

UNIVERSITY OF LIVERPOOL

DEPARTMENT OF ELECTRICAL ENGINEERING AND ELECTRONICS

Radiation Tolerant Acquisition System for Beam Loss Monitoring at High–Luminosity LHC

CERN-THESIS-2022-333
//2022



Thesis submitted in accordance with the requirements of the
University of Liverpool for the degree of Doctor in Philosophy by

Francesco Martina

October 2022

To

Valeria, light of my life,
my beloved parents Lea and Gigi,
my mates Andrea and Emiliano.

Abstract

At CERN's Large Hadron Collider, the strategy for machine protection and the quench prevention of superconductive magnets heavily relies on the Beam Loss Monitoring system. For each turn of the beam train, a large amount of data is acquired from approximately 4000 detectors by 750 digital conversion cards. Within a strict deadline, the readout electronics compare the measured loss radiation with unique per detector thresholds, thus deciding whether the particle beams are permitted to continue circulating or trigger an interlock. In view of the accelerator High-Luminosity upgrade, this protection device needs to be improved in terms of resolution and response speed, enhancing its ability to detect potential hazardous scenarios promptly. Such targets require an entirely renewed electronic front-end that withstands a TID of at least 3 kGy and can be positioned as close as possible to the detection spots. Compared to the currently installed equipment, it will also provide a four times faster sampling rate of 100 kHz. Here, the research activities and the development of such upgrades are reported. They involve the engineering, implementation and validation of a real-time system that digitises the output current of the detectors, forwards the results to a non-radiation area and performs an optimised signal processing, in order to obtain precise and accurate beam loss current measurements. Amongst them, the support for the design and the characterisation of an application-specific integrated circuit is also included, which plays a crucial role in the analogue to digital conversion. A frequency domain analysis of the produced data stream is conducted, demonstrating the advantages of this approach over conventional methods. In this way, a detailed evaluation of the digitiser's behaviour has been carried out, in response to a wide range of inputs, and appropriate filtering techniques have been identified. Hence, the assembly of an overall system test-bench and the standardised measurement procedures have led to the validation of the readout firmware, the signal treatment chains and the system conversion performance. The devices successfully complied with the design specifications, including the input dynamic range of more than 160 dB and the tolerance to radiation. In addition, the novel processing has provided optimal measurement precisions or, complementary, fast step responses. Finally, the flexibility of the developed instruments allows them to be easily adapted beyond the original scope, such as for medical and technological applications or in other particle accelerators.

Acknowledgements

The present research was conducted at the Beam Loss Section at CERN (SY-BI-BL), in the context of the Radiation to Electronics (R2E) and the High-Luminosity LHC (HL-LHC) projects. I thank the CERN, the University of Liverpool and the Cockcroft Institute, especially the QUASAR Group, for their scientific and financial support. In particular, I acknowledge my CERN supervisor Eng. Christos Zamantzas for his constant help and guidance, the Professors James W. Bradley and Carsten P. Welsch for believing in my project and accompanying me during the academic course. I am delighted by the warm hospitality of the Beam Loss Section and the Beam Instrumentation Group. Moreover, I want to give special thanks to my former supervisor Dr. Francesco Zappon, whose trust in me made it possible to undertake this experience at CERN. I sincerely thank all colleagues who were closely involved in the research at various stages, in particular Mathieu Saccani, Ewald Effinger, William Viganò, Eva Calvo, Jean Michel Meyer, Sara Morales Vigo, Erika Lima, Lorenzo Stefanini, Fernando S. Domingues Sousa, George Gimas and the BLMASIC designers of the EP-ESE-ME Section at CERN: Jan Kaplon, Luca Giangrande and Pedro V. Leitao. Thank you, Giuseppe Venturini, for your pioneering research into microelectronic systems for BLM and for your excellent contribution to CERN and the entire scientific community. Furthermore, the achieved results would not have been possible without fruitful cooperation with the EP-ESE-BE Section. Notably, I acknowledge Sophie Baron and Daniel Hernández Montesinos for their support in operating the lpGBT device, Jan Troska, Leonardo Marcon and Carmelo Scarcella for the CWDM SM-VTRx optical transceivers. From the BE-CEM-EPR Section, I thank Salvatore Danzeca, Rudy Ferraro and the PIF staff at PSI, who assisted us in everything concerning the radiation tests. Raphael Berberat and the CERN PCB Workshop are acknowledged for manufacturing the prototypes. I thank my colleague Manoel Barros Marin, for his contribution to the project, his trust in my work and the exciting discussions on the design challenges. Finally, I would like to thank all the unmentioned friends and colleagues who made my stay in Geneva more pleasant.

Contents

1 Introduction **1**
1.1 Statement of Intent and Objectives 1
1.2 Thesis Content Organisation 4

2 System Architecture **7**
2.1 Overview 8
2.2 The BLM Detectors 10
2.3 State-of-the-Art Electronics 12
2.3.1 Digitisation Front-End 13
2.3.2 Readout Electronics 14
2.4 Front-End Upgrades 17
2.5 Back-End Upgrades 19
2.6 Technical Specifications 21

3 ASIC Converters **23**
3.1 Current Measurements 23
3.2 Conversion Architectures 25
3.2.1 The Current to Frequency Converter (CFC) 25
3.2.2 The $\Delta\Sigma$ Converter 27
3.3 Converter Integration 31
3.3.1 Specifications 32
3.3.2 ASIC Structure 33
3.3.3 Integrator Circuit 34
3.3.4 Switches 36
3.3.5 Quantisers 37
3.3.6 Calibration DAC 38
3.3.7 Auxiliary ADC 39
3.3.8 Digital Logics 40
3.3.9 Package 44

4 ASIC Test-Bench **47**
4.1 Objectives 47

4.2 Characterisation Board . . . 48
4.2.1 Design of Schematics . . . 49
4.2.2 Board Layout . . . 52
4.3 Setup Architecture . . . 53
4.3.1 Initial Setup . . . 53
4.3.2 Converter Noise Characterisation Setup . . . 55
4.3.3 Post-Irradiation Setup . . . 56
4.4 Test-Bench Methodology . . . 57
4.4.1 Basic Initial Tests . . . 57
4.4.2 Digital Interface Validation . . . 57
4.4.3 Analogue Circuit Characterisation . . . 58
4.5 Data Analysis . . . 59
4.6 Basic Electrical Results . . . 61
4.7 Digital Characterisation Results . . . 63
4.8 Analogue Characterisation Results . . . 66
4.8.1 Conversion Precision . . . 66
4.8.2 Parallelised Data Analysis . . . 67
4.8.3 Maximum Achievable Resolution . . . 67
4.8.4 Integrating Capacitance . . . 68
4.8.5 Integration Asymmetry . . . 69
4.8.6 Calibration DAC . . . 70
4.8.7 Integrator Output Issue . . . 70
4.9 Selection Verdict . . . 72
4.10 Changelog of Versions . . . 72
4.10.1 BLMASIC Version 2 (2020) . . . 72
4.10.2 BLMASIC Version 3 (2021) . . . 73
4.11 X-Ray Irradiation Tests . . . 73

5 Signal Processing 75
5.1 Accuracy and Precision . . . 75
5.2 Measurement Composition . . . 77
5.3 Spectrum of Converter Measurements . . . 79
5.4 Noise Considerations . . . 84
5.5 Processing Techniques . . . 86
5.5.1 Moving Average Filters . . . 86
5.5.2 Infinite Impulse Response Filters . . . 88
5.6 Processing Metrics . . . 94
5.6.1 Equal Loss Charge . . . 94
5.6.2 Equal Rise Time . . . 95
5.7 Comparison of Methods . . . 96
5.8 Prototype Implementation . . . 102

