 25 are required for direct Opto Board communication.

One of the first tests performed with the preproduction longeron following electrical connection was a check of the data transmission quality using BERTs. Successful communication was established immediately upon powering all modules, as indicated by the decoding alignment status in the FELIX system. Figure 7.9 shows the FELIX service web interface displaying decoding alignment status, confirming that the Aurora stream from the FEs is successfully decoded within FELIX. All FEs installed on the preproduction longeron are therefore functional, and their thermal and electrical performance as well as the performance of the sensors is currently under evaluation.

To evaluate data transmission quality, BERTs were performed for each electrically connected FE at both warm and cold temperature. Every module on the preproduction longeron as listed in Table 7.2 and Table 7.3 was tested. At a CO₂ cooling temperature of around 10°C, the module NTCs measured temperatures of around 22°C. A second test set used a CO₂ cooling temperature of around −22°C with module temperatures near −7°C. The nominal current for powering modules was 5.88 A. All modules powered up correctly at cold temperatures as expected with the new FE chip version, with no cold-start issues observed in the preproduction modules in contrast to the problems with the ShuntLDO encountered in the RD53A demonstrator.

For each module a BERT ran for approximately 216 s, transmitting 274 877 906 944 bits using a PRBS7 pattern from the module to the lpGBT on the Opto Boards. At both temperature conditions and for all modules, zero bit errors were detected. This corresponds to an upper limit on the bit error rate (at 95 % confidence level) of 1.09×10^{-11} .

```

$ python3 influxdb2csv.py -h
usage: influxdb2csv.py [-h] [-c CONFIG] [-o OUTDIR]

Dump WinCC OA data from an InfluxDB instance to csv

optional arguments:
  -h, --help            show this help message and exit
  -c CONFIG, --config CONFIG
                        Configuration file (defaults to
                        "datarequest.json")
  -o OUTDIR, --outdir OUTDIR
                        Output directory (defaults to current working
                        directory)

```

Listing 7.1: Help message of the InfluxDB database export tool.

7.3 Software developments toward uniform and automatic LLS testing

As previously noted, testing procedures are being defined to ensure efficient testing with results that can be compared across the five integration clusters that will build OB LLSs [199]. Within the scope of this thesis, several software tools were developed to implement some of these procedures, as described below.

7.3.1 Retrieval of archived DCS data

Every LLS QC setup uses identical infrastructure hardware and the same DCS software project developed in this thesis (and presented in Section 9.3). For the archival of DCS data, an InfluxDB⁴ database is hosted on the DCS server of each setup. Docker images⁵ were adopted for deployment of the necessary software. This was a decision that built on the experiences made with the microservices of the online software framework.

To retrieve archived DCS data from the InfluxDB database, a Python tool⁶ was developed that connects to the local InfluxDB instance and exports the DCS data to Comma-Separated Values (CSV) files. Listing 7.1 shows the help documentation of this tool.

Given a configuration file for the data export tool as illustrated in Listing 7.2, the database is queried and for each “Data” element data is written to a CSV file containing two columns:

⁴An open-source time series database, <https://docs.influxdata.com/influxdb/v1/>

⁵<https://gitlab.cern.ch/atlas-itk-pixel-systemtest/influxdb-and-grafana-for-lls-qc-dcs>

⁶<https://gitlab.cern.ch/atlas-itk-pixel-systemtest/influxdb-tool-for-dcs-data>

```

{
  "InfluxDB": {
    "url": "https://localhost:8086/",
    "database": "winccoa",
    "username": "lls qc1user",
    "password": "",
    "verify": false,
  },
  "Timerange":{
    "start": "1970-01-01 00:00:00",
    "end": ""
  },
  "Data": [
    {
      "alias": "OB/L3/B01/C/SP1/M06/Tmon"
    }
  ]
}

```

Listing 7.2: Example configuration file for the InfluxDB database export tool. One can specify connection details about the InfluxDB instance, a time range for the exported data, and the name of the data to export.

one for the timestamp and one for the recorded value. The script was validated using the LLS QC InfluxDB at CERN.

7.3.2 Programmatic remote control of the DCS

For module tests on LLSs, an orchestration tool called module-qc-tools⁷ was adopted. This framework is already used for systematic module tests during their production. Reusing this tool with identical test parameters ensures comparability between module production tests and LLS testing. Some module-qc-tools tests require DCS data (for example module temperatures) or even power supply control to turn on and off power to the modules. During standard module QC, this is achieved by directly interfacing with lab power supplies, typically via serial connections.

A Python script called DCS adapter⁸ was developed to replace this connection in the LLS QC setup. It leverages the REST API of the LLS QC DCS, that was developed within the context of this thesis and is presented in Section 9.3.12. Figure 7.10 illustrates the connection scheme of the DCS adapter. The adapter requires a configuration file

⁷<https://atlas-itk-pixel-mqt.docs.cern.ch/latest/>

⁸<https://gitlab.cern.ch/atlas-itk-pixel-systemtest/module-qc-winccoa-dcs>




Figure 7.10: Connection scheme of the DCS adapter Python software that is in between the module-qc-tools and the API of the DCS software (WinCC OA) in the LLSs.

specifying the URL of the DCS API and the module serial number of the module under test. Using this serial number, the adapter retrieves module data from the DCS or controls the module’s SP-chain. This requires the DCS to know the module serial numbers.

The HTTP REST API is interfaced through Python classes and functions generated by an API client generator⁹. The DCS software provides the necessary API definition following the OpenAPI 3.0 specification¹⁰. Consequently, should the API change or be updated, Python code can quickly be regenerated.

The DCS adapter was designed as a drop-in replacement for existing adapters to lab equipment in module QC tests. It was validated by comparing its functionality and outputs against established adapters. After successful validation in the empty LLS QC setup, it is now used for testing modules in the preproduction longeron.

Given the significant focus of this thesis on DCS developments, including developments for the LLS QC DCS, the remainder of this thesis is dedicated to that topic.

⁹<https://github.com/openapi-generators/openapi-python-client>

¹⁰<https://spec.openapis.org/oas/v3.0.3.html>

Detector Control Systems

A major part of this thesis is the development of a Detector Control System (DCS) for the ATLAS ITk upgrade. A technical description of the new silicon tracker was presented in [Chapter 4](#). This chapter introduces control systems following closely Ref. [179]. An overview of the DCS of ATLAS is given in [Section 8.1](#). The typical DCS hardware, that is used in ATLAS, is briefly described in [Section 8.2](#). [Section 8.3](#) presents the software, specifically the SCADA software that is employed in the DCS and that was used for many developments carried out with this thesis and discussed in [Chapter 9](#). [Section 8.4](#) details the FSM as the operational model of a control system for a large detector. This forms the basis for [Section 9.4](#), where an FSM model for the ITk Pixel detector is presented.

8.1 ATLAS DCS

The ATLAS experiment is a highly complex detector comprising many subsystems. As such, it requires a sophisticated control system to enable coherent control and supervision while preventing accidental mishandling. The purpose of the DCS is to prepare and maintain the detector in a well-defined state for high-quality physics data taking. The DCS ensures the effective and safe configuration of subsystems and equipment, that control and monitor all services, infrastructures and sub-detectors involved in the experiment [200]. Any abnormal behavior must trigger alarms within the DCS. The type of alarms and the alarm-handling distinguishes the DCS from the Detector Safety System (DSS) and the CERN Safety System (CSS), as the DCS only signals “Level 1” alarms (for example a tripped power supply) as defined in Ref. [201]. Equipment protection during potentially

dangerous conditions is the responsibility of the DSS [202, 203], while the CSS [204, 205] acts when “potential danger to human life or the environment” (Level 3 alarm as defined in Ref. [201]) is present. The DCS thus serves as a more flexible system enabling highly complex detector operations.

From a high-level perspective, ATLAS is operated by two complementary systems: the DCS and TDAQ. During LHC operation, TDAQ handles readout of detector data generated by particle collisions and assumes control of the DCS. With over 150 PCs, the DCS is a highly distributed, hierarchically organized system for detector operation. It supervises individual detector components and supporting infrastructure during both physics data taking and maintenance/calibration periods. This supervision involves reading and processing on-detector or environmental sensor data along with monitoring the state of its active equipment. These parameters can be archived for future reference. The DCS analyzes live data, detects errors, and alerts operators via alarms in the control room, enabling manual or automatic corrective actions. All of this must be done in an accessible manner to the operator in the control room. The DCS also provides the operator with control over the detector to bring it into the desired operational state.

For a large experiment like ATLAS, that is embedded in a whole set of machinery and hardware at CERN, DCS interoperability with external systems such as LHC run control, ATLAS magnets or the DSS is essential. Key requirements therefore include platform-independent communication interfaces and scalability to manage large sets of parameters.

The implemented architecture for the DCS as shown in Fig. 8.1 comprises Front-End (FE) equipment and Back-End (BE) systems. FE equipment includes DCS hardware such as power supplies, environmental sensors and cooling circuits. Commercial systems are generally preferred for the FE equipment to ensure long-term maintainability and efficient development. Custom designs are developed only when commercial solutions are inadequate. For instance, the custom ELMB [200] hardware component serves as a multipurpose I/O device, for example for temperature or humidity sensors. An independent hardware interlock system enables detector shutdown in case of an emergency. The communication between FE and BE happens primarily via CAN/CANopen or Ethernet/OPC-UA protocols. The BE consists of industrial rack-mounted PCs running software to integrate FEs controls. The Supervisory Control and Data Acquisition (SCADA) software [207] SIMATIC WinCC Open Architecture (WinCC OA), formerly PVSS, [208] was selected by all experiments at the LHC following a careful and detailed selection process [209]. In ATLAS, a distributed system of WinCC OA instances ensures high performance and scalability of the control system with sub-detector-specific projects grouped under their respective operational hierarchies. Each sub-detector has a certain level of operational independence for local control.

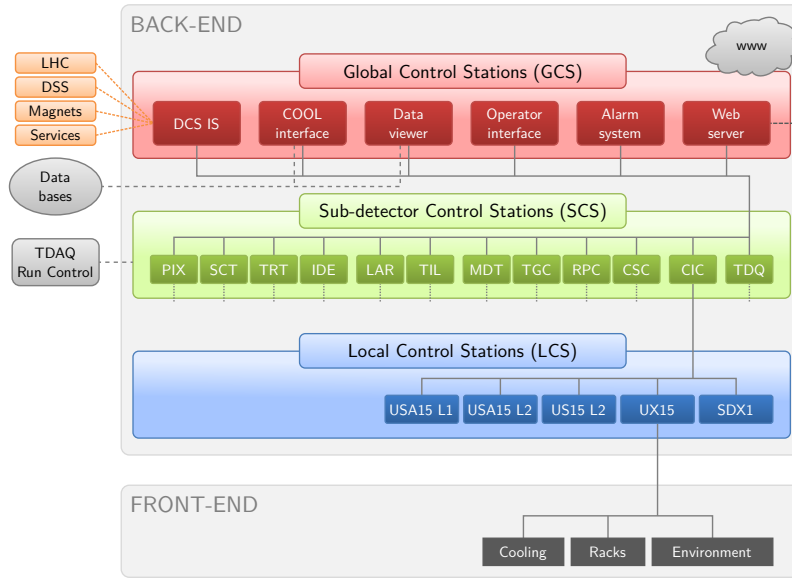


Figure 8.1: Schema of the DCS architecture. The detector hardware is interfaced to the DCS Front-End systems (sensors, PSUs etc.). The data is processed by a hierarchical Back-End system with Local Control Stations for the detector and service caverns and presented to the control room operator using a single User Interface. External systems are interfaced through the DCS Information System (IS). (Reproduced from Ref. [206], licensed under CC BY 4.0)

Given significant commonality in BE development across LHC experiments, the Joint Controls Project (JCOP) collaboration¹ was established [210]. The JCOP collaboration defines standards for the use of DCS hardware and software, provides OPC-UA² servers for power supply integration, and develops the JCOP framework components for WinCC OA and the FSM software. The FSM models the hierarchical structure of the control system. General information about the FSM software system is given in Section 8.4. Reflecting ATLAS’s detector/sub-detector hierarchy with operational independence at multiple levels, the distributed FSM aligns with this structure in the operational model.

Detector and DCS parameters (such as calibration constants, voltage settings and alarm thresholds) are stored in a Configuration Database (an Oracle database). The settings are retrieved as “recipes”.

It was decided to separate this from the detector conditions data (detector condition at a given time), that is stored in a Conditions Database, to enable efficient access to conditions data by offline calibration, reconstruction, and analysis applications [200, 211].

¹<https://jcop.web.cern.ch/>

²Open Platform Communications Unified Architecture, see Section 8.3.3

8.2 The DCS Front-Ends

To safely operate the detector and read out physics data, a complex control system is required. The environment must be continuously monitored, with automatic actions taken to minimize risks to equipment or personnel. The hardware for the DCS FE components is described in the following sections.

8.2.1 ELMB

The Embedded Local Monitor Board (ELMB) features an 8-bit microcontroller with 64 analog input channels multiplexed to a 16-bit ADC [200]. Its general-purpose software provides access to 16 digital output and 16 digital input channels. The analog channels support sensor readout for temperature (for example via NTC sensors), humidity or dew point. Using its built-in CAN controller, the ELMB asynchronously transmits data to a BE system for slow environmental monitoring.

The components were selected to enable operation in high magnetic fields and under ionizing radiation. Power is distributed to the board through the CANbus cable, with two wires for communication/monitoring data and two additional wires for power.

For the HL-LHC upgrade, where radiation levels will exceed ELMB specifications, a radiation-hardened drop-in replacement board was developed, referred to as ELMB2 [212]. Irradiation tests confirm that the ELMB2 withstands a TID of up to 200 Gy. The ITk Common DCS will employ the ELMB2 for environmental monitoring (environmental temperature, humidity and radiation) within the ITk volume [213].

8.2.2 Power supplies

ATLAS employs commercial VME crates that house power supplies in rooms not affected by radiation for ELMBs and detector electronics. As a critical safety requirement, power supplies must provide over-voltage and over-current protection.

The DCS BE monitors and controls these power supplies, including setting and reading of voltages and currents, via OPC-UA servers acting as middleware. The power supplies commonly communicate via CAN or SNMP³ with the OPC-UA servers.

³Simple Network Management Protocol

8.2.3 Interlock

While the DCS software manages slow control of the detector, a hardware-based fast interlock system addresses all detector safety-critical scenarios [214]. ATLAS' subsystems implement independent interlock solutions that are able to switch off necessary power supplies while operating completely independently of the control system. For the current Pixel detector, this system protects sensitive detector elements, as especially overheating is a major risk and can cause irreparable damage. Temperature and other sensors feed directly into the interlock system, where an FPGA evaluates all inputs to trigger appropriate interlock alerts.

Though entirely hardware-based, the interlock system must communicate and explain its actions to the DCS. For instance, the ATLAS IBL [62] uses ELMBs to monitor interlock signals [181]. These ELMBs can transmit test signals to the interlock system via digital output channels but cannot alter the decision logic of the FPGA. The IBL's Interlock Matrix Crate (IMC) served with minor modifications as the demonstrator's interlock solution, as already detailed in [Section 5.5](#).

For the ITk upgrade development focuses on a new interlock solution for its pixel and strip detectors based on a set of Local Interlock Safety Systems (LISSYs) and Main Interlock Crates (MICs) [215]. This solution is detailed in [Section 9.1](#).

8.3 The DCS Back-Ends

The hardware platform for the Back-End (BE) system consists of industrial, rack-mounted server machines that feature redundant, hot-swappable power supplies and HDDs in a RAID system. The machines run different software components to receive information from the Front-Ends, process it, and present the information to operators in the control room. They also enable control of the active DCS hardware.

The OPC-UA servers that facilitate communication between the FE and BE run on these machines, as well as the SCADA software. The JCOP framework provides guidelines, common components and tools for easy integration of hardware into the SCADA system. The SCADA software WinCC OA parametrizes all control and monitoring information and provides GUIs to operators and experts. The SCADA software, the JCOP framework and the middleware solutions are briefly described below. For operation of this complex detector system, WinCC OA is extended with a Finite State Machine (FSM), as described in [Section 8.4](#).

8.3.1 SCADA software

The SCADA software WinCC OA⁴ [208] was selected by all LHC experiments among other factors due to its scalability, platform independence, adaptability and support for industrial standards. This section describes key concepts of WinCC OA and their advantages for high-energy physics SCADA applications, as well as terminology for the ITk Pixel DCS technical implementations as discussed in Chapter 9.

WinCC OA is a software package designed for automation technology to monitor and control systems. Its architecture uses a client-server model, where servers provide information and clients consume it. Each WinCC OA project is functionally separated into several processes called *managers*, which communicate with each other following the client-server principle. Figure 8.2 shows the managers in a typical WinCC OA project. The only mandatory managers are the database manager (DB) and the event manager (EV). Running only required managers in a project reduces resource requirements. Additional managers/processes can be started when needed, utilizing available cores of multi-core processors. Data processing and inter-process communication in WinCC OA are normally event-oriented, meaning only value changes are processed or forwarded immediately, making communication very efficient. Managers exchange data via TCP messages which allows projects to span over multiple computers across IP networks. If required, this enables to connect a WinCC OA project across multiple computers to distribute load. The distribution manager (DIST in Fig. 8.2) further connects several autonomous WinCC OA projects into *distributed* systems. This architecture offers scalability from small projects to large multi-server deployments.

WinCC OA uses an object-oriented data model based on *Data Points (DPs)*. DPs are hierarchically structured pieces of information (following a template called *Data Point Type (DPT)*), often associated with devices, containing a defined number of process variables. All process variables and their values are stored as DPs in a database (handled by the database manager), with value changes handled by the event manager in an event-based approach using multi-threaded callback routines. In a distributed system, projects can access data from all connected systems.

The software also provides Graphical User Interfaces (GUIs). For users, an interface to the database is provided with the Para program⁵. Para displays all DPs available in the distributed system. Process information contained within a DP is provided by *Data Point Elements (DPEs)*. Figure 8.3 illustrates how the DPs and DPEs are represented in Para. DPEs support types like integers, floats or strings, and associate metadata via so-called *configs*. Every DPE includes configs for the “original” value and “common” information

⁴<https://www.winccoa.com/>

⁵Para stands for *parameterization*.

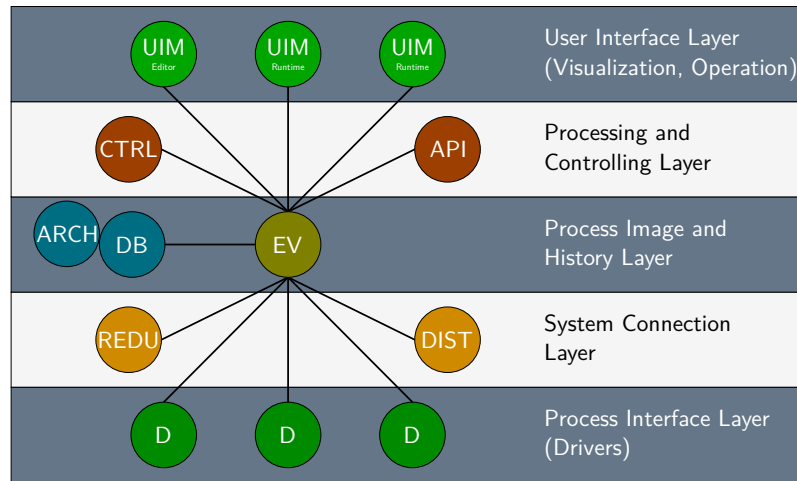


Figure 8.2: The manager (function-specific units) of a WinCC OA project. The event manager (EV) keeps a current image of all process variables in memory and is the communication center of WinCC OA. The database manager (DB) is the link to a database and handles the configuration data of the application. The event and database manager are the two only *mandatory* managers. The distribution manager (DIST) handles the connections in a distributed system and exchanges data point information. The redundancy manager (REDU) can be used to set up redundant projects and guarantee high availability of the system. The drivers (D) can be used to connect to hardware. For example, WinCC OA offers a built-in OPC-UA client. Control managers (CTRL) are used to run control scripts in the background and UI managers can display panels. Via the API, one can extend WinCC OA with custom functionality as an independent additional manager. (Adapted from the WinCC OA documentation)

(e.g. additional descriptions or unique aliases for alternative addressing). Aliases enable alternative data views beyond DPE names. Other configs of a DPE manage archiving or alert handling. Thresholds can be defined per DPE such that values above (or below) trigger alerts. The control flow for the alert handling is shown in [Fig. 8.4](#).

For the development of scripts and design of panels, an Integrated Development Environment (IDE) called Gedi is provided. Panels are run in a User Interface Manager (UIM) called VISION and can display system states, parameter values, trends and control options. Scripts use a C-like scripting language called “Control” and are executed in the background of panels or by Control managers (CTRL). The Control language offers functions to access and change data in the database via `dpGet()` and `dpSet()`. A Control script or a GUI panel can use the `dpConnect()` function to subscribe to changes of a DPE. The performance of a project correlates strongly with the total number of DPs (affecting mainly memory) and the number of `dpConnect()` callback functions that need processing.

The drivers (D) as shown in [Fig. 8.2](#) enable communication with hardware. WinCC OA includes built-in drivers like OPC-UA clients to connect to OPC-UA servers (see

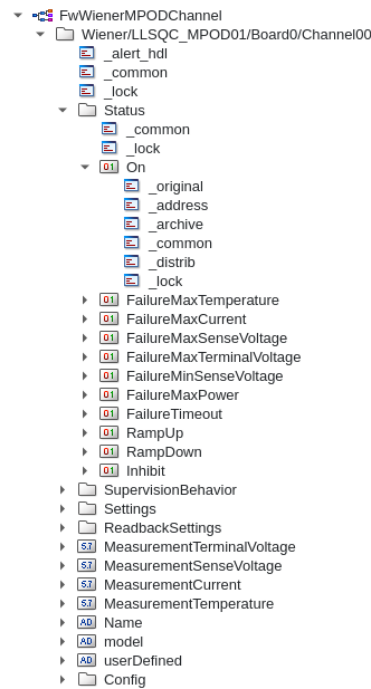


Figure 8.3: Example for data points in Para. At the very top is the data point type. Below, one data point of this type is shown with its data point elements. The data point name is usually derived from the hardware it represents. For a better understanding in which way the hardware will be used, one can assign aliases to data points or data point elements.

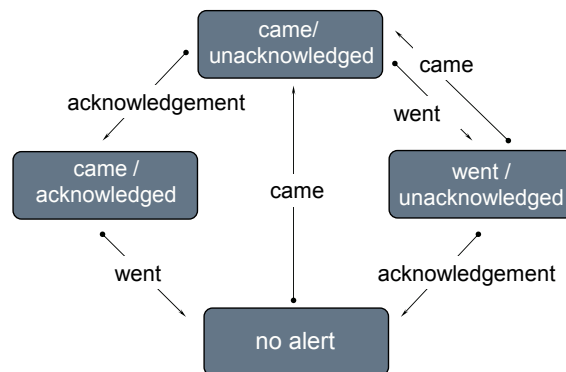


Figure 8.4: Alert handling for data point elements in WinCC OA. The state of the alert is shown in rectangles. “Came” and “went” transitions happen, when the condition for the alert is fulfilled or not anymore. Implementing hysteresis is possible. A user has to acknowledge the alert. Several severity levels for the alert can be defined.

Section 8.3.3). The ITk Pixel DCS uses these almost exclusively. The FASER DCS also employs the built-in MODBUS driver to connect to some of their hardware as explained in Section 10.2. Hardware connections are established via peripheral address configurations (`_address`).

The “OA” in the software package’s name denotes Open Architecture, meaning that WinCC OA provides a C++ API to extend functionality with custom managers. The C++ API is used to enable communication via protocols like DIP and DIM (see Section 8.3.4). It was also used to build the WinCC OA HTTP REST API manager as described in Section 9.3.12.

In spring 2024, the community upgraded from WinCC OA v3.16 to v3.19, coinciding with an operating system migration to AlmaLinux 9 for all computers in the lab. WinCC OA can archive values in the project’s database into an external database for persistent storage using an archive manager (ARCH in Fig. 8.2). The external database used in production systems is an ORACLE⁶ database. Version 3.19 saw an improved support for InfluxDB⁷ as an external database. Archiving configurations (e.g. smoothing settings) reside in DPE `_archive` configs.

8.3.2 JCOP framework

To streamline the development of WinCC OA panels and scripts, and because many required functionalities are common across LHC experiments, the JCOP⁸ framework provides standardized tools for WinCC OA. This ensures high conformity and simplified maintenance. The framework offers components and libraries for hardware access/control (especially for CERN-standard devices), access control, alarm tools, trending tools, and an FSM library to model the hierarchy of the control system. These JCOP Framework components serve as extensions to WinCC OA. When connecting to hardware, components interface with JCOP-provided OPC-UA servers.

8.3.3 OPC-UA servers

The Open Platform Communications Unified Architecture (OPC-UA) protocol is the next-generation evolution of the Open Platform Communications standard [216], designed for exchanging complex data between distributed industrial automation systems. Unlike its predecessor, OPC-UA is platform-independent and supports binary communication over TCP/IP.

⁶A relational database management system, <https://www.oracle.com/database/>

⁷An open-source time series database, <https://docs.influxdata.com/influxdb/v1/>

⁸<https://jcop.web.cern.ch/>

OPC-UA provides extensive capabilities for modeling shared data. The client can query specific portions of the data model via discovery requests, receiving information about a subset of the available structure. To receive data, the client can either poll data from the server or subscribe to data changes. Connection failures are automatically recognized by both the client and the server.

While the dependency on Ethernet renders OPC-UA unsuitable for environments with strong magnetic fields, it has been adopted as the new standard middleware outside the detector for the ATLAS DCS [217]. The quasar framework [218, 219] enables efficient development of OPC-UA servers: By defining an object-oriented information model in an XML server design file, the quasar framework can generate an executable OPC-UA server application.

The OPC-UA server usually connects to the hardware via CAN or SNMP. When CAN is used, a CAN module, which establishes connection to the hardware, needs to be connected to the host machine. The server may also perform conversion from raw to physical values.

A configuration file that is read by the server at start up defines the names of all data sources which a connected OPC-UA client can use to read from and write data to. For example, the Wiener OPC-UA server interfaces Wiener power supplies⁹ with the SCADA system: Power supplies connect via Ethernet using SNMP to the Wiener OPC-UA server, which then communicates with WinCC OA through OPC-UA.

8.3.4 Other communication protocols

DIM

The Distributed Information Manager (DIM)¹⁰ protocol provides efficient and reliable inter-process communication across different platforms [220]. Its communication mechanism is based on the publish/subscribe method and allows for asynchronous communication and updates to multiple destinations. The use of a central name server allows clients to remain unaware of the data source prior to the first connection. Additionally, the protocol supports sending commands from subscribing clients to servers. The underlying network communication is implemented via TCP/IP sockets. DIM serves as the communication protocol for the processes of the FSM.

⁹power supplies by the Kontron HARTMANN-WIENER GmbH, 51399 Burscheid, Germany

¹⁰<http://dim.web.cern.ch/>, <https://lhcb-online.web.cern.ch/ecs/fw/FwDim.html>

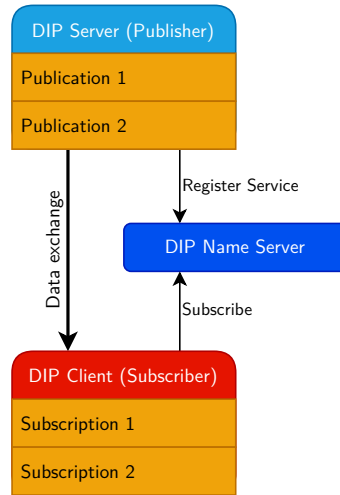


Figure 8.5: Data exchange between two systems using the DIP protocol.

DIP

The Data Interchange Protocol (DIP)¹¹ protocol [221] is used by ATLAS to interface with external control systems, such as those for magnets or cooling [200]. Built on top of DIM, DIP is a message-based communication system enabling reliable exchange of small volumes of soft real-time data between loosely coupled, heterogeneous systems. Publishers register publications with a central name server, while clients query the server for available publications and subscribe to one or more. After subscription, clients continuously listen to any data changes in the publication. Figure 8.5 schematically illustrates this data exchange. A key feature is the system’s ability to detect and report connection failures to clients, with automatic reconnection once network issues are resolved. WinCC OA can be extended via a JCOP Framework component to function as both a DIP server and a DIP client.

MODBUS

The MODBUS protocol [222] is an industrial automation standard commonly used for PLC communication. MODBUS can run over a serial connection or, for the use cases in this thesis, over TCP/IP. MODBUS follows a client/server model where only the client initiates communication by a request, either for information (reading) or sending a command (writing). The MODBUS message payload consists of two fields: a code field and a data field [223]. The client specifies a function code in the code field and includes request-specific data in the data field. Upon successful communication, the server responds with the same function code and response data in the data field. Servers address data items

¹¹<https://dip.web.cern.ch/>

using unsigned integers starting at 0. The possible data item types are called “discretes” (read-only single bits), “coils” (read/write single bits), “input registers” (read-only 16-bit words) and “holding registers” (read/write 16-bit words). Function codes define requested services and must match the referenced data item type (e.g. a specific function code reads holding registers). WinCC OA includes a built-in MODBUS driver, as used in the FASER DCS to communicate with the interlock system or in the ITk LLS QC setups to communicate with the MARTA cooling plant.

CAN/CANopen

The Controller Area Network (CAN) standard is a serial bus protocol enabling node-to-node communication without a master [224]. It is a message-based protocol, where data is transmitted sequentially. Every device on a bus is identified with a unique node-ID and can send and receive data. If two devices want to send data at the same time, the device with the lower priority will automatically stop sending data on the bus.

CANopen [171] is a high-level protocol layered over CAN. The ELMB general-purpose software conforms to the CANopen DS-401 Device Profile for I/O modules [200]. The bus is physically realized over a twisted pair of wires. Advantages of CAN/CANopen include low cost, material savings, and insensitivity to magnetic fields due to differential signaling, making it ideal for slow monitoring data readout from detector components. For ITk Pixel, CAN/CANopen is the selected communication method for the MOPS chip. To interface WinCC OA with devices connected via CAN, an OPC-UA server is put in between as middleware.

A summary of protocols used by the DCS in this thesis is given in [Table 8.1](#).

Table 8.1: Summary of the protocols as used by the DCS in this thesis and some examples of their use cases.

Protocol	Underlying layer	Communication model	Use cases
OPC-UA	TCP/IP	client/server (synchronous and asynchronous)	Communication between WinCC OA and servers to connect to hardware (power supplies, ITk interlock, MOPSHUB4Beginners, ...)
SNMP	UDP/IP	client/server (synchronous)	Communication between server and power supplies
HTTP	TCP/IP	client/server (synchronous)	Control of online software and DCS for orchestration
DIM	TCP/IP	publish/subscribe (asynchronous)	Communication between FSM nodes
DIP	TCP/IP	publish/subscribe (asynchronous)	Communication with cooling plant in SR1
MODBUS	TCP/IP	client/server (synchronous)	Communication with FASER interlock or MARTA
CANopen	CAN	client/server (synchronous)	Communication with MOPS

8.4 Operation: Finite State Machines for Detector Control Systems

Finite State Machines (FSMs) are widely used in modeling of systems that can be in a finite number of pre-defined states. The ATLAS detector also follows this paradigm, with individual detector components mapped to FSMs and grouped across multiple hierarchical levels to present the state of the whole experiment. In this model, the current state of the experiment is determined by many input parameters. Additionally, depending on the current state, different control actions are allowed to enable safe transitions to other states. A key advantage of FSMs is their ability to simplify complex systems and enable decision-making processes. These are essential requirements for the safe operation of the ATLAS detector by a small shift crew of non-experts in the control room.

This section forms the basis for the discussion of an FSM model for the ITk Pixel detector in [Section 9.4](#). First, a brief introduction to the theory of FSMs is provided in [Section 8.4.1](#). Since WinCC OA lacks native FSM functionality, this is implemented using the SMI++ framework detailed in [Section 8.4.2](#). The connection between WinCC OA and SMI++ is

explained in [Section 8.4.3](#). ATLAS-specific extensions to this connection are then described in [Section 8.4.4](#).

8.4.1 Theory of Finite State Machines

An FSM is a formalism for specifying an object's behavior based on its state [225]. It is an abstract machine that is in exactly one of a finite number of states at any given time and can change from one state to another via a transition activated because of some input. Inputs can be external actions and/or reactions to some changes of conditions. Common applications include verifying input sequences in code locks, managing elevator stop sequences or controlling traffic lights. Control systems are typically modeled as FSMs with output, where the current state is determined by many input parameters, and system-dependent control actions can be automatically executed or offered to an operator. This enables non-experts to control large systems.

An FSM is defined by its set of possible states, the initial state, the inputs which the FSM recognizes, the permitted transitions, and the available state-dependent actions. Graphically, this is represented using a directed graph called *state diagram*. [Figure 8.6](#) illustrates such a diagram, where the two states of a turnstile and the possible transitions are shown.

Two fundamental types of FSMs with output exist [226]. In the Moore machine, the output depends solely on the current state. In the Mealy machine, the output depends on both the current state and an input. These models are functionally equivalent (one can be transformed into the other), but the Moore machine is conceptually simpler due to its independence from input variables, though it typically requires more states than an equivalent Mealy machine. In practice, often a mix of both is used.

For complex systems, basic FSMs may be nested or set up in a hierarchy. If higher-level FSMs can react to state changes from lower objects, the FSM is called *event-driven*. The control systems of complex detectors use hierarchical, event-driven state machines as explained in the following sections. Since parent FSMs must handle all possible states of child FSMs, the use of Mealy machines, which require fewer states, is preferred. Mealy machines have been implemented in the ITk Pixel and FASER Preshower FSMs within the context of this thesis.




Figure 8.6: State diagram for an FSM modeling a turnstile. The turnstile has two states, “Locked” and “Unlocked”. Inserting a coin when in the Locked state transitions the turnstile into the Unlocked state, where the bar can turn. A person pushing through the turnstile changes the state back to Locked. The next person must insert a coin again, a simple push will not work. The arrow into the Locked node from the black dot indicates the initial state. (Chetvorno, CC0, via Wikimedia Commons)

8.4.2 State Manager Interface framework

At the core of the FSM functionality of the control system lies the SMI++ framework [227, 228], an extended version of the State Manager Interface (SMI) [229], originally developed for the DELPHI experiment [230] at LEP but rewritten in C++. SMI¹² is a framework for designing and implementing distributed control systems [227]. It models physical systems as a collection of hierarchically structured objects that behave as FSMs. This hierarchy enables independence of subtrees and decomposes complex systems into manageable entities, aligning with the operational model of large experiments.

The FSM logic of objects is defined using the State Manager Language (SML), which supports event-driven and rule-based logic allowing for the creation of output-generating FSMs: Objects can react to state changes of other objects and automatic actions can be executed to handle errors or automate processes. SML permits the execution of automatic actions by implementing either a Moore machine (where actions depend only on the current state) or a Mealy machine (where actions depend on both the current state and inputs). The language fully specifies the class of an object such as its states, actions and rules, with actions described as sequences of instructions typically executed asynchronously. Common instructions include the `do` instruction to send commands to other objects, the `if` instruction to test the state of an object, the `wait` instruction to synchronize the execution of actions, and the `move_to` instruction to terminate an action and transition to a new state.

¹²<https://smi.web.cern.ch/>

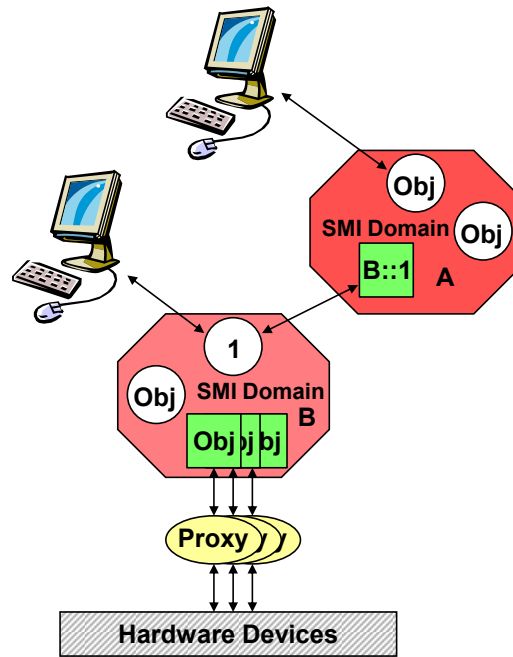


Figure 8.7: SMI++ runtime environment. The quadratic objects are of type associated. These are different in that they cannot execute any action: a command received by an associated object is passed through. Conversely, a state change occurring in the entity controlled by the associated object is faithfully reflected. To address an object from another domain, “domain::object” is used. (Adapted from Ref. [231])

Objects may represent concrete entities (then they are called *associated* objects), such as power supplies or temperature sensors, or abstract entities, such as groups of power supply channels. For state determination, associated objects interface with hardware through a proxy (normally WinCC OA), whereby the associated object only reflects the state of its proxy and relays actions to it. Abstract objects implement their logic entirely in SML.

Logically related objects can be grouped together into SMI domains, which run as separate (potentially distributed) processes. Associated objects from one domain may link to other domains, commonly the top-level object of the other domain. Inter-domain communication is handled by the DIM protocol. A State Manager executes the SML code of a domain and becomes the “State Manager Process” of that domain. Figure 8.7 illustrates an SMI++ runtime environment containing two SMI domains with abstract and associated objects as explained above.

Objects (both associated and abstract) can be grouped into “object sets” to simplify the manipulation of many objects. Objects may be dynamically added to or removed from these sets.

The generation of the SML code for experiments is typically done through WinCC OA tools provided by the Controls Hierarchy framework component.

8.4.3 The Controls Hierarchy framework component

The Controls Hierarchy framework component (FwFsm)¹³ is an FSM toolkit for WinCC OA [231], providing a generic, platform-independent implementation of a hierarchical state machine utilizing SMI++. It offers tools to define FSM objects for experiments, structure them in hierarchy trees, generate the SML descriptions, start and stop the SMI++ runtime and interact with the FSM objects through WinCC OA Control scripts. WinCC OA serves as a proxy for associated FSM objects, a capability that is enabled by its C++ API.

Not only associated FSM objects are interfaced, however. The attributes of all instantiated FSM objects within an SMI domain are made persistent as data points in the associated WinCC OA project. This facilitates archiving of the FSM states and transitions, communicating the experiment status to operators and providing operators control over the FSM via GUIs.

Figure 8.8 illustrates a generic FSM tree generated by the FwFsm component. Commands propagate downward through the tree, while statuses propagate upward. The FwFsm component uses the following terminology for the different types of FSM object classes:

1. Device Unit (DU): These are the leave nodes of the FSM tree and correspond to concrete physical objects (e.g. a power supply channel) represented as data points in WinCC OA. The state logic runs as a Control script in a Control Manager in WinCC OA, which is also used as the proxy to interface the hardware with the associated objects in the SMI++ runtime.
2. Logical Unit (LU): These are abstract objects (e.g. a system that groups several power supply channels) that cannot be controlled independently. LUs cannot have CUs as children and run within an SMI domain process.
3. Control Unit (CU): These are also abstract objects, with the main difference to LUs being that different CUs can be controlled by different operators. One CU corresponds to one SMI domain process (it is the top-level abstract object in an SMI domain).