5.9 Measurement Error Forecasting . . . 106

6 Front-End Development 107
6.1 Board Introduction . . . 108
6.2 Technical Specifications . . . 108
6.3 Board Architecture . . . 109
6.3.1 Component Selection . . . 110
6.4 Circuit Design . . . 111
6.4.1 Back-Panel Connection . . . 111
6.4.2 Converter Inputs . . . 112
6.4.3 Measurement Data Path . . . 114
6.4.4 Fast Lanes . . . 117
6.4.5 Board Control and Redundancy . . . 119
6.4.6 Hardwired Configuration . . . 121
6.4.7 Reset . . . 121
6.4.8 Diagnostics . . . 122
6.4.9 Power Distribution . . . 124
6.5 Layout Design . . . 127
6.5.1 Board Geometry . . . 127
6.5.2 PCB Stack-Up . . . 128
6.5.3 Component Placement . . . 130
6.5.4 Routing . . . 131
6.6 Board Start-Up and Configuration . . . 132

7 Back-End Development 137
7.1 Key Purposes . . . 137
7.2 Data Readout Architecture Overview . . . 138
7.3 FPGA Transceivers . . . 139
7.4 The lpGBT IP Core . . . 140
7.5 Phase Alignment . . . 143
7.6 Redundancy . . . 145
7.7 BLMASIC Data Retriever . . . 145
7.7.1 Physical Coding Sublayer . . . 146
7.7.2 Frame Interpreter . . . 150
7.8 Processing Implementation . . . 151
7.9 Diagnostics . . . 154
7.10 Test-Bench Data Transmission . . . 155
7.11 Performance Considerations . . . 157

8 Results 159
8.1 Overall System Assembly . . . 159
8.2 Front-End Power Consumption . . . 161

8.3 Analysis Strategy . . . 161
8.4 Data Calibration . . . 164
8.5 Conversion Performance . . . 166
8.6 Merit Factors . . . 173
8.7 Optimal Parameters . . . 175
8.8 Accuracy Considerations . . . 180
8.9 Full-Scale Reliability . . . 182
8.10 Radiation Test . . . 183
8.10.1 Irradiation Facilities . . . 184
8.10.2 Test-Bench Installation . . . 185
8.10.3 Irradiation Session . . . 186
8.10.4 Radiation Test Results . . . 188

9 Conclusions **193**

Appendices **199**

A Radiation Effects in Electronics **201**
A.1 Total Ionising Dose (TID) . . . 201
A.2 Single Event Effects (SEE) . . . 202
A.3 Displacement Damage (DD) . . . 203

B Ionisation Chamber Assembly **205**

C BLMASIC Registers **207**

D BLMASIC Pinout **213**

E Front-End Schematics **217**

F Front-End Layout **229**

References **233**

Chapter 1

Introduction

1.1 Statement of Intent and Objectives

CERN [S1], the European Organisation for Nuclear Research, represents the largest centre in particle physics, one of the most respected in the world. Designers at CERN develop and employ groundbreaking instruments to discover the primary constituents of matter and study its subatomic interactions. This allows us to provide insights into the fundamental laws of Nature and advance the boundaries of human knowledge by delving into the smallest building blocks of the Universe. Furthermore, the multitude of subjects and the breadth of challenges addressed at the centre extend the CERN mission to quality education for tomorrow's scientists, constructive cooperation amongst countries and significant enrichment of the technological know-how.

In order to accomplish these goals, CERN includes a massive complex of cascaded particle accelerators and detectors. Accelerators boost bunches of particles up to near the speed of light and produce controlled collisions within each other or towards fixed targets. The detectors observe all the types of scattered particles produced in the events

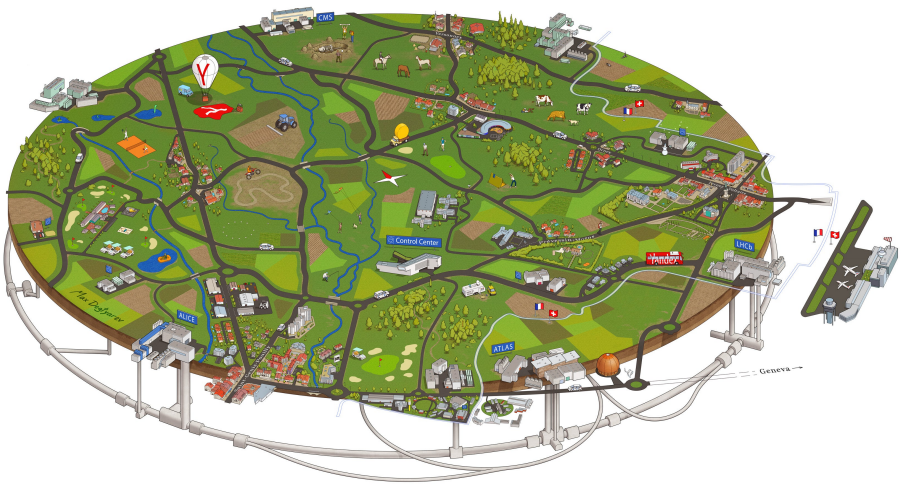

Figure 1.1: Illustration of the LHC area (by Max Degtyarev, from the Reddit page of CERN).

and they record the results of interest. Thus, a worldwide team of researchers, having different expertises, analyse such data to validate new theories, investigate unknown phenomena and improve the precision of measurements of already known Physics.

The most powerful accelerator built at CERN is the Large Hadron Collider (LHC) [S2], which has been active since 10th September 2008. The LHC is a synchrotron with circular layout and provides collisions of opposite moving proton beams, with a centre of mass energy up to 14 TeV. Following the discovery of the long-awaited Higgs boson [1], [2], achieved thanks to the active cooperation amongst ATLAS and CMS experiments, LHC continues to represent a worldwide reference for fundamental science research. Therefore, it will most likely retain its leadership as the world's highest energy accelerator for at least the next two decades.

Even though the accelerator optics are constantly optimised to impose strict confinement of the beam, a variety of phenomena cause particles to escape from the desired trajectory and spread outside the beam pipes. Therefore, it is very likely that they will interact with the surrounding materials, producing extensive secondary particle showers and electromagnetic radiation. These events lead to unwanted energy deposition in the various machine elements. Their probability of occurrence and harmfulness grow as more powerful accelerators are built.

The Beam Loss Monitoring system [3], hereafter referred to as the *BLM*, is responsible for keeping the level of radiation surrounding the beam pipes under control, estimating the magnitude of energy depositions and acting promptly in view of protecting the accelerator components. The system is extended throughout the particle path and it involves thousands of measurement spots, which are located according to the expected radiation levels and the criticality of the installed equipment. At each turn of the LHC beam train of bunches, a large amount of data are acquired from approximately 4000 detectors by 750 readout cards. Then, within a period of a few tens of microseconds, the back-end electronics compare the measured loss radiation with unique per detector thresholds, so they decide whether the particle beams are permitted to continue circulating, or an interlock should be triggered. The major challenge faced by the BLM concerns the need to carry out the sequence of operations within a strict deadline, allowing a reliable machine protection system.

In addition, the BLM provides an efficient diagnostic tool for the optimisation and inspection of the accelerator, since it can concurrently detect both local and extended anomalies. This does not only concern the spatial multi-spots monitoring, but it also includes the possibility of recording a long time history of measurements, paving the way for a more high-level assessment of the system behaviour.

The exploitation of the full performance of LHC has been the foremost priority stated by the European Strategy for Particle Physics Update (ESPPU) in 2013 [4]. Later, it was recognised that a further improvement of its components would allow an enhancement of its discovery capabilities. The upgrade design study began in 2011 and it has set the goals of increasing the instantaneous collision rate by a factor of 5 and

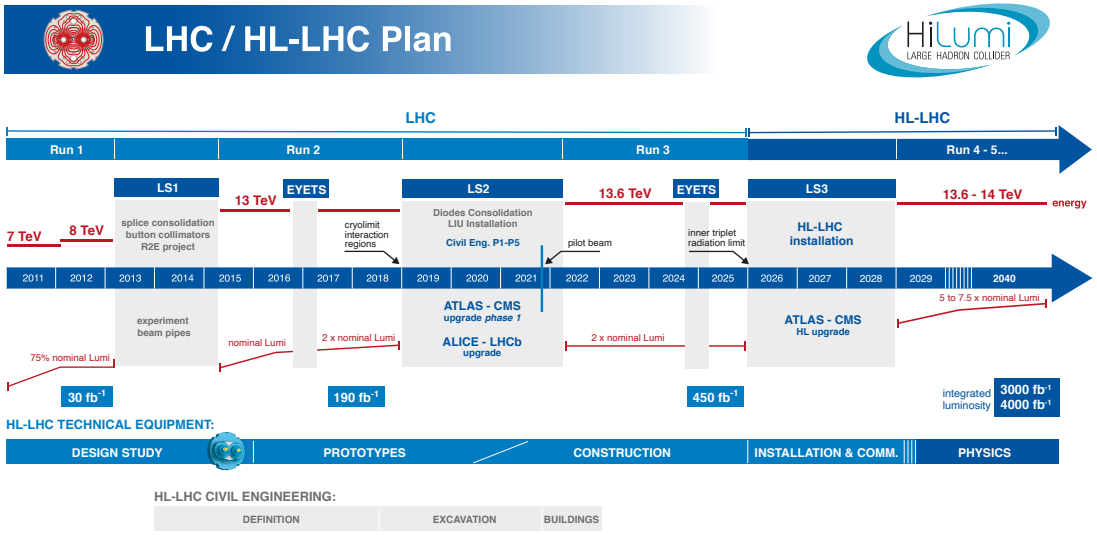

Figure 1.2: Work plan for LHC and HL-LHC accelerators in January 2022, including the centre of mass collision energy and the integrated luminosity forecasts.

the integrated luminosity by a factor of 10, with respect to the nominal design values. The renewed accelerator will be named High-Luminosity LHC (HL-LHC) [5], [S3].

Fig. 1.2 shows the work plan for developing and commissioning the HL-LHC, updated in 2022. Design efforts, which have been going on for ten years, have currently achieved relevant challenging features [5], amongst which: novel tuning of accelerator dynamic optics, groundbreaking superconducting magnet working up to 12 T, innovative design of magnets (such as Canted Cosine Theta and super-ferric), compact superconducting RF cavities for beam rotation and ultra-precise phase control, new technologies and materials for beam collimation and high-current superconducting links with almost null dissipation. Furthermore, in the context of such a modern accelerator, the Beam Loss Section (SY-BI-BL [S4]) is providing its expertise to develop a cutting-edge version of the BLM, exploiting completely renewed electronics. This choice is configured as a natural upgrade of the system, which has proven to be very robust and reliable during its entire lifetime, but is getting outdated and difficult to maintain.

In view of the HL-LHC commissioning, during the Long Shut-Down 3 (see Fig. 1.2), the renewed BLM system will be installed. Its mission-critical electronics shall provide an increased radiation tolerance of factor 6 at least [S5], a step response 4 times faster, a lower conversion noise and improved diagnostic features. The present work reports the design activities, the implementation and the validation of a real-time system, which 1) digitises the currents produced by the ionisation chamber detectors, minimising the lengths of analogue transmission lines, 2) propagates the results to a non-radiation area and 3) performs an optimised signal processing to obtain precise and accurate beam loss current measurements. Particular attention has been devoted to the choice of the

design solutions and the components, in order to guarantee the commissioning of a reliable system whilst reducing the costs and facilitating maintainability during its life cycle. Furthermore, a new analysis technique in the frequency domain has been implemented. In particular, by illustrating the theoretical foundations and specific numerical procedures, it can be proved how operatively assist the design of such instruments and improve the methodologies adopted for their characterisation. Actually, this approach not only provides a formalisation of the signal processing and an effective way of investigating the noise sources, but also makes it possible to accurately estimate the precision of the outcomes for any desired linear filtering. Thus, instead of performing extensive iterations of measurements and averaging, the validation of the conversion precision can now be replaced by a more efficient and flexible computation.

Finally, the device performance has been exhaustively validated both under the Section laboratory conditions and the radiation-controlled environments provided by the CERN X-Ray facilities and the proton accelerator at PSI (Villigen, Switzerland) [S6]. The results showed that all the design specifications had been fulfilled, which enabled us to overcome the prototyping stage successfully. Since the project challenges have involved, in an integrated fashion, pure hardware solutions, efficient processing tool developments and theoretical studies, the delivered product and the acquired expertise provide a class of accurate real-time measurement instruments. Their applications go beyond the specific commissioning target, contributing to the scientific enrichment of the topic.

1.2 Thesis Content Organisation

Before presenting the specific contents, we recall that part of the whole material has already been published in the papers [6], [7] and presented as posters at the *International Beam Instrumentation Conference* 2020 & 2021 by the author of this work.

In Chapter 2, the work begins with the architectural description of the state of the art BLM, the involved devices and the grounds justifying an update of the system. Thus, the design specifications for such renewal will be established, highlighting those demanding the development of improved electronics and processing. Later, integrated circuit solutions to perform the conversion of beam loss currents into digital data streams are introduced in Chapter 3. In particular, they are focused on the *BLMASIC* integrated circuit, which represents a key component for the new BLM front-end. The preliminary validation of its performance, by means of developed test-benches, and the data analysis software will be reported in Chapter 4. The theoretical explanations of the novelty signal processing methods and their prototype implementation are described in Chapter 5. The same includes examples on real measurements and the design of the filter bank. The development of the new radiation tolerant front-end is reported in Chapter 6. Its technical specifications, the electronic schematics and the *PCB* layout are there displayed in detail. Following the hardware implementation,

the back-end architecture and the firmware, which is in charge of carrying out the data acquisition and the diagnostic monitoring, are described in Chapter 7. It also includes the feasibility study of the novel signal processing implementation into an FPGA device, demonstrating the possibility of using a floating-point representation logic. Within Chapter 8, the overall system test-bench is presented together with the undertaken validation sessions. The data calibration procedures will be exposed and all the precision characterisation results will be reported and discussed. It will follow an analysis of the converter accuracy and the behaviour above the full-scale. The proton irradiation test carried out at PSI will be presented, showing the performance drifts due to the sustained TID and the single event effects (*SEE*) occurred in the front-end electronics. Conclusions of research will be discussed in the final Chapter 9. Finally, the appendices include: an introduction to the effects of radiation in electronics, technical details of the employed devices and the complete schematics and layout of the front-end board.

Chapter 2

System Architecture

An introduction to the CERN LHC accelerator is provided, including the description of its operational characteristics during the most recent runs. Then, the beam loss phenomenon and the risks that it poses to the plant are explained, highlighting the need for a monitoring system. The architecture of the operational BLM will be presented, including the details concerning the employed detectors, the digitisation electronics and the readout. Afterwards, the reasons for an upgrade of the electronics in view of HL-LHC are outlined, emphasising the necessity of developing a completely renewed front-end board, based on radiation-tolerant components. Finally, the technical specifications for the new system are reported.