Control Units can be distributed over several machines and used in various partitioning modes to enable monitoring or control of only a subsystem independently of the rest. The possible partitioning modes are shown in Fig. 8.9. Logical Units and Device Units only implement the modes “enabled” and “disabled”.

¹³<https://lhcb-online.web.cern.ch/ecs/fw/FwFsm.html>

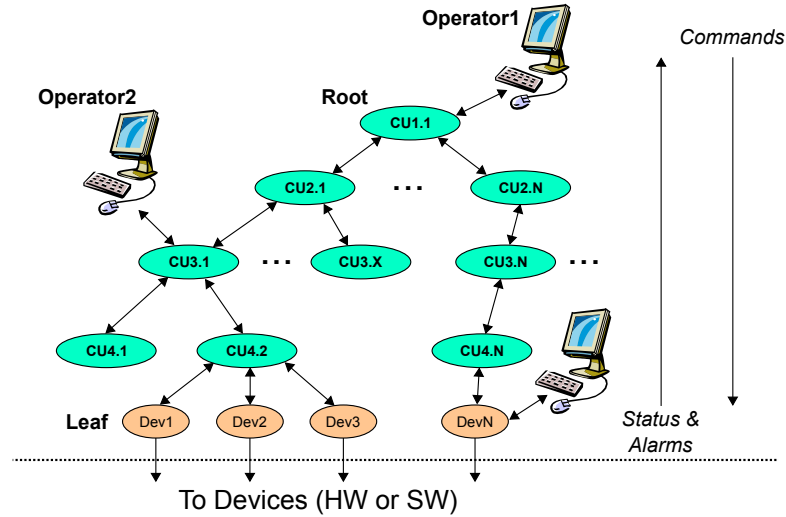


Figure 8.8: Generic architecture of an FSM tree. (Adapted from Ref. [231])

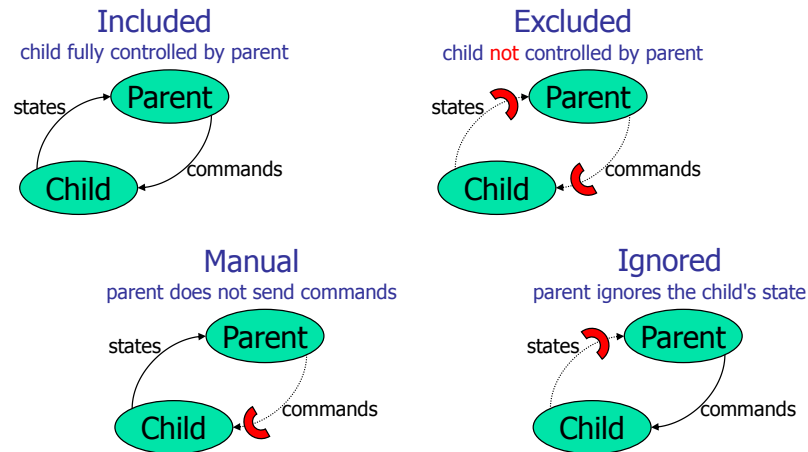


Figure 8.9: Partitioning Modes of Control Units. For each of the four partitioning modes it is shown how states and commands can propagate between a parent and a child Control Unit.

8.4.4 ATLAS-specific extensions

ATLAS decided to extend the functionality of FwFsm with the fwFsmAtlas component. The partitioning modes obtain an additional attribute: ownership. Only the owner of an FSM (sub-)tree can decide whether ownership is exclusive or shared, allowing control by others.

The ATLAS FSM component also introduces a second attribute for FSM objects. Each object thus has two attributes: the *state* and the *status*. The *state* represents the operational

OK	WARNING	ERROR	FATAL
-----------	----------------	--------------	--------------

Figure 8.10: Possible status attributes for FSM objects and their color coding as implemented by the fwFsmAtlas component.

state of the object, while the *status* represents the safety status using alerts. The *status* attribute can have one of the following values or severity levels: “OK”, “WARNING”, “ERROR” or “FATAL”. Each value is associated with a specific color shown in Fig. 8.10. Both the state and the status are propagated upwards independently, meaning alerts do not overshadow the operational state.

For operators in the ATLAS control room, a dedicated user interface is available. Each FSM object has a panel associated with it. Fig. 8.11 shows an example of the FSM screen and its modules. On the left-hand side are the navigation module and FSM module which enable to freely navigate through the FSM tree and display the current object with its state and status, as well as all children with their respective states and statuses. Here, operators can send commands and perform partitioning actions. Below is the secondary module, which displays a second FSM object independently of the main module in the center of the panel, which supplies the operator with detailed information about the selected object. At the top right is the problem list: a table that collects all alerts (bad statuses) of the FSM and their timestamps. Clicking on an alert navigates the main module to the origin of the alert to enable investigation of the problem.

For conventions to be used when creating a sub-system, an FSM integration guideline [233] is available. It defines the state diagram in Fig. 8.12 for the highest sub-detector node, ensuring the sub-detector can be embedded in the general ATLAS FSM tree as a reference below the Global Control Station (GCS). In the control room, shifters operate the detector only through the FSM and an alarm screen, a table that collects all active alerts in the system. Thus, all necessary commands for normal operation must be implemented through the FSM.

The fwFsmAtlas component, which includes ATLAS-specific extensions for the Controls Hierarchy framework component for WinCC OA, is also used by the FASER DCS as described in Section 10.2.




Figure 8.11: ATLAS DCS user interface during a regular LHC fill for luminosity production with pp -collisions. It shows an overview of the state of all detector components. (Reproduced from Ref. [232], © ATLAS)

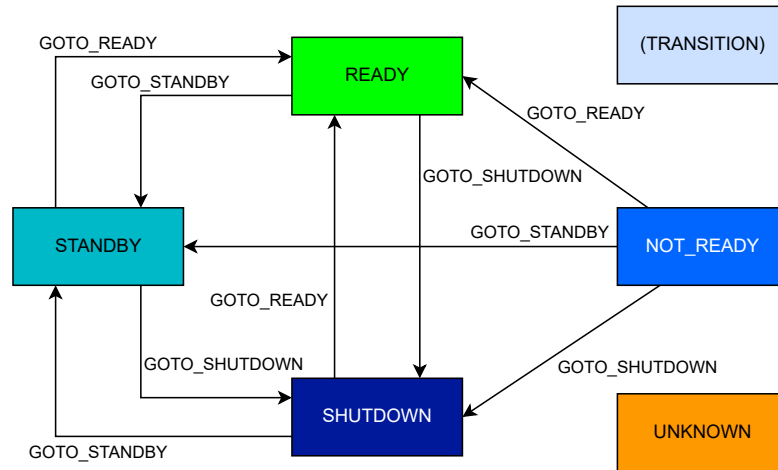


Figure 8.12: State diagram for an ATLAS sub-detector. This state diagram is required for all sub-detectors to be able to integrate into the Global Control Station (GCS) (see Fig. 8.1).

The Detector Control System for the ITk Pixel Detector

A large part of the thesis work focused on the development of a control system for OB LLS QC setups (as presented in [Chapter 7](#)), with the intention of eventually scaling this solution to a larger control system. The design is oriented toward what a final system could look like. From a detector design perspective, the architecture of the ITk Pixel Detector Control System (DCS) is already defined.

This chapter explains the overall architecture of the ITk Pixel detector DCS as implemented in the detector design in [Section 9.1](#) and the general requirements for the ITk Pixel DCS in [Section 9.2](#). [Section 9.3](#) then presents developments for a DCS specific to OB LLS QC setups, which build upon previous developments presented in Ref. [\[179\]](#). These developments leverage the SCADA software (see [Section 8.3.1](#)) and include the general system design, monitoring tools, implementation of automatic testing procedures, configuration tools for the DCS, and the creation of remote interfaces. Throughout these developments, the final design for a full detector DCS was kept in mind. Therefore, the design of the FSM follows a scheme suitable for a full detector. The design of the FSM is detailed and discussed in [Section 9.4](#) and performance tests provide an idea of the current system when scaled to a full detector. [Section 9.5](#) outlines potential next steps in developing the final detector DCS.

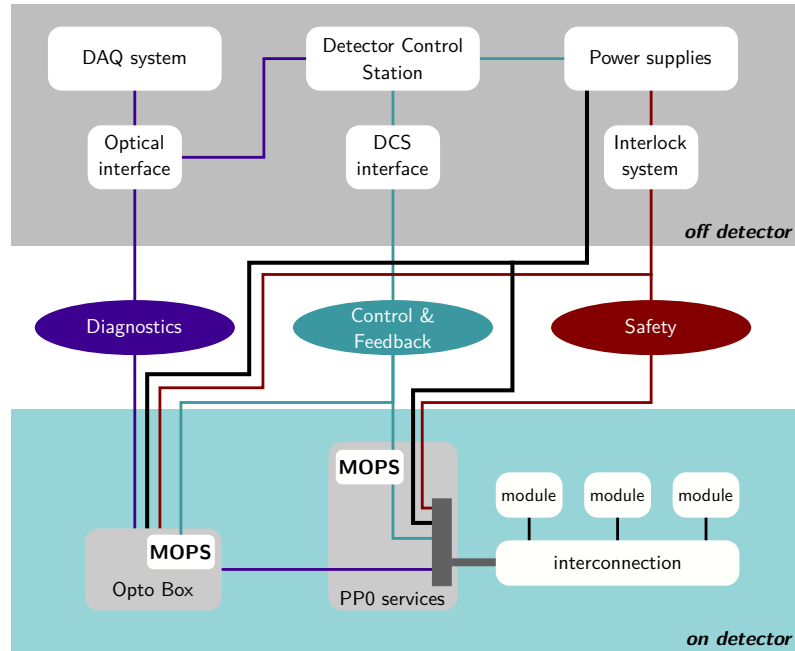


Figure 9.1: Overview of the paths for the ITk Pixel DCS system. Modules are connected into an SP-chain. One module per chain has its NTC connected to the interlock system. NTCs within the Opto Box are also linked to the Safety path. (Adapted from Ref. [234])

9.1 The architecture of the ITk Pixel DCS

The concept of the ITk Pixel DCS is driven by the needs of the serial powering. Four to fourteen modules are connected into an SP-chain and powered through a single supply line (LV) by a current source, forming the smallest logical power unit. In total there will be 1012 SP-chains, with 448 located in the OB region alone. Additionally, the Opto Boxes, which house the Opto Boards with optical transceivers and ASICs that handle the data transfer, require monitoring and control.

To ensure safe operation and reliable data acquisition, the Pixel DCS comprises three independent paths: Diagnostics, Control and Feedback, and Safety. This concept is shown in Fig. 9.1. The three paths are described in more detail below.

9.1.1 Control and feedback path

The Control and Feedback path manages the power supplies for the SP-chains and the Opto Boxes. Under serial powering, individual modules cannot be turned on or off, but only an entire SP-chain can be controlled as a unit. Detailed connection schemes for the SP-chain are provided in the Appendix in Fig. A.1. Each SP-chain is controlled via one LV channel

from a current source to power the FE chips and two to four HV channels for depletion of the silicon sensors. FE chip power consumption varies based on the chip activity and therefore on the position of the chip in the detector, with higher activity expected in inner layers. The current required by a single SP-chain is expected to vary between 5.55 A and 6.6 A, with a total voltage drop across an SP-chain (including the LV cable) from 10 V to 29 V [235, 236]. The bias voltage is distributed in parallel to the sensors in an SP-chain. Only in layer 0, a single HV line per module will be used, while in the outer layers up to 7 modules will be biased in parallel by one HV channel resulting in 2 to 4 HV channels per SP-chain. For the 3D sensors in the inner layers, a depletion voltage of up to 250 V at the end of HL-LHC is anticipated, while the expectation for planar sensors is up to 600 V at the end of life. A total of 2228 HV channels will be needed for detector operation [235].

Approximately 300 Opto Boxes will be installed just outside of the ITk detector volume for detector readout. Within each Opto Box, all Opto Boards associated to one SP-chain are powered by one supply line (OPTO). The Opto Boxes have up to five OPTO supply lines and require an input voltage between 8 V and 12 V for the first stage of the two-stage DC-DC conversion performed internally.

For the monitoring of temperatures and voltages of modules in an SP-chain as well as temperatures and voltages inside an Opto Box, the DCS connects to the Monitoring Of Pixel System (MOPS). The MOPS is a DCS ASIC developed specifically for the Pixel detector [169, 173]. It is used for aggregation of on-detector monitoring data. One MOPS is assigned to each SP-chain (see Section 4.5.2). For external communication, the MOPS uses a CAN bus. The data is received at PP3 by the MOPSHUB [237], aggregated and transmitted to the counting room via an Embedded Monitoring and Control Interface (EMCI) and an Embedded Monitoring and control Processor (EMP) as shown in Fig. 9.2 [238]. The power for the MOPSHUB (VPP3) is independent of the power for the MOPS (VCAN). In total, 86 VPP3 channels will be required for the operation of the ITk Pixel detector [235]. All power supply crates will be installed in standard 19 inch racks in the services caverns, connected to the detector in the detector cavern via long cables.

The Control and Feedback path serves as the slow control path and is required during all types of operation. This path must be functional for the FSM to determine the state of the Pixel detector.

9.1.2 Safety path

The safety path comprises the hardwired interlock system to protect the detector against risks. It demands the highest reliability and runs permanently. To prevent overheating of the detector, temperature sensors on the SP-chain and in the Opto Boxes monitor

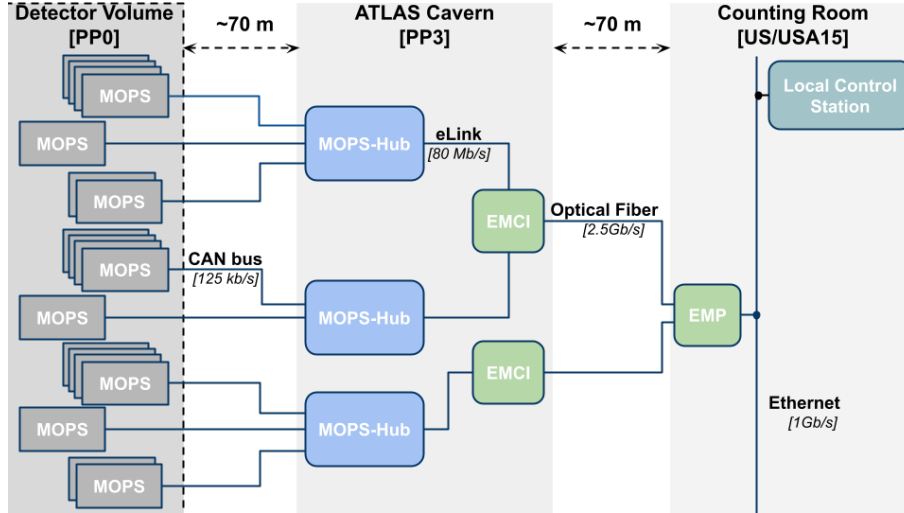


Figure 9.2: The MOPSHUB topology foreseen at PP3. Up to 8 MOPSHUBs connect to one EMCI, which contains an lpGBT and a VTRx⁺ and sends the aggregated data to an EMP in the counting room. (Reproduced from Ref. [239], licensed under CC BY 4.0)

temperature conditions, and are connected directly to the LISSY crates of the interlock system to enable switch-off actions independent of the Control and Feedback path [240]. Up to 16 LISSY crates can interface with a MIC to receive signals from systems such as the DSS, the cooling infrastructure or the beam interface system [213, 215]. During incidents, a LISSY crate acts with high segmentation on the power supplies to minimize the number of detector units which are being disabled.

9.1.3 Diagnostics path

Each FE chip collects internal monitoring data via an on-chip ADC. The data contains, for example, internal voltages as well as temperature information from additional NTCs on the module PCB beyond the one used by the MOPS chip. This data is transmitted via the standard (physics) data path to FELIX, where it is separated and forwarded to the DCS. No detector-safety critical data will be transferred via this path.

9.2 Requirements for the ITk Pixel DCS

High-level functional and performance requirements for the ITk Pixel DCS are documented in Ref. [234]. The ITk Pixel DCS must interface with all active ITk DCS detector hardware components, which is the on-detector MOPS and the MOPSHUB on PP3 to monitor FE electronics like FE chips and Opto Boxes. Additionally, the DCS must interface with power

supplies using an OPC-UA server as middleware. When selecting DCS or DCS-related hardware or software components, standardized and commonly used solutions should be preferred in order to allow for centralized support (to be provided by the respective central teams) and to minimize the maintenance effort needed to run the control system over several years.

For SP-chain control, an important task of the ITk Pixel DCS, the following control parameters are required [234]:

- LV on/off,
- LV current setting,
- LV voltage limit setting,
- HV on/off,
- HV voltage setting,
- HV current limit setting.

Monitoring parameters for an SP-chain include:

- LV voltage monitoring,
- LV current monitoring,
- LV status (on, off, interlock),
- LV monitoring of voltage limit setting,
- HV voltage monitoring,
- HV current monitoring,
- HV status (on, off, interlock),
- HV monitoring of current limit setting.

Additionally, the DCS must also monitor the voltage drops across modules in the SP-chain and the module temperatures.

Regarding monitoring of the infrastructure, the DCS must be able to monitor the interlock system and other auxiliary parts, for example cooling systems or MOPSHUB crates. Concerning availability, the design of the DCS must be such that it is able to operate during periods when monitoring and control functions of detector or infrastructure elements are required, and not only during periods of data taking.

Finally, the DCS will interface with external control systems such as the DSS and perform actions related for example to beam safety. During ATLAS physics runs, communication with DAQ is necessary to power and configure the detector into a state for data-taking.

9.3 Developments for the DCS for ITk Pixel OB LLS QC setups

Not all requirements outlined in [Section 9.2](#) apply to the OB LLS QC DCS development. Although these setups test small, complete detector units with SP-chains, they are not full detectors as installed at the LHC and they are not embedded in a final infrastructure like in Point 1. Consequently, beam safety protocols or operational procedures for physics data taking are currently irrelevant. Nevertheless, while these developments do not target the final detector, the design principles of a full-scale DCS were consistently integrated to enable seamless evolution from test setups to the final system using all gained experience.

The users of the LLS QC test setups will be interested in all the details of the detector hardware and infrastructure needed to operate the setup and test the LLS, particularly if there are errors on the control side which need debugging outside normal operation. Thus, the LLS QC DCS may provide more accessible low-level control than a real detector DCS. During standard operation, however, the highest interest lies in knowing the operational state of the detector unit, monitoring performance parameters during tests and exporting information from the DCS for offline analysis.

Accordingly, the LLS QC DCS was created with panels that were built with the detector properties in mind and with state calculation of the detector components via an FSM, extending the work established in Ref. [179]. To address QC-specific needs, the system was significantly enhanced with monitoring capabilities for new setup components, automated testing procedures and remote programmatic interfaces.

Greater emphasis was placed on lower-hierarchy panels, since users test individual LLSs. The FSM implementation ensures a consistent graphical interface while simplifying operation and preventing operator errors during LLS production for the final detector.

The following subsections detail these developments for the LLS QC DCS.

9.3.1 System overview and project setup

As already described in [Chapter 7](#), there are five clusters that will set up OB LLS QC test stands. They will use the same DCS hardware as presented in [Fig. 7.2](#). For connecting the DCS to the hardware, three OPC-UA servers (see [Section 8.3.3](#)) run on a dedicated DCS machine with the AlmaLinux 9 operating system: one server is used for the Wiener current source in PL512 format to provide LV for the SP-chain; one for the Wiener MPOD Mini Crate that houses a LV module for OPTO, VPP3 and VCAN and a HV module for the SP-chain HV; and one server for the connection to the MHFB. All OPC-UA servers

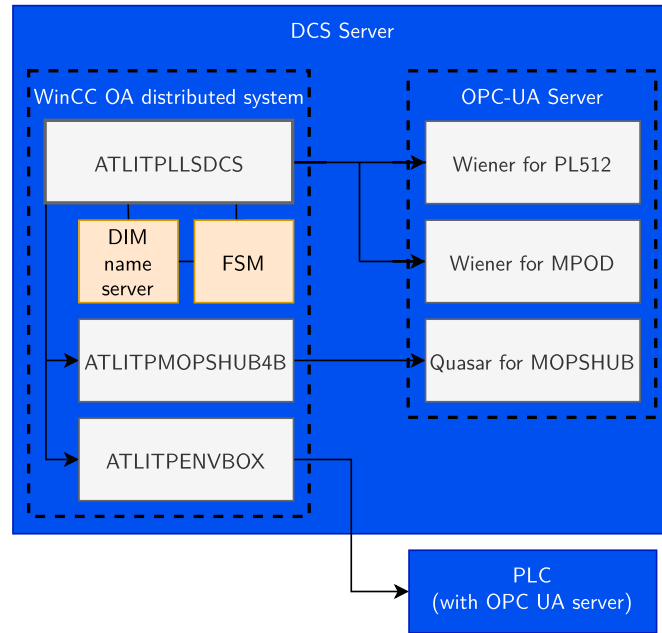


Figure 9.3: Schematic overview of the ITk OB LLS QC DCS software setup.

run as `systemd`¹ services to enable automatic startup. This is useful after a power cut, for example, when the computer is configured to automatically start as soon as power is restored. Consequently, all OPC-UA servers resume operation after such scenarios without manual intervention. A fourth OPC-UA server runs directly on the Siemens S7-1500 PLC.

On the same DCS machine, WinCC OA (see [Section 8.3.1](#)) in version 3.19 is installed and three WinCC OA projects, connected in a distributed system with unique system IDs, interface with the OPC-UA servers to enable control and monitoring of the setup. [Figure 9.3](#) shows a schematic overview of the DCS software components on the DCS machine. Also the WinCC OA projects and a DIM name server, required for the communication between SMI domains of an FSM, run as `systemd` services to enable automatic startup. An FSM is set up in the ATLITPLLSDCS WinCC OA project, which serves as the main project and entry point for the operator GUIs. During development of the WinCC OA project and the FSM, the general ATLAS guidelines for DCS [241] and FSM [233] were followed. Additionally, requirements from [Section 9.2](#) were followed where appropriate for the setups with QC infrastructure hardware.

The main WinCC OA project contains created panels, libraries and scripts, and is used to configure the system for a specific LLS to be tested. The system was designed modularly to add additional hardware connections through integration of subprojects. Two subprojects

¹System and Service Manager for Linux operating systems



Figure 9.4: The Device Editor and Navigator showing the three different trees of the project. The hardware tree in (a) displays DP names, the logical tree in (b) displays alias names. Single operation panels can be opened from the hardware and logical view. The FSM follows closely the names of the logical tree. It can be started from the FSM view as shown in (c) and the corresponding GUI can be opened from there.

are added: one to establish connection to the MOPSHUB4Beginners and one to connect to the environment and interlock PLC. Both subprojects retain the information of their hardware via DPs and include detailed panels for the MOPSHUB4Beginners or PLC, respectively. The main project has access to both the DPs and the panels provided by the subprojects.

The main project heavily uses JCOP framework components to integrate hardware information into WinCC OA using predefined DPTs and correct peripheral address settings in the DPEs. Components exist for the Wiener PL512 power supplies and the Wiener MPODs, each adding their own drivers to the WinCC OA project.

The primary development tool for hardware integration is the Device Editor and Navigator (DEN). This tool is used to integrate hardware devices like power supplies and create a “hardware” tree of the setup displaying the available DPs as shown in Fig. 9.4(a). This procedure typically needs to be done only once for the LLS QC setups, as their infrastructure hardware does not change.

To configure the system for a specific LLS, aliases are written to DPs (or DPEs) describing the connected hardware. For example, an alias for a power supply channel describes which SP-chain it is connected to, or an alias for a MOPS ADC channel describes the module

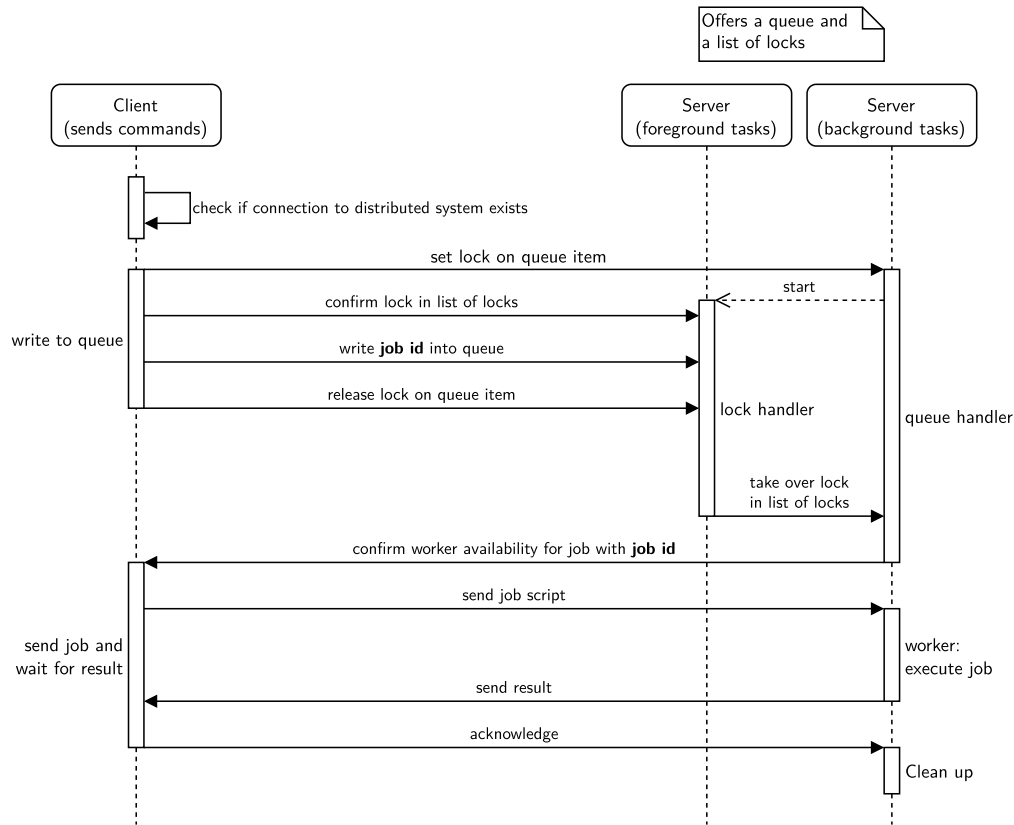


Figure 9.5: Communication scheme to execute scripts in a remote WinCC OA project.

it is monitoring. This configuration is executed from the main project, but aliases are written to DPs in other projects as well. A restriction exists from the JCOP components and WinCC OA: some of the commands for alias creation can only run in the local project. To overcome this, functionality was implemented to execute scripts in remote projects by a CTRL manager and returning the result of the script to the calling project.

Therefore, every subproject runs a CTRL manager hosting a “Distributed Commands Server”. This server creates DPs that are used to handle communication, e.g. transfer of the script code and the script result. The main project acts as a client, sending scripts to the server in the subproject. [Figure 9.5](#) illustrates the client-server interactions for executing scripts in a remote project. If multiple clients request remote execution, a First In, First Out (FIFO) queue is implemented on the server. The queue stores job IDs (short hash values derived from the job submission time and script content), not full scripts. When the server writes a job ID into a DPE to confirm the next job, the client sends the script and waits for the job result.

The Distributed Commands Server allows the system setup to be performed from one central point, the main project. [Figure 9.6](#) shows the system setup panel in the main

project. The connectivity file is a CSV file that contains the mapping of aliases to DPEs (across all projects). It was decided to impose conventions on alias names and enforce that they are unique across all projects. The alias-to-DPE mapping represents the physical cabling, for instance, the information about which detector component is connected to which power supply channel. If a different type of LLS is tested and cabling changes, the system can be reconfigured. Aliases in a WinCC OA project form the “logical” tree in the DEN, as shown in [Fig. 9.4\(b\)](#).

The creation of the logical tree and common actions concerning aliases are supported by an alias handling library, which consolidates frequently used alias functions. The hierarchical design, both in DP names and aliases, is realized by using the slash (“/”) as hierarchy separator per JCOP convention. For example, the first HV channel of an SP-chain (from a Wiener MPOD, represented as a DP in [Fig. 8.3](#)) can have an alias such as:

OB/L3/B01/A/SP1/HV/CH1 .

This indicates the SP-chain is the first on the A-side of the first longeron in layer 3 in the OB region. The naming scheme loosely follows the official names in Ref. [\[242\]](#), but adapts to the limitations and requirements inside a WinCC OA project.

When writing an alias to a DP, additional alias elements are also written to its DPEs. For instance, the “MeasurementCurrent” DPE of the first HV channel from the previous example receives the alias

!OB/L3/B01/A/SP1/HV/CH1.Imon .

The preceding “!” symbol prevents display of this DPE in the logical tree, allowing one to focus only on the full aliases. Assigning these additional aliases is useful when displaying the value of these DPEs and when archiving the values together with the alias. The logical tree in the DEN shows then a “detector view” that can be already used for simple (channel-wise) control of the detector.

In a second step of the system setup as shown in [Fig. 9.6](#), the interlock matrix of an interlock crate is configured to display the correct association between input signals and output signals. For more information see Ref. [\[179\]](#) and [Section 9.3.5](#).

In the final step of the system setup, an FSM control layer is created using the available aliases in the system. The hierarchical structure of the FSM is displayed in the FSM tree of the DEN as shown in [Fig. 9.4\(c\)](#). From there the GUI of the FSM can be opened. Setting appropriate alerts to DPEs is done in the start up scripts of the FSM. This allows to quickly adjust thresholds by restarting the FSM.

SR1SystemSetup: System Setup <@itkpix-llsqc1-dcs-01>

log in as root! root

Setup the aliases and FSM of the Project Configured

1. Delete configuration of current project (all aliases and FSM)

Adjust the paths (if necessary) to the configuration files for project setup:

Connectivity file (mapping alias to dpe):
 Connectivity config file for creating alias

Interlock Matrix Crate (IMC) configurations (one config file per IMC):

Read mode of IMC files:
☐ only use given DPE (and already existing alias)
☐ use given DPE and overwrite existing alias with supplied
☒ only use supplied alias (default)

Archiving config file (optional - can be changed in running project):
 NGA) InfluxDB - EVENT

Metadata config file (optional - can be changed in running project):

☐ Skip creation of FSM

2. Configure current project (read config files and create FSM)

Figure 9.6: The system setup panel to read the connectivity information from a CSV file and to build the FSM by using the available aliases in the system. Additionally, interlock matrices can be configured for example for the monitoring of the PLC interlock logic, where all necessary information is read from an Extensible Markup Language (XML) file.

Optionally, the system setup panel allows the configuration of how the data of DPEs is archived into a persistent database, i.e. the configuration of smoothing parameters, by providing a configuration file (see also [Section 9.3.3](#)). It also enables saving serial numbers and other metadata of the device under test in the DCS to better associate DCS test results with the tested hardware (see also [Section 9.3.10](#)).

For all configuration steps, simple text files can be used. This allows distributing configuration files for different LLS types to testing sites, enabling DCS setup without DCS software expertise.

9.3.2 Identifying information using aliases

The creation of the hierarchical FSM as described above is realized by checking all existing aliases in the current system and matching them to general patterns describing components in the setup (see also [Section 9.4.1](#)). Therefore, an alias written to a DPE describes where the DPE's data is used in GUIs and in the FSM. This mapping of alias to DPE was unique and thus came with one limitation: when a DPE holding detector data receives an alias, the data now belongs to a specific branch in the alias or FSM tree, but cannot be easily associated with another branch. It was decided to introduce functionality to allow for multiple aliases to be mapped to one DPE to overcome this limitation.

The functionality is implemented by creating auxiliary DPs for aliases that map to the same DPE. These auxiliary DPs receive the pure alias and contain the DPE name (to which the alias should be mapped) as their value. Mapping from aliases to DPs remains unique. Conversely, a DPE has multiple aliases assigned as a list of single aliases. To retrieve a specific alias, the user must provide the starting name of the branch for which the alias should be returned. The scheme for assigning multiple aliases to a single DPE is shown in [Fig. 9.7](#). The advantage of storing all information in DPs is that this information is easily accessible, enabling the connectivity to be retraced at any time.

An example of this functionality is the DP that describes the powering channel for the Opto Boards belonging to the OB/L3/B01/A/SP1 SP-chain. This DP is assigned two aliases:

$$\begin{aligned} & \text{OB/L3/B01/A/SP1/OPTO/CH} \quad , \\ & \text{and INFRA/OPTOSYS/BOX01-L/PWRBOARD/bPol12V-E/CH} \quad . \end{aligned}$$

The first alias describes which SP-chain is affected by the powering channel, while the second alias specifies the Opto Box and the bPOL12V E. With this solution, the information described by this DP can be used in the branch for the OB detector and in the branch

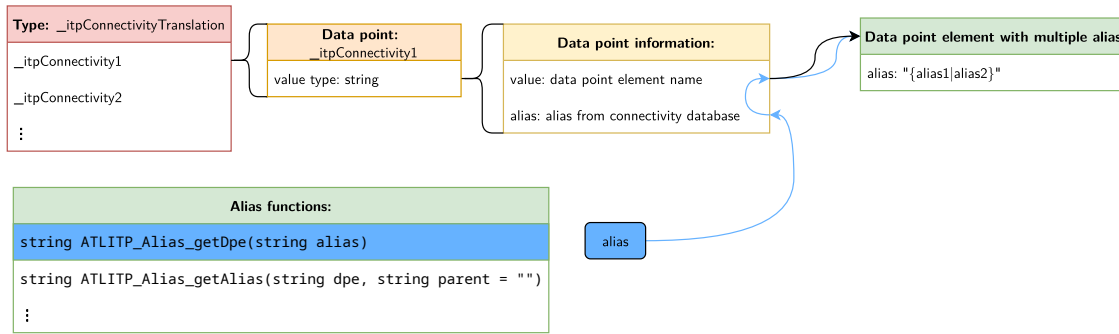


Figure 9.7: Schema to support multiple aliases per DPE in WinCC OA. If multiple aliases are assigned to one DPE, auxiliary DPs (of type `__itpConnectivityTranslation`) are created to allow for the translation of a single alias to its corresponding DPE. The final DPE stores all its associated aliases in a formatted list.

describing detector infrastructure (INFRA). The standard WinCC OA function for alias-to-DPE conversion has been replaced in all self-written code by a custom wrapper function.

The performance impact of this design is that, instead of looking up the DPE for an alias (a fast process executed within a CTRL manager), the system must additionally read the value of a DP, which requires sending a message to the event manager of WinCC OA. However, this does not significantly affect the FSM: The lookup is performed once at the start of the FSM, while all callbacks for state calculations are connected to DPEs and not to aliases.

9.3.3 Configuration of archive settings

The storage of DCS parameters (like temperatures or voltages) and associated timestamps into persistent database is important for reconstructing the state of the setup at a later time (see also [Section 7.3.1](#)). To save space, archived data can be smoothed. Smoothing settings must be configured for every DPE, but several groups of DPEs (such as all module temperatures) will likely require identical configuration. To simplify mass configuration of these archiving settings, configurations can be associated to alias patterns. A Tab-Separated Values (TSV) configuration file for archiving settings as shown in [Listing 9.1](#) can be used during system setup (see [Fig. 9.6](#)) or at a later time to perform mass configuration. The alias pattern in the first column of the configuration file uses wildcards like `*` for matching any sequence of characters and `?` for matching any single character to apply the same setting to a group of DPEs.

For manual configuration or verification that DPEs have been configured for archiving, a GUI as shown in [Fig. 9.8](#) is provided for the operator in the lab.

```

!OB/*/HV/CH?.Imon    true    TIME_AND_VALUE_SMOOTH    0.00000004    3600
!OB/*/HV/CH?.Vmon    true    TIME_AND_VALUE_SMOOTH    0.5    3600
OB/*/M??/Tmon        true    TIME_AND_VALUE_SMOOTH    0.5    3600
OB/*/M??/Vmon        true    TIME_AND_VALUE_SMOOTH    0.01    3600
...

```

Listing 9.1: Example TSV configuration file for mass configuration of archive settings of DPEs. The first column contains the pattern to match DPEs that should be configured, the second column enables smoothing, the third column specifies the smoothing procedure, the fourth column contains the value deadband and the fifth column the time deadband for smoothing.

ArchiveConfig: Archiving Configuration <@itkpix-llsqc1-dcs-01>

Configure Archiving

By data point elements:

The search supports wildcards, i.e. "?" to represent any single character, and "*" to represent any string of characters. However, when searching for data point elements, "*" does not cross hierarchy levels. Consider using "***" in that case. Don't forget to specify a system or use ":" at the beginning to search in all connected systems when searching by data point elements.

By alias:

By config file:

Archive group:

Interval of Validity (IOV)

Get the IOV of an alias stored in the archive:

alias

Start time	End time	Data point element

Currently assigned Dpes for archiving

Data point element	Alias	Config

Figure 9.8: A panel to configure archive settings of DPEs and confirm that data is archived. The configuration can be done by DPE name, alias or configuration file as given in Listing 9.1. The interval of validity shows from which start time to which end time an alias was associated to a DPE. The table at the bottom lists all DPEs that are configured for archiving.

9.3.4 Automatic procedures for HV and LV scans

Some electrical tests of the LLSs are performed by the DCS, specifically “HV scans” and “LV scans” (see [Section 6.2.1](#) and [Section 6.2.2](#)). During production and QC tests, particular interest is placed on $I(V)$ curves to check for sensor degradation.

User Interfaces (UIs) and automatic procedures were developed to perform these tests. The UI for HV scans is shown in [Fig. 9.9](#). A scan can be configured using start and stop values along with a step size. In each step, after setting a new value, the scan routine waits for a specified duration before performing multiple measurements and averaging the results. For HV scans, the voltage range is scanned with the HV channel always on. For LV scans, in each step the channel is cycled (turned off, a new value is set, and turned on again). All control functions use FSM commands to ensure consistent operation. Scan results are written to disk as CSV files and can be imported into analysis tools.

With the development of a remote interface, the REST API for WinCC OA, a scan functionality has also been implemented in the REST API manager and can be accessed via HTTP endpoints (see [Section 9.3.12](#)).

9.3.5 Implementation of additional interlock strategies

Although the interlock system is purely hardware based, it is required to communicate the interlock actions to the rest of the DCS. For the RD53A demonstrator, ELMBs were used to read out the incoming and outgoing interlock signals and report them to the DCS (see [Section 5.5](#)). For the OB LLS QC sites, a PLC contains the interlock logic and performs interlock actions. In the final detector, LISSY crates will be used. All monitoring signals are latched, meaning the information about a value exceeding a threshold is kept until manually reset, to allow the user to reconstruct the behavior of the system even in case of very short interlock signals. The panel for monitoring the interlock displays the interlock logic matrix, which connects input and output signals. This allows the operator to retrace the origin of interlock actions. In addition, the state of every output and every input signal is shown. Note, that it is not possible from the DCS to change the interlock logic.