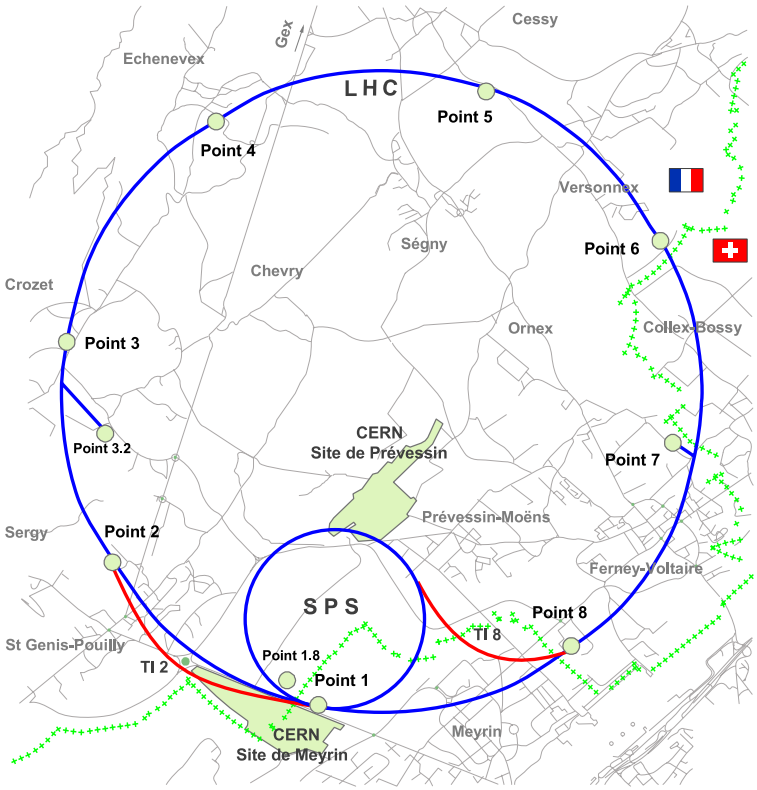

Figure 2.1: Geographical location of LHC and SPS ring accelerators.

2.1 Overview

Fig. 2.1 shows a map of the Franco-Swiss border area with the nearby city of Geneva. The ring paths of LHC and the linked accelerator SPS ($450GeV$) are represented in blue, while red is used for the injection transfer lines. Eight interaction points are present alongside the large accelerator, each including a variety of surface and underground infrastructures. The points P.1, P.2, P.5 and P.8 concern the main user experiments of LHC, which are named: *ATLAS* [S7], *ALICE* [S8], *CMS* [S9] and *LHCb* [S10], respectively. The main beam cleaning collimators are installed in P.3 and P.7, to reduce the off-momentum and the beam halo particles. Finally, the accelerating radio frequency cavities are installed in P.4 and the beam dumping system in P.6.

The LHC accelerator is made of eight arcs and eight insertions, which are 2.45 km and 545 m long, respectively, and it is extended for almost 27 km in total circumference. Each arc consists of 23 cells, which integrate a variety of magnets to bend and tune the geometry of the particle beams. On the other hand, the accelerator insertions are made of straight sections, which end with proper transition regions called dispersion suppressors. The layout of each insertion depends on the specific role that it plays in the system, such as experimental area access, injection, beam dumping or cleaning. The bending arc dipoles can currently produce a magnetic field of a nominal magnitude of 8.33 T, which is 10^5 times greater than the natural geomagnetism. Such a performance is only feasible by exploiting superconducting electromagnets of special Niobium-Titanium alloy coils, cooled down to the cryogenic temperature of $1.9K$.

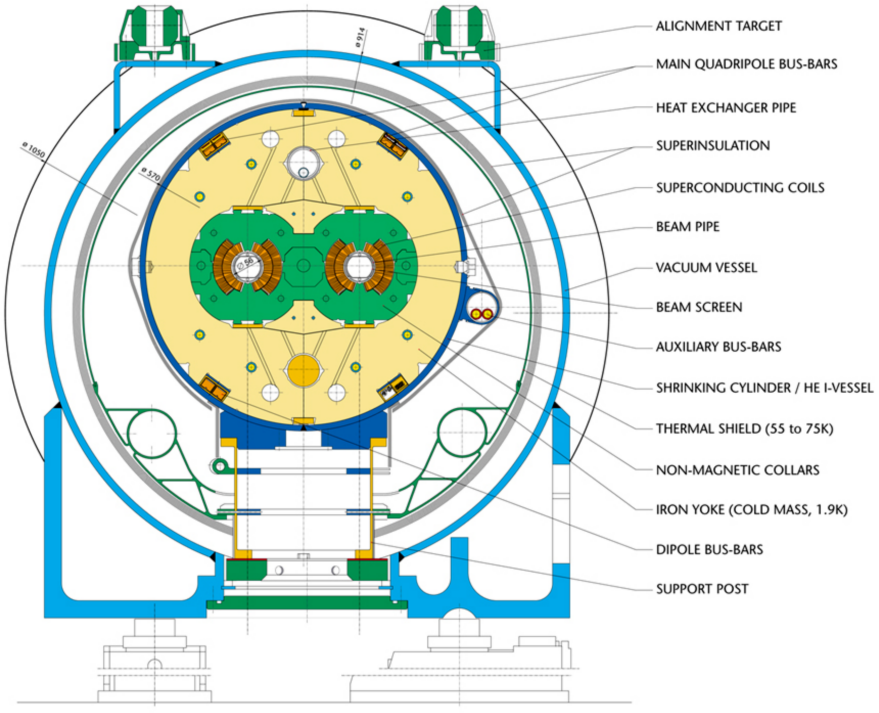

Figure 2.2: Cross section of an LHC dipole.

This minimises the power losses and allows current flows up to $11,850\ A$. In Fig. 2.2, it is shown the cross-section of an LHC dipole, where it is possible to distinguish the structural mechanics, the vacuum vessel, the two beam pipes surrounded by the superconductive coils and the related cooling plant.

The working parameters of LHC and the relative beam features during *Run 1* and *Run 2* are reported in Tab. 2.1 [8]. During the latest period, LHC operated at $6.5\ TeV$ with 2556 bunches of $\sim 1.15 \times 10^{11}$ protons for each of them. Thus, it achieved the peak luminosity of $2.01 \times 10^{34}\ cm^{-2}s^{-1}$. In these conditions, the kinetic energy of the beam alone is about $340\ MJ$, hence even minor particle losses may have a severe impact on the systems. The most representative example concerns the radiation-induced quenching of the superconductive magnets [9]–[13]. This occurs when a coil section mutates from the superconducting to the normal resistive state, due to the energy locally deposited by a beam loss event. The consequent Joule heating causes an abrupt temperature rise, which quickly extends the process to the surrounding areas and develops in a chain reaction. At the LHC nominal beam energy of $7\ TeV$, the deposited energy density of $1\ mJcm^{-3}$ is estimated as the threshold for the quenching process. Higher energy values, in the order of $1\ Jcm^{-3}$, can cause permanent damage to the magnets. While a quench leads to hours of accelerator downtime, the damage often requires delays of months and high repairing costs [14]. Therefore, the presence of instruments that promptly detects such a type of event and perform the right actions is mandatory in an accelerator of the LHC scale.

The BLM is the main system [15]–[17] in charge of preventing the quenches induced by events lasting from $356\ \mu s$, corresponding to 4 LHC bunch turns, up to $100s$. Namely, these are included in the categories: *fast*, *intermediate* and *slow* losses. The system must perform the loss measurements all along the beam path, with the positioning of detectors depending on the specific location. Areas with a higher expected dose rate, or

Parameter	**Design**	**2012**	**2016**	**2017**	**2018**
Beam Energy $[TeV]$	7	4	6.5	6.5	6.5
Bunch Spacing $[ns]$	25	50	25	25	25
$\beta^\star$ CMS/ATLAS $[cm]$	55	60	40	40 or 30	30-25
Crossing Angle $[\mu rad]$	285	290	370 or 280	300 - 240	320 - 260
Bunch Population N $\left[10^{11}ppb\right]$	1.15	1.65	1.1	1.15	1.15
Normalised Emittance ϵ $[\mu m]$	3.75	2.5	2.2	2.2	2.0
Number of Bunches per Ring k	2808	1374	2220	2556	2556
Peak Luminosity L $\left[10^{34}cm^{-2}s^{-1}\right]$	1	0.75	1.4	2.05	2.01
Peak Average Event Pile-Up μ	20	35	50	55	60
Peak Stored Energy $[MJ]$	360	145	270	320	340

Table 2.1: Comparison amongst LHC machine parameters across years.

involving more sensitive equipment, are monitored with denser detection spots. Some relevant examples are represented by: the beam cleaning collimators, the intersection points and the Triplet Magnets close to ATLAS and CMS experiments.

The massive hardware volume and a large amount of collected data by the BLM system make the accelerator operations dependent on its performance, including diagnostics not strictly circumscribed to the machine protection [18], [19]. Indeed, longevity and high reliability [20], [21] are required in the electronics and the firmware for detection and data processing, in order to ensure a failure rate below $10^{-7}/hour$ for each channel [22], [23]. Furthermore, proper long-term maintenance and calibration allow to successfully fulfil the goals of machine protection throughout its decades of operation [24]–[27]. Since the loss radiations must not exceed the limits in all the single detectors, careful adjustment of thresholds is mandatory and it is the outcome of detailed simulations and measurements [28]. If the system is configured with too demanding safety margins, it will dramatically affect the accelerator availability [29]. On the contrary, potentially dangerous situations could be missed if the loss protection is not strict enough. Generally, in the LHC, a relevant percentage of the interlock beam dumps are triggered by the BLM: about 18% during 2010/2011.

2.2 The BLM Detectors

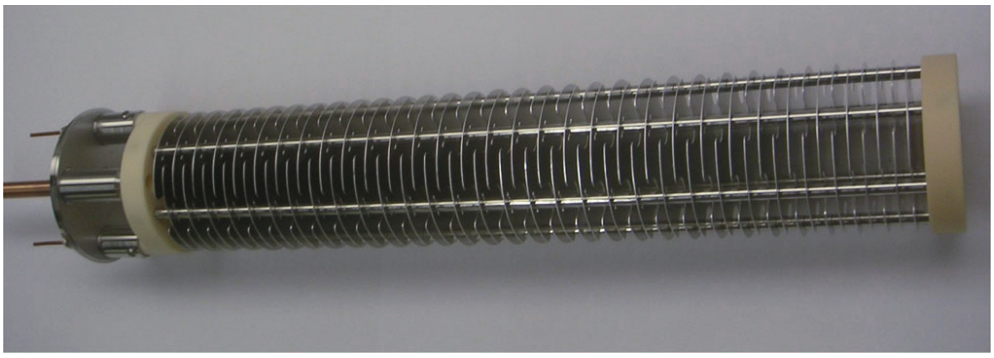

Figure 2.3: Interior view of an LHC ionisation chamber detector.

The main detector employed in the BLM system at CERN is the parallel plates gas Ionisation Chamber (IC, see Fig. 2.3 and the technical drawing in Appendix B). This is composed of a set of 61 aluminium electrodes, alternately connected (evens and odd) to two external nodes and spaced about 5.71 mm to 5.75 mm. Each electrode diameter is 9 cm and the total length of the set, including the mechanical structure, is about 50 cm. The assembly can be electrically modelled as the parallel of 60 flat-faced capacitors, for a total value of $\approx 500\ pF$ (measured at 1 kHz by an LC analyser). The whole is installed inside a stainless-steel tube, to which it is secured by means of two very high resistivity ceramics joints (Al_2O_3). Finally, the containing tube is filled with high-purity Nitrogen at 1.1 bar of pressure, before its definitive sealing.

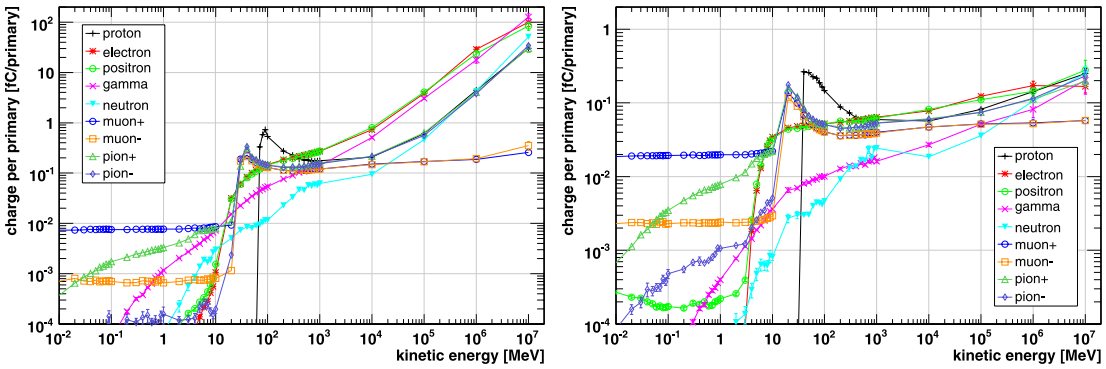

Figure 2.4: Ionisation chamber detector response to longitudinal (left) and transversal (right) interaction with different components of the radiation field.

During the measurements, a bias voltage of about 1500 V is applied to the detector nodes, establishing an electric field across the plates of about 262 kVm^{-1}. Whenever ionising radiation goes beyond the outer walls, it releases part of its energy to the gas atoms, producing ion-electron couples. Hence, the electric field accelerates the particles towards the opposite electrodes, which collect their charge. The required drift times for the electrons and ions are in the order of 300 ns and 80 μs, respectively.

Outside the sealed volume, a BNC output signal connector is joined to one of the detector nodes, allowing the charge to drain. The other node is used for the high-voltage biasing, which is provided by a second BNC HT connector and stabilised by an RC filter. This latter is composed of a resistor of 10 $M\Omega$ and a radiation-tolerant capacitor of 470 μF.

Depending on the monitoring location, the detectors are exposed to a wide range of radiations, with kinetic energies spread in the range of 10 keV to 10 TeV. In order to determine the chamber response, *GEANT4* [S11] simulations have been carried out with different particle types and incidence angles [30]. The relation between the particle energy and the amount of charge produced by the gas interaction is reported in Fig. 2.4. By comparing the two plots, it can be observed that the detector is more sensitive to

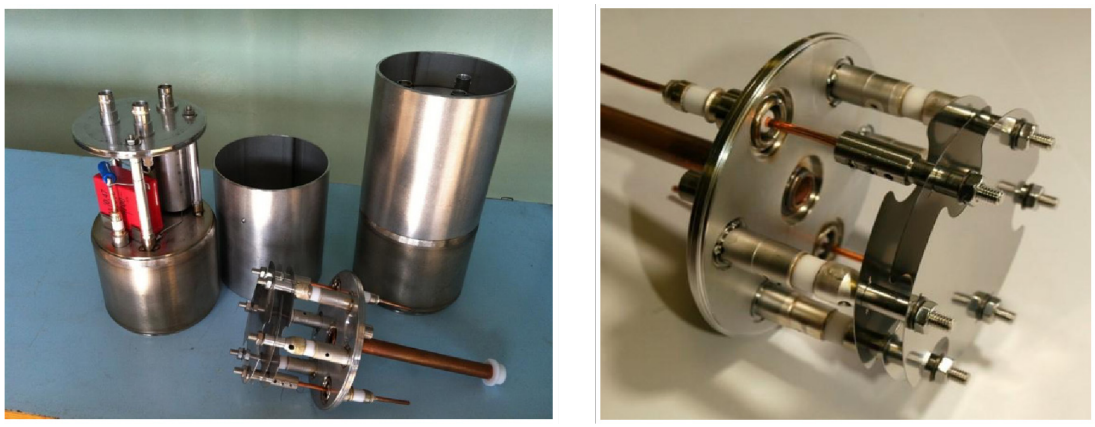

Figure 2.5: Assembly parts of a *LIC* (left) and a *SEM* (right) detectors.

high energy impacts if they occur longitudinally to its axis, because of the larger material traversed. In general, smaller detection cut-offs are reported in case of transversal interactions, which represent the most common case in the current BLM installations. In this situation, the study reported minimum observable energies of: 30 MeV for protons, 2 MeV for neutrons, 3 MeV for electrons and 300 keV for photons.