To make the adoption of other interlock strategies easier, the code from Ref. [\[179\]](#) has been made more generic. With this, support for PLC, LISSY and a custom interlock solution for EC LLSs testing [\[243\]](#) has been implemented.

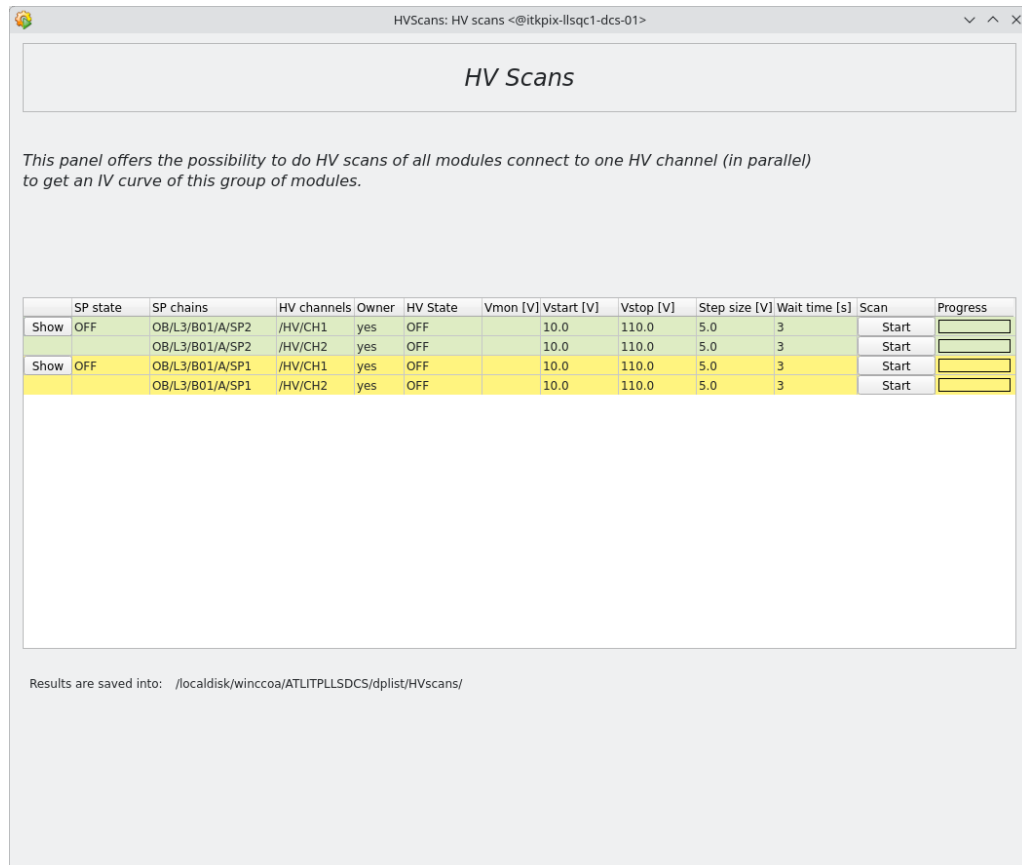


Figure 9.9: A panel to perform automatic HV scans. A scan can be started for a specific HV channel. Several scans can run in parallel and the results are written to disk as CSV files. Clicking on “Show” navigates to the corresponding SP-chain in the FSM. The “Owner” column indicates if the current user has control over the SP-chain in the FSM.

9.3.6 Opto Box monitoring

An Opto Box comes in two flavors: a “normal” Opto Box and a “mirrored” Opto Box. The two flavors differ in which direction electrical data cables enter the box and fibers leave the box. Five bPOL12Vs are installed on the power board in the box to provide power to up to 8 Opto Boards. A connector board defines which Opto Boards are associated to which bPOL12V. The monitoring of the Opto Box then consists of the information that is sent to the interlock system, and information read out by a MOPS in the Opto Box. The interlock monitors the temperature on the power board and on the connector board. The MOPS chip can monitor the temperatures on the Opto Boards but also voltage levels provided by the two different DC-DC converter types in the Opto Box (see also [Section 5.7](#)).

The monitoring panel for an Opto Box is shown in [Fig. 9.10](#). The interlock information is displayed with LED graphics that indicate the interlock status. The MOPS monitoring

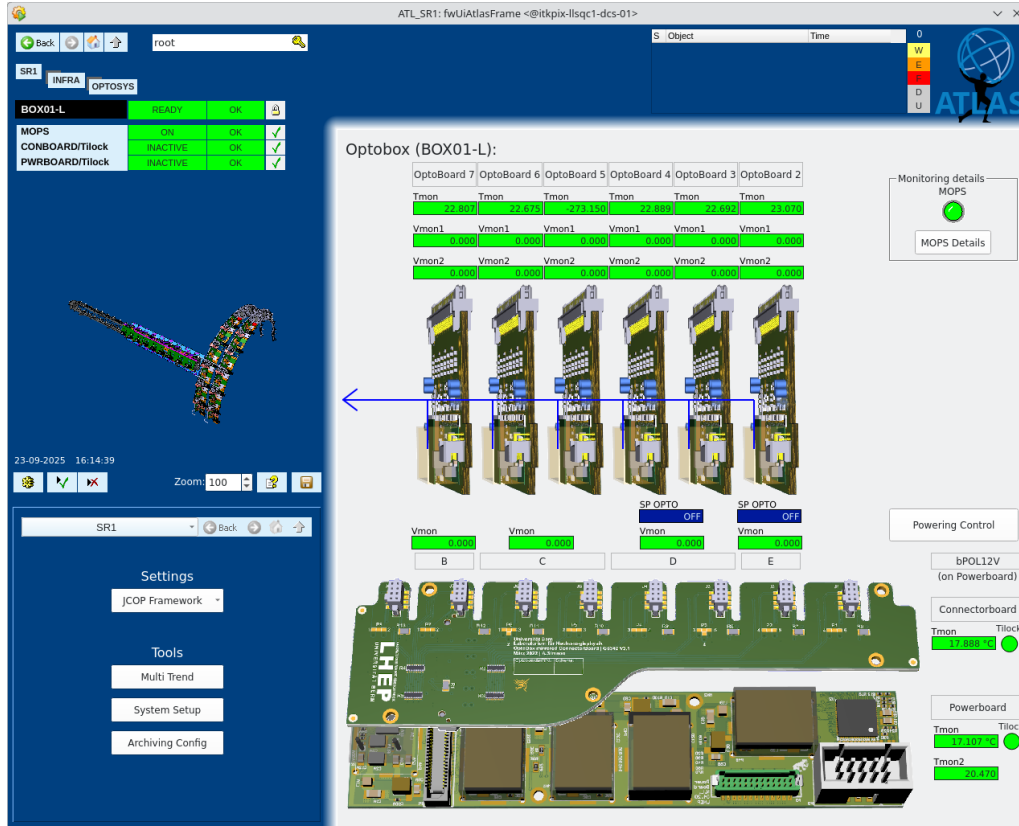


Figure 9.10: The panel to monitor an Opto Box embedded in the general FSM UI framework. Shown is the Opto Box as used in the CERN LLS QC setup. The fibers are leaving the Opto Box to the left as indicated by the blue arrow. Four power channels are connected to the bPOL12Vs B to E, where B and E are powering a single Opto Board and C and D power two Opto Boards. This mapping is defined by the connector board in the Opto Box. Monitoring via a MOPS chip of the bPOL12Vs on the power board and bPOL2V5s on the Opto Boards show their output voltages.

information is shown next to the monitored hardware. In addition, the state of the OPTO power of the associated SP-chains is shown.

The monitoring panel dynamically shows the correct Opto Box (or connector board) configuration as defined by the alias structure in the system. Every possible configuration can be displayed.

9.3.7 MARTA monitoring

The LLS QC setups utilize a MARTA cooling plant for cooling the modules of an LLS (see also [Section 4.5.3](#)). MARTA features a touch display for local control (turning it on/off,

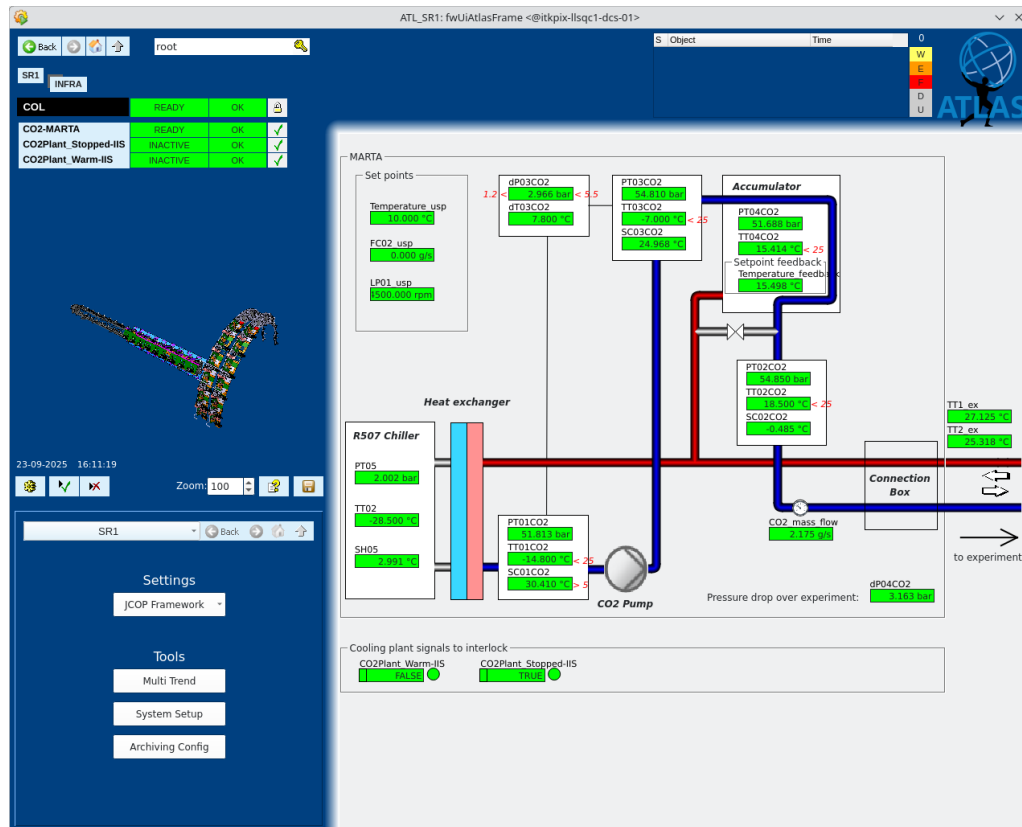


Figure 9.11: The panel to monitor the MARTA cooling plant embedded in the general FSM UI framework. Shown is monitoring information that can be received via a MODBUS connection. No control is possible via this connection.

changing the temperature set point). Additionally, it can be connected to the local network to provide monitoring data via a MODBUS connection.

To integrate MARTA monitoring into the DCS, for example to display the set point of the cooling plant, a framework was developed that takes the MODBUS register map and creates corresponding DPs which connect to these registers. New monitoring values are periodically pulled from MARTA at a refresh rate of 0.5 s.

Figure 9.11 shows the monitoring panel for the MARTA cooling plant. The design closely mirrors the touch screen interface, displaying both the temperature set point and feedback temperature. Additional operational parameters are shown, and alerts are raised, if these parameters exceed thresholds that could compromise normal cooling plant operation.

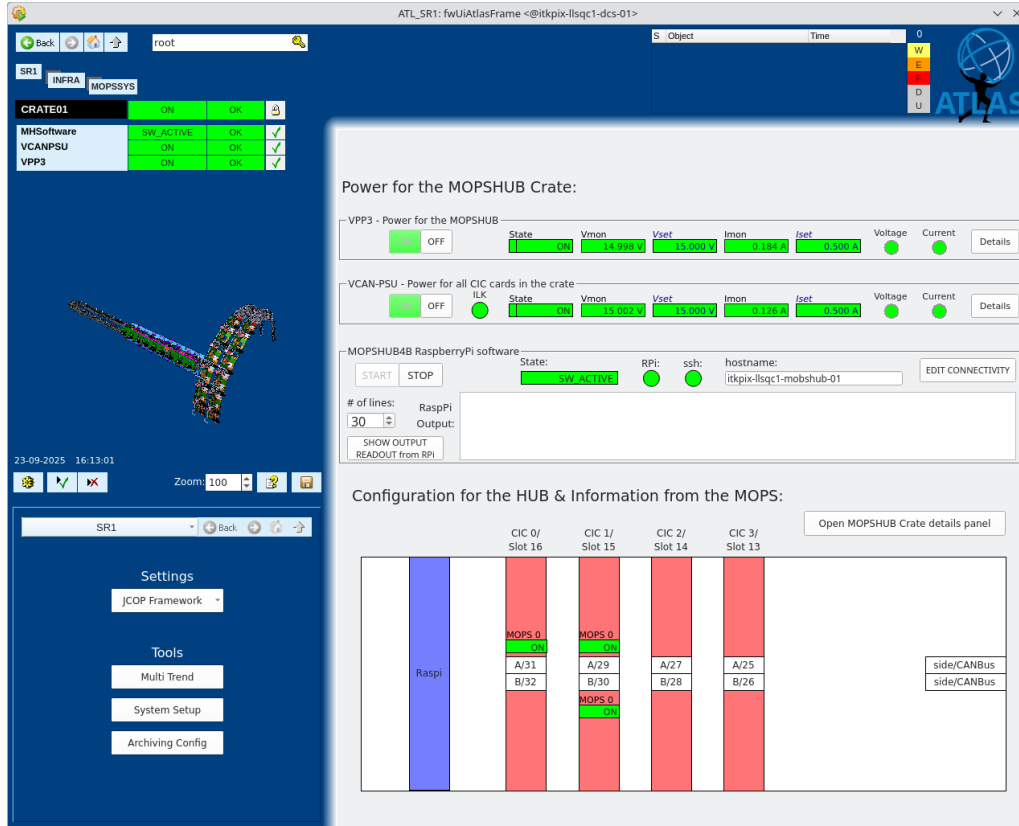


Figure 9.12: The panel to monitor the MHFB crate embedded in the general FSM UI framework. On top are shown the power supply channels for the Raspberry Pi and the CICs. In the middle is a control panel for the readout software running on the Raspberry Pi. The bottom shows the assignment of MOPS chips to CICs.

9.3.8 MHFB monitoring and control

Every OB LLS QC testing site is equipped with a MHFB to read out the MOPS chips in the system. The MHFB and the CAN Interface Cards (CANs) as shown in Fig. 5.6 need power. Power is provided by two LV MPOD channels. The power channel for the Raspberry Pi is labeled INFRA/MOPSSYS/CRATE01/VPP3, while the power channel for the CICs is labeled INFRA/MOPSSYS/CRATE01/VCANPSU. This naming scheme follows closely the official one for an FPGA in the final MOPSHUB.

Figure 9.12 shows the monitoring panel for a MHFB crate. The power supply channels are shown on the top and the allocation of the CICs is shown on the bottom. For remotely monitoring the state of the readout software on the Raspberry Pi and the possibility to start and stop it, a remote control functionality was implemented, that is displayed in the middle of the panel [244].

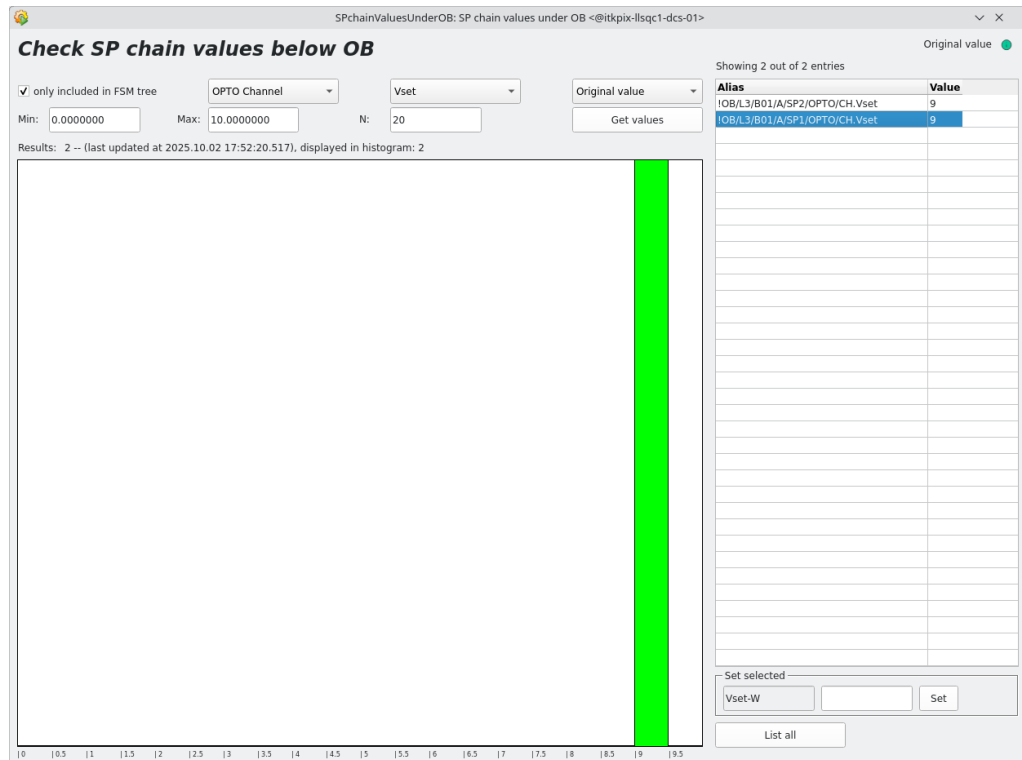


Figure 9.13: Histogram of SP-chain values. Shown are the voltage settings for power supply channels powering Opto Boards for readout of the preproduction longeron at CERN. The range and the number of bins of the histogram can be set at the top. The histogram in this example shows two SP-chains configured both with the same voltage setting for OPTO.

9.3.9 Mass configuration and consistency checks of SP-chains

A key requirement for the DCS is to simplify the identification of abnormalities in the setup and to simplify the verification of consistency in the setup settings. To address this, a tool was developed that displays the distribution of critical parameters across multiple SP-chains as histograms to enable rapid detection of any outliers. This tool is shown in [Fig. 9.13](#). Operators can select an SP-chain parameter along with the histogram range and binning. Values outside the selected range are indicated in the result text. Selecting a histogram bin lists all included values and aliases in the right-hand table, where operators can correct any undesired settings.

This panel is useful for verifying the consistency of SP-chain settings before powering an LLS or identifying abnormal monitoring values during operation.

Alias	Type	Serial-Nr.	Serial-Nr. of Cell	Additional Info	Note(s)
OB/L3/B01/A/SP1/M01	Module	20UPGM22601160			
OB/L3/B01/A/SP1/M02	Module	20UPGM22601182			
OB/L3/B01/A/SP1/M03	Module	20UPGM22601181			
OB/L3/B01/A/SP1/M04	Module	20UPGM22601146			
OB/L3/B01/A/SP1/M05	Module	20UPGM22601158			
OB/L3/B01/A/SP1/M06	Module	20UPGM22601171			

Figure 9.14: The metadata panel of an SP-chain. Shown are the metadata for an SP-chain (e.g. the serial number of a PP0), the metadata for the two HV groups of an SP-chain, and the metadata for all modules in an SP-chain. Each module can have a module serial number and an additional serial number as a loaded cell.

9.3.10 Metadata for detector hardware

During production, multiple LLSs will be tested in the setups. Electrical tests will be performed by the DCS, with the possibility to export results. To unambiguously associate test results with the tested detector hardware, a framework was developed to attach metadata like serial numbers to certain aliases (for example module aliases). Storing the serial number of a connected module in the test setup enables saving it alongside test results, enabling unambiguous association of results to tested hardware.

During system setup, the framework creates a new DP for each alias, using a name composed of a constant prefix and the alias. This DP stores metadata associated with the alias. Metadata of an SP-chain can be displayed in a panel that is shown in [Fig. 9.14](#). Possible metadata are serial numbers, “additional information” and “notes”. Modules can additionally have a “cell” serial number since they are equipped with a bare cell in OB (see [Section 4.5](#)). The “additional information” field is foreseen to be used to store formatted information like HV settings. The “notes” field can be used for manual notes, for example about non-working module temperature readings. Operators can export all system metadata to a CSV file or import metadata as needed.

The framework also enables lookup of aliases for given serial numbers. This functionality is used in the DCS adapter (see [Section 7.3.2](#)) when a module's serial number is provided to identify the associated SP-chain and control said SP-chain.

9.3.11 Deployment and usage

Since the LLS QC test setups are identical across the five clusters, a copy of the CERN DCS software package was distributed to all sites. Development has continued, with updates pushed to a Git repository on CERN's GitLab instance. A Continuous Integration (CI) pipeline was configured, such that when a tag is created, the pipeline job collects all relevant files and packages them into a release that can be installed (and updated) as a JCOP framework component following the specifications in Ref. [245]. Updates to the DCS software are distributed through this channel. Additionally, extensive documentation [246] was made available for all colleagues at the testing sites.

9.3.12 WinCC OA HTTP REST API manager

The main motivation for this development is the reduction of testing time during QC tests while in production. Given limited time allocated in the project schedule for testing LLSs after assembly but before integration into the final detector structure, automated testing without manual intervention is essential to maintain the schedule while ensuring quality control. Such testing identifies defects in built detector components and enables corrective actions. To enable automated testing, a software interface to the DCS of the LLS QC setup is required. This interface allows testing routines to control the detector unit under test, to bring it into any desired state (and then run tests from the DAQ) or to perform tests directly in the DCS.

Mirroring their use in the online software framework (see [Section 5.7](#)), the feasibility of implementing an HTTP REST API for the DCS was explored. WinCC OA lacks built-in functionality to easily access its parameterization, which is required for certain type of data for example power supply channel information (see [Fig. 8.3](#)). However, WinCC OA provides a C++ API for developing custom managers to extend its functionality. Therefore, a HTTP REST API manager was developed for WinCC OA using this C++ API.

The internal message queue between front-end and back-end

The developed HTTP REST API manager adheres to REST principles by exposing the resources of the distributed WinCC OA system, which can be manipulated using HTTP methods (“verbs”) [247]:

- GET: Retrieve the representation of the targets resource
- POST: Process the representation enclosed in the request
- PUT: Create or update the target resource’s state using the enclosed representation
- DELETE: Request deletion of the target resource’s state
- PATCH: Apply partial modifications to a resource

The HTTP REST API front-end of the manager uses the Oat++ framework² to expose HTTP endpoints representing the system’s resources. The back-end of the manager utilizes WinCC OA’s C++ API to communicate with the event manager of the main WinCC OA project.

The Oat++ framework creates a new thread for each connection to process the request and return a response. However, the functions offered by the WinCC OA C++ API are not thread-safe, and a WinCC OA manager must run a single main thread to manage communication with the event manager of WinCC OA. Therefore, a message queue was implemented to serialize requests from connection threads for processing by the main thread.

The message queue has multiple “producers” of messages (the connection threads), but only one “consumer”, the main thread. Access must be regulated to prioritize the main thread (high priority) over connection threads (low priority), ensuring the main thread can regularly read messages without blocking. This access pattern to the message queue is implemented using three mutexes:

1. queue mutex (Q)
2. next-access mutex (N)
3. low-priority mutex (L)

A connection thread must acquire first the L mutex, then the N mutex and finally the Q mutex. Only then can it release the N lock (and allow only the main thread to acquire it next) and append messages to the queue. Afterward, the connection thread releases Q and L. The main thread waits only on N. Once it acquires a lock on N, it attempts to lock Q, unlocks N and reads messages from the queue. Finally, the main thread releases Q. This access pattern is illustrated for two connection threads in Fig. 9.15. It guarantees that

²<https://github.com/oatpp/oatpp>, an open-source web framework written in C++

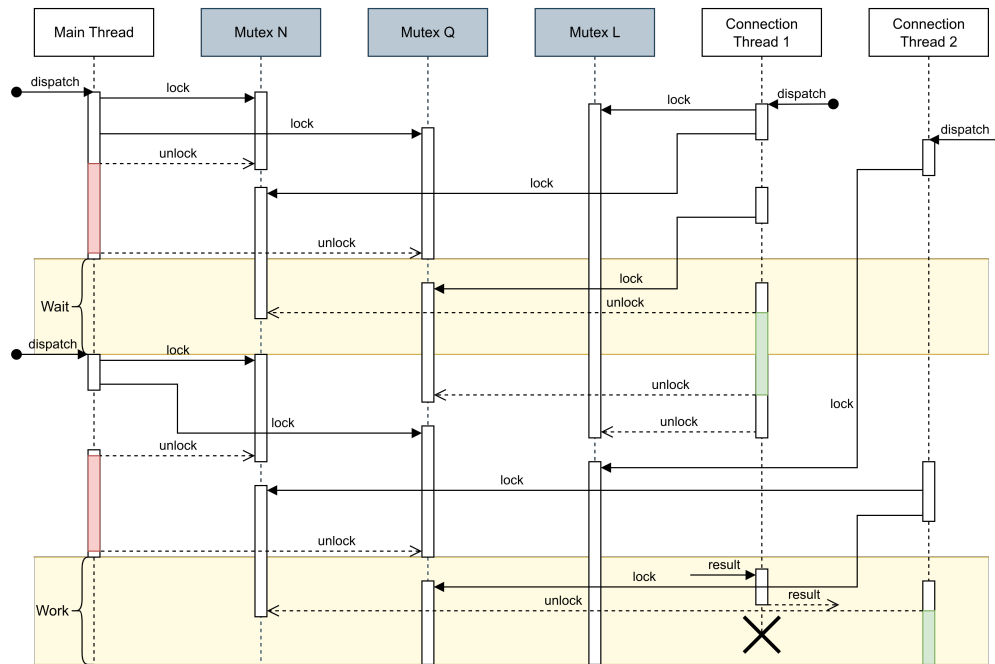


Figure 9.15: Access to the message queue from connection threads and the main thread in the WinCC OA HTTP REST API manager. The three mutexes governing the access to the queue are highlighted in grey. Phases with red indicate times in which a thread is reading from the message queue. Phases in green indicate times in which a thread is appending messages to the queue. In yellow is highlighted the time in which the main thread waits for 2 ms (on purpose, because there are no messages to process) or performs work (sending or receiving messages from the WinCC OA event manager). “Dispatch” in the main thread here indicates the start of a new loop in the main thread. This queue access pattern minimizes the time for the main thread to wait for access to the message queue.

when connections exist, at least one connection thread can add messages between accesses of the main thread. Also while the main thread *processes* messages, connection threads can add to the queue.

To simplify development, an interface class for WinCC OA functions, similar to the functions in the CTRL scripting language, was created. This interface class is used to send messages to the main thread as described above by adding the message queue as a shared static member variable. A drawback of this “global” variable is, that replacing the queue for testing would be difficult. An alternative approach would be implementing the queue via the singleton pattern [248].

Implemented functionality

The REST API manager has access to the parameterization of the WinCC OA project by exchanging information with the event manager of the project. It, however, does not have access to already written CTRL code that is used in many places in a WinCC OA project, for example to configure and run the FSM or to handle the alias functionality. An advantage of the interface class mentioned before and the similarity of CTRL code to C code is, that most of the CTRL code can be copied and used as C++ source code in the REST API manager with only minimal changes necessary. The HTTP REST API manager therefore does not only expose low-level WinCC OA functionality but can also provide high-level functionality for controlling the FSM or other aspects of the system, often based on the underlying alias structure. The Oat++ framework includes the functionality to add a Swagger UI³, which is shown in Fig. 9.16 for the REST API manager in the CERN LLS QC setup.

Implemented functionality includes reading all ADC channels of a MOPS chip and getting a set of environmental data over a period of time. For this, the manager also accesses the archive database of the system. Additionally, HV and LV scans are directly implemented in the manager, meaning they can be started and stopped and their results can be retrieved.

Both scan types are accessible as resources under their respective endpoints. Internally, a scan is modeled as an FSM with the state diagram given in Fig. 9.17. Several scans can run in parallel. At start time, the configuration parameters of the scan are sent, including start, stop and step size values. On a successful run of a scan, the scan results can be retrieved in JavaScript Object Notation (JSON) format, resetting the scan. The scan results contain voltage and current monitoring values from the power supply, MOPS monitoring data and serial numbers of the tested hardware, as introduced in Section 9.3.10. The scan functionality has been successfully tested and used in the CERN LLS QC setup with the preproduction longeron. A pipeline for automatic analysis of the data is being developed.

Measurement tools used during module QC tests can use the LLS DCS using this API manager (see Section 7.3.2).

Benchmarks of the API manager

To understand how the REST API manager performs under stress, benchmark tests have been conducted. For these tests, the development environment for the DCS system was used, which is hosted at CERN on a rack-mounted server from the year 2011 with the following specifications:

³<https://swagger.io/tools/swagger-ui/>, a REST API documentation tool

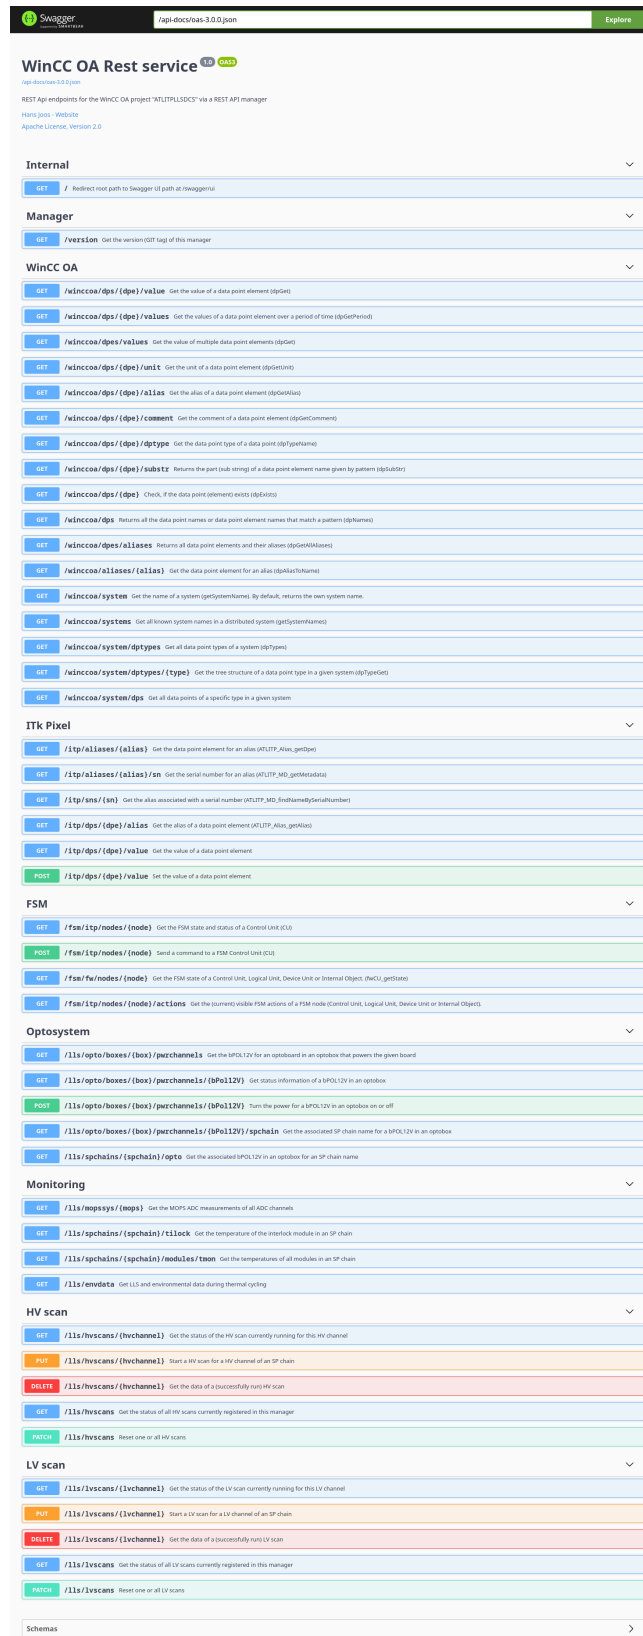


Figure 9.16: Swagger UI of the WinCC OA HTTP REST API manager.

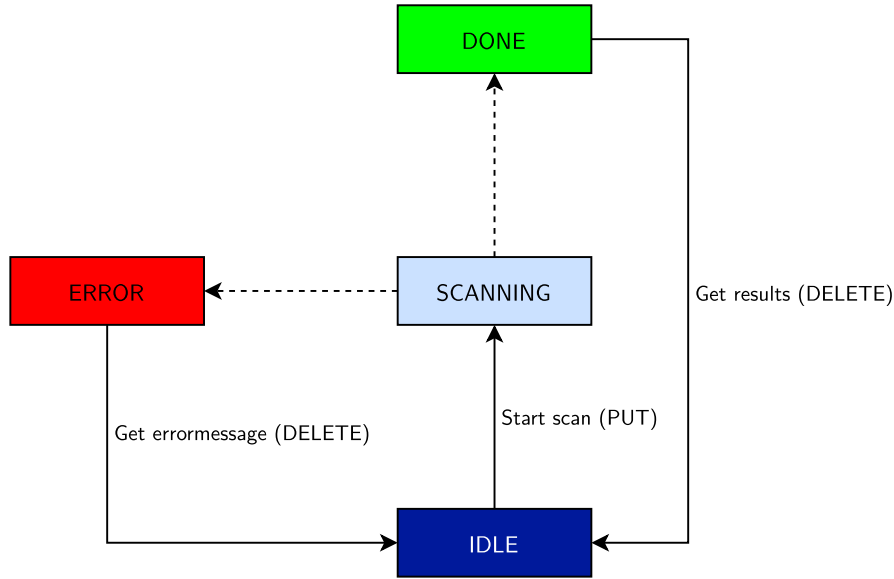


Figure 9.17: State diagram of the FSM for LV or HV scans performed with the WinCC OA HTTP REST API manager.

- Model: Dell Inc. PowerEdge R610
- CPU: 8x Intel(R) Xeon(R) CPU E5620 @ 2.40 GHz
- Memory: 7.5 GiB: 4x DDR3 2 GiB single rank 1333 MT/s
- Storage: 250 GB HDD

The benchmarks were conducted using the Apache benchmarking tool⁴. Both a high number of sequential and concurrent requests for a DPE value were sent and the response times measured. Regarding concurrent requests, the underlying Oat++ framework is limited in the number of concurrent requests it can accept by the number of file descriptors it can open. With a default soft limit of 1024, this represents the maximum number of concurrent requests that the API manager can handle. This limit could theoretically be raised to 524 288 on the test machine. All tests used the HTTP KeepAlive feature, where multiple requests were executed within one HTTP session to stress the WinCC OA back-end of the manager.

One benchmark set compared response times when connections originated locally versus from within the CERN campus network. Figure 9.18(a) shows the response times for both request types as a function of the percentiles of 900 000 requests. The results show that response times remain consistent across many requests, with only minor outliers. The mean response time when connecting from the local machine with 1000 concurrent connections is approximately 40 ms for the request of a DPE value. When connecting from the CERN

⁴<https://httpd.apache.org/docs/2.4/programs/ab.html>




Figure 9.18: Benchmarks of the WinCC OA HTTP REST API manager. The request for a DPE value was tested. (a) shows a comparison of the response time when running the benchmark locally or from a device connected in the same network with 1000 concurrent connections. The x -axis shows the percentiles of 900 000 requests. (b) compares the response time for various amount of concurrent connections during the benchmark. The x -axis shows the percentiles of 10 000 requests.

campus network, the response time rises to 150 ms, with the added latency attributable to the network transmission.

Figure 9.18(b) shows response times for different level of concurrencies as a function of percentiles of 10 000 requests. This test was performed from another machine within the CERN campus network. The response time increases linearly with the number of concurrent requests.

Generally, the API manager provides slow control and is not required to meet strict response-time thresholds. Instead, it must provide a reliable way to interact with or retrieve information from the DCS. The previous tests focused on a heavy load scenario. For the more common case of single requests, the typical response time is an interesting performance indicator. Table 9.1 gives response times for various types of requests that allow a grading of the performance of the API manager. The typical response time for natively accessing DP values (using a CTRL manager) is given and can be compared to the response time when the API manager is used. Distinct response times for requests coming from the same machine and from a machine in the same network are given. It is visible that the network adds latency. The average response time of 2.263 ms for requests via the API manager is significantly slower than native access. This time is heavily influenced by the waiting time of the main thread, which is set to 2 ms (see also Fig. 9.15). A fastest achievable response time of the API manager is given by testing the HTTP endpoint that simply returns the manager's version without going through the message queue. This time is given as 0.256 ms. The provided benchmarks confirm that the API manager fulfills the requirements for a slow control interface.

Table 9.1: Given are various performance indicators that enable a comparison of the performance of the API manager. Each value was obtained as an average of 10 000 repeated tests.

Indicator	Average time
Request of the API manager version from same machine (without KeepAlive)	0.256 ms
Request of a DPE value via the API manager from same machine (without KeepAlive)	2.263 ms
Request of a DPE value via the API manager from remote machine (without KeepAlive)	7.803 ms
Request of a DPE value from a CTRL manager	0.175 ms

Communication between the API manager and the event manager of the WinCC OA project is also organized by a message queue. The event manager accepts (by default) up to 100 000 messages in its queue in the receiving direction. Being limited to approximately 1000 concurrent connections, this limit cannot be exceeded by the API manager. Every request from an API connection needs to wait until it receives an answer from the event manager allowing it to clear its queue.

9.4 The Finite State Machine for ITk Pixel setups

To ensure safe operation of the LLS QC setup and enable supervision of detector units by non-experts, an FSM control hierarchy has been implemented for the QC DCS. It acts as a layer between the low-level hardware-oriented components of the DCS and the setup operation and supervision. This FSM hierarchy enables control of parts as well as the entire LLS through simple commands. Combined with the alert system, it simplifies debugging of problems.

9.4.1 FSM control hierarchy

The FSM control hierarchy is structured as a tree as shown in [Fig. 9.19](#). Commands propagate downward through the tree, while states propagate upward. The leaves of the tree are Device Units (DUs), shown in orange. White nodes represent Logical Units (LUs) and blue colored nodes in the tree denote Control Units (CUs) as defined in [Section 8.4.3](#). The tree has been refined from Ref. [\[179\]](#) and consists of two main subtrees: the infrastructure tree and the detector tree. In the detector tree, the DUs are mainly power supply channels for powering the detector parts, with one powering channel for the

Opto Boards of an SP-chain as in the final design. The DUs in the infrastructure tree are mainly environment sensors and status signals from infrastructure hardware.