Finally, scaled detector geometries are employed in the accelerator areas characterised by very high dose rates (like injectors), to allow faster response times and lower saturation levels. In particular, the Secondary Emission Monitors (SEM) [31] and the Little Ionisation Chambers (LIC), shown in Fig. 2.5, have a structure resembling the standard IC one, but they include three electrodes only (made of Titanium for the SEM). Compared to the IC, the LIC is characterised by a lower sensitivity of a factor of 60, while the SEM is 7×10^{-4} times less, since it operates in the vacuum environment at 10^{-4} $mbar$ [32], [33].

2.3 State-of-the-Art Electronics

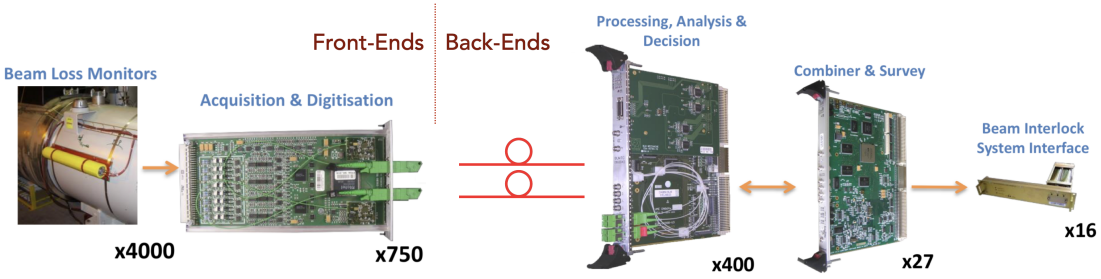

Figure 2.6: BLM signal chain within the LHC installation.

Fig. 2.6 shows a simplified view of the overall signals chain and the electronics that it is included in the BLM system. In particular, the components and the reported quantities refer to the plant that is currently operational in the LHC accelerator.

About 4000 detectors are installed throughout the beam path, producing current signals that are proportional to the loss radiation in each specific spot. A common example of their placement is shown in Fig. 2.7. Afterwards, a chain of coaxial and multi-conductor cables allows the propagation of the analogue signals toward the front-ends. In the ARCs, those are located underneath the quadrupole magnets, within dedicated compact crates. Instead, in the straight sections, they are concentrated at two specific positions at each of the LHC interaction regions due to the higher radiation levels. Finally, the readout electronics are installed at the surface facilities and they process exclusively digital signals.



Figure 2.7: Three IC detectors installed on an LHC quadrupole magnet. The other three are placed on the opposite cryostat side, to optimise the monitoring of both beams.

2.3.1 Digitisation Front-End

The front-end board [34], [35] provides all the circuitry aimed to carry out the digitisation of the current measurements and their transmission to the surface electronics via an optical link. In particular, the LHC installation employs the custom designed board named *BLECFC* (internal ref. EDA-00593). This device allows the interfacing to 8 IC detectors, the currents of which are concurrently acquired by input channels characterised by a high dynamic range.

The analogue circuitry employs commercial components and it is based on the current-to-frequency converter (CFC) architecture, the details of which will be set out in Par. 3.2.1. Due to the required Total Ionising Dose (TID) tolerance up to 500 Gy, every component has been extensively tested and their behaviour drifts have been taken into account during the development. Furthermore, to achieve a safety integrity level compliant with the *IEC 61508 SIL3* standard (10^{-8} to 10^{-7} $failures/h$) [22], several diagnostic features and circuit protections have been implemented.

The data processing and serialisation are performed by a flash-based Actel *A54SX72A* FPGA, embedded in the board. This hosts the firmware dedicated to: the retrieval of measuring currents, the monitoring of the IC biasing status, the generation of transmission packets and their related identification and validation keywords. These latter include: a unique code per card, which is used to ensure the data to detector assignments, a report of the internal logic status, an incremental counter value to identify the packet and a CRC code. The FPGA is interfaced with two redundant Gigabit Optical Links (GOL), equipped with a data serialiser and a laser driver. Thanks to these, the front-end is able to transmit the measurements to the readout electronics, which are installed at the surface facilities. The overall computation is synchronous to a system clock of 40 MHz and the measurement packets are propagated every 40 μs, which corresponds to an input sampling frequency of 25 kHz. Finally, the FPGA acquisition and control logics have been tripled, to guarantee reliability and redundancy.

2.3.2 Readout Electronics

The BLM readout electronics represents the decision unit of the entire machine protection system and, in the information chain, it sits between the raw loss measurements stream and the beam interlock interface. Two specific card types are involved in this process: 1) the *BLETC* or *DAB64x* card [36], [37], which performs the data acquisition and thresholding, 2) the *BLECS* (internal ref. EDA-01660), which combines the results and triggers the interlock signal. The installation of LHC includes about 400 BLETC and 27 BLECS, accommodated within VME64 based crates (ANSI/VITA 1-1994 standard), for a total of 25. Three of these are located in each of the interaction points except P.7, where an additional one allows better monitoring of the collimator plants. The crate is usually equipped with 16 acquisition cards, allowing the simultaneous tracking of 256 detectors, and a unique combiner that always occupies the rightmost slot. Furthermore, an off-the-shelf CPU card (PowerPC based) acts as the VME bus master and it is installed on the leftmost slot. It hosts a LynxOS real-time operating system that can be used for all the non-critical processes, such as the system configuration and data logging. This last task is crucial at every dump event, to find the triggering detector and to inquire which specific threshold has been exceeded. Thus, the collected information assists the operators investigating the causes of the event.

The following paragraphs will describe in more detail the tasks performed by the BLETC and BLECS devices.

The BLETC Processing Back-End

The processing of BLM data is assigned to the BLETC programmable-logic based card, which has been developed in cooperation with the whole beam instrumentation group at CERN (internal ref. EDA-00998). It exploits an Altera [S12] Stratix as the main computational device, an Altera MAX CPLD for the boot configuration and the VME bus management, and it is equipped with three 512 KiB SRAM memories. In addition, specifically designed for the BLM, a mezzanine card has been installed on it, enabling the reception of four gigabit optical links. Since double redundancy is employed for the data coming from the tunnel electronics, two front-ends can be connected to the BLETC.

The card hosts custom firmware, designated for the BLM operations [37], [38]. As the front-end packets are acquired by the FPGA interface, a transmission check module is used to validate their integrity and choose which stream will be employed for the computation between the redundant. It preliminary performs the 8b/10b decoding, then it computes the CRC codes and flags any detected errors. Finally, it compares the two CRC result values, to verify the equivalence of data. By means of these tests, the system can exploit the link redundancy to improve its availability and be protected against coder, transmission and decoder misbehaviours. Tab. 2.2 reports the actions programmed in the module according to the validation results of channels A and B.

Case	CRC Check A	CRC Check B	CRCs Comparison	Action
1	Error	Error	Mismatch	**Dump**
2	Error	Error	Equal	**Dump**
3	Error	OK	Mismatch	**Use B**
4	Error	OK	Equal	**Use B**
5	OK	Error	Mismatch	**Use A**
6	OK	Error	Equal	**Use A**
7	OK	OK	Mismatch	**Dump**
8	OK	OK	Equal	**Use A**

Table 2.2: Actions of the redundant data receiver.