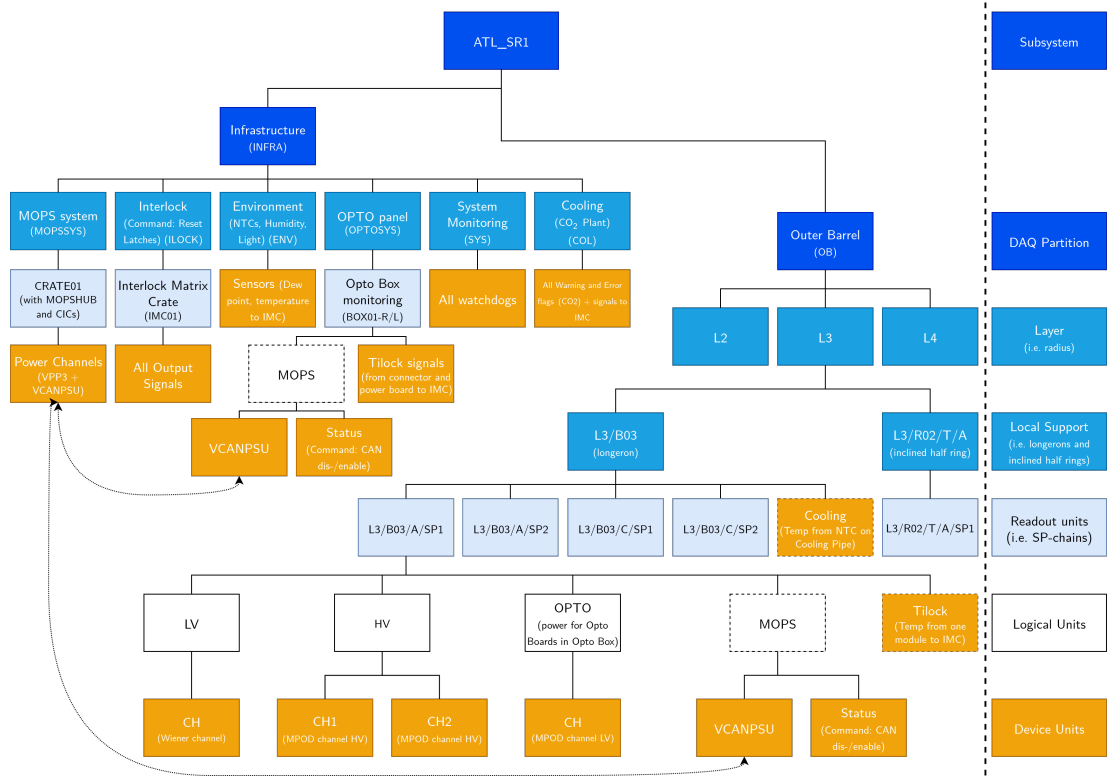


Figure 9.19: FSM hierarchy tree of Control Units, Logical Units and Device Units for the ITk Pixel setups. The tree consists of two main subtrees, the infrastructure tree and the detector tree. Due to lack of space only one SP-chain CU is shown in full detail. FSM objects shown with dashed rectangles have only their *status* propagated upwards in the tree. A MOPS that is switched off raises an alert without changing the *state* of the parent CU. The three VCANPSU DUs that are connected with arrows connect to the same power supply channel.

The non-expert operator of the LLS QC setup will usually use the lowest CU to gather detailed information about the state of the detector parts. The lowest CU in the detector tree is an SP-chain. It is the smallest unit that can be controlled independently from other parts of the detector. A detailed description of the SP-chain CU as an example for the implementation of FSM logic is given in [Section 9.4.2](#).

The general creation of CUs, LUs and DUs follows the conventions in ATLAS [233]. For the DUs, their operational state is computed inside a CTRL script in WinCC OA. The status of a DU is usually obtained from the alert handling of the associated data points. All state, status and command labels have upper case letters. The ATLAS specific extension for the FSM components for WinCC OA predefines a color coding scheme of 8 different

colors for the state of FSM objects. An orange or red color indicates a problem. A light blue color indicates a transition state that is expected to transition into a static state after some time.

For the time being, the FSM logic implemented in the objects does not contain a software interlock. If an error is detected, the FSM or alert system will signal the error to the operator and allow the operator to take action. This concerns mainly the status received from the power supplies. In case of dangerous conditions in the environment, the interlock will act directly on the power supplies, independent of the FSM. However, the FSM will reflect the changed state immediately. If, for instance, a power supply trips, the FSM reports the trip and then moves to represent the state with the power supply channel being off. This might be an error state and the FSM could be programmed to take automatic actions such as turning other power supplies off to move to a save state.

As described previously, the FSM tree for the project is created by using the aliases in the system. A library of functions automates the creation of the FSM tree in the main system by traversing down the logical tree in all connected, distributed systems and pattern matches found aliases to FSM object types. This means that the FSM only runs on the main system, but the DUs will access data points on remote systems if so instructed. Using this approach, the FSM tree can be easily adjusted to the detector hardware available in the system. This includes, for instance, if a different LLS type is being tested, a new connectivity can be read and a new FSM tree can be generated. The DU's name is then the alias of the associated data point and already includes the parent tree structure. In the names for the LUs and CUs the “/” symbol is replaced by an underscore (`_`), since the slash is a reserved character in these cases.

9.4.2 SP-chain Control Unit

The modules in an SP-chain require a supply line (LV) in order to power their readout chips by a constant current source, and a sensor bias line (HV) in order to deplete their silicon sensors. Data from the FE chips travels electrically over a distance of approximately 6 m to Opto Boards, where the electrical signals from different chips are serialized and transmitted to the FELIX readout system via an optical link. The active components on the Opto Boards require a separate powering line (OPTO), where the design of the detector system guarantees always exactly one OPTO line for all Opto Boards belonging to an SP-chain.

The operational model and hardware design dictate a strict power-on sequence of an SP-chain: To protect the inputs of the GBCRs on an Opto Board, OPTO power must be provided before turning on the modules in the connected SP-chain (see [Section 5.7](#)).

Table 9.2: FSM states of a power supply channel DU.

State	Description
FSM_TEST	Warning state: The DU is forcibly set to a testing state
ON	Static state: The channel is switched on and outputting constant voltage/current.
OFF	Static state: The channel is switched off.
RAMPING_UP/DOWN	Transition state: The channel is ramping up or down its output.
UNKNOWN	Error state: Connection to the power supply is lost.
INHIBIT	Static state: Power channel receives external inhibit signal. The channel is switched off and cannot be turned on in this state.
TRIPPED	Severe error state: The channel has tripped after exceeding operational limits.

To have well-defined leakage current paths in an SP-chain, LV must be turned on before turning on HV (see [Section 4.5.1](#)). To enforce this sequence, the SP-chain CU is the lowest CU in the FSM detector tree design and contains all three power line types (OPTO, LV and HV). The content of the SP-chain CU is also shown in [Fig. 9.19](#). The OPTO, LV and HV units are added as LUs. The DUs are then the associated channels of the power supplies. The HV unit groups together two power channels. Currently, no individual modules are added as children, since no single-module control is possible in the serial powering scheme.

The Device Units and Logical Units of an SP-chain

The FSM logic for the DUs of the power supply channels is implemented using scripts in WinCC OA. The behavior of the DUs for different power supply types is almost identical. The possible *states* of these DUs are listed in [Table 9.2](#). The decision flow implemented in the DU script for state determination is shown in [Appendix C](#) in [Fig. C.2](#). During FSM startup, an auxiliary DP with a structure shown in [Fig. 9.20](#) is created for each power supply channel DU to provide the means to set the DU into the FSM_TEST state and test the state propagation of the FSM tree up to the highest level.

The *status* of the power supply DUs is derived from the alert handling of the associated data points (for channel-tripped information) and an additional consistency computation using the auxiliary DP in [Fig. 9.20](#). The consistency check takes into account a final state that has been requested from the parent SP-chain CU when sending a command, and

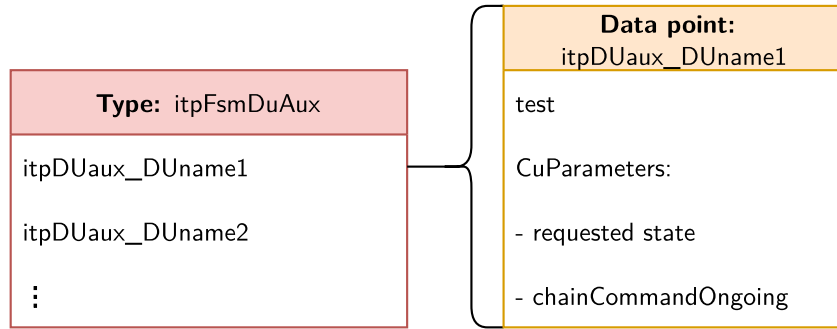


Figure 9.20: Auxiliary DPs for each power supply channel DU to introduce FSM testing functionality and reporting functionality in case of inconsistencies in the SP-chain CU. The “CuParameters” elements are copies of parameters from the SP-chain CU and can be used to compare the own DUs state to the requested state from the last SP-chain command. In case the requested state was not reached after finishing the command, an alert is raised.

compares this requested state with the DU’s own state to raise an alert (and move to the ERROR status) in case the DU’s state is inconsistent with the requested state from the CU. The convention implemented here is that the *state* of the FSM describes the state of the detector without taking into account if this state is safe, useful or currently wanted. The *status* of the FSM is used to convey this type of information and to raise alerts if something needs to be checked by the operator.

The DU auxiliary DP shown in Fig. 9.20 could additionally be used in the future to implement dynamic thresholds for raising alerts depending on the mode of the detector, for example during stable beams situations versus calibration mode.

The higher-level LV and OPTO LU units replicate the state behavior of their single DU, while the HV LU unit adds an additional MIXED state, in case that the two HV power channels are in different states.

The operational model of the SP-chain Control Unit and chain commands

The SP-chain CU consists of its own set of states, the transition between these states, and actions available in each state. The state of each SP-chain CU summarizes the states of the child LUs in a way that highlights the operational procedure. The number of states should be kept relatively small to simplify the integration of the SP-chain CU in higher-level CUs. The states available for the SP-chain CU are described in Table 9.3.

The normal states are: OFF (where OPTO, LV and HV are off), OPTO_ON (where the OPTO power is first supplied to the Opto Boards), LV_ON (where additionally LV power turns on the FE chips) and ON (where additionally HV depletes the sensors). Concerning

Table 9.3: FSM states of the SP-chain CU ordered by priority in the state calculation flow with the highest priority on top.

State	Description
FSM_TEST	Warning state: Any child in FSM_TEST state
UNKNOWN	Error state: One of OPTO, LV or HV is in state UNKNOWN or some children are disabled
TRANSITION	Transition state: One of OPTO, LV or HV is in state TRANSITION
TRIPPED	Error state: One of OPTO, LV or HV is in state TRIPPED. Tries to CLEAR the tripped flag (but not the alarm).
ON	Static state: OPTO, LV and HV are ON.
LV_ON	Static state: OPTO and LV are ON while HV is OFF or INHIBIT.
OPTO_ON	Static state: OPTO is ON while LV and HV are OFF or INHIBIT.
OFF	Static state: OPTO is OFF while LV and HV are OFF or INHIBIT.
OPTO_RAMP	Transition state: OPTO is ramping up or down while LV and HV are OFF or INHIBIT.
LV_RAMP	Transition state: OPTO is ON, LV is ramping up or down while HV is OFF or INHIBIT
HV_RAMP_UP/DOWN	Transition state: HV is ramping up or down while OPTO and LV are ON.
OPTO_CHANGE	Transition state: OPTO is ramping up or down while LV is ON or ramping up or down.
NO_OPTO	Error state: OPTO is not ON while LV is ON or ramping up or down.
HV_MIXED	Warning state: OPTO and LV are ON, but only one HV channel is ON.
HV_ON	Error state: LV is OFF or INHIBIT, while HV is on or ramping up or down.
INHIBIT	Error state: OPTO is INHIBIT while LV and HV are OFF or INHIBIT.
HV_RESET	Forced transition state: Action to turn off HV.
OPTO_RESET	Forced transition state: Actions to turn off first HV and then LV.

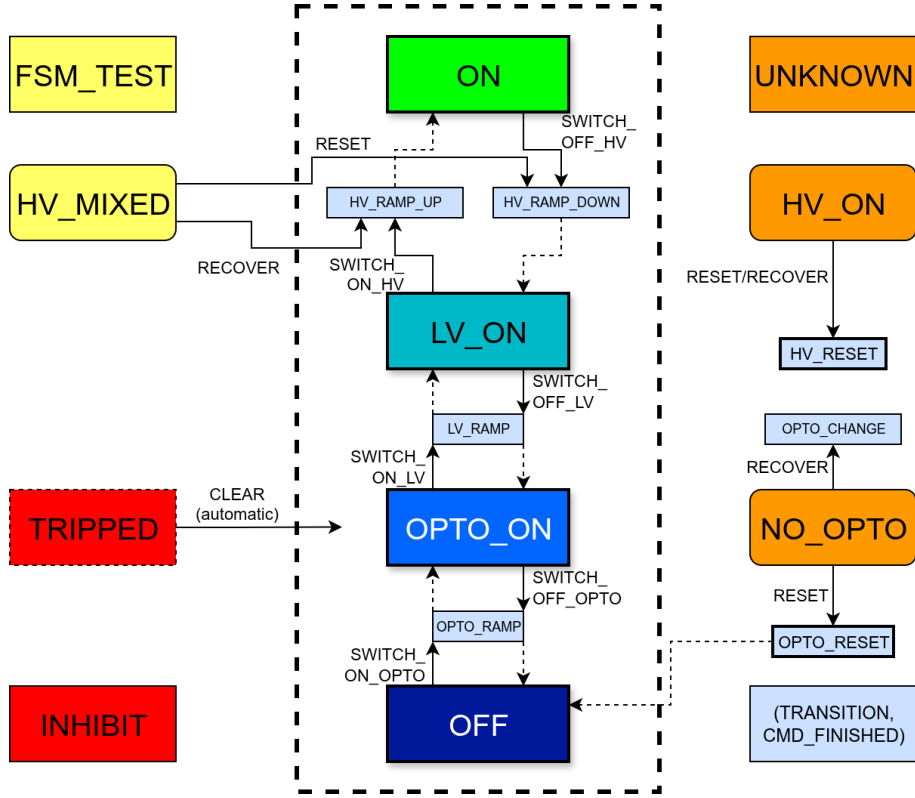


Figure 9.21: State diagram of the SP-chain Control Unit. The normal operational states are OFF, OPTO_ON, LV_ON and ON. The named arrows indicate commands. The HV_RESET and OPTO_RESET states are “forced” transition states that can only be reached through commands from the HV_ON or NO_OPTO states, respectively.

the number of states, it was decided to create all necessary states such that to move from one normal state to the next, only one command has to be sent to exactly one child LU. Two error states exist with HV_ON and NO_OPTO that need fast reaction from the operator. A HV_MIXED error state indicates that one of the two HV channels in an SP-chain failed to turn on. These three error states are grouped together on the next higher level of the FSM tree. The state transitions that are foreseen during a normal operation are shown in the state diagram in Fig. 9.21. The arrows in the diagram indicate the direction of commands.

The TRIPPED state is a temporary error state entered while attempting to clear the tripped flag from the associated DU (without clearing the alert that is raised at the same time, preserving this information). This ensures the tripped power supply channel is shown as OFF promptly, allowing the SP-chain CU to reflect the actual state of the SP-chain again. INHIBIT and OFF states of DUs are treated mostly the same in the SP-chain CU. For example, if only HV is INHIBIT, which can be the case in LLS QC setups if the light sensor disables HV (see Table 7.1), or in the final detector during unstable beams, it is treated as

OFF and control of OPTO and LV is retained while reflecting the actual SP-chain state. If a request to turn on HV is sent but cannot be fulfilled due to INHIBIT, an alert is raised in the HV DU as previously explained. Similarly, if HV control is undesired and the HV LU is disabled, this is treated as if HV is OFF to allow control of the rest of the SP-chain.

The TRANSITION state is entered immediately after sending a command and waiting for the DU to acknowledge receipt. While in TRANSITION, the CU does not accept new commands, but typically moves shortly after into one of the RAMPING states, which allow sending commands again.

It is foreseen to operate the SP-chain from its CU. To prevent incorrect operation, the commands directly onto the LUs that are available from the FSM UI are marked as “expert” commands.

The three different LUs have to be handled cautiously, since they are not independent and need commands sent in the correct order. To be able to go with one command through several states in the state diagram of the CU as shown in [Fig. 9.21](#), a few options are available.

1. The SML language offers a WAIT_FOR instruction that can be used when declaring an action. A second command can then only be sent if a certain condition is fulfilled. During execution of the chain command however, the state propagation is blocked and intermediate states are not captured.
2. One can add an extra DU under the CU and let it handle all the commands. A reliable connection then has to be guaranteed. This method was chosen for the smallest CU in the FSM of the current ATLAS Pixel detector.
3. One can add additional states for each possible next command in a command chain. These additional states have the same conditions for entering than the stable states. The difference is that on reaching the additional state, they automatically execute the next command in the command chain. Each additional state has one command associated with it. This is akin to the Moore machine discussed in [Section 8.4.1](#) and would result in many more states.
4. One can use object parameters that persist over state changes. These parameters can be used to trigger the next action. This is akin to the Mealy machine.

For the implementation of chain commands for the SP-chain CU, option 4 with object parameters was chosen.

There are four parameters, the previous state, the final state, the requested state and the command type. When a chain command is issued, these four parameters are set with the requested state being the same as the final state but not being deleted at the end of the chain command. Then the first action is executed. Once the CU is in the next static state,

it checks the parameters again. By knowing the final state, the object knows where to proceed next. By knowing the previous state and comparing to the current one, the object knows, whether it is proceeding in the right direction or should stop the chain commands because it is stuck in a state. The command type for this scenario is either RECOVER or GOTO. RECOVER commands do not need a final state parameter and stop as soon as a normal state is reached. GOTO commands can be used to continue when a transition from an irregular state to a normal state occurs. The command type parameter can be used to check if a chain command is ongoing, i.e. if its value is not empty. At the end of a chain of commands, the unit goes into a transition state called CMD_FINISHED, in which it resets all parameters except for the requested state parameter. This state will also appear in logs and can help to understand and retrace the transitions of a unit. A parent of the unit will also be informed of the state changes. The requested state parameter (used by power supply channel DUs to detect inconsistencies) can be manually reset via the GUI.

The simple SWITCH_ON_* commands for the CU are then mostly hidden in the FSM UI and GOTO_* commands are used (for example GOTO_ON), where the name of the command is independent of the current state.

Currently, no delay is implemented between two actions in a chain of commands. This might pose a problem if a device (FE chip or Opto Board) is not immediately ready, once the supply voltage is provided, but the next action has already started and the device cannot be brought into a ready state afterwards. If more information for the DCS is needed about the state of an FE chip or Opto Board configuration, it can be added under the LV LU or OPTO LU respectively, and then be taken into account during chain commands.

In addition to the previously discussed LV, HV and OPTO children, the SP-chain CU also has a Tilock DU and a MOPS DU. These two units, however, do not play a role in the state calculation for the SP-chain. They only propagate their *status* up to the CU. The Tilock raises an alarm if the module temperature, that is routed to the interlock system, is too high. The MOPS LU will raise an alarm if the DCS monitoring of the SP-chain is defective.

The GUI for the SP-chain Control Unit

The control of the FSM is provided via an ATLAS customized UI framework in which created panels are embedded. The FSM UI frame shows information about the FSM tree and offers control of and navigation through the tree. [Figure 9.22](#) shows the SP-chain panel embedded in the FSM UI for the M12 SP-chain of the preproduction longeron at CERN. One can navigate through the tree by clicking on the names of the different units on the left. The alarm table at the top of the panel collects all bad statuses in the FSM

tree. By clicking on the alarm, the FSM UI navigates directly to the unit where the alarm was raised. With this, one can quickly investigate the cause of an alarm.

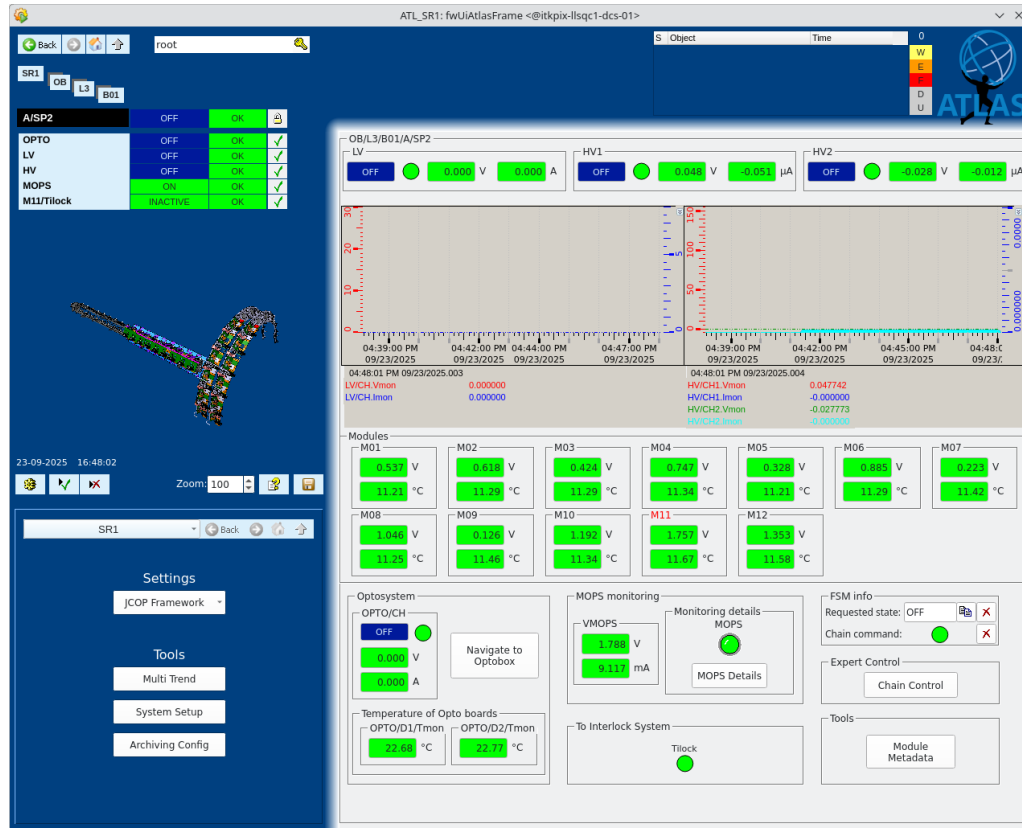


Figure 9.22: The SP-chain monitoring panel, here for the M12 of the preproduction longeron at CERN. The panel displays all required DCS parameters as listed in [Section 9.2](#).

The monitoring panel for the SP-chain shows the alias of the SP-chain at the top. Then there are the states of the LV channel and the two HV channels. The text box shows status information like ON, OFF, RAMPING, TRIPPED or INHIBIT. This information comes directly from the power supplies. The Light-Emitting Diode (LED) graphic next to it shows the corresponding interlock signal that is sent from the interlock system to the power supply channel. A click on it opens the interlock panel that is mentioned in [Section 9.3.5](#). The currently measured output voltage and output current of the power supply channels are shown.

The trending plots below show the voltage and current of the LV channel over time on the left side and the voltage and current of the two HV channels over time on the right side. Below the plots is a field that shows the temperature and voltage drops of the modules as measured by the MOPS chip of the SP-chain with one exception: The red-highlighted module M11 is the module with its NTC connected to and read by the interlock system.

Since different SP-chains can have different number of modules, the panel adds them dynamically on load by checking which aliases exists in the database.

At the bottom left, there is the power channel for the Opto Boards assigned to the SP-chain. The temperatures of the Opto Boards are also shown. These temperatures are read out by another MOPS chip inside the Opto Box.

The supplied power on the CAN bus cable to which the MOPS on the SP-chain is connected is shown in the middle. The state of this MOPS chip is indicated with an LED graphic to the right of it. Below, the generated interlock (input) signal from the interlock module NTC (here M11) is shown with label “Tilock”. In case this temperature sensor generates an interlock signal, the LED graphic turns red and an alert is triggered. A right-click on the LED graphic opens a panel to acknowledge the alert. On the right is information about the SP-chain CU. The current value of the requested state parameter is shown with the possibility to set it to the current state or clear it. Additionally, an LED indicates whether a chain command is ongoing and a button offers the possibility to cancel the ongoing chain command.

An expert panel can be opened which allows channel wise control of the SP-chain, bypassing the FSM control logic. Lastly, the panel displaying the metadata of all detector hardware of the current SP-chain as shown in [Fig. 9.14](#) can be opened.

9.4.3 Testing, operation and performance considerations

The DCS was developed on a dedicated development machine, tested successfully for the first time in the empty OB demonstrator setup and subsequently used for demonstrator tests. It underwent continuous development and the updated version has operated at the LLS QC setup at CERN after prior validation in the empty setup. Tests of the FSM control path and state calculation in an empty setup were performed using open circuits on voltage sources and a high-power resistor on the current source. The DCS proved reliable and useful during the operation of the OB demonstrator and the preproduction longeron. Throughout these operations, the DCS exhibited no performance degradation, and the computer never ran out of available memory.

Regarding system performance and scalability, the setup was designed to enable quick addition of extra hardware by reading a CSV file. If the alias of the hardware follows the standard naming scheme, it will automatically be included in the generation of the FSM tree. The FSM library was written in a way such that by default all connected distributed systems are scanned for aliases in order to check if they should be included in the FSM tree. During the system setup, it is not important if the time needed to

create the FSM tree is short. Once the system is in operation, however, look-ups of data points in the local database by their alias should proceed quickly. The panels are built in a way where they utilize aliases to display information. During initialization of a panel, different parts might only be added after a data point look-up was performed, which can delay the moment until the complete panel is shown. If the full alias of a data point is known, the look-up of the corresponding data point name is instantaneous, even if it is in a remote system (on the same computer). The following performance considerations were already given in [179] and are repeated here for completeness. With native time measurement functions, the time needed to execute the `dpAliasToName()` function that performs this look-up is less than the measurement precision. If only a part of the alias is known and the look-up is performed with wildcards, e.g. `OB/L3/B01/A/SP1/HV/CH*` with the `dpGetAllAliases()` function, the time needed for finding the corresponding aliases increases drastically. Native time measurement functions give an execution time for one function call of this type of approximately 70 ms. In this simple example, where one can guess the range of possible aliases that will match the given search pattern, it is much more efficient to simply check these possible aliases (here CH1 and CH2) individually. Executing 1000 calls to `dpAliasToName()` currently costs approximately the same time as one call to `dpGetAllAliases()`, as was seen in a quick test for the system setup. The usage of `dpGetAllAliases()` should therefore be kept to a minimum and kept as a last resort. In some situations, for example for a function checking for the module in an SP-chain with the interlock connection, the usage of a cache system was realized, where after the first successful search, the results are reused. A general cache system, especially for the alias to DPE mapping, could be implemented for each newly opened UI manager, if it should turn out that querying the internal Data Point identification container, that contains the whole information of DPs and aliases, is too resource intensive in a larger system.

As the system was designed for scalability, it is planned to use it during testing of larger detector structures. One such structure is an inclined unit as shown in [Fig. 4.16](#). The largest inclined unit comprises 9 IHRs and therefore 18 SP-chains. To verify the current design of the DCS can handle this scale, a performance test has been conducted for an FSM with 18 configured SP-chains (but no detector hardware available). Since no power supply emulator for WinCC OA was available at the time of these tests, an OPC-UA server has been set up that connected to one spare Wiener PL512 voltage source with 12 channels (without any load), but exposed 6 copies of each channel via OPC-UA to WinCC OA. This configuration replicated real power supply channel ramping behavior in the FSM. More precisely, changing one physical channel's state propagated to 6 FSM channels.

The test evaluated whether state changes in an FSM of this scale could be processed by the computer hardware. The same rack-mounted server as in the benchmark tests in [Section 9.3.12](#) was used (specifications repeated for convenience):

- Model: Dell Inc. PowerEdge R610
- CPU: 8x Intel(R) Xeon(R) CPU E5620 @ 2.40 GHz
- Memory: 7.5 GiB: 4x DDR3 2 GiB single rank 1333 MT/s
- Storage: 250 GB HDD


The 18 SP-chains in the FSM corresponded, together with the infrastructure tree, to 42 CUs or SMI processes. The system resources of the machine were monitored using the prometheus⁵ monitoring tool. The CPU utilization over time during testing is shown in Fig. 9.23. Figure 9.23(a) shows the CPU times of DCS-related processes. Testing periods were from 17:13:00 to 17:16:00 and from 17:22:00 to 17:25:00 with commands being sent to induce state changes. Notably, no SMI process shows significant CPU time during these tests. The CTRL manager, that includes handling of all DUs, reached 15 % utilization of one CPU core during these tests. The UI manager spiked when opening the FSM UI, but stabilized afterwards. The event manager showed only minor increases in CPU time. Prometheus exhibited significant fluctuations. Overall CPU utilization is shown in Fig. 9.23(b). The time waiting on input/output (I/O) shows spikes that coincide with the opening of the FSM UI and could be caused by loading data from the Hard Disk Drive (HDD). In general, Fig. 9.23(b) confirms that FSM state calculations impose negligible load on the 8 core system, with ample headroom remaining.

The scaling tests, due to the way they were set up, revealed one issue: In rare cases, the FSM was forwarding a command to a DU whose state recently changed in which the command was not valid anymore. This can be mitigated by checking in the DU script itself, if a received command is valid.


9.4.4 Discussion of the chosen implementation

The current implementation of the Pixel FSM has proven well-suited for system tests and QC setups, supporting the intended testing procedures. The design of the SP-chain CU ensures control of the detector unit even when incomplete (for example missing HV channels). Compared to Fig. 9.1, both the control and feedback path and the safety path are implemented, but the diagnostics path is still missing, as it depends on solutions for the FELIX readout which were not available at the time of writing. Depending on the diagnostics information that will be available, modifications to the design of the SP-chain CU may prove beneficial. New states might be helpful when incorporating DAQ information like the configuration state of a module.

⁵<https://prometheus.io/>, an open-source monitoring system



(a)



(b)

Figure 9.23: CPU performance under an FSM with 18 SP-chains. SP-chain activity started at 17:13:00 by sending commands to cause state changes in the FSM until 17:16:00 and then again from 17:22:00 on. The tests continued until 17:25:00. (a) shows the CPU usage of different DCS-related processes (or groups of processes) as CPU seconds per real seconds, meaning that a value of 1 s implies that the process is using one CPU core fully. Note, that only processes are listed in the legend, that had at some point a value above 50 ms. All SMI processes stay well below 1 ms. (b) shows the overall CPU utilization of the 8 core system.

Nevertheless, even with the current setup, an alternative FSM tree structure could have been implemented. Given that the DCS system was designed from the outset for the OB demonstrator and subsequent LLS QC tests, it was decided to make the system configurable via configuration files and to align the FSM tree with the mechanical structure of the detector. Although LLS testing sites test only one LLS at a time, the FSM tree incorporates all detector levels required by the official LLS naming scheme, i.e. the official naming scheme follows the mechanical structure. This also allows to accommodate for any number or combination of LLSs in a setup. The smallest detector unit that can be controlled independently is a SP-chain, which is reflected in the FSM tree by making the SP-chain the lowest CU.

However, alternative tree structures could group SP-chain CUs by cooling groups (to quickly act on all SP-chains within one cooling group), by readout paths (to be able to turn off all SP-chains connected to one Opto Box in case the box needs to be replaced) or by off-detector services (following the detector’s cabling structure). Ultimately, many different “views” of the detector may be necessary and should be implemented, with one tree being the main tree used to signal the detector status to the central ATLAS DCS.

Experience from current setups will clarify which commands to detector hardware are required for experts versus operators. The control of the setup should mainly be realized by using FSM commands as the FSM logic prevents errors and simplifies handling of the hardware. Access control is currently not implemented, but could restrict permissions for non-experts through rights management.

Additionally, no automatic recovery actions currently exist and an operator should always supervise the setup. While only power supply channel trips are automatically recoverable by the DCS, proactive measures could prevent interlock triggers (for example gracefully powering down detector units before interlock conditions are met if cooling degrades and module temperatures rise).

The current implementation has proven its usefulness for the operation of the OB demonstrator and in LLS QC setups. A detector slice test at CERN with several LLSs from different subsystems is currently in preparation, offering further opportunities to refine the system.

9.5 Towards a final detector DCS

Following the preceding discussion, several general improvements to the current DCS emerge. The connectivity information (the mapping of aliases to DPEs) is currently derived from a CSV file. The community has initiated work to integrate connectivity data into a

database. Importing this information into the DCS with consistency checks on the provided structure would eliminate errors when creating the CSV file manually.

The solution presented here will support an upcoming slice test at CERN with LLSs from the OB and EC and eventually IS subsystems. This setup will progressively incorporate production hardware, approximating real operational scenarios at a smaller scale. While not made for final conditions so far, this DCS implementation provides a robust foundation and could develop in the slice test to the DCS of the final detector.

During the slice test, it will be important to observe the DCS performance and reliability under real world conditions, in order to evaluate coverage of common use cases and identify areas that require revision. Attention was paid to make the DCS system compatible with different configurations. Currently, the DCS project focused mostly on the control and feedback path of the ITk Pixel DCS. But it is already possible to add monitoring of a LISSY crate to the DCS. Once a final MOPSHUB crate is available in the slice setup, integration should prove straightforward, as minimal differences to the MHFB in the control and monitoring functionality are expected.

The diagnostics path is the path with the most undetermined future. As the slice setup closely mirrors the final detector in hardware and functionality, it offers an ideal platform to explore solutions for the final DCS and DAQ. Concurrent DAQ system development toward a production version enables progress on the DCS diagnostics path.

Also other developments in the DCS can go hand in hand with the evolution of the DAQ system. In order to synchronize the state of the detector with the DAQ during operation, bi-directional communication between DCS and DAQ control must be provided [234]. How tight this integration is required to be, is not yet clear. Sharing information about the state of an SP-chain from the DCS to DAQ may be necessary for the DAQ to determine when a FE chip can be configured.

Within the context of this thesis, potential implementations were partially explored, yielding a small Proof of Principle (PoP). A possible implementation of a communication architecture between DCS and DAQ is shown in Fig. 9.24. An OPC-UA server would provide state and status information of all relevant FSM objects. It would receive this information via DIM directly from the SMI processes of the FSM. The DAQ could connect to this OPC-UA server and subscribe to the state information of a certain SP-chain. Additionally, it could contribute information about the configuration state of the modules associated to an SP-chain to the OPC-UA server. This information, on the other hand, could be read by the DCS.

A PoP was developed using the quasar framework to implement the path labeled DIM in Fig. 9.24. This prototype scans all available FSM objects connected to a DIM name server




Figure 9.24: Possible implementation of a communication system between DCS and DAQ.

(DNS) and connects to their state, status and parameter information. This information is then made available via the OPC-UA protocol.

It is an exciting time, as many developments for DCS targeting operation are currently gaining momentum and a picture is emerging on the operational model of the detector during calibration and data taking. The aim is to have a fully functional and well-developed system ready for operation in Run 4. Only then can the potential of the HL-LHC be fully exploited. The presented DCS project is a step in this direction.

This concludes the work on the ATLAS ITk project for this thesis. [Chapter 10](#) presents contributions to the FASER experiment, a small forward experiment at the LHC not far from the ATLAS IP.

Controlling the FASER Experiment

The two general purpose detectors at the LHC are the ATLAS and the CMS experiments. Their cylindrical design around the beam pipe allows them to detect heavy particles with large transverse momenta (p_T). With the HL-LHC, significant efforts are being made to increase the experimental data set and enable precision measurements to confirm or find inconsistencies in the current theory.

While the Standard Model (SM) has proven hugely successful in providing experimental predictions for collider experiments, there are observations that the SM cannot explain. The SM, for example, does not contain any viable dark matter candidate that is needed to explain observed gravitational effects. The search for new particles at ATLAS and CMS usually targets high- p_T signatures. There is, however, a complementary class of viable new dark matter candidates that are much lighter, with masses in the MeV to GeV range, and much more weakly coupled to the particles of the SM [249]. At the LHC, searching for such weakly coupled particles is ineffective in the high- p_T region due to insufficient cross-sections for producing them in sufficient numbers. The situation differs significantly at low p_T where the inelastic cross-section is many orders of magnitude larger with most of it in the very forward direction [250, 251]. Such weakly coupled particles are typically long lived and travel macroscopic distances without interacting before decaying to SM particles.

In 2021, a new experiment was installed at the LHC, called ForwArd Search ExpeRiment (FASER) [252–254]. FASER searches for light, very weakly-interacting particles which may be produced in the collisions at the ATLAS interaction point (IP1), travel long distances through concrete and rock without interacting and then decay to visible particles in the

detector decay volume [255]. These Long-Lived Particles (LLPs) are a suitable candidate for dark matter.

FASER started taking LHC collision data in 2022 and will run throughout LHC Run 3. Within the context of this thesis, contributions were made to the Detector Control System (DCS) of the FASER experiment with the commissioning and integration of a Water Leak Detection (WALD) system, as well as a development of a DCS for the new FASER Preshower upgrade installed in February 2025. This chapter therefore first presents the detector and the new Preshower in Section 10.1, and then briefly describes the FASER DCS in Section 10.2 as it was developed before the installation of the Preshower upgrade. In Section 10.3 the WALD system is introduced and its installation, calibration and integration into the existing DCS are detailed. Lastly, Section 10.4 presents the newly developed DCS for the Preshower upgrade.

10.1 The detector

The FASER detector is positioned along the beam collision axis line-of-sight, 480 m from IP1 in a service tunnel, TI12, that was originally used for the transfer lines of LEP and not needed anymore for the LHC [6]. Figure 10.1 shows the location of the detector in TI12. To position the detector on the line-of-sight, a shallow trench was excavated in the upward-sloping tunnel. The size of this trench sets the overall detector dimensions, including the 10 cm radius transverse-size of the active part of the detector and the total length of about 7 m. FASER acceptance covers a pseudorapidity of $\eta > 9.2$. The FASER ν sub-detector [256, 257] has a slightly larger acceptance of $\eta > 8.5$ [6].

Detailed simulation studies showed that the only SM particles reaching the FASER detector at an appreciable rate are muons and neutrinos [254]. Most neutrinos traversing FASER originate directly from collisions in IP1, whereas muons primarily originate from showers produced when particles from IP1 collisions interact with the LHC machine elements.

These simulations also estimated the radiation level in the TI12 tunnel, yielding a TID of less than 5 mGy and NIEL of less than 5×10^7 n_{eq} per year. These values are estimated to be comparable to or lower than those in the ATLAS service cavern (USA15), where non-radiation-hard commercial electronics are used.

One of the detector's original goals was to enable searches for LLPs decaying into pairs of charged particles, such as electron-positron decays from so-called dark photons [255]. Requirements for the FASER detector before its installation included [6]:



Figure 10.1: The location of the FASER detector at the LHC. The detector is located in the unused TI12 service tunnel 480 m from the ATLAS interaction point (IP1) along the beam collision axis line-of-sight with the LHC machine, concrete and rock in between. (Reproduced from Ref. [6], licensed under CC BY 4.0)

- Implementing a veto system for charged particles that enter the detector to reject background events from high-energy muons (10^9 muons are expected to traverse FASER during operation in Run 3),
- Precisely tracking high-energy charged particles (the magnetic field in the decay-volume and spectrometer must separate the decay products, such as electron-positron pairs, to measurable distances within the detector),
- Measuring the energy of decay products in the electromagnetic calorimeter up to $\mathcal{O}(\text{TeV})$ with a few percent resolution.