Case 1 represents the worst possible since both the links report errors and the logic has no choice but to trigger a dump request. Case 2 and 7 are related to situations of disagreement between the individual checkers and this is likely due to errors in the CRC generation or validation. It is not possible to determine which of the signals contains the valid measurements, so the selection terminates again with a dump. In all the remaining cases (3, 4, 5, 6, 8), the module can make an aware choice of the signal to be actually used and this is propagated to the successive processing. However, while 8 is the ideal condition for everything operating correctly, those cases highlighted in yellow still involve faulty conditions, which will be signalled for subsequent maintenance.

Later, the correct signal is processed by a chain of digital modules aimed to retrieve the loss measurements from the front-end raw data. In fact, it should be pointed out that the digitisers described in Par. 2.3.1 do not produce numeric values which can be directly attributed to the input currents, as it would be usually expected by employing standard ADC devices. Instead, they generate modulated digital signals, from which the information of interest can be extracted if a suitable back-end processing is performed. This solution has been adopted to simplify the circuitry placed in the harsh radiation environment, improving the system reliability and enabling more flexible solutions in terms of response speed and precision characteristics.

The 25 kHz frequency at which the BLM measurements are updated allows the detection of fast loss events. Nevertheless, the monitoring of intermediate and slow losses is likewise essential for the complete protection of the accelerator machine. Hence, a repertoire of temporal measurement integrations is employed to keep track of the losses over different time scales.

The required computation is implemented into the FPGA as a set of 11 running sums. This is essentially a convenient method to realise digital moving average filters, the results of which lack the final division by the width of the window. In particular, the logics include: a shift register that is fed with the new samples, a subtractor operating on the newest and the oldest stored samples, and an accumulator. A cascading

organisation of the running sum logics allows the exploitation of shorter integration windows for the computation of the wider ones. This makes it possible not to store a huge amount of samples, which would otherwise be unachievable by the BLETC hardware, especially considering the longest integrations. However, such a solution has the disadvantage of featuring an inherent decimation of data, meaning that longer integration outputs will be updated with lower rates. This limitation is usually tolerable, because the filtering process still attenuates the high-frequency components of the measurement signals. Tab. 2.3 reports the variety of windows implemented, together with the respective output periods and tunings nomenclature.

By the study of how the running sums affect the stochastic processes of measurements, it is noticeable that longer integrations allow smaller errors and higher precision, at the expense of slower responses. In addition, such a computation carries a physical significance for the BLM system because the integral of the losses can be directly related to the energy deposited by them. In fact, the quench of magnets can be caused either by spikes of high magnitudes and short durations or by weak but persistent losses. The effects of signal processing on the measurements will be examined in detail in Chapter 5.

The latest module of the signal treatment chain ensures that the beam loss currents are constrained within proper configured boundaries [39], [40]. In particular, the output of each running sum is compared with a designated threshold and a dump request is triggered whenever one of these is exceeded. Finally, since the damaging effects of losses depend on the beam energy [19], different thresholding tables are employed by

Integration Window		**Output Periods**		**Signal Name**
Samples	Time $[ms]$	Samples	Time $[ms]$	
1	0.04	1	0.04	RS0
2	0.08	1	0.04	RS1
8	0.32	1	0.04	RS2
16	0.64	1	0.04	RS3
64	2.56	2	0.08	RS4
256	10.24	2	0.08	RS5
2048	81.92	64	2.56	RS6
8192	327.68	64	2.56	RS7
32768	1310.72	2048	81.92	RS8
131072	5242.88	2048	81.92	RS9
524288	20971.52	32768	1310.72	RS10
2097152	83886.08	32768	1310.72	RS11

Table 2.3: Processing integration time windows and periods of the corresponding decimated outputs.

the system, accordingly to the actual accelerator status. The tables are matched with 32 beam energy levels (from 0.45 to 7 TeV) and they are stored on non-volatile memory. The updating capability of this is provided either by the front panel connector or the crate CPU card. Concerning the diagnostics, the BLETC is also able to generate non-critical warnings and reports, in case a programmed percentage of the loss limits is breached.

The BLECS Combiner

Dump requests are emitted by the processing cards acting upon a specific digital flag, which is named *"beam permit"*. This is daisy chained through each BLETC by using a custom back-panel connector and a simple cascaded AND logic, which combines it with the flags locally generated by each processing card. In particular, they will be asserted if the loss measurements lie within the set constraints. Then, the propagated flag is finally collected by the combiner, representing the main interface between the BLM crate and the Beam Interlock Interface (BIS) [41], [42]. In case the *"beam permit"* is de-asserted, a dump of the beams is carried out within a few bunch turns.

Such an implemented architecture provides high protection reliability to the accelerator, prioritising the safe beam extraction even if faulty conditions occur in the BLM electronics. For instance, in case one or more cards is affected by a power supply issue, the daisy chain is cut and the *"beam permit"* is de-asserted consequently. The same situation applies whether the BLECS is faulty itself.

Finally, the BLECS card is also employed to distribute the beam energy information to all the BLETCs installed within the crate. These latter shall select their thresholding tables accordingly.

2.4 Front-End Upgrades

The digitiser card represents the most critical electronic component of the whole BLM system. Their development involved significant challenges concerning conversion performance, reliability qualifications and compliance with the harsh working environment. In fact, the high radiation dose rate and the TID, which the device must withstand during its lifecycle, are definitely the key factors to be taken into account for the success of the project

In the operating system, the BLECFC has performed well on the goals, however precautionary measures have been demanded to protect the electronics against excessive radiation. In particular, in the proximity of the intersection points (P.1, P.2, P.5, P.8) and of the main cleaning collimators (P.3, P.7), the digitisers have been concentrated in crates located far from the beam pipe. This has allowed not exceeding the BLECFC specification margins, but it has required the employment of long analogue cables (up to $\sim 800\ m$) to connect the IC detectors.

In LHC, the beam loss magnitudes are spread over a very high dynamic range, as well as the detector currents that the BLM system has to measure. The ICs produce currents below 10 pA as a result of the weakest losses, but values above 1 mA can be measured for the most intense ones. This implies that the electronics must operate over an input dynamic range of more than 160 dB, while ensuring the reliability of measurements and accuracy. While higher currents are easily distinguished by the operational BLM complex, weak ones often have the same power as the background noise and their detection proves difficult. This is accentuated by the fact that the LHC environment includes high-power equipment producing strong electromagnetic interference (EMI), which may influence the BLM operations. Thus, all the analogue transmissions should be limited to short distances, in order to avoid degradation of the characteristics. Furthermore, signal transient issues must also be taken into account since cables introduce an intrinsic attenuation and a low-pass filtering effect, increasing with the transmission length. This can reduce the BLM system response speed, thus inhibiting its ability to detect fast losses.

On account of the stated issues, the first goal of this work is the development of an upgraded front-end, named BLEIC, which allows the digitisation of the detector currents as close as possible to the measuring spots. Hence, the most demanding challenge is posed by the enhancement of the card radiation tolerance specifications. In addition, the new system will operate in the HL-LHC environment, for which even higher radiation levels are expected in comparison to the former accelerator, due to the increased beam luminosity [S5]. Thus, for the new digitiser, a TID increased by a factor > 6 has been assessed (3 kGy), to cope with such new requirements and maintain the same operational lifetime expectation at 20 years.

The HL-LHC upgrade will also involve the replacement of critical equipment, relying on the BLM protection against excessive loss events, as the new triplet magnets of ATLAS and CMS experiments. Thus, a fast measurement response speed will play a key role in the ability to promptly detect potentially dangerous situations. Indeed, design efforts have been made in the whole detection chain to minimise the latency between the loss events and the interlock triggering. Amongst these, it has been decided to increase the digitisers sampling frequency by a factor of 4, producing new samples every 10μs ($f_s = 100\ kHz$).