The detector was designed to fulfill the above requirements. Other factors, such as the limited time available for design, construction, commissioning, and installation, motivated the use of existing detector components.

10.1.1 The original detector layout

Figure 10.2 shows a sketch of the FASER detector, with particles originating from collisions in IP1 entering from the right. The different detector systems are briefly explained below. For a thorough overview, see Ref. [6].

The magnet system

Three permanent dipole magnets with a magnetic field of 0.57 T are used in the FASER detector. The first dipole (1.5 m long) surrounds the decay volume, while the other two




Figure 10.2: A sketch of the original FASER detector (without the Preshower upgrade), showing the different sub-detector systems. Particles originating from collisions in IP1 enter the detector from the right. The FASER coordinate system is also shown. (Reproduced from Ref. [6], licensed under CC BY 4.0)

(each 1 m long) are part of the tracking spectrometer. Their primary purpose is to separate closely spaced charged particles from LLP decays and measure the charge of muons arising from neutrino interactions.

The tracking system

The tracking system consists of two distinct components: the tracking spectrometer (see Ref. [258]) and the Interface Tracker (IFT). The tracking spectrometer allows the trajectories of charged particles traversing the detector to be reconstructed, and their position, momentum, and charge to be measured. The IFT, placed immediately after the FASER ν emulsion detector, enables tracks reconstructed in the emulsion to be matched with those in the active detectors, allowing, for example, charge identifications of muons from neutrino interactions.

Both systems use identical hardware components assembled in *stations*. The tracking spectrometer comprises three stations, while the IFT uses one. Each tracking station contains three double-layers of single-sided silicon microstrip detectors. A layer consists of eight silicon strip modules arranged in two columns of four modules. The modules used are spares from the ATLAS SCT barrel [259]. The SCT modules have a strip pitch of $80\text{ }\mu\text{m}$ and a stereo angle between the two sides of 40 mrad . They are aligned such that the resulting track resolution of approximately $20\text{ }\mu\text{m}$ in the precision coordinate is aligned

with the magnet bending plane (y), while the lower resolution of $800\text{ }\mu\text{m}$ is along x . The spectrometer is able to resolve closely spaced, charged particle tracks separated by 100 to $200\text{ }\mu\text{m}$.

The SCT modules operate at 150 V bias voltage and are cooled with two water chillers located in TI12 set to $15\text{ }^{\circ}\text{C}$.

The calorimeter and scintillator systems

Four scintillator systems, used for vetoing and trigger purposes, and an electromagnetic calorimeter, designed to measure the energy of high-energy electrons and photons, are installed in FASER. Each scintillator system consists of at least two scintillator counters read out by PMTs.

The first two scintillator systems veto charged particles entering the detector. They exceed the active transverse size of FASER to veto muons that could enter FASER at an angle with respect to the detector axis. The third system is a timing scintillator to provide a trigger for charged particles exiting the decay volume and to give a precise time for triggered events with sub-nanosecond resolution ($<1\text{ ns}$). The fourth system is the original preshower scintillator system, comprising two scintillator counters, each read out by a single PMT, interleaved with tungsten absorbers, and graphite blocks installed to reduce back-splash. This configuration aimed to distinguish neutrino events from photon events in the calorimeter.

The calorimeter system comprises four spare modules from the LHCb experiment's outer electromagnetic calorimeter (ECAL) [260] with a depth of $25 X_0$ to provide energy measurements. Each module measures $12 \times 12\text{ cm}^2$ in the transverse plane and contains 66 layers of lead and plastic scintillator. Wavelength-shifting fibers transport the collected light to a PMT at the back of the module. The modules are arranged in a 2×2 configuration and therefore provide very coarse transverse segmentation. The calorimeter cannot resolve two closely spaced electromagnetic showers but measures only the total electromagnetic energy per event.

The FASER ν emulsion detector

The FASER ν emulsion detector [256, 257] was added at the very front of the FASER detector to study high-energy neutrinos produced in colliders. A schematic illustration of the emulsion detector is shown in Fig. 10.3. The detector consists of 730 1 mm-thick tungsten plates interleaved with emulsion films and weighs approximately 1.1 t. The tungsten acts




Figure 10.3: A schematic illustration of the FASER ν emulsion detector at the very front of FASER. Particles originating from collisions in IP1 enter the detector from the left. (Reproduced from Ref. [261], licensed under CC BY 4.0)

as the target for the neutrino to interact with, while emulsion films record trajectories of charged particles from the interaction.

Unlike other detector components, the emulsion system is purely passive. For data readout, the emulsion detector must be extracted, and the emulsions developed and scanned. Once particle trajectories become visible, neutrino interaction vertices can be reconstructed. Given that the detector records all charged particle trajectories, it must be replaced before the track multiplicity in the emulsion becomes too high. The emulsion detector is replaced after exposure to 30 to 50 fb $^{-1}$ of collision data, typically three times per year.

10.1.2 The Preshower upgrade

In 2022, the collaboration proposed an upgrade of the simple tungsten-scintillator preshower to a tungsten-silicon high-precision Preshower [113] to significantly extend FASER’s sensitivity to LLPs that decay into photon pairs, such as Axion-Like Particles (ALPs) [262, 263]. This upgrade allows the experiment to distinguish single-photon events from those with multiple, very closely spaced photons by identifying the cores of electromagnetic showers initiated in the preshower’s tungsten layers. Achieving this requires the silicon detector to possess high granularity and a high dynamic range to detect particles over a wide energy range [264].

The original design specified six layers of 1 X_0 (3.5 mm) tungsten absorber alternating with planes of monolithic silicon pixel detectors (see Section 4.2.4). Refined simulations

evaluated three layouts featuring six layers of silicon pixel detectors and tungsten absorbers with varying thickness [265]. During assembly and qualification, the design was reduced to four pixel layers due to the yield in module production. It is not expected that this will have a significant negative impact on the achievement of the physics goals. Simulations showed that 4 detector planes are enough to identify with high certainty the cores of electromagnetic showers; the last two planes were originally added for redundancy and rejection of fake cores [266]. Installed in February 2025, the final Preshower configuration comprises 6 mm of tungsten in front of the first two pixel layers each and 4 mm in front of pixel layers 3 and 4. The two thicker tungsten layers ensure near-complete photon conversion before pixel layer 2, generating hits in at least three detector planes.

The following section briefly presents the design of the new Preshower detector. Further details are available in Refs. [113, 266].

Monolithic pixel detector (ASIC)

The design of the monolithic detector ASIC builds on research initiated at the University of Geneva in 2015 [143, 267]. The design of the production ASIC was finalized in May 2023 and delivered to the University's laboratories in July 2024 [266]. It is a 150 μm -thick monolithic silicon detector with a total size of $22.15 \times 15.35 \text{ mm}^2$. The hexagonal pixels have a side length of 65 μm , yielding an effective pitch of approximately 100 μm in both horizontal and vertical directions. The ASIC's active area contains 26 624 pixels arranged in a matrix with 208 columns and 128 rows.

The power consumption of the ASIC is dominated by the analog current in the FE amplifier electronics. A single ASIC consumes approximately 0.470 W [266], translating to 2.8 W for a module consisting of 6 ASICs and 33.5 W for a detector plane consisting of 12 modules. These values are compatible with FASER's existing water cooling system with a coolant temperature of 15 °C.

Detector modules

The readout module consists of an assembly of six ASICs arranged in a 2×3 array. Each module measures around $31 \times 67 \text{ mm}^2$. Figure 10.4 shows an exploded CAD view of a Preshower module. Electrical connection for communication with and powering of the six ASICs is made through a flexible PCB and 134 wire-bonds between the PCB and each ASIC. The module flex features two connectors: one for digital signals, clock and command and the second for power delivery. The power connector provides two Low Voltage (LV)




Figure 10.4: Exploded CAD view of a Preshower module. A module is composed of six ASICs glued to an aluminum base-plate. The module flex with the electrical interconnection and the SMD components is glued on the top of the ASICs. The bottom layer is the thermal interface sheet that is added when integrating a module to the cooling plate of a plane. (Reproduced from Ref. [113], licensed under CC BY 4.0)

lines, V_{ccA} for the analog circuitry of the ASIC and V_{ddD} for its digital circuitry, as well as a High Voltage (HV) line for sensor depletion.

Detector planes

A Preshower detector plane comprises 12 modules arranged in two columns of 6 modules. Each column connects to an Active Patch Panel (APP), forming an electrically independent half-plane. The left half-plane (when looking towards IP1) is called A-side, the right half-plane B-side of a Preshower plane. [Figure 10.5\(a\)](#) shows a fully assembled Preshower plane with all 12 modules mounted on a cooling plate within an aluminum frame. The APP for each half-plane is mounted above in an aluminum support frame and consists of a Patch Panel Board (PPB) and a Logic Board (LB). Connections are made between the modules and the PPB for communication (flat orange cable) and for power (twisted red and white cable). Also visible in [Fig. 10.5\(a\)](#) are a temperature and a humidity sensor mounted top-right on the cooling plate. The volume with the modules is flushed with dry air to avoid condensation during operation.

The PPB contains only passive electronics. It provides connectors for the modules and the LB. Additional connectors are for humidity and temperature sensors, the LV and HV cables from the power supplies, as well as for the PIM board. The PIM board monitors voltages




Figure 10.5: (a) Image of a fully assembled Preshower plane. (Reproduced from Ref. [266], licensed under CC BY 4.0) (b) The upgraded Preshower during installation in TI12. For interlock and monitoring, a crate hosts Preshower Interlock and Monitoring (PIM) boards (see Section 10.4).

and temperatures of the modules and provides interlock signals to the power supplies in case of undesirable operating conditions (see Section 10.4 for more).

The LB hosts a Cyclone 5 FPGA that aggregates data from the 36 ASICs in one half-plane and sends it to the readout board mounted (together with the PIM boards) in the PIM crate. The FPGA is also powered via the PPB. Figure 10.5(b) shows the Preshower with its four planes during installation in the TI12 tunnel. In this figure, the following is highlighted: the position of the detector planes (the scintillators of the original Preshower were kept) and the location of the PIM crate.

The DCS for the new Preshower was developed within this thesis and is presented in Section 10.4. Section 10.2 therefore provides a brief overview of the FASER DCS before the Preshower upgrade.

10.2 The Detector Control System of FASER

The Detector Control System (DCS) of the FASER experiment controls and monitors the parameters of the experiment's power system, monitors the environmental conditions in the detector and the tunnel, monitors the hardware-interlocks on the detector, and implements automatic procedures and alerts to ensure the safe operation of the experiment [6]. The components handled by the DCS can be grouped into the following categories:

- HV and LV detector power systems with PSUs for the tracker and the PMTs
- a Versa Module Eurocard (VME) crate for housing readout electronics of the calorimeter and scintillator detectors
- Power Distribution Units (PDUs) that distribute power to the PSUs of the detector power system, the VME crate and, via an additional small power supply, the readout electronics and trigger system for the tracker
- a cooling system
- a hardware interlock for the tracker (via Tracker Interlock and Monitoring (TIM) boards)

For a detailed explanation of these components, see Ref. [6]. A brief summary of each component shall be given here.

The power supplies for the detector systems are stored in three 19 inch rack mountable crates called the MPOD LV/HV Power Supply System¹. The MPOD crates house LV and HV power supplies for powering the FE chips in the SCT modules and for biasing the sensors of the tracker. An additional HV power supply is needed to power the PMTs for the calorimeter, preshower, timing, and veto scintillator stations. The readout electronics of the calorimeter and scintillator detectors are housed in a 19 inch rack mountable VME crate¹. Several PDUs are used to power the infrastructure and allow for remote control. The PDUs can be interfaced via the MODBUS protocol, while a Wiener OPC-UA server is available for the MPOD crates and VME crate to allow for control via the OPC-UA protocol.

Two air-cooled water chillers are installed in TI12 as the cooling system of FASER. One is always in use to cool the detector at 15 °C, while the other acts as a hot spare ready to be switched into operation should the first one fail. The cooling system is provided by the CERN cryogenic and ventilation group (EN-CV). The EN-CV group allows users to subscribe to monitoring information of the cooling system via a DIP publication (see [Section 8.3.4](#)). If neither chiller is operational, the power supply system is forced to be turned off through a hardware interlock signal acting directly on a full MPOD crate.

¹manufactured by W-IE-NE-R Power Electronics GmbH

An additional interlock system acting channel-wise on the power supplies is provided by the TIM boards. A TIM board monitors temperatures (via NTCs) and humidity levels at a tracker station. It is equipped with an AM335X microcontroller² and a comparator circuit which generates a hardware interlock signal to the LV and HV power supplies if the frame temperature of a tracker station exceeds 25 °C. Additionally, the TIM board digitizes the monitoring data and provides it to the DCS system via an Ethernet connection and the MODBUS protocol. The DCS controls all of the power supplies and centrally monitors the temperature and humidity measurements provided by the TIM boards as well as the voltage and current of the HV and LV power supplies. An automatic software action in the DCS shuts down the tracker in case the temperature measured by the NTCs on the module reaches 31 °C.

As the SCADA system, the LHC-wide DCS software WinCC OA is used (see also [Section 8.3.1](#)). The DCS machines running the software are located above ground in the SR1 rack area and are connected to the hardware in the TI12 tunnel via standard 1 Gbit/s Ethernet in a network secured from the general CERN network. The DCS archives all operational and monitoring data into a CERN-wide ORACLE database. This data is duplicated to an InfluxDB instance and displayed as time series via a Grafana³ interface for shifters and experts. The DCS also implements an FSM to define operational states and transitions between these states and to allow for easy control of the detector. FASER uses the same software components for the implementation of the FSM as ATLAS (see [Section 8.4.4](#)). [Figure 10.6](#) shows the UI, provided by the FSM, which can be used by the FASER run manager or experts to stay informed about the state of the detector and to control it.

The entry point to the FSM as shown in [Fig. 10.6](#) is provided by a supervisory control station (SCS). This station connects to lower-level WinCC OA projects that establish the connection to hardware. The structure of the FASER DCS software components is shown in [Fig. 10.7](#). The original structure included the SCS project and the local control station (LCS) project. This thesis has added two new projects: the WALD system was integrated into the FASER DCS as presented in [Section 10.3](#) and a new DCS was developed and integrated for the Preshower upgrade as presented in [Section 10.4](#).

²manufactured by Texas Instruments Incorporated

³<https://grafana.com/>, an open-source visualization web application

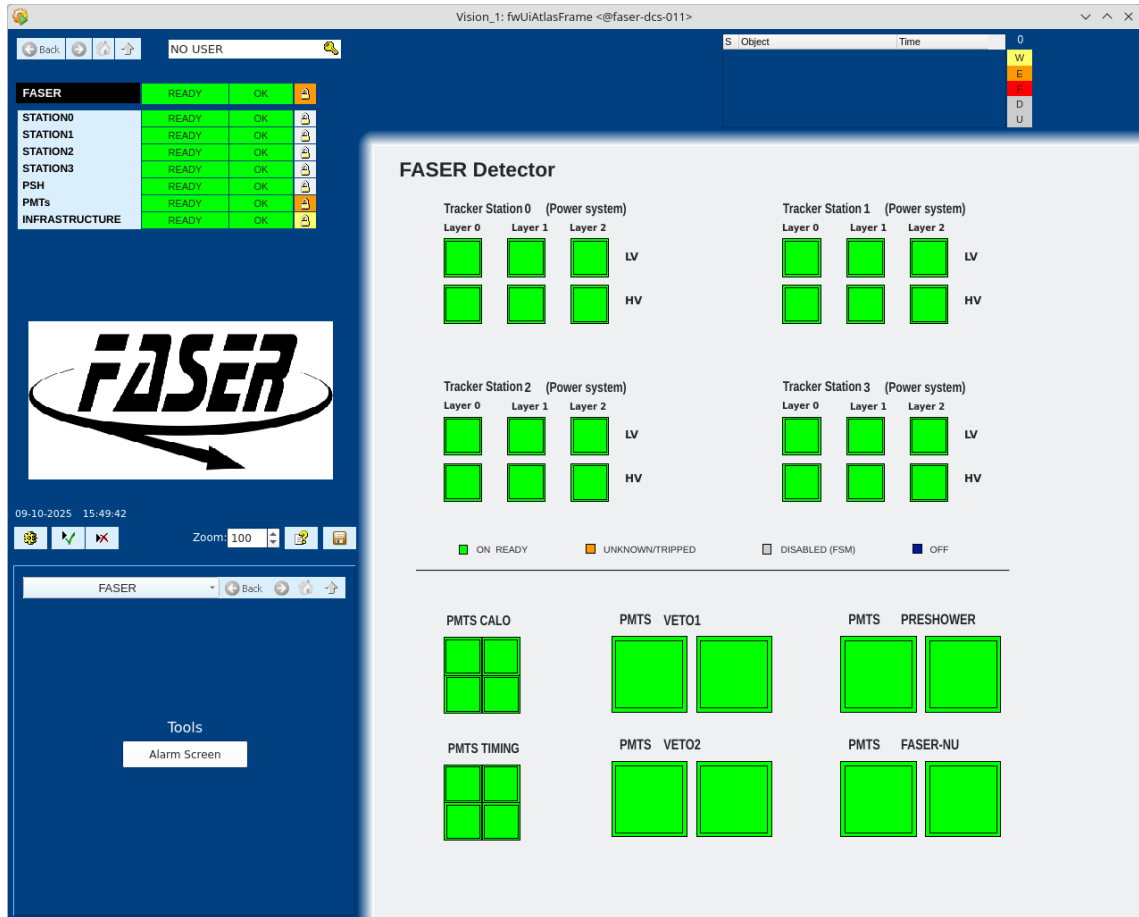


Figure 10.6: The FASER main DCS FSM panel. From here, the whole of FASER can be controlled. The FASER DCS uses the same software components as ATLAS.

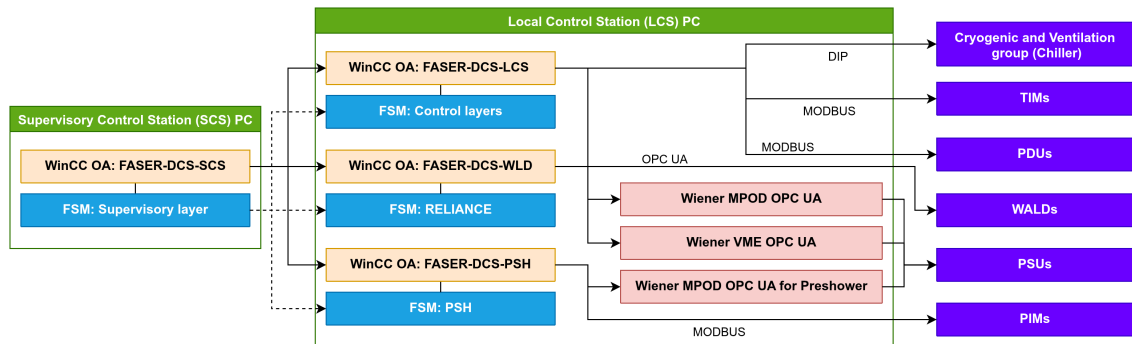


Figure 10.7: Diagram of the software components of FASER's DCS. As high-level entry point (for experts and shifters), a supervisory control station is set up. Lower-level WinCC OA projects run on a separate machine and are connected to the supervisor station in a distributed system. The low-level projects establish the connection to the hardware. The WALD system (FASER-DCS-WLD) and Preshower system (FASER-DCS-PSH) were added with this thesis.

10.3 Commissioning and integration of a Water Leak Detection system

The cooling system of FASER consists of two water chillers installed together with the detector in the TI12 tunnel to provide cooling to the tracking system, as previously discussed. The operational status of the cooling system can be monitored through the provided DIP publication. Additionally, a webcam has been installed in TI12 to visually check the water level in the chillers. However, neither monitoring capability is suited to detect smaller leaks in the cooling system, for instance if a joint starts leaking water and endangers nearby electronics. With the recent verification and installation of a new Water Leak Detection (WALD) system at ATLAS (also called Reliable Liquid Detection for Critical Environments, or RELIANCE) [268], a solution became available for FASER as well. In a collaboration with ATLAS' DSS experts, the WALD system was installed at FASER.

The WALD system is based on Carbon Nano-Tubes (CNTs) materials that act as water sensors. The material used is the Smart Paper [269] developed by the University of Washington. It is a water-sensitive, conductive, paper-like material produced by mixing conductive CNTs as fillers with cellulose nanofibers during the standard paper-making process [268]. This creates a conductive paper with conduction paths through its volume. The conductivity is affected by changes in the fiber dispersion due to external conditions such as temperature, mechanical deformation, and mainly humidity. The key feature of the Smart Paper is that when in contact with water, its conductivity reduces because the water causes the pulp to swell, separating the highly conductive CNTs thereby increasing its resistance. This increase in resistance may reach a factor of 30, making it suitable for water detection [270].

The WALD system thus works by continuously measuring the resistance of the sensor material and triggering an alarm if the resistance suddenly increases. The increase in resistance persists as long as there is water and only returns to normal values after the water naturally evaporates. A key advantage of this paper-like material over the default solution with water leak detection wires is its ability to cover large surfaces.

10.3.1 The RELIANCE box

The monitoring system consists of up to four Smart Papers as sensors connected to a RELIANCE box that houses custom electronics and a Raspberry Pi single board computer. For the resistance measurements, 3.3 V are directly provided by the box. The custom electronics include a conditioning path for a signal using filters and an amplifier before it

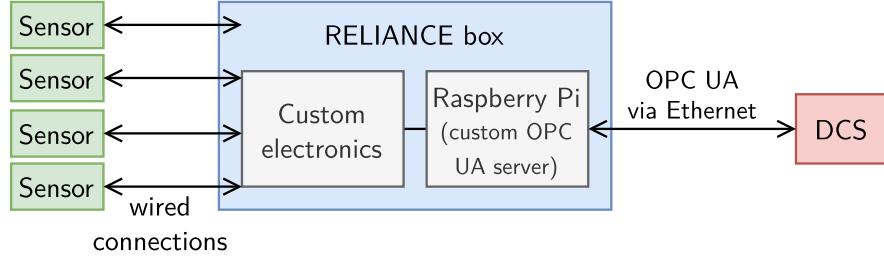


Figure 10.8: Schematic representation of the WALD monitoring system. Up to 4 sensors can be connected to one RELIANCE box. The readout is handled by a Raspberry Pi single board computer, which sends the monitoring data to the DCS.

reaches a commercial 16-bit differential ADC. The ADC is interfaced with the Raspberry Pi, which runs a custom OPC-UA server that performs the readout of the four sensors connected to the box and provides, among others, information about the resistance and water leak alerts to the DCS. A schematic representation of this monitoring system is shown in Fig. 10.8.


In order to detect the presence of water and raise a water leak alert, a signal extraction algorithm is required. It needs to discriminate slowly changing, large fluctuations in the background due to humidity changes and fast noise fluctuations from a water leak signal, that results in a sudden increase in resistance that only slowly returns to normal. The OPC-UA server therefore includes a Chasing Averages (CA) algorithm. In a series of resistance measurements R_1 to R_{C+D} with R_{C+D} being the most recent measurement, the algorithm computes the ratio of the average of the most recent D measurements to the average of the preceding C measurements:

$$\text{CA} = \frac{\overline{R_D}}{\overline{R_C}} = \frac{\frac{1}{D} \sum_{i=C+1}^{C+D} R_i}{\frac{1}{C} \sum_{j=1}^C R_j} \quad (10.1)$$


This number is continuously updated as two sliding windows advance in time with new measurements being made every second. The parameters C and D can be optimized to improve signal extraction; their default values are $C = 900$ and $D = 100$. At constant background, the CA evaluates to 1. Any systematic increase in resistance leads to CA values larger than 1. If the CA exceeds a certain threshold, an alert is raised. All three parameters, C , D and $\text{CA}_{\text{threshold}}$, can be set through the OPC-UA interface by the DCS. At the start of the readout, the algorithm needs to wait for $C + D$ measurements. During this time, the RELIANCE box is in “training” mode. Once enough measurement points are available, it signals its readiness to the DCS.

Table 10.1: Assignments of RELIANCE box channel to WALD sensors.

RELIANCE box channel	Sensor location
WALD1-0	Tracker Station 0 (IFT) (planned)
WALD1-1	Tracker Station 1
WALD1-2	Tracker Station 2
WALD1-3	Tracker Station 3
WALD2-0	Spare
WALD2-1	Spare
WALD2-2	Left chiller
WALD2-3	Right chiller



(a)



(b)

Figure 10.9: Installation of WALD sensors in TI12. (a) shows a Smart Paper installed below a tracking station of the tracking spectrometer. (b) shows one Smart Paper below the piping for each water chiller. All Smart Papers were put on Kapton sheets (in orange).

10.3.2 Installation and calibration of the WALD system

Two RELIANCE boxes were installed in TI12 at the beginning of 2024, providing 8 sensor connections. Five Smart Papers were installed: one paper under each tracking station of the spectrometer and one paper under each water chiller. They were connected to the RELIANCE boxes in the electronics rack in TI12. The space below the IFT was inaccessible at the time of the installation of the other Smart Papers. The assignments of RELIANCE box channels to sensor locations is given in Table 10.1. Figure 10.9 shows photos of the installation of the Smart Paper in TI12.

As mentioned previously, the parameters C and D in the CA algorithm in Eq. (10.1) must be carefully chosen to properly extract spikes in the resistance measurements while

suppressing noise. Additionally, a sensible threshold $CA_{\text{threshold}}$ must be established to detect all water leaks while minimizing the number of false alarms.

The following procedure was used to determine C , D and $CA_{\text{threshold}}$: Resistance measurements for all sensors were recorded over approximately 100 hours. The RELIANCE box provides a measurement point roughly every second. For each sensor, the CA was calculated for various combinations of C and D . A threshold $CA_{\text{threshold}}$ was determined for each combination by calculating the mean (\overline{CA}) and sample standard deviation (σ) of the CA distribution:

$$CA_{\text{threshold}} = \overline{CA} + 15 \cdot \sigma \quad . \quad (10.2)$$

For each sensor, the resulting plots were manually inspected. Regions with extreme noise were excluded from the dataset used to calculate $CA_{\text{threshold}}$ in order to keep sensitivity to real signals. False alarms due to a low alert threshold are more acceptable than missed water leaks. To validate threshold choices, a small water droplet was dropped on every installed Smart Paper, and the algorithm's response was verified.

Figure 10.10 shows the measured resistance of the WALD1-1 sensor over time. Below the resistance plot, CA values for different combinations of C and D are displayed. Horizontal dashed lines indicate the calculated threshold for each combination. Regions, in which the CA exceeds the calculated threshold, have a colored background. Only data before the water droplet test (left of the dashed vertical line) was used for threshold calculation. Figure 10.10 reveals sudden large resistance drops in this channel, the origin of which is unclear. A similar behavior is also visible in the WALD1-3 and WALD2-3 channels. ATLAS reported seeing this behavior when the Smart Paper touches metal [271]. Putting the Smart Paper in FASER on Kapton sheet, however, did not resolve the issue. A resistance drop leads to a CA smaller than 1, which will not cause an alert. However, when the resistance returns to normal values, it will lead to values of CA larger than one that could potentially trigger an alert. In WALD1-3 and WALD2-3, some of these regions have been manually removed for the threshold calculation. As shown in Fig. 10.10, CA curves with $C = 200$ would generate many false alerts when using 15 standard deviations for the threshold determination. Thus, a balance must be found between suppression of noise and reliable leak detection. The 15 standard deviations in Eq. (10.2) were selected (after making a choice for C and D), such that all water droplet tests that were performed resulted in an water leak alert.

The parameters C and D were finally selected after analyzing their influence on signal generation during the water leak test. A close-up of the water droplet test is given in Fig. 10.11, displaying the same data as in Fig. 10.10 but limited to the time of the water




Figure 10.10: Measured resistance of the WALD1-1 sensor over time on the top. Below are the calculated CAs (see Eq. (10.1)) for different values of C and D . The threshold is determined for each combination as the mean plus 15 standard deviations of the data to the left of the black dashed vertical line, and displayed as a horizontal dashed line of the same color. To the right is a test with a water droplet at around 6780 min showing an increase in resistance. Regions where the CA exceeds the threshold are additionally shaded during the time of the alert. A close-up view of the water droplet test is given in Fig. 10.11.

droplet test. It shows, that the water droplet causes an increase in resistance of the Smart Paper. This is registered by the CA algorithm as a spike. The magnitude of the spike primarily depends on C . Larger C values produce larger spikes. Spikes also get larger when D approaches C . However, large values for C delay detection of the water leak. The CA curve exceeds the threshold at a later time. This effect is even more pronounced when D is close to C . In the end, $C = 1000$ and $D = 400$ were chosen to balance between large signals in case of a water leak and minimal delay in reporting the water leak. The calculated thresholds for this combination for each sensor are given in Table 10.2.




Figure 10.11: The same data as in Fig. 10.10 is shown, but limited to the time of the water droplet test to be able to see the effect of different C and D values on the alert generation.

Table 10.2: Calculated thresholds for the CAs of the WALD sensors in the case of $C = 1000$ and $D = 400$.

RELIANCE box channel	Calculated $CA_{\text{threshold}}$
WALD1-1	1.1501
WALD1-2	1.0016
WALD1-3	1.0110
WALD2-2	1.0025
WALD2-3	1.0057

These C , D and $CA_{\text{threshold}}$ values successfully detected leaks in the performed water droplet tests. However, occasional false alarms still occur at a rate of approximately every three to four weeks, which is acceptable for the shifters. When shifters receive an alert, a manual check of the resistance curve allows them to distinguish between a real water leak and a false alarm. A possible design and implementation of a new algorithm with better suppression of noise was beyond the scope of this thesis.

10.3.3 Integration of the WALD system into the FASER DCS

For the integration of the WALD system into FASER's DCS, software components developed for ATLAS have been used. These components include UIs and FSM logic for both individual RELIANCE boxes as well as systems with multiple boxes. A new WinCC OA project was set up on the local control station machine to monitor the WALD system for FASER (see [Fig. 10.7](#)). A RELIANCE FSM tree structure was integrated into the main FASER FSM hierarchy to provide centralized access to UIs. When the CA algorithm detects a water leak, the WinCC OA software sends an alert email to the shifters. This email includes a link to the live monitoring page where operators can manually verify potential false alarms. Shifters must use the FSM UI to reset any triggered alerts. As of this writing, no actual water leaks have occurred at the FASER detector.

10.4 The FASER Preshower DCS

With the installation of the Preshower upgrade (described in [Section 10.1.2](#)) in February 2025, FASER received a new silicon detector, for which a control system became necessary. This system would take the operational model of the new detector design into account, seamlessly integrate into the existing FASER DCS, and provide a simple control interface for shifters without compromising on required control procedures. This section presents the DCS developed for the Preshower upgrade. First, the hardware infrastructure of the new detector system and its corresponding operational model are detailed. Then the FSM control hierarchy tree designed for the upgrade is presented, with a detailed look into the control of a Preshower half-plane. Finally, testing steps and initial experience of the system in operation are described.

10.4.1 The Preshower Interlock and Monitoring board

The Preshower Interlock and Monitoring (PIM) boards are the DCS front-end hardware of the Preshower upgrade, similar to the TIM boards used in the tracker. They control

power to the Preshower modules, monitor their electrical properties and temperatures, and provide a hardware interlock to the power supplies of the Preshower in case of abnormal operating conditions. Each board is equipped with an AM335X microcontroller and a comparator circuit to generate a hardware interlock signal. The board is powered with 24 V by a small power supply that can be controlled through a PDU. One PIM board connects to the A- and B-side of one Preshower plane. Consequently, four PIM boards are installed in TI12.

Figure 10.12 shows the connections between one Preshower plane, a PIM board and the Preshower MPOD crate housing the power supplies. One preshower half-plane is powered by four LV channels transported through one cable to the PPB with their sense lines terminated at the PPB. One channel powers the FPGA on the LB (this channel is called APP). The three other channels each feed a DC/DC converter, which provides power to two module interface (IF) blocks in parallel. An IF block contains LDOs that generate V_{ccA} and V_{ddD} for a detector module. These voltages are then distributed to the modules via the red and white power cable. Using one LV channel per two modules reduces the required channel count, allowing the use of a single MPOD crate in the limited space of the TI12 tunnel. One HV channel is connected to a half-plane, split on the PPB and transported through the IF blocks to each module in parallel. When traversing the IF block, the HV line passes through a Solid State Relay. The PIM board can disabled or enable all LDOs and the Solid State Relay in an IF block to isolate a module from operation (for example in case of failure) or reintegrate it.

The PIM board thus enables/disables single modules, provides monitoring information to the DCS and makes hardware interlock decisions. It monitors the input voltage to the three DC/DC converters, V_{ccA} and V_{ddD} from each IF block to a module, the temperature of each module via an NTC on the module PCB and the HV leakage current of each module via a shunt resistor on the PPB. Additionally, it monitors the temperature on the PPB and, for the B-side, the temperature and humidity in the module frame. For interlock functionality, it connects to the LV power supply through the PPB and to the HV power supply through a HV Interlock Fan-in (HVIF) board. The HV interlock is triggered if the module temperature exceeds a threshold of 30 °C or in case any of the three module LV channels is down. LV interlock occurs only if module temperatures exceed 30 °C and then not immediately, but with a short delay to allow for HV to turn off first. Only module power is cut in this case, while power to the FPGA on the LB is not affected by interlock decisions. The temperature-based interlock decision considers only module 5 in each half-plane.

The PIM digitizes all monitoring data and, along with the state of the module enable signals, provides this data to the DCS via an Ethernet connection and the MODBUS

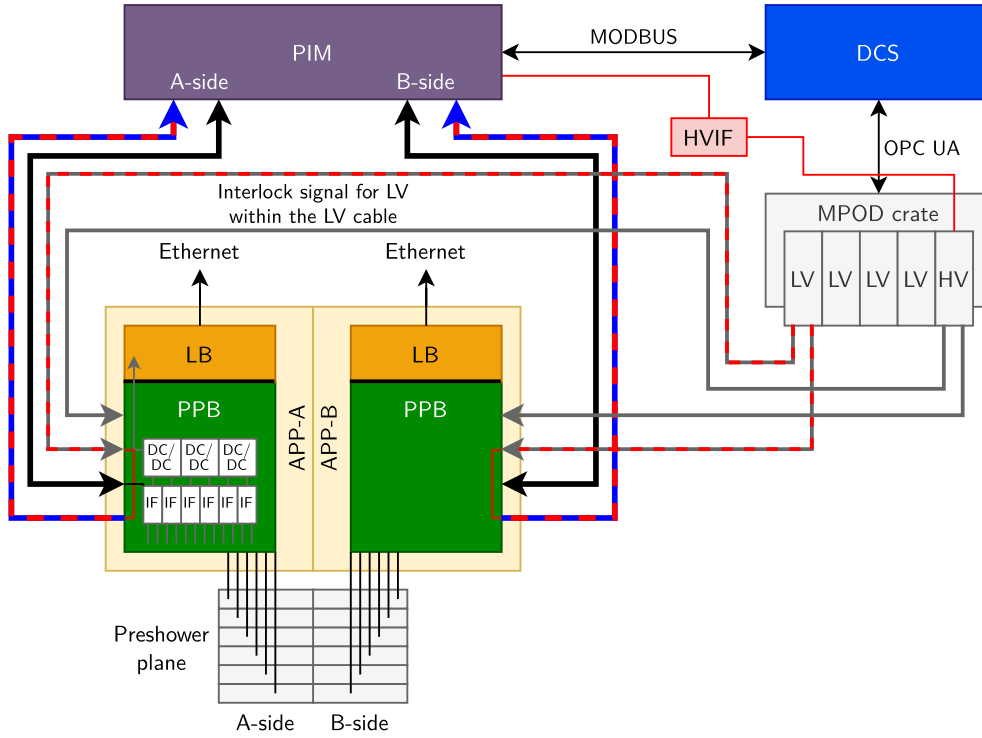


Figure 10.12: Schematic overview of the PIM and power supply connections to a Preshower plane. One PIM is connected to one Preshower plane with two connections per half-plane (A-side and B-side). Details on the PPB are only shown for the A-side. Six modules connect to one PPB.

protocol. To integrate PIM communication into the WinCC OA-based FASER DCS, the framework developed for MARTA cooling (see [Section 9.3.7](#)) was reused. Most PIM MODBUS registers are read-only, except those controlling the module enable signals. The DCS monitors detector components via the PIMs, enables or disables modules and controls power to the modules and LB through the power supplies in an MPOD crate using an OPC-UA interface.

The default setting for the PIM board after a power-cut is that all modules are disabled. Therefore, during power-up, modules must first be enabled.

10.4.2 The FSM control hierarchy of the Preshower

To ensure safe operation of the Preshower and enable supervision of the detector by non-experts, an FSM control hierarchy has been implemented for the Preshower DCS. It acts as a layer between the low-level hardware-oriented components of the DCS and the setup

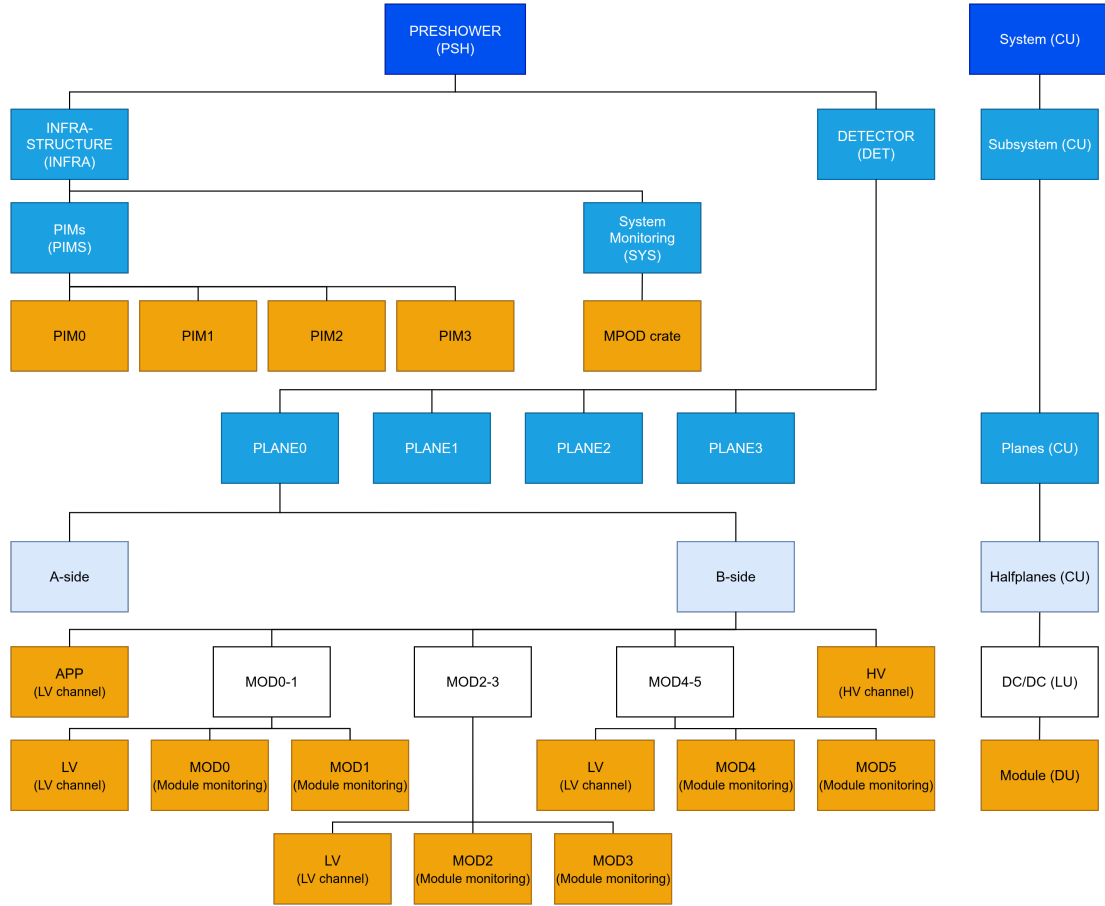


Figure 10.13: The FSM hierarchy tree of CUs, LUs and DUs for the upgraded FASER Preshower. The tree consists of two subtrees, the infrastructure tree and the detector tree. CUs are given with blue background, LUs with white and DUs with an orange background. Due to lack of space only one half-plane CU is shown in full detail. The two modules, powered by one LV channel, are combined into one LU.

operation and supervision. This FSM hierarchy enables control from single Preshower half-planes to the full detector through simple commands.

The FSM control hierarchy is structured as a tree shown in Fig. 10.13. The leaves of the tree are DUs, shown in orange. White nodes represent LUs and blue colored nodes in the tree denote CUs as defined in Section 8.4.3. The tree consists of an infrastructure tree and a detector tree. The infrastructure tree includes the state of the PIM boards and the communication state of the MPOD crate. The states of the detector modules are collected in the detector tree.

The lowest CU in the detector tree is a half-plane. It is the smallest unit that can be controlled independently from other parts of the detector. A detailed description of the half-plane CU is given in the following [Section 10.4.3](#).

The FSM tree for the Preshower is automatically generated (and mirrors) the alias structure in the system similar to the approach used for ITk (see [Section 9.3.1](#)). For example, to safely shut down the detector prior to maintenance, a single command can be issued to the top-level node (PSH). This command propagates down through the hierarchy to each half-plane CU, where the Preshower’s operational model is applied.

10.4.3 The half-plane Control Unit

The half-plane CU implements the operation model of the FASER Preshower. This model prevents dangerous powering configurations that could damage the detector hardware. Specifically, HV must never be enabled when LV is off, as this can destroy the ASICs of a module. This condition is also enforced by the interlock logic of the PIM. Additionally, single modules should only be enabled, if so requested. Following a power cut, malfunctioning modules must remain disabled, and it must not be possible to power on the detector while failing to enable modules marked for operation. Lastly, it is advantageous to provide power to the FPGA on the LB before turning on the modules of a half-plane.

The half-plane CU consists, as shown in [Fig. 10.13](#), of a powering channel for the FPGA on the LB (APP), LUs that group the state of two modules powered by the same LV channel together, and the HV channel for biasing the sensor. For computing the states of the power supply channel DUs (APP, LV, HV), the same logic as in [Section 9.4.2](#) is used. Consequently, the states of a power supply channel DUs are given in [Table 9.2](#), except that the FSM_TEST state is not implemented in the FASER FSM.

The module DU represents the powering state of a module taking the LV and HV channels into account. For the module enable setting, in addition to the actual setting in the PIM board, storage of a target setting is implemented such that the enable setting survives power cycles of the PIM board. The module DU continuously compares target and actual settings and changes the *status* to ERROR in case of inconsistencies. The implemented *states* for the module DU are shown in the module DU state diagram in [Fig. 10.14](#). When no LV and HV are applied and target and actual enable settings agree, the module’s state is either OFF or DISABLED depending on the enable setting. If the target and actual enable settings do not agree, the module’s state is UNSET, and the target setting needs to be sent to the PIM board first. This forces an operator to apply target enable settings first before being able to turn on the detector. The module DU then simply represents the powering state of the module where a distinction is made between an actually enabled

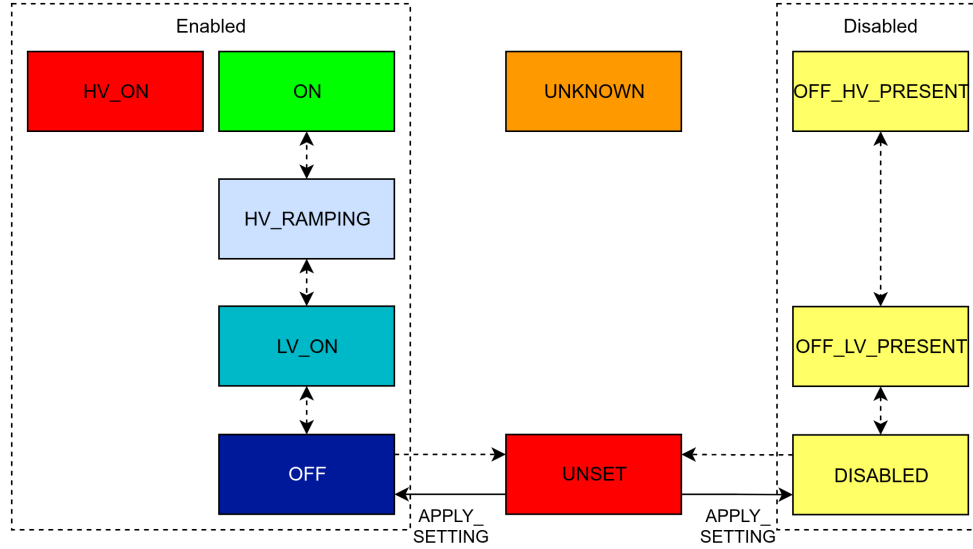


Figure 10.14: The state diagram for the module DU. No control over the powering of the module is possible, instead it passively reflects the powering state of the module. A distinction between an actually enabled and a disabled module is made, since it is possible that the corresponding powering channels are on, but the power to the module is disabled in the IF block on the PPB. This design enforces that changed enable/disable settings are only applied, when no voltage is present on the IF block.

versus a disabled module. The parent LU in the FSM tree combines the state of the one LV channel and the two associated modules. Its state diagram is given in [Appendix D](#) in [Fig. D.1](#).

A state of the half-plane CU summarizes the states of the APP power channel, the states of the LV power channels and their associated modules, and the state of the HV channel. The states available for the half-plane CU are given in the state diagram in [Fig. 10.15](#). The arrows between the states denote transitions. Solid lines indicate transitions initiated by commands with the command name given next to the arrow. Dashed lines indicate automatic transitions. The operation states foreseen during normal operation are within the dashed box. The normal states are: SHUTDOWN (where APP, all LVs and HV off), APP_ON (where APP is on to power the FPGA on the LB), LV_ON (where additionally all three LV channels power the module's ASICs) and ON (where additionally HV depletes the sensors). The UNSET state is automatically taken (before a command can be sent to turn on LV channels) when APP is on and the actual and target module enable settings disagree. This state persists until target settings are applied, reminding operators that correct settings must be restored (for instance after a power-cut) before operation. The implemented procedure also prevents changing module enable states while the modules are powered, which could be a potentially damaging action. The target setting, however, can

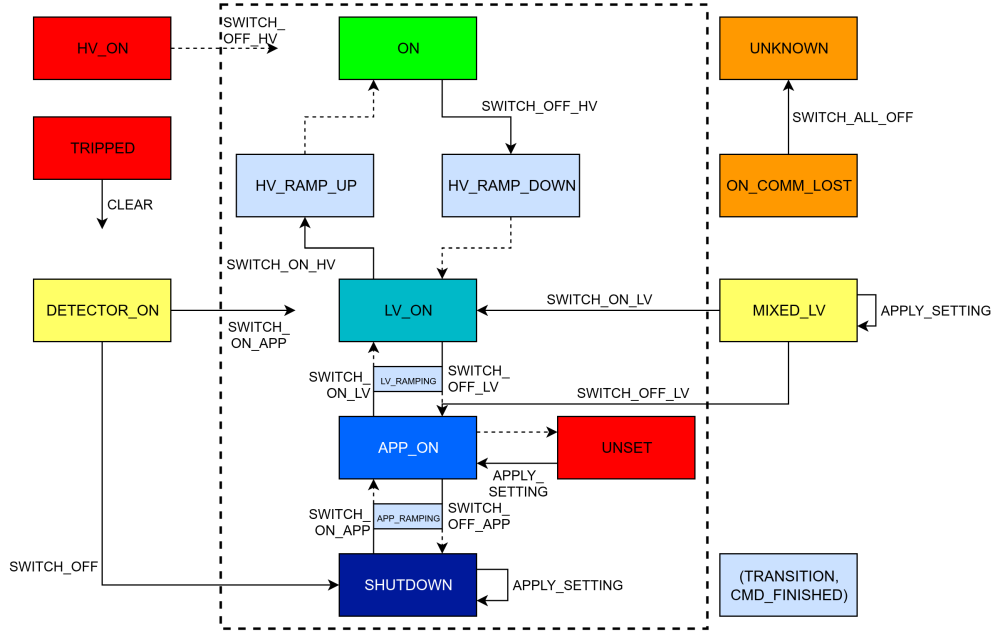


Figure 10.15: The state diagram for the half-plane CU. Normal operation states are inside the dashed box.

be changed at any time, triggering an alert while it differs from the actual setting, and the new setting can only be applied after turning off the power to the corresponding module.

The error states in Fig. 10.15 are HV_ON (where HV is turned on while at least one LV channel is off), TRIPPED (where a power supply channel tripped), DETECTOR_ON (where the LB FPGA is not powered while the modules are turned on), MIXED_LV (where the three LV channels are in different states), ON_COMM_LOST (where communication with the PIM board was lost, but the detector is still powered) and UNKNOWN (where communication with the power supplies was lost). The HV_ON state should be prevented by the interlock logic of the PIM board. However, the FSM implements an automatic action to turn HV off as an additional safety measure.

The normal states were chosen, such that when moving from one state to the next, a command to only one type of child nodes has to be sent. During a power-on sequence, the DCS can turn on the detector completely before the FASER DAQ starts configuring the Preshower modules for data taking. To go from SHUTDOWN to the ON state in the order given by the state diagram in Fig. 10.15, the half-plane CU implements chain commands using FSM object parameters as described in Section 9.4.2. This allows to send a single GOTO_ON command when in the SHUTDOWN state and the FSM takes care of moving through the intermediate APP_ON and LV_ON states to reach the ON state, which greatly simplifies operation for a shifter.

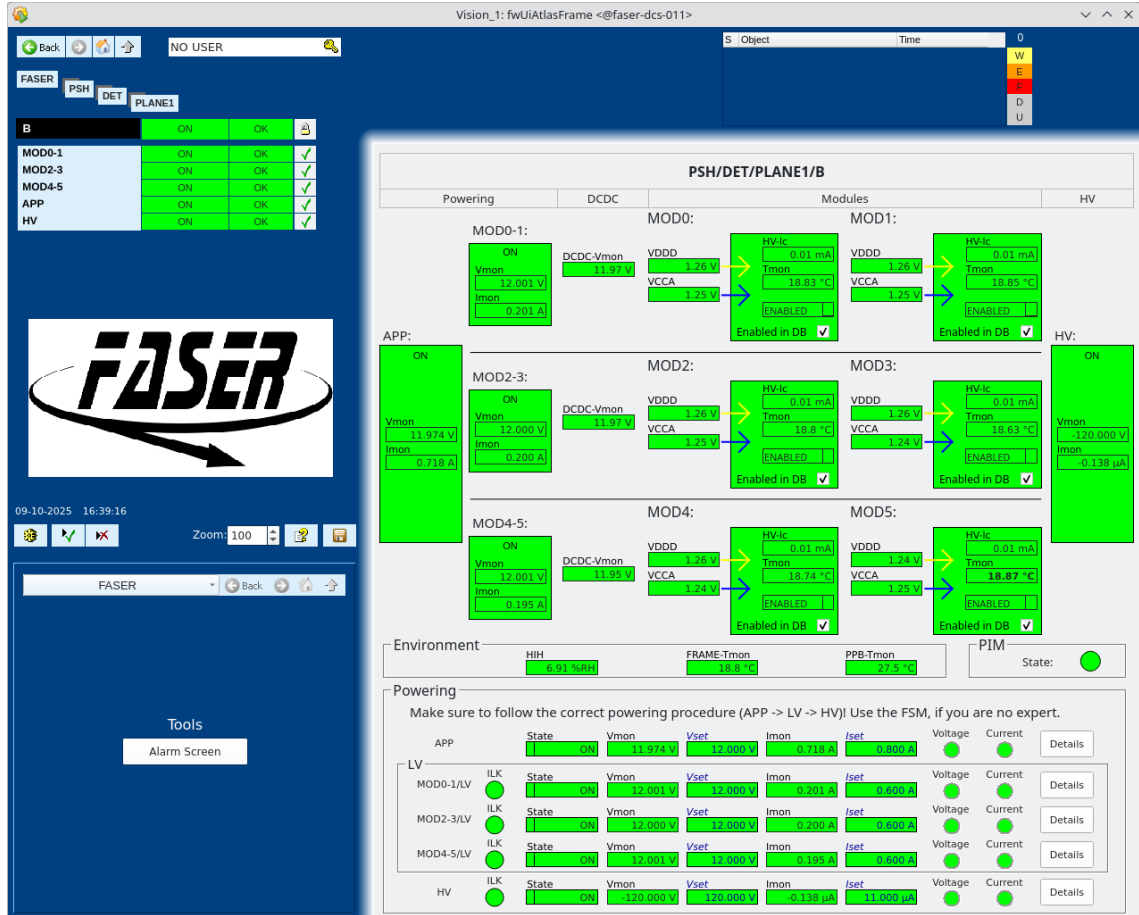


Figure 10.16: GUI for the Preshower half-plane. In the bottom, configurations for the different power channels for a half-plane are shown. The main part displays information about the powering state of the modules and different monitoring values obtained from the PIM board. On the very left is first shown the state of the APP power channel and then the three LV power channel. The PIM board monitors the voltage on the DC/DC converter inputs. For each module, the supplied V_{CCA} and V_{DDD} are shown, as well as a leakage current measurement, temperature measurement and the actual enable setting of a module, all provided by the PIM board. The target module enable setting is also displayed here and can be changed. To the right is the status of the HV channel for sensor biasing.

The GUI for the half-plane shown in Fig. 10.16 is embedded in the general FSM framework for FASER. It displays all monitoring information received from the PIM board for the half-plane as well as the state of the associated powering channels. Via the menu on the left, an operator can issue commands while the FSM ensures correct procedures. The target module enable settings can be changed in the UI. Since disabled modules are not reported as ERROR, a table in the top-level Preshower FSM node lists all disabled modules in the detector for a quick overview. No modules are currently disabled in operation.

10.4.4 Testing, commissioning and initial operational experience

The new Preshower DCS was developed using a test setup comprising an MPOD crate with power supplies for APP, LV and HV, a PPB, a module PCB connected to the PPB and a PIM board. This allowed development of the half-plane control, which was subsequently scaled to the full Preshower detector. The DCS was also used in two test beam campaigns for Preshower planes to gather feedback from operators. Following successful validation, it was integrated into the general FASER DCS in February 2025 for operation of the Preshower in TI12 as shown in [Fig. 10.7](#).

Since commissioning, shifters have successfully operated FASER with the new Preshower. Each week, a designated run manager is in charge of all manual actions affecting the physics run and power state of the detector. The run manager can connect remotely to the supervisory control station (SCS) to monitor or adjust the state of the detector. From there, also control of the Preshower upgrade is possible. The DCS additionally monitors all operational and environmental parameters of the Preshower and archives them into the CERN-wide ORACLE database. This data is published to the Grafana monitoring pages of FASER and presented in weekly operation meetings. [Figure 10.17](#) shows such a Grafana page for the environmental data. The new DCS runs reliably, supporting the operation of the new Preshower.

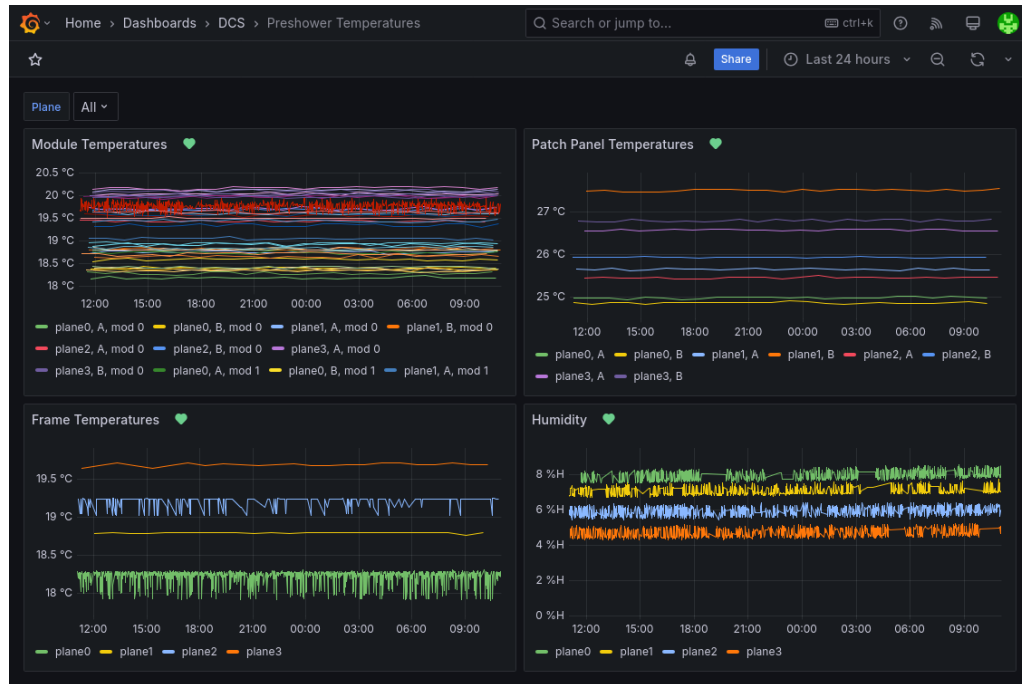


Figure 10.17: Environmental monitoring of the Preshower upgrade detector. Time series of module and environmental temperatures as well as humidity at Preshower planes are visualized for shifters using the Grafana application. The displayed data is received by the DCS from the PIM boards.

Summary and Conclusions

The work described in this thesis contributed significantly to various aspects of the hybrid pixel detector of ATLAS's planned new inner detector for the HL-LHC, the all-silicon ITk. In the new pixel detector, modules are combined into SP-chains and loaded on ring and stave shaped low mass carbon fiber local supports (LLS). The design of these LLSs undergoes a strict verification process with external reviews between subsequent milestones. During the prototyping phase, an Outer Barrel (OB) detector demonstrator with RD53A prototype modules was set up at CERN to evaluate the performance of the modules after loading and thereby verify the design of the LLSs.

The test setup for the demonstrator tests included significant infrastructure to study the detector within a complete detector system close to the final design. The test setup was commissioned and infrastructure components characterized in the context of this thesis.

A half-filled, stave shaped, layer 3 longeron demonstrator with layer 2 electrical services was subsequently tested. In total, 18 RD53A modules were mounted in two SP-chains. Produced $I(V)$ curves showed no sensor performance degradation when compared to $I(V)$ curves of the modules before installation on the longeron. $V(I)$ curves of the SP-chains characterized the ShuntLDOs and confirmed their linear behavior. Electrical scans were performed on all 18 modules and their test results were compared to results from previous stages in the assembly process to find potential performance degradation in a particular step.

Each module's performance was characterized using digital scans, analog scans, threshold scans, Time over Threshold scans and disconnected bump scans. In summary, no systematic

issues were found and no FE showed degradation off its electrical performance. Disconnected bump scans showed no increase in disconnected bumps across all FEs. These results were validated on 7 FEs with a radioactive source scan showing no additional bump delaminations. Additional BERTs and threshold scans evaluated the performance of the longeron under multi-SP-chain operation, showing effectively no difference between one and two SP-chains operated at the same time on the longeron.

The results of these tests demonstrated that the detector units consistently fulfill the detector requirements. The test results underwent external review, which the OB community successfully passed, allowing the project to move to the preproduction phase.

During preproduction, the LLS QC test setups will be qualified. One OB LLS QC setup was built at CERN. Before the first preproduction longeron was tested, the ADCs of the MOPS chips that were used to later monitor temperatures and voltages of the preproduction longeron modules were calibrated. This improved the accuracy of the temperature measurements significantly, especially at cold temperatures.

A preproduction longeron with ITkPix-V1.1 modules was installed at the CERN LLS QC setup. The initial BERTs were performed on all modules of the preproduction longeron showing perfect communication at both warm and cold temperatures. Two Python software packages were developed to support the testing procedures of the LLSs in all QC setups. One package can be used to retrieve DCS data from a database, the other can be used to connect testing tools used in module production to the DCS of the LLS setup allowing for reuse of common tools.

The OB LLS QC setups are complex test stands for testing the smallest fully-functional detector units. A DCS is important for the safe operation and high-level control of a detector unit. In the course of this thesis, a comprehensive DCS was developed for usage by LLS testers to operate the detector units. A SCADA system records all live data relevant to the operation and provides graphical interfaces to access this data. Special attention was paid to designing a system that is easily configurable for different types of LLSs and that enables remote control allowing automation of testing procedures. To this effect, an HTTP REST API manager for the SCADA system was developed and benchmarked. Requirements for operational procedures were identified and the operational model translated into the state model of an FSM. A scaled-up FSM was found to have acceptable performance for potential use in larger detector tests.

In summary, the contributions to the ATLAS ITk Pixel project enabled the OB to move to the next phase for the production of LLSs and laid the groundwork for a control system for the final detector. Only with a detector built from well-performing LLSs and a control

system which ensures its safe and successful operation can the potential of the HL-LHC be fully exploited.

Another part of this thesis was dedicated to contributions to the DCS of the FASER experiment. A new Water Leak Detection system was installed for the detector to alert in case of water leaks in the detector's cooling system and to protect sensitive electronics. The system was commissioned and calibrated. Thresholds for water sensors were found that reduce the amount of false alerts due to noise while guaranteeing the detection of water. The system was integrated into the general FASER DCS.

Additionally, a new DCS was developed for the Preshower upgrade of FASER, a monolithic pixel detector with four layers and tungsten absorbers in between. Also for the Preshower DCS, the requirements for operational procedures were identified and the operational model translated into the state model of an FSM. The Preshower FSM logic, together with graphical UIs, enables shifters to control the Preshower using safe procedures. The new DCS was validated and installed for operation of the new Preshower in the TI12 tunnel. The new DCS runs reliably, supporting the operation of the new Preshower, helping FASER to extend its physics reach and probe new areas in parameter space for new physics.

Acknowledgments

This work would not have been possible without the help and support of many people. As a CERN doctoral student, I consider myself extremely lucky to have two incredible advisors to whom I owe a great debt of gratitude. To Stan Lai, my professor in Göttingen, thank you very much for letting me be part of your group. Thanks to your organizational talent, I always felt in good hands when it came to the not-so-interesting parts (bureaucracy), and thanks to our many research chats, I equally felt in good hands when it came to the interesting parts (research and keeping a connection to Göttingen). During your stay at CERN, I experienced the advantages of being able to quickly chat with you and got to know good restaurants in Geneva. To Benedikt Vormwald, I am so lucky to have you as my CERN supervisor. I should feel bad about asking you so many questions, but you always took your time to answer them allowing me to learn so much during my studies. You guided me through my work, asking *me* the right questions whenever I came up with ideas. These were the talks I enjoyed the most. You both supported me when I applied for the CERN doctoral program. Thanks to the Wolfgang Gentner Program of the German Federal Ministry of Education and Research (grant no. 13E18CHA), I could conduct my studies in the EP-ADE-TK section at CERN. Thank you to Brian and Leyre, who made the start of my studies so enjoyable. Thank you to Jenny, Anna and Sebastian, for the best office and cakes and the most memorable adventures. Anastasia, I hope your dreams come true. I appreciate your feedback and persistent push to do things right.

I received much help from people outside of my section. Many thanks to Krista Morassutti for helping me with formalities and to Jörn Große-Knetter for being part of my thesis committee. I learned a lot from Kerstin Lantzsch about control systems of large detectors. Thank you very much for being available to answer my questions and introducing me to this immense scale of ATLAS. Also many thanks to Matthias Hamer. I really appreciate your help when learning about serial powering. To Simon Florian Koch, it was always great working with you and discussing ideas. Thank you very much.

Acknowledgments

A big Thank You to the colleagues at FASER Jamie Boyd, Brian Petersen, Hidetoshi Otono, Tomohira Inada and Stefano Zambito for the trusting cooperation, for always being available and showing me many different aspects of installing and/or operating a detector. I will not forget the walk along the LHC to get to FASER.

I would like to thank my family for their support over the years. Meine lieben Brüder, danke für die „lehrreichen“ Unterhaltungen. Mama und Papa, euch habe ich zu verdanken, dass ich meinen Weg gehen konnte. Dear Dasha, thank you very much for sharing the last few years with me. Thank you for being a caring girlfriend and my best friend. I think, it is time to get a dog together.

Bibliography

- [1] ATLAS Collaboration, *Observation of a New Particle in the Search for the Standard Model Higgs Boson with the ATLAS Detector at the LHC*, [Phys. Lett. B **716** \(2012\) 1](#).
- [2] CMS Collaboration, *Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC*, [Phys. Lett. B **716** \(2012\) 30](#).
- [3] O. Aberle et al., *High-Luminosity Large Hadron Collider (HL-LHC): Technical Design Report*, [CERN Yellow Rep.: Monogr. **CERN-2020-010** \(2020\)](#).
- [4] ATLAS Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, [J. Instrum. **3** \(2008\) S08003](#).
- [5] ATLAS Collaboration, *ATLAS Phase-II Upgrade Scoping Document*, CERN-LHCC-2015-020, 2015, CDS: [2055248](#).
- [6] FASER Collaboration, *The FASER Detector*, [J. Instrum. **19** \(2024\) P05066](#).
- [7] M. K. Gaillard, P. D. Grannis, and F. J. Sciulli, *The Standard Model of Particle Physics*, [Rev. Mod. Phys. **71** \(1999\) S96](#).
- [8] L. Evans and P. Bryant, *LHC Machine*, [J. Instrum. **3** \(2008\) S08001](#).
- [9] A. Einstein, *Die Grundlage der allgemeinen Relativitätstheorie*, [Ann. Phys. \(Berl.\) **354** \(1916\) 769](#).
- [10] A. Zee, *Einstein Gravity in a Nutshell*, Princeton University Press, 2013, ISBN: 978-0-691-14558-7.
- [11] R. Mann, *An Introduction to Particle Physics and the Standard Model*, [CRC Press](#), 2010, ISBN: 978-0-429-14122-5.
- [12] A. Pich, *The Standard Model of Electroweak Interactions*, 2012, arXiv: [1201.0537 \[hep-ph\]](#).

- [13] M. Thomson, *Modern Particle Physics*, [Cambridge University Press](#), 2013, ISBN: 978-1-107-03426-6.
- [14] M. D. Schwartz, *Quantum Field Theory and the Standard Model*, [Cambridge University Press](#), 2014, ISBN: 978-1-107-03473-0.
- [15] S. P. Martin and J. D. Wells, *Elementary Particles and Their Interactions*, [Springer](#), 2022, ISBN: 978-3-031-14368-7.
- [16] W. N. Cottingham and D. A. Greenwood, *An Introduction to the Standard Model of Particle Physics*, [Cambridge University Press](#), 2023, ISBN: 978-1-009-40168-5.
- [17] H. Fritzsch, M. Gell-Mann, and H. Leutwyler, *Advantages of the Color Octet Gluon Picture*, [Phys. Lett. B](#) **47** (1973) 365.
- [18] H. D. Politzer, *Reliable Perturbative Results for Strong Interactions?*, [Phys. Rev. Lett.](#) **30** (1973) 1346.
- [19] D. J. Gross and F. Wilczek, *Ultraviolet Behavior of Non-Abelian Gauge Theories*, [Phys. Rev. Lett.](#) **30** (1973) 1343.
- [20] S. L. Glashow, *Partial-Symmetries of Weak Interactions*, [Nucl. Phys.](#) **22** (1961) 579.
- [21] S. Weinberg, *A Model of Leptons*, [Phys. Rev. Lett.](#) **19** (1967) 1264.
- [22] A. Salam, *Weak and Electromagnetic Interactions*, [Conf. Proc. C](#) **680519** (1968) 367.
- [23] R. P. Feynman, *Space-Time Approach to Quantum Electrodynamics*, [Phys. Rev.](#) **76** (1949) 769.
- [24] P. A. Zyla et al. (Particle Data Group), *Review of Particle Physics*, [Prog. Theor. Exp. Phys.](#) **2020** (2020) 083C01.
- [25] G. 't Hooft, *Renormalization of Massless Yang-Mills Fields*, [Nucl. Phys. B](#) **33** (1971) 173.
- [26] G. 't Hooft, *Renormalizable Lagrangians for Massive Yang-Mills Fields*, [Nucl. Phys. B](#) **35** (1971) 167.
- [27] P. W. Higgs, *Broken Symmetries and the Masses of Gauge Bosons*, [Phys. Rev. Lett.](#) **13** (1964) 508.
- [28] F. Englert and R. Brout, *Broken Symmetry and the Mass of Gauge Vector Mesons*, [Phys. Rev. Lett.](#) **13** (1964) 321.
- [29] P. W. Higgs, *Spontaneous Symmetry Breakdown without Massless Bosons*, [Phys. Rev.](#) **145** (1966) 1156.
- [30] Super-Kamiokande Collaboration, *Evidence for Oscillation of Atmospheric Neutrinos*, [Phys. Rev. Lett.](#) **81** (1998) 1562.

- [31] Planck Collaboration, *Planck 2018 Results: VI. Cosmological Parameters*, [Astron. Astrophys.](#) **641** (2020) A6.
- [32] M. B. Gavela, P. Hernández, J. Orloff, and O. Pène, *Standard Model CP-Violation and Baryon Asymmetry*, [Mod. Phys. Lett. A](#) **09** (1994) 795.
- [33] R. D. Peccei and H. R. Quinn, *Constraints Imposed by CP Conservation in the Presence of Pseudoparticles*, [Phys. Rev. D](#) **16** (1977) 1791.
- [34] S. Myers and E. Picasso, *The Design, Construction and Commissioning of the CERN Large Electron–Positron Collider*, [Contemp. Phys.](#) **31** (1990) 387.
- [35] B. Povh, K. Rith, C. Scholz, F. Zetsche, and W. Rodejohann, *Teilchen und Kerne: Eine Einführung in die physikalischen Konzepte*, [Springer](#), 2014, ISBN: 978-3-642-37821-8.
- [36] F. Landua, *The CERN Accelerator Complex Layout in 2022*, 2022, CDS: [2813716](#).
- [37] J. Vollaie et al., *Linac4 Design Report*, [CERN Yellow Rep.: Monogr.](#) **6** (2020).
- [38] CMS Collaboration, *The CMS Experiment at the CERN LHC*, [J. Instrum.](#) **3** (2008) S08004.
- [39] LHCb Collaboration, *The LHCb Detector at the LHC*, [J. Instrum.](#) **3** (2008) S08005.
- [40] ALICE Collaboration, *The ALICE Experiment at the CERN LHC*, [J. Instrum.](#) **3** (2008) S08002.
- [41] S. Myers and H. Schopper, eds., *Particle Physics Reference Library: Volume 3: Accelerators and Colliders*, [Springer](#), 2020, ISBN: 978-3-030-34244-9.
- [42] O. Brüning and L. Rossi, *The High Luminosity Large Hadron Collider*, Advanced Series on Directions in High Energy Physics, [World Scientific](#), 2015, ISBN: 978-981-4675-46-8.
- [43] *LHC / HL-LHC Plan*, URL: https://hilumilhc.web.cern.ch/sites/default/files/2025-01/HL-LHC_Plan_January2025.pdf (visited on 05/07/2025).
- [44] R. Garoby et al., “Upgrade Plans for the LHC Injector Complex”, [Proc. IPAC’12](#), TUXA02, New Orleans, LA, USA, 2012, 1010.
- [45] ATLAS Collaboration, *Combination of Searches for Higgs Boson Pair Production in pp Collisions at $\sqrt{s} = 13$ TeV with the ATLAS Detector*, [Phys. Rev. Lett.](#) **133** (2024) 101801.
- [46] ATLAS Collaboration, *Projected sensitivity of Higgs boson pair production combining the $b\bar{b}\gamma\gamma$ and $b\bar{b}\tau^+\tau^-$ final states with the ATLAS detector at the HL-LHC*, ATL-PHYS-PUB-2022-005, CERN, 2022, CDS: [2802127](#).
- [47] A. Dainese et al., *Report on the Physics at the HL-LHC, and Perspectives for the HE-LHC*, CERN-2019-007, CERN, 2019, CDS: [2703572](#).

- [48] ATLAS Collaboration, *Snowmass White Paper Contribution: Physics with the Phase-2 ATLAS and CMS Detectors*, ATL-PHYS-PUB-2022-018, 2022, CDS: [2805993](#).
- [49] S. Navas et al. (Particle Data Group), *Review of Particle Physics*, [Phys. Rev. D **110** \(2024\) 030001](#).
- [50] S. Mehlhase (ATLAS Collaboration), *ATLAS Detector Slice (and Particle Visualisations)*, 2021, CDS: [2770815](#).
- [51] B. Andersson, G. Gustafson, G. Ingelman, and T. Sjöstrand, *Parton Fragmentation and String Dynamics*, [Phys. Rep. **97** \(1983\) 31](#).
- [52] ATLAS Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider: A Description of the Detector Configuration for Run 3*, [J. Instrum. **19** \(2024\) P05063](#).
- [53] R. M. Bianchi, *ATLAS Experiment Schematic or Layout Illustration*, 2022, CDS: [2837191](#).
- [54] ATLAS Collaboration, *ATLAS Inner Detector: Technical Design Report, Volume 1*, tech. rep. CERN-LHCC-97-016, 1997, CDS: [331063](#).
- [55] ATLAS Collaboration, *ATLAS Pixel Detector Electronics and Sensors*, [J. Instrum. **3** \(2008\) P07007](#).
- [56] ATLAS Collaboration, *Operation and Performance of the ATLAS Semiconductor Tracker*, [J. Instrum. **9** \(2014\) P08009](#).
- [57] ATLAS TRT Collaboration, *The ATLAS Transition Radiation Tracker (TRT) Proportional Drift Tube: Design and Performance*, [J. Instrum. **3** \(2008\) P02013](#).
- [58] I. Perić et al., *The FEI3 Readout Chip for the ATLAS Pixel Detector*, [Nucl. Instrum. Methods Phys. Res. A **565** \(2006\) 178](#).
- [59] D. Attree et al., *The Evaporative Cooling System for the ATLAS Inner Detector*, [J. Instrum. **3** \(2008\) P07003](#).
- [60] ATLAS Collaboration, *Experiment Briefing: Keeping the ATLAS Inner Detector in Perfect Alignment*, 2020, CDS: [2723878](#).
- [61] ATLAS Collaboration, *ATLAS Insertable B-layer: Technical Design Report*, tech. rep. CERN-LHCC-2010-013, 2010, CDS: [1291633](#).
- [62] ATLAS Collaboration, *Production and Integration of the ATLAS Insertable B-layer*, [J. Instrum. **13** \(2018\) T05008](#).
- [63] M. Garcia-Sciveres et al., *The FE-I4 Pixel Readout Integrated Circuit*, [Nucl. Instrum. Methods Phys. Res. A **636** \(2011\) S155](#).
- [64] A. Ahmad et al., *The Silicon Microstrip Sensors of the ATLAS Semiconductor Tracker*, [Nucl. Instrum. Methods Phys. Res. A **578** \(2007\) 98](#).

- [65] ATLAS Collaboration, *Operation and Performance of the ATLAS Semiconductor Tracker in LHC Run 2*, *J. Instrum.* **17** (2022) P01013.
- [66] ATLAS Collaboration, *Performance of the ATLAS Transition Radiation Tracker in Run 1 of the LHC: Tracker Properties*, *J. Instrum.* **12** (2017) P05002.
- [67] ATLAS Collaboration, *The ATLAS Inner Detector Commissioning and Calibration*, *Eur. Phys. J. C* **70** (2010) 787.
- [68] ATLAS Collaboration, *ATLAS Liquid Argon Calorimeter: Technical Design Report*, tech. rep. CERN-LHCC-96-041, 1996, CDS: [331061](#).
- [69] ATLAS Collaboration, *ATLAS Calorimeter Performance: Technical Design Report*, tech. rep. CERN-LHCC-96-040, 1996, CDS: [331059](#).
- [70] ATLAS Collaboration, *Operation and Performance of the ATLAS Tile Calorimeter in Run 1*, *Eur. Phys. J. C* **78** (2018) 987.
- [71] ATLAS Collaboration, *ATLAS Muon Spectrometer: Technical Design Report*, tech. rep. CERN-LHCC-97-022, 1997, CDS: [331068](#).
- [72] ATLAS Collaboration, *ATLAS New Small Wheel: Technical Design Report*, tech. rep. CERN-LHCC-2013-006, 2013, CDS: [1552862](#).
- [73] ATLAS Collaboration, *The ATLAS Trigger System for LHC Run 3 and Trigger Performance in 2022*, *J. Instrum.* **19** (2024) P06029.
- [74] ATLAS Collaboration, *Athena*, version 21.0.127, Zenodo, 2021, URL: <https://zenodo.org/record/4772550>.
- [75] J. Anderson et al., *FELIX: A High-Throughput Network Approach for Interfacing to Front End Electronics for ATLAS Upgrades*, *J. Phys.: Conf. Ser.* **664** (2015) 082050.
- [76] G. Avoni et al., *The New LUCID-2 Detector for Luminosity Measurement and Monitoring in ATLAS*, *J. Instrum.* **13** (2018) P07017.
- [77] ATLAS Collaboration, *Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC*, *Eur. Phys. J. C* **83** (2023) 982.
- [78] ATLAS Collaboration, *Letter of Intent for the Phase-I Upgrade of the ATLAS Experiment*, tech. rep. CERN-LHCC-2011-012, 2011, CDS: [1402470](#).
- [79] ATLAS Collaboration, *Letter of Intent for the Phase-II Upgrade of the ATLAS Experiment*, tech. rep. CERN-LHCC-2012-022, 2012, CDS: [1502664](#).
- [80] ATLAS Collaboration, *Technical Design Report for the Phase-II Upgrade of the ATLAS Trigger and Data Acquisition System - Event Filter Tracking Amendment*, tech. rep. CERN-LHCC-2022-004, 2022, CDS: [2802799](#).
- [81] ATLAS Collaboration, *ATLAS TDAQ Phase-II Upgrade: Technical Design Report*, tech. rep. CERN-LHCC-2017-020, 2017, CDS: [2285584](#).

- [82] ATLAS Collaboration, *ATLAS Liquid Argon Calorimeter Phase-II Upgrade: Technical Design Report*, tech. rep. CERN-LHCC-2017-018, 2017, CDS: [2285582](#).
- [83] ATLAS Collaboration, *Technical Design Report for the Phase-II Upgrade of the ATLAS Tile Calorimeter*, tech. rep. CERN-LHCC-2017-019, 2017, CDS: [2285583](#).
- [84] ATLAS Collaboration, *Technical Design Report for the Phase-II Upgrade of the ATLAS Muon Spectrometer*, tech. rep. CERN-LHCC-2017-017, 2017, CDS: [2285580](#).
- [85] ATLAS Collaboration, *Technical Design Report: A High-Granularity Timing Detector for the ATLAS Phase-II Upgrade*, tech. rep. CERN-LHCC-2020-007, 2020, CDS: [2719855](#).
- [86] G. Pellegrini et al., *Technology Developments and First Measurements of Low Gain Avalanche Detectors (LGAD) for High Energy Physics Applications*, [Nucl. Instrum. Methods Phys. Res. A](#) **765** (2014) 12.
- [87] ATLAS Collaboration, *The LUCID 3 Detector for the ATLAS Phase-II Upgrade*, CERN-LHCC-2021-016, CERN, 2021, CDS: [2780604](#).
- [88] B. Giacobbe (ATLAS LUCID Group), *LUCID-3: The Upgrade of the ATLAS Luminosity Detector for High-Luminosity LHC*, [J. Instrum.](#) **19** (2024) C03053.
- [89] ATLAS Collaboration, *Modelling Radiation Damage to Pixel Sensors in the ATLAS Detector*, [J. Instrum.](#) **14** (2019) P06012.
- [90] ATLAS Collaboration, *Measurements of Sensor Radiation Damage in the ATLAS Inner Detector Using Leakage Currents*, [J. Instrum.](#) **16** (2021) P08025.
- [91] ATLAS Collaboration, *Technical Design Report for the ATLAS Inner Tracker Pixel Detector*, tech. rep. CERN-LHCC-2017-021, 2017, CDS: [2285585](#).
- [92] ATLAS Collaboration, *Technical Design Report for the ATLAS Inner Tracker Strip Detector*, tech. rep. CERN-LHCC-2017-005, CERN, 2017, CDS: [2257755](#).
- [93] ATLAS Collaboration, *Expected Tracking Performance of the ATLAS Inner Tracker at the High-Luminosity LHC*, [J. Instrum.](#) **20** (2025) P02018.
- [94] H. Bethe, *Zur Theorie des Durchgangs schneller Korpuskularstrahlen durch Materie*, [Ann. Phys. \(Berl.\)](#) **397** (1930) 325.
- [95] H. Bethe, *Bremsformel für Elektronen relativistischer Geschwindigkeit*, [Z. Phys.](#) **76** (1932) 293.
- [96] F. Bloch, *Bremsvermögen von Atomen mit mehreren Elektronen*, [Z. Phys.](#) **81** (1933) 363.
- [97] L. D. Landau, “On the Energy Loss of Fast Particles by Ionisation”, *Collected Papers of L.D. Landau*, Elsevier, 1965, 417, ISBN: 978-0-08-010586-4.
- [98] H. Bichsel, *Straggling in Thin Silicon Detectors*, [Rev. Mod. Phys.](#) **60** (1988) 663.

- [99] F. Hartmann, *Evolution of Silicon Sensor Technology in Particle Physics: Basics and Applications*, [Springer](#), 2024, ISBN: 978-3-031-59719-0.
- [100] H. Bethe and W. Heitler, *On the Stopping of Fast Particles and on the Creation of Positive Electrons*, [Proc. R. Soc. Lond. A](#) **146** (1934) 83.
- [101] Y.-S. Tsai, *Pair Production and Bremsstrahlung of Charged Leptons*, [Rev. Mod. Phys.](#) **46** (1974) 815.
- [102] H. Kolanoski, *Particle Detectors: Fundamentals and Applications*, [Oxford University Press](#), 2020, ISBN: 978-0-19-885836-2.
- [103] H. A. Bethe, *Molière’s Theory of Multiple Scattering*, [Phys. Rev.](#) **89** (1953) 1256.
- [104] G. R. Lynch and O. I. Dahl, *Approximations to Multiple Coulomb Scattering*, [Nucl. Instrum. Methods Phys. Res. B](#) **58** (1991) 6.
- [105] S. F. Koch et al., *Measuring the ATLAS ITk Pixel Detector Material via Multiple Scattering of Positrons at the CERN PS*, [Eur. Phys. J. C](#) **85** (2025) 381.
- [106] M. Turala, *Silicon Tracking Detectors—Historical Overview*, [Nucl. Instrum. Methods Phys. Res. A](#) **541** (2005) 1.
- [107] S. Seidel, *Silicon Strip and Pixel Detectors for Particle Physics Experiments*, [Phys. Rep.](#) **828** (2019) 1.
- [108] R. C. Alig and S. Bloom, *Electron-Hole-Pair Creation Energies in Semiconductors*, [Phys. Rev. Lett.](#) **35** (1975) 1522.
- [109] P. Sellin and J. Vaitkus, *New Materials for Radiation Hard Semiconductor Detectors*, [Nucl. Instrum. Methods Phys. Res. A](#) **557** (2006) 479.
- [110] M. De Napoli, *SiC Detectors: A Review on the Use of Silicon Carbide as Radiation Detection Material*, [Front. Phys.](#) **10** (2022).
- [111] L. Rossi, P. Fischer, T. Rohe, and N. Wermes, *Pixel Detectors: From Fundamentals to Applications*, [Springer](#), 2006, ISBN: 978-3-540-28332-4.
- [112] P. P. Altermatt, A. Schenk, F. Geelhaar, and G. Heiser, *Reassessment of the Intrinsic Carrier Density in Crystalline Silicon in View of Band-Gap Narrowing*, [J. Appl. Phys.](#) **93** (2003) 1598.
- [113] FASER Collaboration, *The FASER W-Si High Precision Preshower Technical Proposal*, CERN-LHCC-2022-006, 2022, CDS: [2803084](#).
- [114] S. Sze and K. K. Ng, *Physics of Semiconductor Devices*, [Wiley](#), 2006, ISBN: 978-0-471-14323-9.
- [115] S. Parker, C. Kenney, and J. Segal, *3D — A Proposed New Architecture for Solid-State Radiation Detectors*, [Nucl. Instrum. Methods Phys. Res. A](#) **395** (1997) 328.

- [116] C. Da Via et al., *3D Silicon Sensors: Design, Large Area Production and Quality Assurance for the ATLAS IBL Pixel Detector Upgrade*, [Nucl. Instrum. Methods Phys. Res. A](#) **694** (2012) 321.
- [117] W. Shockley, *The Theory of P-n Junctions in Semiconductors and p-n Junction Transistors*, [Bell Syst. Tech. J.](#) **28** (1949) 435.
- [118] M. Moll, *Radiation Damage in Silicon Particle Detectors: Microscopic Defects and Macroscopic Properties*, PhD thesis: Universität Hamburg, 1999.
- [119] A. Chilingarov, *Generation Current Temperature Scaling*, tech. rep. PH-EP-Tech-Note-2013-001, 2013, CDS: [1511886](#).
- [120] A. Chilingarov, *Temperature Dependence of the Current Generated in Si Bulk*, [J. Instrum.](#) **8** (2013) P10003.
- [121] C. Leroy and P.-G. Rancoita, *Silicon Solid State Devices and Radiation Detection*, [World Scientific](#), 2012, ISBN: 978-981-4390-04-0.
- [122] P. Drude, *Zur Elektronentheorie der Metalle*, [Ann. Phys. \(Berl.\)](#) **306** (1900) 566.
- [123] D. Caughey and R. Thomas, *Carrier Mobilities in Silicon Empirically Related to Doping and Field*, [Proc. IEEE](#) **55** (1967) 2192.
- [124] C. Jacoboni, C. Canali, G. Ottaviani, and A. Alberigi Quaranta, *A Review of Some Charge Transport Properties of Silicon*, [Solid-State Electron.](#) **20** (1977) 77.
- [125] W. Shockley, *Currents to Conductors Induced by a Moving Point Charge*, [J. Appl. Phys.](#) **9** (1938) 635.
- [126] S. Ramo, *Currents Induced by Electron Motion*, [Proc. IRE](#) **27** (1939) 584.
- [127] *Radiation Effects in the LHC Experiments: Impact on Detector Performance and Operation*, [CERN Yellow Rep.: Monogr. CERN-2021-001](#) (2021), ed. by I. Dawson.
- [128] ATLAS Collaboration, *Sensor Response and Radiation Damage Effects for 3D Pixels in the ATLAS IBL Detector*, [J. Instrum.](#) **19** (2024) P10008.
- [129] G. Lindström et al., *Radiation Hard Silicon Detectors—Developments by the RD48 (ROSE) Collaboration*, [Nucl. Instrum. Methods Phys. Res. A](#) **466** (2001) 308.
- [130] J. D. Cressler and H. A. Mantooth, eds., *Extreme Environment Electronics*, CRC Press, 2017, ISBN: 978-1-4398-7430-1.
- [131] M. Boscardin et al., *Advances in 3D Sensor Technology by Using Stepper Lithography*, [Front. Phys.](#) **8** (2021) 625275.
- [132] S. Grinstein, *Technical Specification and Acceptance Criteria for 3D Sensors for the ATLAS Pixel Tracker Upgrade*, tech. rep. AT2-IP-EP-0002, ATLAS Collaboration, 2018, EDMS: [1817524 v.1](#) (internal documentation).

- [133] F. Hügging, *Technical Specification and Acceptance Criteria for the Bump Bonding of the ITk Pixel Modules*, tech. rep. AT2-IP-EP-0003, ATLAS Collaboration, 2020, EDMS: [1817526 v.3](#) (internal documentation).
- [134] A. Heggelund et al., *Radiation Hard 3D Silicon Pixel Sensors for Use in the ATLAS Detector at the HL-LHC*, [J. Instrum.](#) **17** (2022) P08003.
- [135] A. Macchiolo, *Technical Specification and Acceptance Criteria for the Planar Pixel Sensors for the ITk Project*, tech. rep. AT2-IP-EP-0006, ATLAS Collaboration, 2018, EDMS: [1817538 v.1](#) (internal documentation).
- [136] W. Snoeys, *Monolithic CMOS Sensors for High Energy Physics — Challenges and Perspectives*, [Nucl. Instrum. Methods Phys. Res. A](#) **1056** (2023) 168678.
- [137] G. Aglieri et al., *Monolithic Active Pixel Sensor Development for the Upgrade of the ALICE Inner Tracking System*, [J. Instrum.](#) **8** (2013) C12041.
- [138] F. Reidt, *Upgrade of the ALICE ITS Detector*, [Nucl. Instrum. Methods Phys. Res. A](#) **1032** (2022) 166632.
- [139] H. Pernegger et al., *First Tests of a Novel Radiation Hard CMOS Sensor Process for Depleted Monolithic Active Pixel Sensors*, [J. Instrum.](#) **12** (2017) P06008.
- [140] M. Dyndal et al., *Mini-MALTA: Radiation Hard Pixel Designs for Small-Electrode Monolithic CMOS Sensors for the High Luminosity LHC*, [J. Instrum.](#) **15** (2020) P02005.
- [141] ALICE Collaboration, *Technical Design Report for the ALICE Inner Tracking System 3 - ITS3*, tech. rep. CERN-LHCC-2024-003, 2024, CDS: [2890181](#).
- [142] J. Liu, *ALICE ITS3: A Truly Cylindrical Vertex Detector Based on Bent, Wafer-Scale Stitched CMOS Sensors*, [Nucl. Instrum. Methods Phys. Res. A](#) **1064** (2024) 169355.
- [143] G. Iacobucci et al., *Efficiency and Time Resolution of Monolithic Silicon Pixel Detectors in SiGe BiCMOS Technology*, [J. Instrum.](#) **17** (2022) P02019.
- [144] L. Rossi, *Pixel Detectors Hybridisation*, [Nucl. Instrum. Methods Phys. Res. A](#) **501** (2003) 239.
- [145] M. Garcia-Sciveres and J. Christiansen, *RD Collaboration Proposal: Development of Pixel Readout Integrated Circuits for Extreme Rate and Radiation*, CERN-LHCC-2013-008, CERN, 2013, CDS: [1553467](#).
- [146] J. Christiansen and F. Loddo, *Extension of RD53*, CERN-LHCC-2018-028, CERN, 2018, CDS: [2637453](#).
- [147] F. Loddo et al., *RD53 Pixel Chips for the ATLAS and CMS Phase-2 Upgrades at HL-LHC*, [Nucl. Instrum. Methods Phys. Res. A](#) **1067** (2024) 169682.

- [148] RD53 Collaboration, *RD53A Integrated Circuit Specifications*, CERN-RD53-PUB-15-001, CERN, 2015, CDS: [2113263](#).
- [149] RD53 Collaboration, *The RD53A Integrated Circuit*, CERN-RD53-PUB-17-001, version 3.51, CERN, 2019, CDS: [2287593](#).
- [150] RD53 Collaboration, *RD53B Design Requirements*, CERN-RD53-PUB-19-001, CERN, 2019, CDS: [2663161](#).
- [151] RD53 Collaboration, *RD53B Manual*, CERN-RD53-PUB-19-002, version 2.21, CERN, 2023, CDS: [2665301](#).
- [152] RD53 Collaboration, *RD53B Users Guide*, CERN-RD53-PUB-21-001, CERN, 2021, CDS: [2754251](#).
- [153] RD53 Collaboration, *RD53C Chip Manual*, CERN-RD53-PUB-24-001, version 1.92, CERN, 2024, CDS: [2890222](#).
- [154] RD53 Collaboration, *RD53 Pixel Readout Integrated Circuits for ATLAS and CMS HL-LHC Upgrades*, *J. Instrum.* **20** (2025) P03024.
- [155] L. Gaioni, *Test Results and Prospects for RD53A, a Large Scale 65 Nm CMOS Chip for Pixel Readout at the HL-LHC*, *Nucl. Instrum. Methods Phys. Res. A* **936** (2019) 282.
- [156] XILINX, *Aurora 64B/66B: Protocol Specification*, version SP011 (v1.3), 2014.
- [157] D. Ta et al., *Serial Powering: Proof of Principle Demonstration of a Scheme for the Operation of a Large Pixel Detector at the LHC*, *Nucl. Instrum. Methods Phys. Res. A* **557** (2006) 445.
- [158] M. Karagounis et al., “An Integrated Shunt-LDO Regulator for Serial Powered Systems”, *2009 Proceedings of ESSCIRC*, Athens, Greece, 2009, 276.
- [159] J. Kampkötter et al., *Characterization and Verification of the Shunt-LDO Regulator and Its Protection Circuits for Serial Powering of the ATLAS and CMS Pixel Detectors*, *J. Phys.: Conf. Ser.* **2374** (2022) 012071.
- [160] S. Möbius, *Pixel Detector Studies for the ATLAS ITk Upgrade for the HL-LHC*, PhD thesis: Georg-August-Universität Göttingen, 2023.
- [161] R. Bates, *Technical Specification for ITk Pixel Modules*, tech. rep. AT2-IP-ES-0009, ATLAS Collaboration, 2024, EDMS: [2019657 v.3](#) (internal documentation).
- [162] D. Álvarez, A. Cueto, S. Kuehn, and B. Vormwald, *Design Overview of the Loaded Local Supports for the ITk Pixel Outer Barrel*, tech. rep. AT2-IP-ER-0046, ATLAS Collaboration, 2023, EDMS: [2822664 v.1](#) (internal documentation).

- [163] D. Giugni et al., *Specifications for the ATLAS ITk Pixel Services. Technical Requirements*, tech. rep. AT2-IP-EP-0007, ATLAS Collaboration, 2023, EDMS: [1817540 v.7](#) (internal documentation).
- [164] M. Garcia-Sciveres and N. Wermes, *A Review of Advances in Pixel Detectors for Experiments with High Rate and Radiation*, [Rep. Prog. Phys.](#) **81** (2018) 066101.
- [165] V. Perovic, *Serial Powering in Four-Chip Prototype RD53A Modules for Phase 2 Upgrade of the CMS Pixel Detector*, [Nucl. Instrum. Methods Phys. Res. A](#) **978** (2020) 164436.
- [166] S. Kersten, L. Püllen, and C. Zeitnitz, *Ongoing Studies for the Control System of a Serially Powered ATLAS Pixel Detector at the HL-LHC*, [J. Instrum.](#) **11** (2016) C02070.
- [167] V. Filimonov et al., *A Serial Powering Pixel Stave Prototype for the ATLAS ITk Upgrade*, [J. Instrum.](#) **12** (2017) C03045.
- [168] F. Hinterkeuser, *Evaluation of a Serial Powering Scheme and Its Building Blocks for the ATLAS ITk Pixel Detector*, PhD thesis: Rheinische Friedrich-Wilhelms-Universität Bonn, 2022, HDL: [20.500.11811/10392](#).
- [169] R. Ahmad, M. Karagounis, S. Kersten, N. Lehmann, and C. Zeitnitz, *Specification for the Pixel DCS ASIC: Monitoring of Pixel System*, tech. rep. AT2-IP-ES-0001, ATLAS Collaboration, 2021, EDMS: [1909505 v.3](#) (internal documentation).
- [170] International Organization for Standardization, *Road Vehicles — Controller Area Network (CAN) Part 2: High-speed Medium Access Unit*, ISO 11898-2:2016, 2016.
- [171] CAN in Automation (CiA), *CANopen Application Layer and Communication Profile*, CiA 301, version 4.2.0, 2011.
- [172] A. Walsemann et al., *A CANopen Based Prototype Chip for the Detector Control System of the ATLAS ITk Pixel Detector*, [PoS TWEPP2019](#) (2020) 013.
- [173] R. Ahmad, *Development and Characterization of the Monitoring of Pixel System (MOPS) Chip to Monitor the ATLAS ITk Pixel Detector*, PhD thesis: Bergische Universität Wuppertal, 2023.
- [174] P. Tropea et al., *CO₂ Evaporative Cooling: The Future for Tracking Detector Thermal Management*, [Nucl. Instrum. Methods Phys. Res. A](#) **824** (2016) 473.
- [175] P. Tropea et al., *Advancements and Plans for the LHC Upgrade Detector Thermal Management with CO₂ Evaporative Cooling*, [Nucl. Instrum. Methods Phys. Res. A](#) **936** (2019) 644.
- [176] L. Zwalinski et al., “Progress towards the Commissioning and Installation of the 2PACL CO₂ Cooling Control Systems for Phase II Upgrade of the ATLAS and CMS Experiments”, [Proc. ICALEPCS’23](#), TUPDP104, Cape Town, South Africa, 2024, 802.

- [177] B. Verlaat, M. Van Beuzekom, and A. Van Lysebetten, “CO2 Cooling for HEP Experiments”, *Proc. TWEPP2008*, Naxos, Greece, 2008, 328.
- [178] C. Amelung, H. Chen, and F. Lanni, *ATLAS Review Office and Review Strategy for Phase-II Detector Upgrades*, tech. rep. ATC-R-MN-0001, ATLAS Collaboration, 2021, EDMS: [1979229 v.3](#) (internal documentation).
- [179] H. L. Joos, *Development of the Detector Control System of the ITk Outer Barrel Demonstrator for the ATLAS Experiment*, II.Physik-UniGö-MSc-2022/01, MA thesis: Georg-August-Universität Göttingen, 2022.
- [180] D. Ecker, *Implementation of the Communication and Data Visualization for the ITk Pixel Monitoring Chip of the ATLAS Experiment at CERN*, MA thesis: Bergische Universität Wuppertal, 2021.
- [181] S. Kersten, P. Kind, C. Riegel, and N. Lehmann, *The Interlock Matrix Crate of the ATLAS IBL Detector*, tech. rep. ATL-IP-ES-0203, ATLAS Collaboration, 2021, EDMS: [1302654 v.1](#) (internal documentation).
- [182] J. N. Pettersen, *ATLAS ITK miniDCS Operation Manual*, manual, CERN, 2023, EDMS: [2815466 v.1](#) (internal documentation).
- [183] L. Franconi, *The Opto-Electrical Conversion System for the Data Transmission Chain of the ATLAS ITk Pixel Detector Upgrade for the HL-LHC*, *J. Phys.: Conf. Ser.* **2374** (2022) 012105.
- [184] L. Zhang et al., *The Design and Test Results of A Giga-Bit Cable Receiver (GBCR) for the ATLAS Inner Tracker Pixel Detector*, *J. Instrum.* **18** (2023) C03005.
- [185] P. Moreira, *The lpGBT: A Radiation Tolerant ASIC for Data, Timing, Trigger and Control Applications in HL-LHC*, Conference talk, Topical Workshop on Electronics for Particle Physics (Santiago de Compostela, Spain), 2019.
- [186] J. Troska et al., *The VTRx+, an Optical Link Module for Data Transmission at HL-LHC*, *PoS TWEPP2017* (2018) 048.
- [187] C. Soos, *The VTRX+ Story*, Seminar talk, EP-ESE Electronics Seminars (CERN), 2024.
- [188] F. Faccio, S. Michelis, G. Blanchot, G. Ripamonti, and A. Cristiano, *The bPOL12V DCDC Converter for HL-LHC Trackers: Towards Production Readiness*, *PoS TWEPP2019* (2020) 070.
- [189] G. Ripamonti et al., *2.5V Step-down DCDCs: A Radiation-Hard Solution for Power Conversion*, *PoS TWEPP2019* (2020) 071.
- [190] T. Heim, *YARR - a PCIe Based Readout Concept for Current and Future ATLAS Pixel Modules*, *J. Phys. Conf. Ser.* **898** (2017) 32053.

- [191] D. Harada, *Evaluation of the Performance and Development of the Production Procedures for the ITk Pixel Detector of the ATLAS Experiment*, PhD thesis: Université de Genève, 2024.
- [192] A. Gaa, *Development of the Detector Control System of the ITk Outer Barrel Demonstrator for the ATLAS Experiment*, II.Physik-UniGö-MSc-2022/07, MA thesis: Georg-August-Universität Göttingen, 2022.
- [193] J. S. Steinhart and S. R. Hart, *Calibration Curves for Thermistors*, *Deep Sea Res. Oceanogr. Abstr.* **15** (1968) 497.
- [194] R. Müller, *Detector Development and Analysis Techniques for Finding Leptoquarks with the ATLAS Detector at the LHC*, PhD thesis: Universität Bern, 2023.
- [195] L. Meng, *RD53A Module Testing Document*, tech. rep. ATL-COM-ITK-2019-045, CERN, 2019, CDS: [2702738](#).
- [196] Ł. Klimczyk and R. Kantor, *User's Manual for MARTA*, manual, 2019, EDMS: [2015836 v.1](#) (internal documentation).
- [197] B. Vormwald, *Outer Barrel Interlock Matrix for LLS QC Setup*, tech. rep. AT2-IP-ER-0055, ATLAS Collaboration, 2023, EDMS: [2962556 v.1](#) (internal documentation).
- [198] R. Ahmad, *Description and Manual to Use the MOPS Chip*, manual, private communication, 2022.
- [199] B. Vormwald, *Electrical QC Testing Procedure of ITk Pixel Loaded Local Supports*, tech. rep. AT2-IP-QC-0005, ATLAS Collaboration, 2023, EDMS: [2601084 v.1](#) (internal documentation).
- [200] A. B. Poy et al., *The Detector Control System of the ATLAS Experiment*, *J. Instrum.* **3** (2008) P05006.
- [201] Technical Inspection and Safety Commission, *Alarms and Alarm Systems*, tech. rep. SAFETY INSTRUCTION 37 (IS37), CERN, 2003, EDMS: [335802 v.4](#).
- [202] S. Lüders, R. B. Flockhart, G. Morpurgo, and S. M. Schmeling, “The CERN Detector Safety System for LHC Experiments”, *Proc. ICALEPCS'03*, TH214, Gyeongju, Korea, 2003, 569.
- [203] O. Beltramello et al., *The Detector Safety System of the ATLAS Experiment*, *J. Instrum.* **4** (2009) P09012.
- [204] E. Manola-Poggioli and L. Scibile, “CERN Safety Alarm Monitoring System”, *Proc. EPAC'06*, WEPOCH156, Edinburgh, UK, 2006, 2293.
- [205] H. Nissen and S. Grau, “CERN Safety Alarm Monitoring”, *Proc. IPAC'11*, TUPS061, San Sebastian, Spain, 2011, 1674.

- [206] K. Lantzsch et al., *The ATLAS Detector Control System*, *J. Phys.: Conf. Ser.* **396** (2012) 12028.
- [207] A. Daneels and W. Salter, “What Is SCADA?”, *Proc. ICALEPCS’99*, MC1I01, Trieste, Italy, 1999, 339.
- [208] *SIMATIC WinCC Open Architecture*, ETM professional control GmbH, URL: <https://www.winccoa.com/>.
- [209] A. Daneels and W. Salter, “Selection and Evaluation of Commercial SCADA Systems for the Controls of the CERN LHC Experiments”, *Proc. ICALEPCS’99*, TA2O01, Trieste, Italy, 1999, 353.
- [210] O. Holme, M. González-Berges, P. Golonka, and S. Schmeling, “The JCOP Framework”, *Proc. ICALEPCS’05*, O3_005, Geneva, Switzerland, 2005.
- [211] ATLAS Collaboration, *Software and Computing for Run 3 of the ATLAS Experiment at the LHC*, *Eur. Phys. J. C* **85** (2025) 234.
- [212] K. S. Nicpon et al., *The Embedded Local Monitor Board Upgrade Proposals*, *PoS TWEPP2018* (2019) 34.
- [213] S. Kersten et al., “The ITk Common Monitoring and Interlock System”, *Proc. ICALEPCS’19*, THCPR06, New York, NY, USA, 2020, 1634.
- [214] T. Henss et al., *The Hardware of the ATLAS Pixel Detector Control System*, *J. Instrum.* **2** (2007) P05006.
- [215] S. Kersten et al., *The ITk Interlock Hardware Protection System*, *Nucl. Instrum. Methods Phys. Res. A* **1045** (2023) 167613.
- [216] W. Mahnke, S.-H. Leitner, and M. Damm, *OPC Unified Architecture*, Springer, 2009, ISBN: 978-3-540-68898-3.
- [217] P. P. Nikiel, B. Farnham, S. Schlenker, H. Boterenbrood, and V. Filimonov, “OPC Unified Architecture within the Control System of the ATLAS Experiment”, *Proc. ICALEPCS’13*, MOPPC032, San Francisco, CA, USA, 2014, 143.
- [218] P. P. Nikiel, B. Farnham, V. Filimonov, and S. Schlenker, *Generic OPC UA Server Framework*, *J. Phys. Conf. Ser.* **664** (2015) 082039.
- [219] P. P. Nikiel, P. Moschovakos, and S. Schlenker, “Quasar: The Full-Stack Solution for Creation of OPC-UA Middleware”, *Proc. ICALEPCS’19*, MOPHA100, New York, NY, USA, 2020, 453.
- [220] C. Gaspar, M. Dönszelmann, and Ph. Charpentier, *DIM, a Portable, Light Weight Package for Information Publishing, Data Transfer and Inter-Process Communication*, *Comput. Phys. Commun.* **140** (2001) 102.

- [221] B. Copy, E. Mandilara, I. P. Barreiro, and F. Varela, “Monitoring of CERN’s Data Interchange Protocol (DIP) System”, *Proc. ICALEPCS’17*, THPHA162, Barcelona, Spain, 2018, 1797.
- [222] Modbus Organization, *MODBUS PROTOCOL*, 2025, URL: <https://modbus.org/specs.php>.
- [223] R. Zurawski, ed., *Industrial Communication Technology Handbook*, CRC Press, 2015, ISBN: 978-1-138-07181-0.
- [224] G. Schnell and B. Wiedemann, *Bussysteme in der Automatisierungs- und Prozesstechnik*, Springer, 2019, ISBN: 978-3-658-23687-8.
- [225] J. Wang and W. Tepfenhart, *Formal Methods in Computer Science*, Chapman and Hall/CRC, 2019, ISBN: 978-0-429-06368-8.
- [226] J. H. Conway, *Regular Algebra and Finite Machines*, Dover Publications, 2012, ISBN: 978-0-486-48583-6.
- [227] B. Franek and C. Gaspar, *SMI++ Object Oriented Framework for Designing and Implementing Distributed Control Systems*, *IEEE Trans. Nucl. Sci.* **45** (1998) 1946.
- [228] B. Franek and C. Gaspar, *SMI++ Object Oriented Framework Used for Automation and Error Recovery in the LHC Experiments*, *J. Phys.: Conf. Ser.* **219** (2010) 22031.
- [229] J. Barlow et al., *Run Control in MODEL: The State Manager*, *IEEE Trans. Nucl. Sci.* **36** (1989) 1549.
- [230] DELPHI Collaboration, *The DELPHI Detector at LEP*, *Nucl. Instrum. Methods Phys. Res. A* **303** (1991) 233.
- [231] C. Gaspar and B. Franek, *Tools for the Automation of Large Distributed Control Systems*, *IEEE Trans. Nucl. Sci.* **53** (2006) 974.
- [232] ATLAS Collaboration, *Detector Control System (DCS) Figures*, URL: <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/ApprovedPlotsDCS> (visited on 05/07/2025).
- [233] ATLAS DCS, *FSM Integration Guideline*, ATL-DQ-ON-0010, 2021, EDMS: [685114 v.3.0](#) (internal documentation).
- [234] R. Kopeliansky and S. Schlenker, *DCS: Requirements Document for HL-LHC*, tech. rep. ATU-GE-ES-0004, ATLAS Collaboration, 2020, EDMS: [2276493 v.1](#) (internal documentation).
- [235] M. Hamer, *Requirements for the Power Supply System of the ATLAS ITk Pixel Detector*, tech. rep. AT2-IP-ES-0014, ATLAS Collaboration, 2019, EDMS: [2192745 v.2](#) (internal documentation).

- [236] M. Marjanović, L. Meng, T. Heim, and E. Pianori, *Electrical Specification and QC Procedures for ITkPixV1.1 Modules*, tech. rep. AT2-IP-QA-0025, ATLAS Collaboration, 2023, EDMS: [2786321 v.1](#) (internal documentation).
- [237] A. Qamesh, *An FPGA-based Data Aggregator for the New ATLAS ITK Pixel DCS*, PhD thesis: Bergische Universität Wuppertal, 2024.
- [238] D. Ecker et al., *EMCI-EMP: Developments and Experience with the Novel Detector Control Solution*, *J. Instrum.* **20** (2025) C01020.
- [239] A. Qamesh et al., *An FPGA-based Data Aggregator for the New ATLAS ITk Pixel DCS*, 2024, arXiv: [2410.24057 \[physics\]](#), URL: <http://arxiv.org/abs/2410.24057>, pre-published.
- [240] C. Gemme et al., *ITk Interlock System Strategy*, tech. rep. AT2-IE-ES-0005, ATLAS Collaboration, 2023, EDMS: [2387552 v.8](#) (internal documentation).
- [241] ATLAS DCS, *ATLAS DCS Integration Guidelines*, tech. rep. ATLAS-DQ-ON-0013, 2007, EDMS: [685451 v.3](#) (internal documentation).
- [242] M. Hamer, *ITk Pixel Connectivity Spreadsheet for Electrical Services*, tech. rep. AT2-IP-ES-0016, ATLAS Collaboration, 2021, EDMS: [2423241 v.1](#) (internal documentation).
- [243] S. F. Koch, *Measurements of ATLAS, Measurements with ATLAS: Construction and Characterisation of ITk Pixel Detector Structures, and a Search for Leptoquarks in Events with Di-Tau Final States*, PhD thesis: University of Oxford, 2025.
- [244] R. Jirásek, *Improvements on the DCS and Monitoring System for the ITk Pixel Loaded Local Supports*, tech. rep. CERN-STUDENTS-Note-2024-168, CERN, 2024, CDS: [2911508](#).
- [245] S. Schmeling, *Installation Tool User’s and Developer’s Guide*, manual, JCOP, 2003, EDMS: [1085148 v.1](#).
- [246] H. L. Joos, *Documentation for the DCS for Outer Barrel LLS*, tech. rep. AT2-IP-AP-0015, ATLAS Collaboration, 2024, EDMS: [2884949 v.1](#) (internal documentation).
- [247] D. Gourley and B. Totty, *HTTP: The Definitive Guide*, O’Reilly, 2002, ISBN: 978-1-56592-509-0.
- [248] K. Iglberger, *C++ Software Design: Design Principles and Patterns for High-Quality Software*, O’Reilly, 2022, ISBN: 978-1-0981-1316-2.
- [249] M. Battaglieri et al., *US Cosmic Visions: New Ideas in Dark Matter 2017: Community Report*, 2017, arXiv: [1707.04591 \[hep-ph\]](#), pre-published.
- [250] ATLAS Collaboration, *Measurement of the Inelastic Proton-Proton Cross Section at $s = 13$ TeV with the ATLAS Detector at the LHC*, *Phys. Rev. Lett.* **117** (2016) 182002.

- [251] H. Van Haevermaet (CMS Collaboration), *Measurement of the inelastic proton-proton cross section at $\sqrt{s} = 13$ TeV*, 2016, arXiv: [1607.02033 \[hep-ex\]](#), pre-published.
- [252] J. L. Feng, I. Galon, F. Kling, and S. Trojanowski, *ForwArd Search ExpeRiment at the LHC*, [Phys. Rev. D **97** \(2018\) 035001](#).
- [253] FASER Collaboration, *Letter of Intent for FASER: ForwArd Search ExpeRiment at the LHC*, tech. rep. CERN-LHCC-2018-030, 2018, CDS: [2642351](#).
- [254] FASER Collaboration, *Technical Proposal for FASER: ForwArd Search ExpeRiment at the LHC*, CERN-LHCC-2018-036, 2018, CDS: [2651328](#).
- [255] FASER Collaboration, *FASER’s Physics Reach for Long-Lived Particles*, [Phys. Rev. D **99** \(2019\) 95011](#).
- [256] FASER Collaboration, *Detecting and Studying High-Energy Collider Neutrinos with FASER at the LHC: FASER Collaboration*, [Eur. Phys. J. C **80** \(2020\) 61](#).
- [257] FASER Collaboration, *Technical Proposal: FASERnu*, CERN-LHCC-2019-017, 2020, CDS: [2702868](#).
- [258] H. Abreu et al., *The Tracking Detector of the FASER Experiment*, [Nucl. Instrum. Methods Phys. Res. A **1034** \(2022\) 166825](#).
- [259] A. Abdesselam et al., *The Barrel Modules of the ATLAS Semiconductor Tracker*, [Nucl. Instrum. Methods Phys. Res. A **568** \(2006\) 642](#).
- [260] LHCb Collaboration, *LHCb Calorimeters: Technical Design Report*, tech. rep. CERN-LHCC-2000-036, 2000, CDS: [494264](#).
- [261] FASER Collaboration, *Reconstruction and Performance Evaluation of FASER’s Emulsion Detector at the LHC*, [J. Instrum. **20** \(2025\) P12018](#).
- [262] J. L. Feng, I. Galon, F. Kling, and S. Trojanowski, *Axionlike Particles at FASER: The LHC as a Photon Beam Dump*, [Phys. Rev. D **98** \(2018\) 55021](#).
- [263] F. Kling and P. Quílez, *ALP Searches at the LHC: FASER as a Light-Shining-through-Walls Experiment*, [Phys. Rev. D **106** \(2022\) 055036](#).
- [264] D. Hayakawa, *The W-Si High Precision Preshower Detector of the FASER Experiment at the LHC*, [EPJ Web Conf. **320** \(2025\) 00056](#).
- [265] R. E. Kotitsa, *Modeling, Simulation Development, and Testing of the High-Resolution FASER Preshower Detector for Dark Matter Searches at CERN*, PhD thesis: Université de Genève, 2024.
- [266] T. Moretti, *The FASER Pre-Shower Upgrade: A High Spatial Resolution W-Si Pre-Shower Detector Enabling Sensitivity to Long-Lived Particle Decays into Photons*, PhD thesis: Université de Genève, 2025.

- [267] L. Paolozzi et al., *Time Resolution and Power Consumption of a Monolithic Silicon Pixel Prototype in SiGe BiCMOS Technology*, [J. Instrum.](#) **15** (2020) P11025.
- [268] I. Asensi Tortajada, *Design and Development of Safety and Control Systems in ATLAS*, PhD thesis: Universitat de València, 2021, HDL: [10550/80793](#).
- [269] A. B. Dichiara, A. Song, S. M. Goodman, D. He, and J. Bai, *Smart Papers Comprising Carbon Nanotubes and Cellulose Microfibers for Multifunctional Sensing Applications*, [J. Mater. Chem. A](#) **5** (2017) 20161.
- [270] C. Solans, *RELIANCE Water Leak Detection System in ATLAS*, tech. rep. ATC-TY-EY-0825, ATLAS Collaboration, 2023, EDMS: [2423056 v.0.6](#) (internal documentation).
- [271] F. Haslbeck, private communication.

APPENDIX A

Connection diagrams

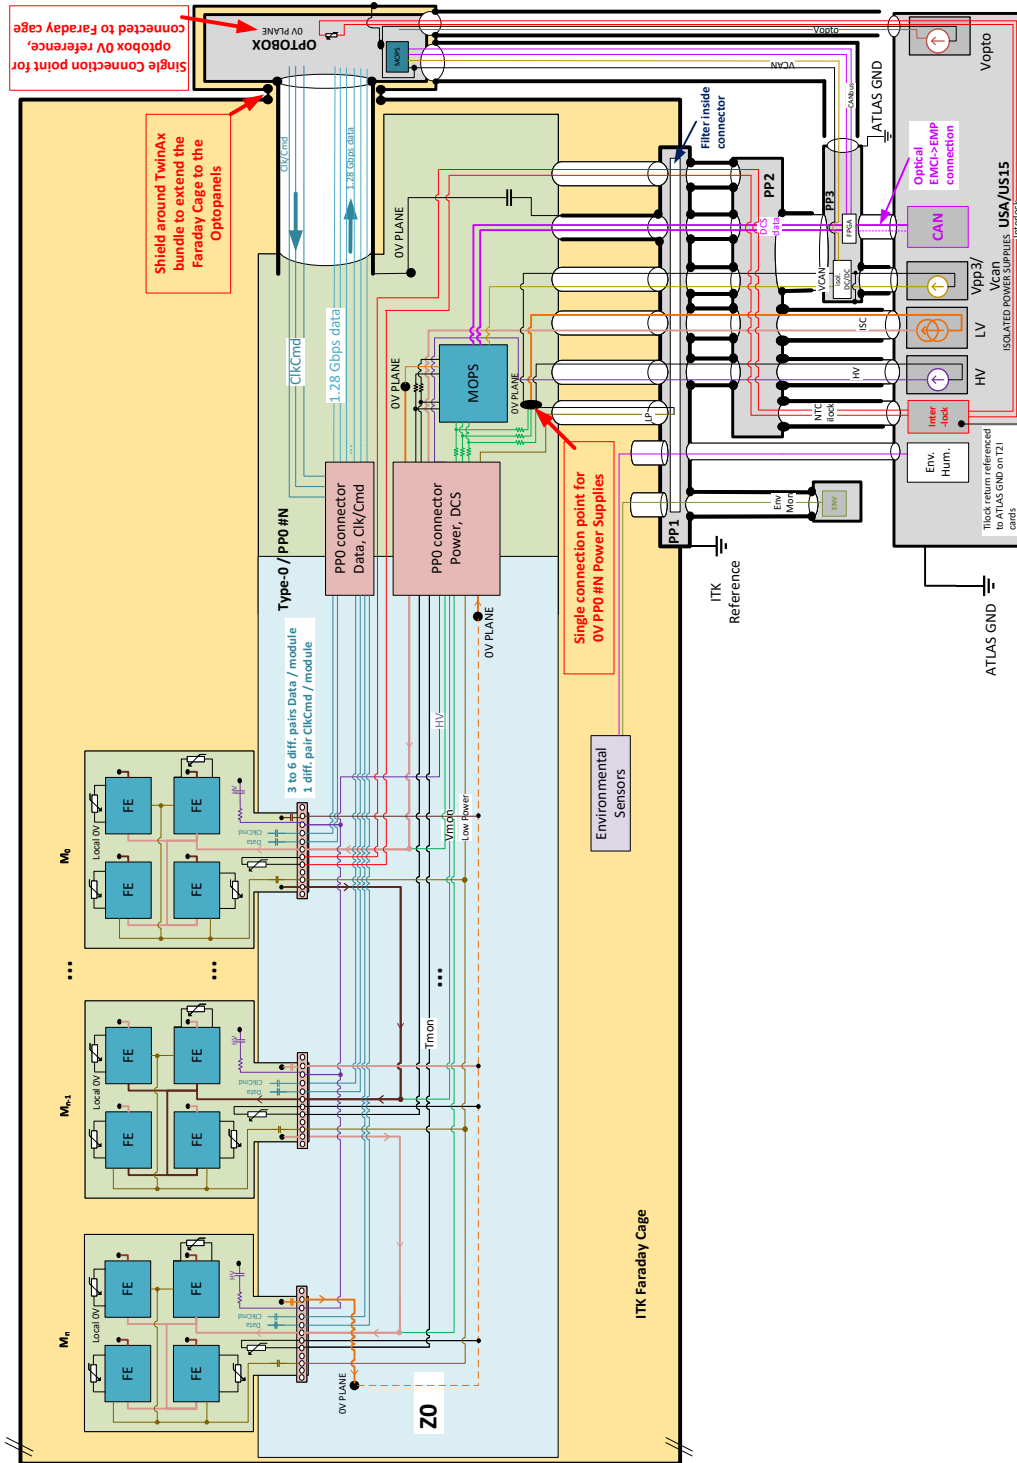


Figure A.1: Overview of the lines per SP-chain of the OB from the detector up to the electronics caverns US(A)15. [163]

Additional material for the results from the Outer Barrel demonstrator

Table B.1: Module temperatures of the M6 SP-chain (A-SP-1) on the demonstrator longeron during performance tests. The temperature at the CO₂ inlet was measured to be 10.7 °C and at the outlet 10.5 °C.

Position in Longeron	Name in DCS	Module	Temperature
A-M-B-1	M01	ThinQ10	19.8 °C
A-M-T-1	M02	Paris10	21.8 °C
A-M-B-2	M03	Paris3	20.2 °C
A-M-T-2	M04	SiegenQ1	broken connection
A-M-B-3	M05 (interlock)	Paris12	25.0 °C
A-M-T-3	M06	SiegenQ2	22.9 °C

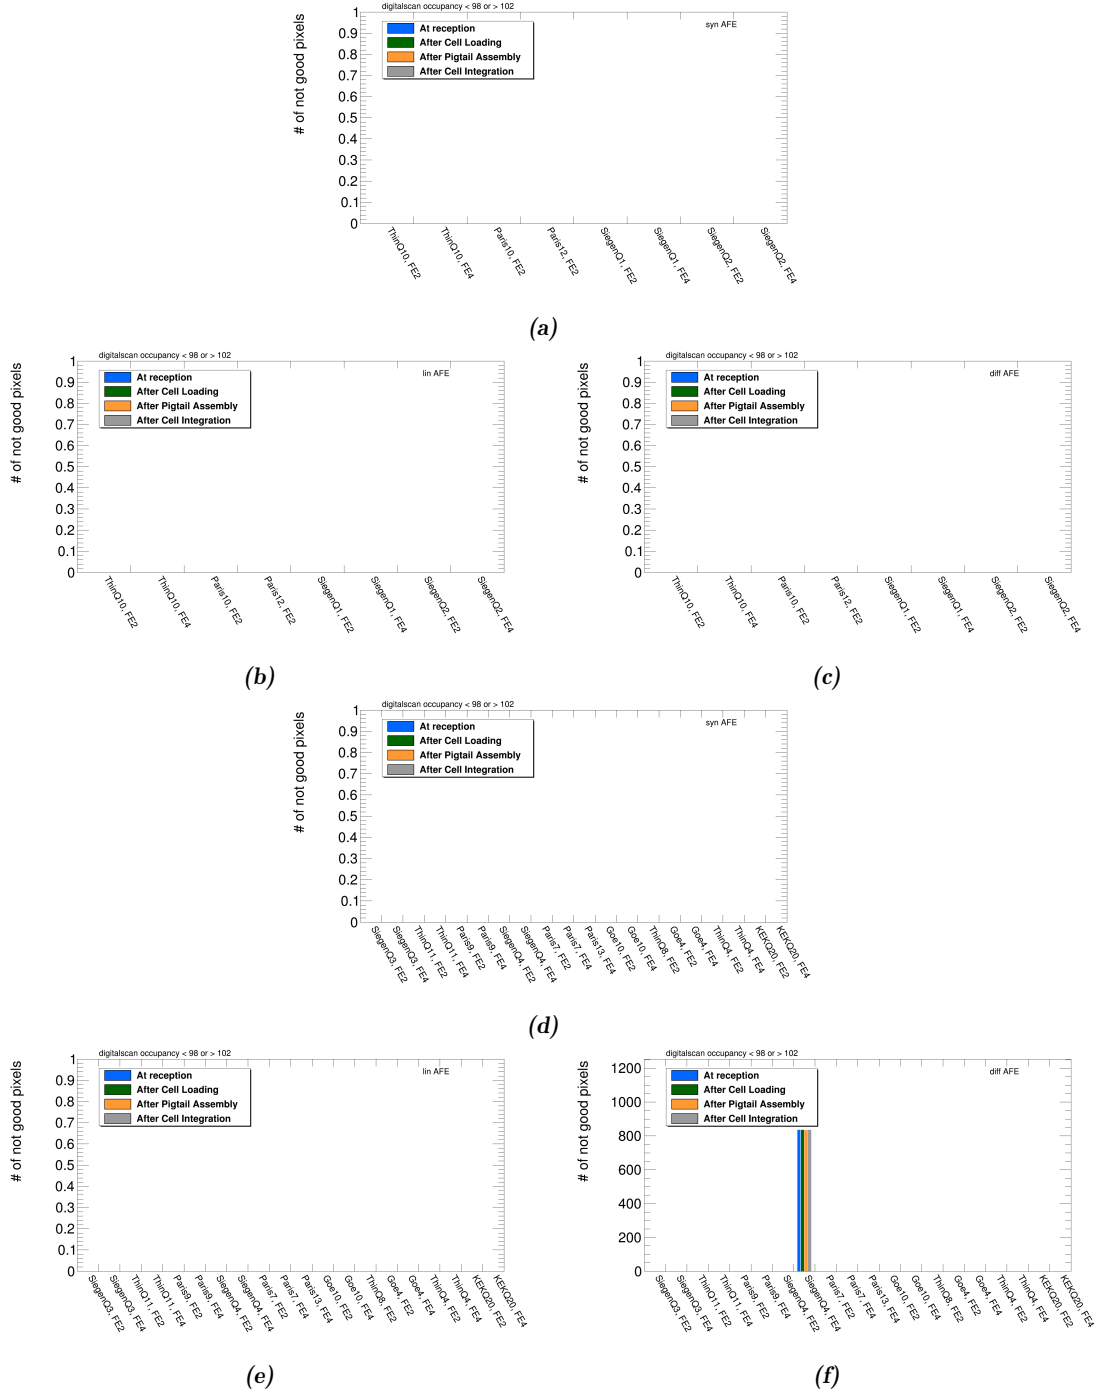


Figure B.1: Comparison of the number of not-good pixels in a digital scan over the four assembly stages. The results for the modules of the M6 are separated into the (a) synchronous, (b) linear and (c) differential AFE. The results for the modules of the M12 are also separated into the (d) synchronous, (e) linear and (f) differential AFE. A not-good pixel when 100 triggers are sent is defined as having less than 98 or more than 102 registered hits. The not-good pixels in the SiegenQ4, FE4 module are visible in the hit map shown in Fig. B.2.

Table B.2: Modules of the M12 SP-chain (A-SP-2) on the demonstrator longeron during performance tests. The temperature at the CO₂ inlet was measured to be 10.6 °C and at the outlet 10.5 °C.

Position in Longeron	Name in DCS	Module	Temperature
A-M-B-4	M01	SiegenQ3	21.8 °C
A-M-T-4	M02	ThinQ11	20.8 °C
A-M-B-5	M03	Paris9	25.3 °C
A-M-T-5	M04	SiegenQ4	21.4 °C
A-M-B-6	M05	Paris7	20.2 °C
A-M-T-6	M06	Paris13	22.2 °C
A-M-B-7	M07	Goe10	22.2 °C
A-M-T-7	M08	ThinQ8	22.2 °C
A-M-B-8	M09	Goe4	20.6 °C
A-M-T-8	M10	ThinQ4	22.4 °C
A-M-B-9	M11	KEKQ20	21.4 °C
A-M-T-9	M12 (interlock)	Liv5	22.3 °C




Figure B.2: Hit map from a digital scan of the SiegenQ1 module, FE4 in the (a) 1. stage of assembly and (b) stage of assembly. The not-good pixels are visible in both stages.

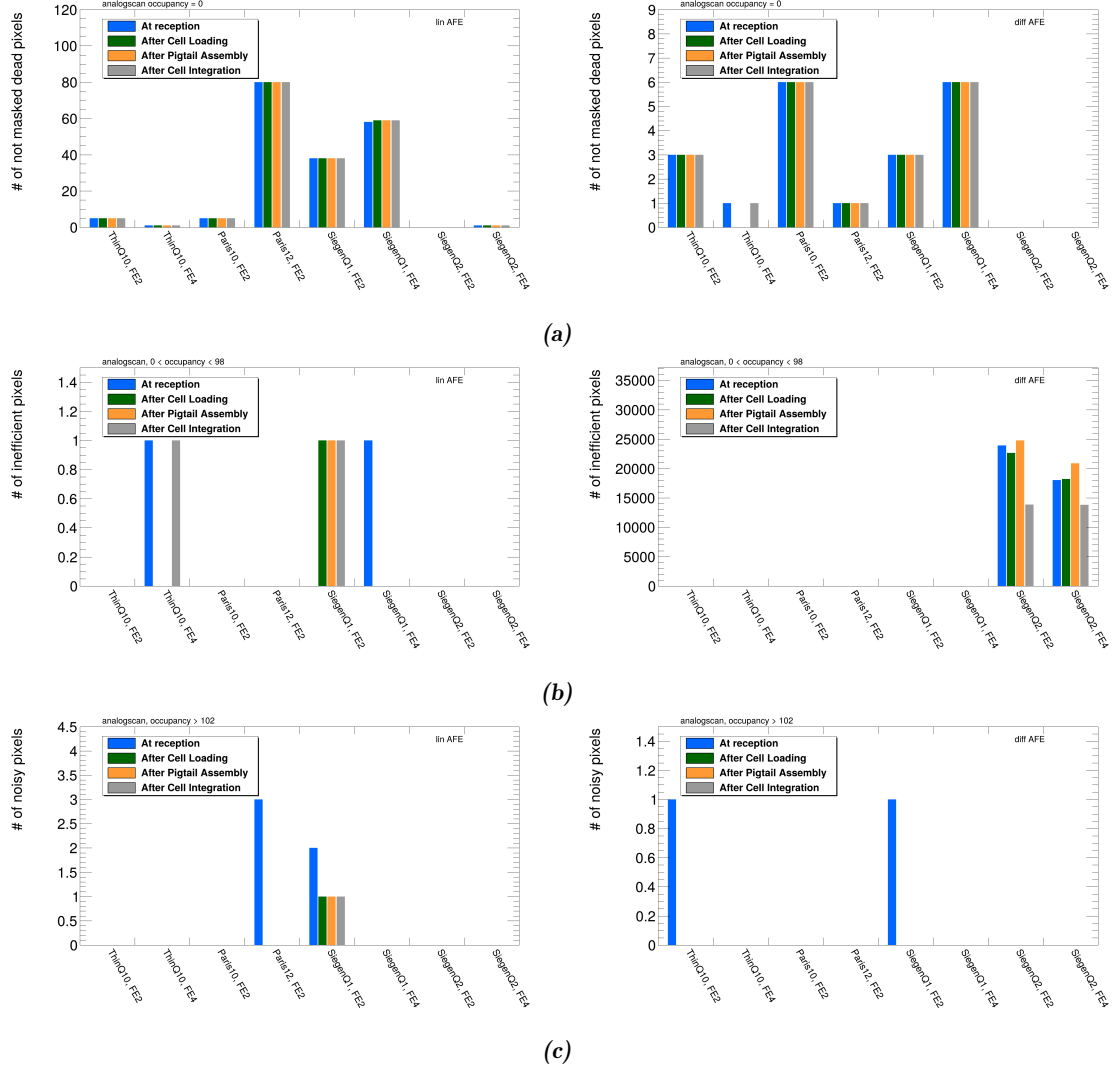


Figure B.3: Comparison of the number of (a) dead, (b) inefficient and (c) noisy pixels in an analog scan of the M6 over the four assembly stages. The evolution is shown for the linear (left) and differential (right) AFEs. See Fig. B.5, which shows the inefficient pixels of the differential AFEs of the SiegemQ2 module already in stage 1.

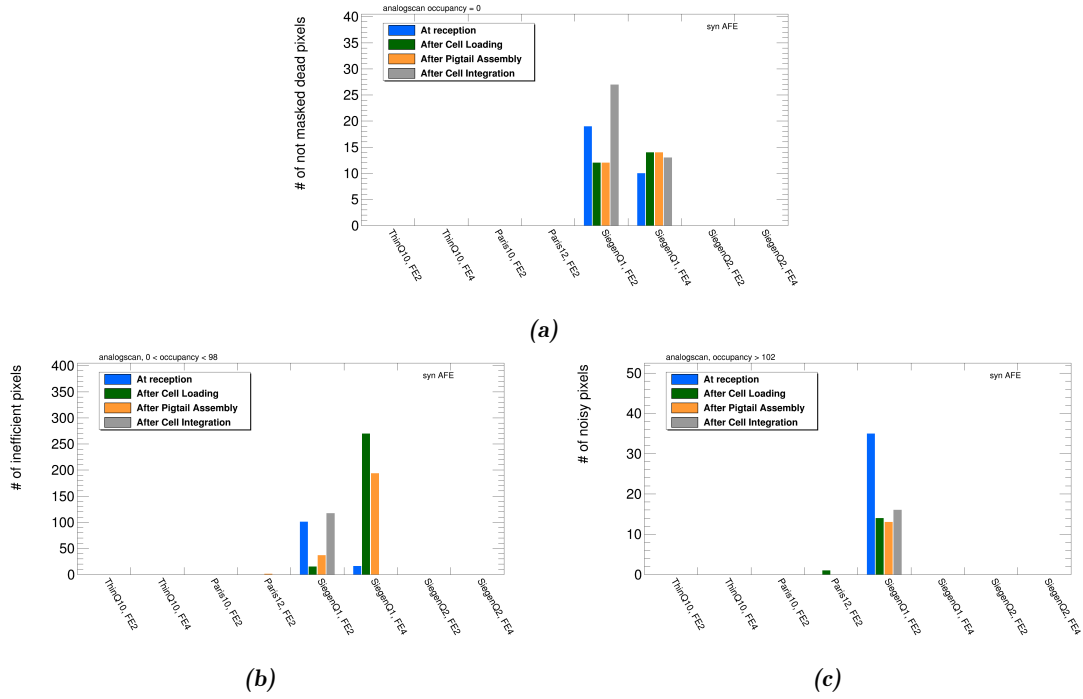


Figure B.4: Comparison of the number of (a) dead, (b) inefficient and (c) noisy pixels in an analog scan of the M6 over the four assembly stages. The evolution is shown for the synchronous AFEs.




Figure B.5: Occupancy map of an analog scan with 100 injections of the full SiegenQ2 module in stage 1. Already in this first stage, the module has many inefficient pixels in the differential AFEs.

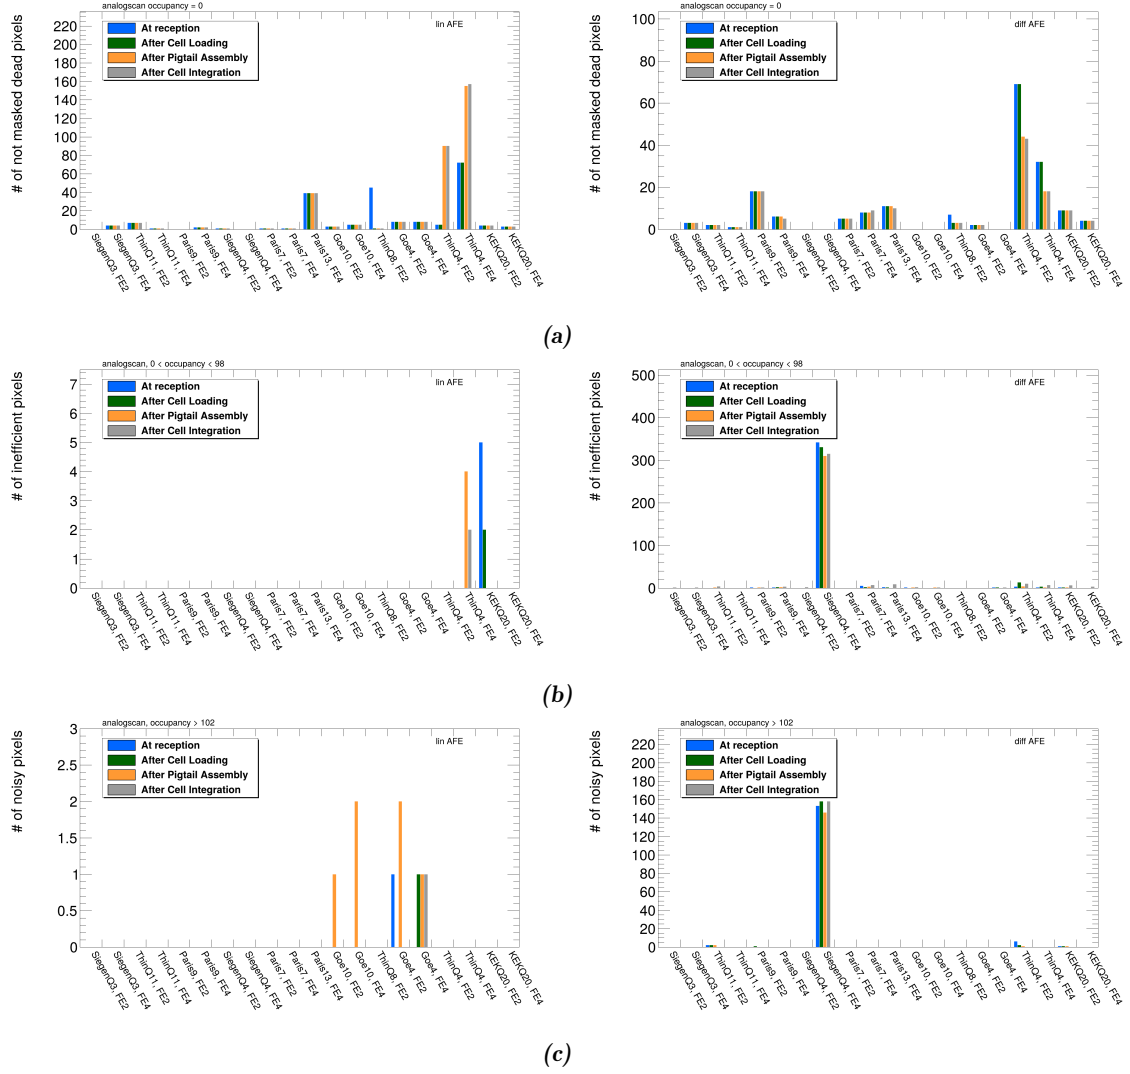


Figure B.6: Comparison of the number of (a) dead, (b) inefficient and (c) noisy pixels in an analog scan of the M12 over the four assembly stages. The evolution is shown for the linear (left) and differential (right) AFEs.


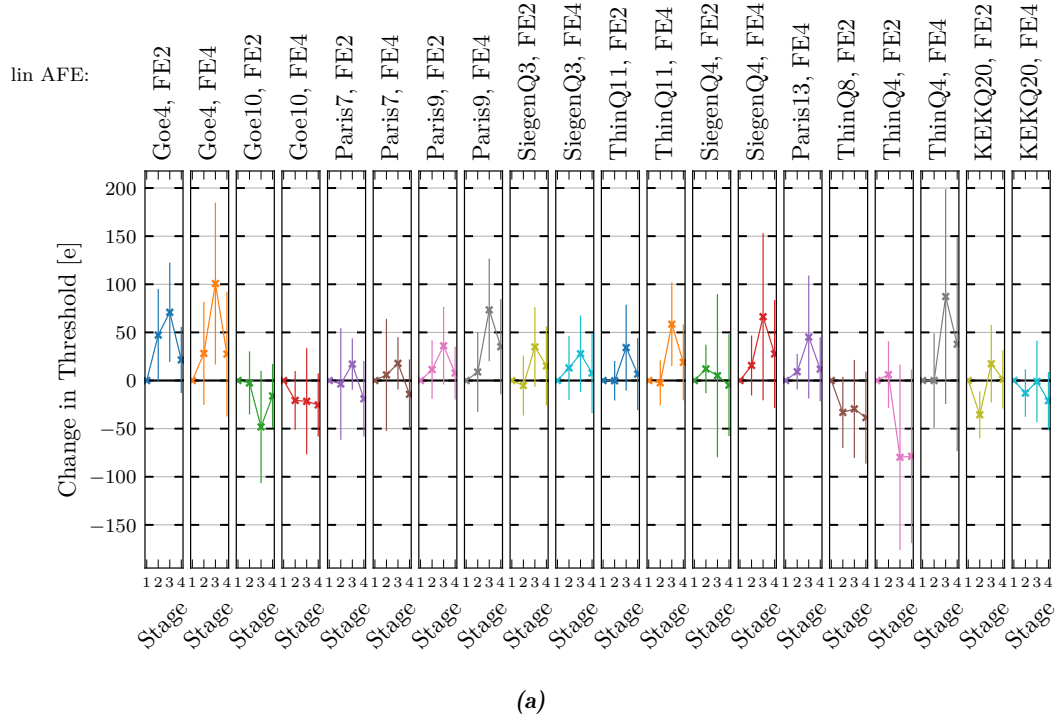


Figure B.7: Per pixel threshold difference evolution for the (a) linear AFEs of the M12 and (b) differential AFEs. Plotted are the means of the distributions of the per pixel changes (always with respect to stage 1) of the threshold. The error bars represent the standard deviation of this distribution. Stages are defined in Table 6.3.

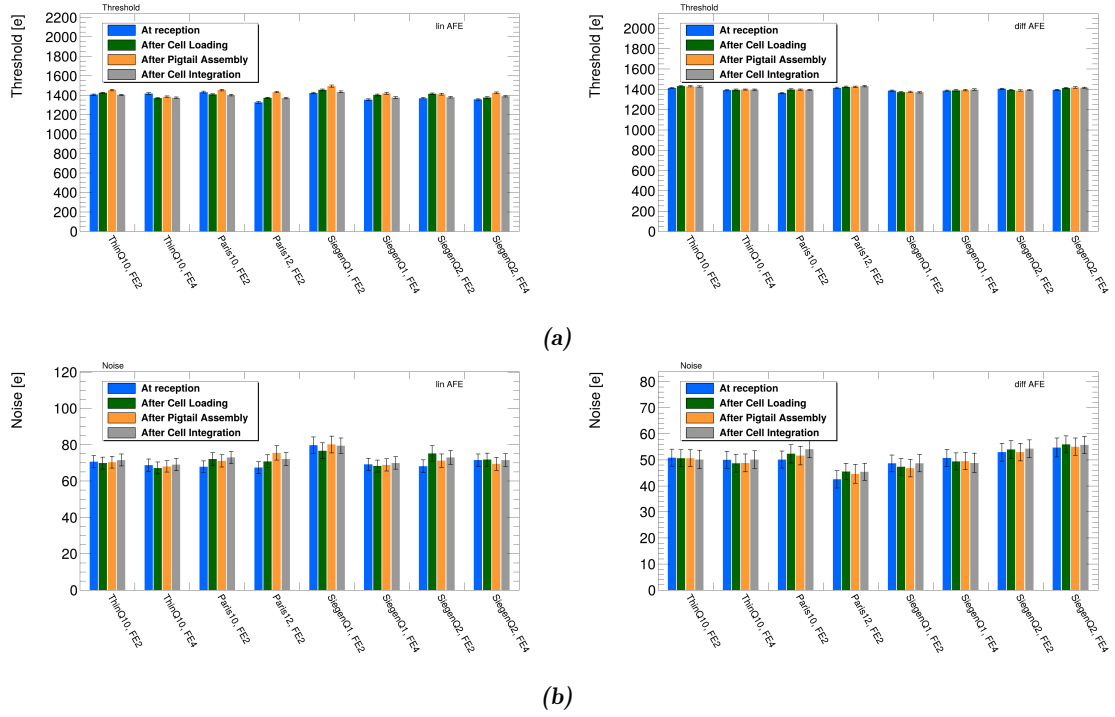


Figure B.8: Evolution of (a) threshold and (b) discriminator noise for linear (left) and differential (right) AFEs of the M6.

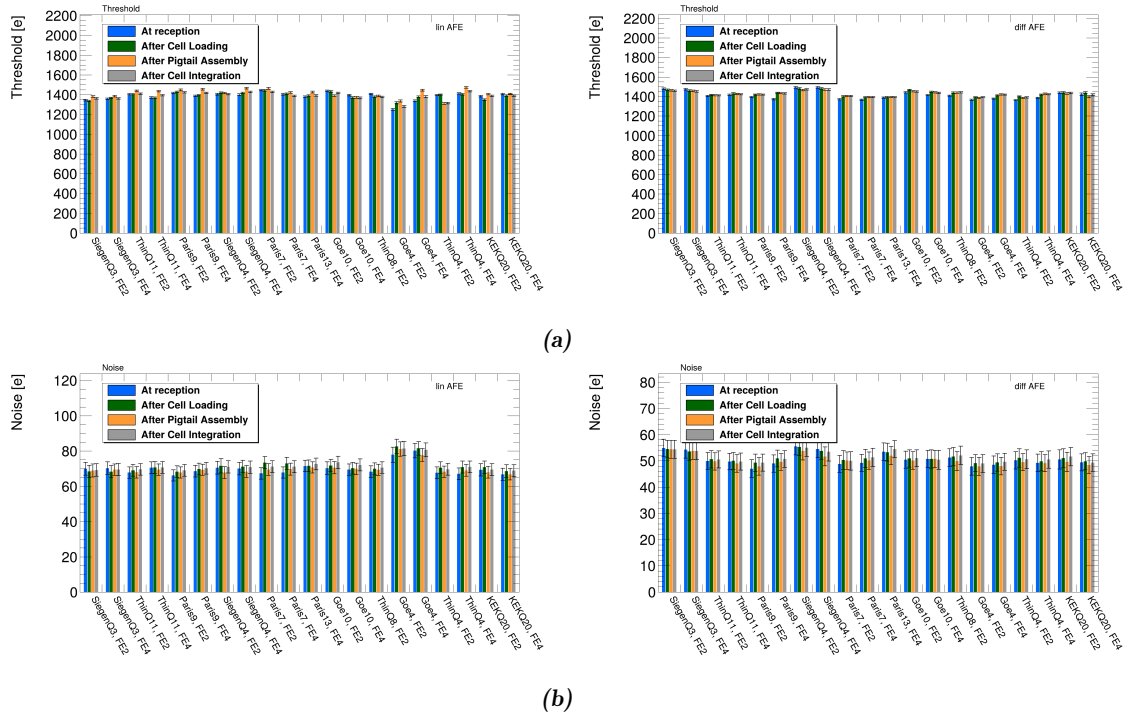


Figure B.9: Evolution of (a) threshold and (b) discriminator noise for linear (left) and differential (right) AFEs of the M12.

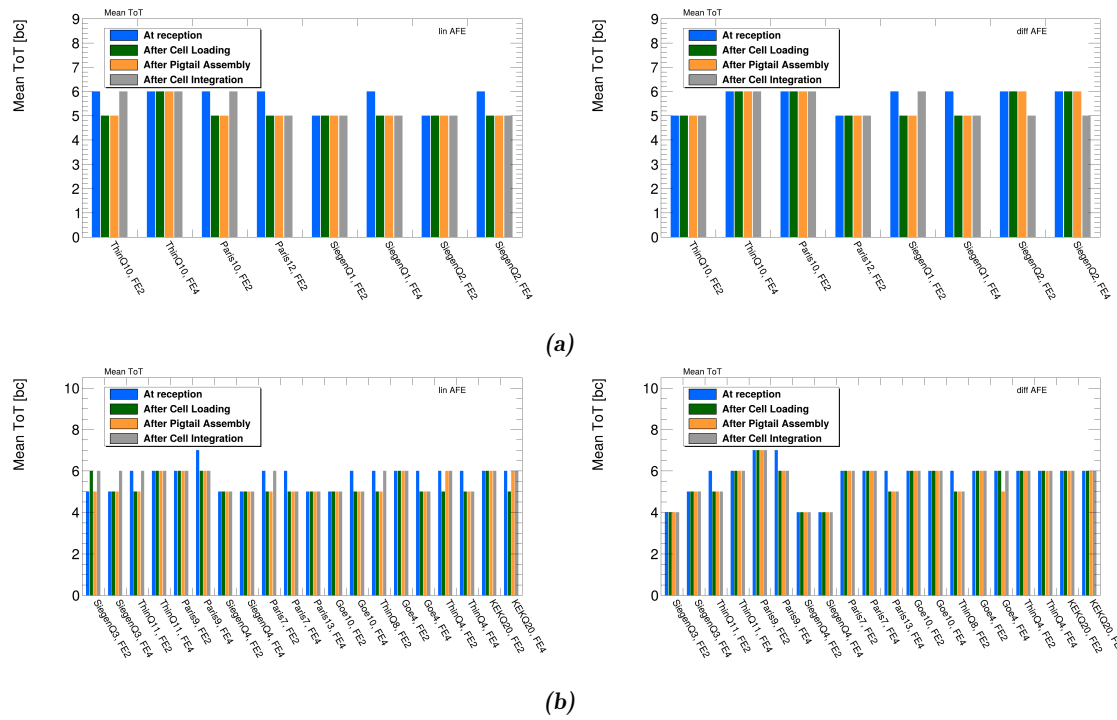


Figure B.10: Evolution of the most common ToT value for (a) the M6 and (b) the M12 for linear (left) and differential (right) AFEs.

Additional material for the ITk DCS developments




Figure C.1: Naming conventions and general structure of aliases used in the ITk DCS.

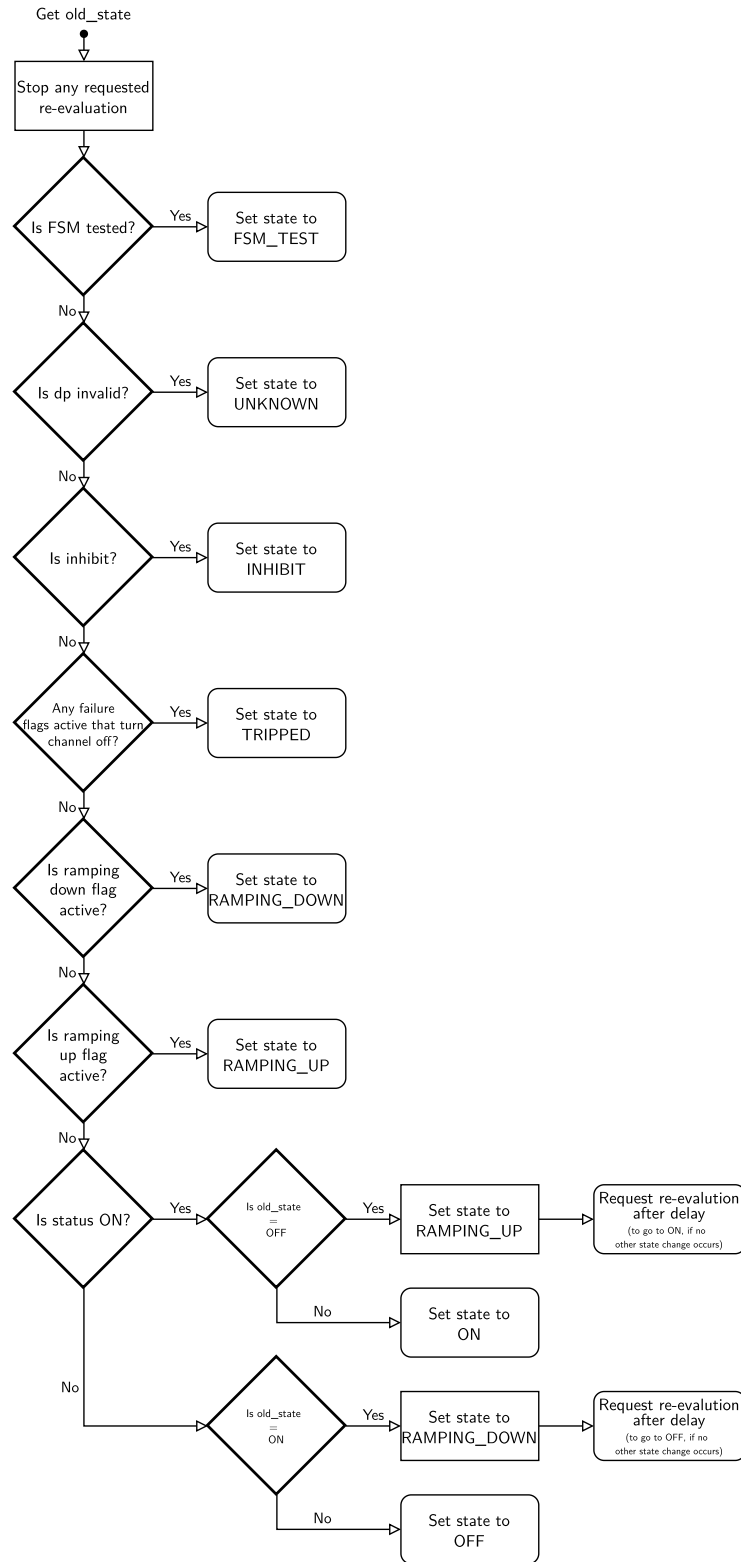


Figure C.2: Decision flow for state determination of power supply channel DUs. This decision flow guarantees a RAMPING state between ON and OFF states, even if the ramping state is not correctly reported by the power supply.

APPENDIX D

Additional material for the FASER Preshower DCS

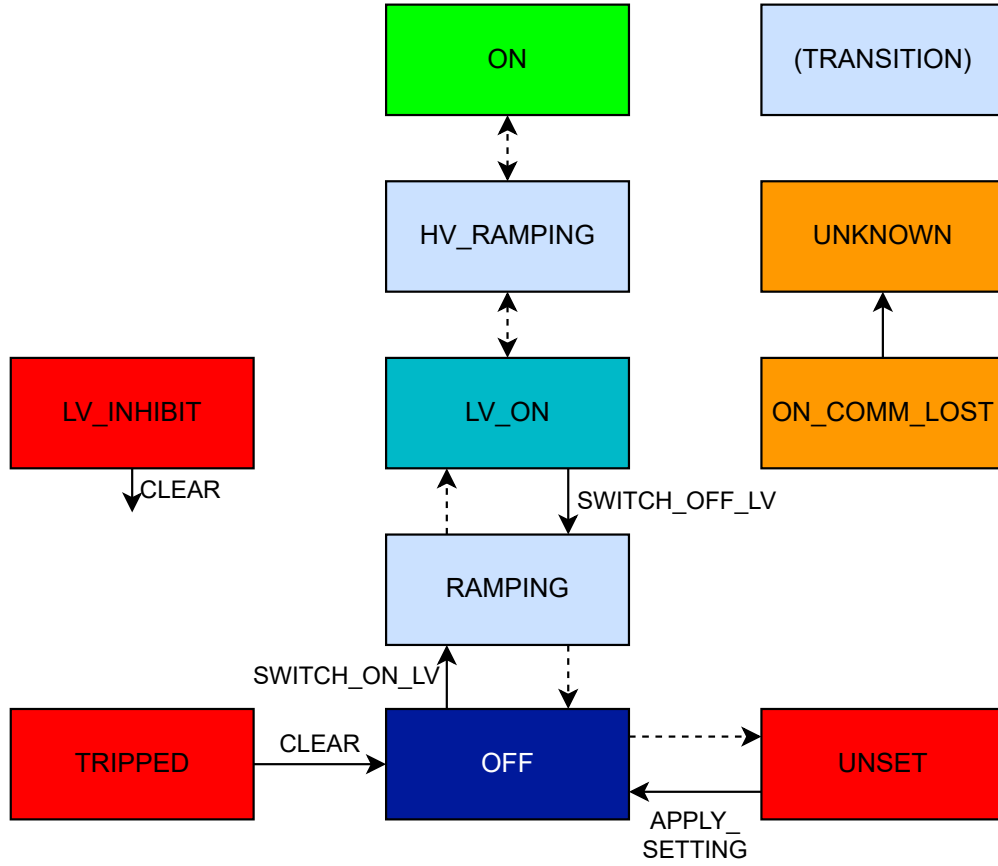


Figure D.1: The state diagram for the Preshower DC/DC LU that summarizes the states of an LV power channel and two modules powered by the LV power channel through a DC/DC converter.