A market survey, carried out during the preliminary development stages, concluded that such requirements are not achievable by using commercial products, so it is necessary to deploy specific micro-electronic components for the purpose. This approach is widely adopted at CERN (as in similar applications [43]) and different devices are already available from the internal supply, to address a variety of needs. This is the case of the low-power Giga Bit Transceiver (lpGBT) [44], [S13], which will be employed to handle the communication with the surface electronics and the linPOL12V device, used for the power regulation. However, an integrated current digitiser compliant with the required range was not yet produced, so it has been developed and validated in

the context of this work. It is named *BLMASIC* and it will represent the key component of the newer front-ends.

The ASIC-based architecture, planned for this upgrade, will allow a dramatic reduction in the number of components populating the card. This translates into benefits concerning reliability, power consumption and costs. Furthermore, in-house development of the integrated circuits enables more effective planning of the spare procurement and device obsolescence, which represents a potential issue for the future maintenance of the BLECFC.

2.5 Back-End Upgrades

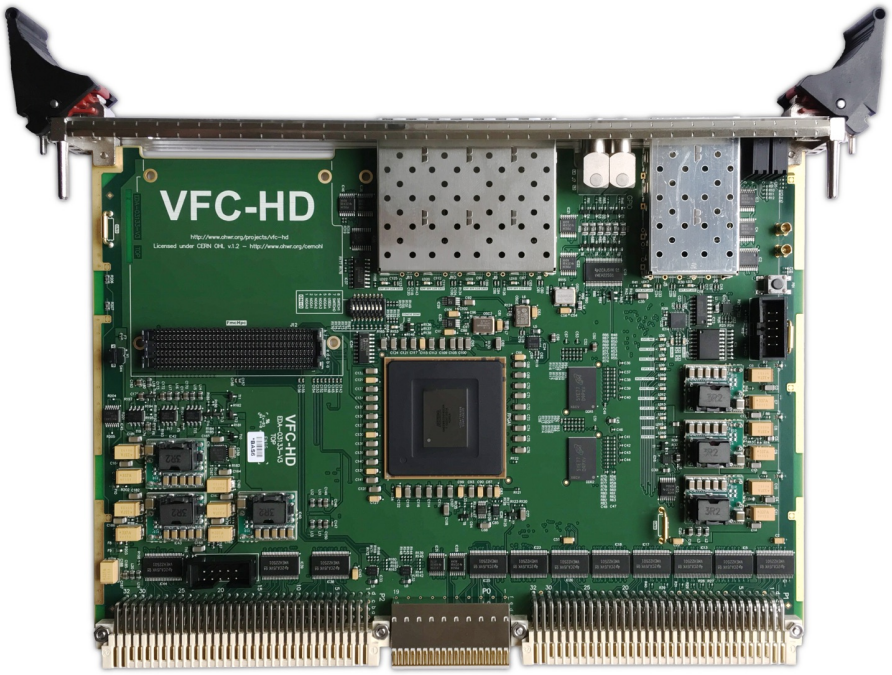

Figure 2.8: The VFC-HD carrier board.

In the context of this work, a simultaneous improvement of the readout electronics has been carried out in order to accommodate the new BLM features. The BLETC card was developed in 2006, exploiting a 130 *nm* Stratix FPGA technology, which was declared obsolete by the manufacturer in 2017. During the various updates of the firmware, it has been equipped with more sophisticated functionalities. In particular, during the latest long shutdown (*LS2*), considerable efforts have been spent by the developers, who pushed the device logic up to its limits. Thus, the available resources have been fully exploited during the operation of the BLM system in LHC and further improvements can only be achieved by a major upgrade of the hardware components.

In view of the continuous evolution of the beam instrumentation, the group has begun the design of a new standardised readout card for some time already. After several prototypes testing, the achieved maturity has allowed to release a newer VME64x based board: the *VFC-HD* [45], [S14] shown in Fig. 2.8, which includes considerable advances in comparison to the previous solution. The board is equipped with an Arria V GX FPGA as the main computation device. It provides a number of logic array blocks (LABs) 6 times larger than the BLETC counterpart and an embedded memory that is 10 times bigger. Furthermore, the extended provision of 6.55 *Gbps* transceivers will allow faster communication with the upgraded front-ends (rate > 5) and the integration of a Gigabit Ethernet interface, which can be employed in test-benches for the acquisition and control. Finally, a 32 *bit* DDR3 memory is included, allowing a total storage capability of 2 *GB*.

The tasks performed by the new firmware are inspired by those already implemented in the BLETC (cf. Par. 2.3.2). However, each digital module will be optimised for the new hosting FPGA device, allowing for manage the higher measurement sampling rate and faster communication links. In this work, a significant amount of the low-level code has been rewritten or extended with updated libraries to cope with the new hardware resources, the BLMASIC features, the novel signal processing and the improved diagnostics provided by the front-ends. A planned firmware upgrade will also enable backward compatibility with the previous system, implementing a transparent commutation between the logic necessary to interface either the BLECFC or the newer BLEIC boards. Amongst the upgrades, the running sums technique (see Par. 2.3.2) will be compared with a more advanced Infinite Impulse Response (IIR) filtering, evaluating its implementation feasibility and which methodology maximises the measurement precision performance.

Finally, in the present work, additional firmware modules have been developed and employed for the system test-bench to certify the specification compliance, before proceeding with the commissioning into the accelerators. Including specific acquisition and profiling features, these modules are able to achieve a complete front-end conversion characterisation, data communication integrity checks and the signal processing validation. These analyses will be carried out both in the section laboratories and in the electronic irradiation facilities.

2.6 Technical Specifications

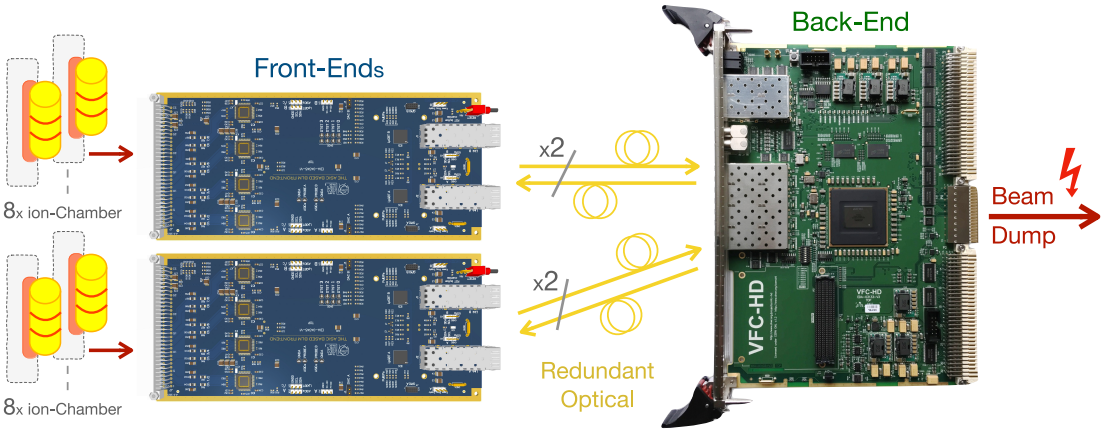

Figure 2.9: Unit of the new BLM electronics, composed of a single VFC-HD back-end and a pair of BLEIC front-ends.

The key technical specifications issued for the BLM electronics upgrade are reported in Tab. 2.4, focusing on the requirements dealt with in this work. These are referred to as a Unit, including the single back-end VFC-HD board that acquires the data streams from two independent BLEIC front-ends, as shown in Fig. 2.9. Clearly, this architecture will be replicated in order to build the whole BLM plant. Additional VFC-HD specifications can be found at the designated repository [S14].

The radiation tolerance features are referred to the BLEIC cards only since the back-end is not affected by this issue, being installed within the surface crates. The indicated TID value is obtained by measurements carried out on the LHC accelerator and estimations of the dose absorbed at HL-LHC during the device lifetime [S5]. Coherently with the summarising table reported in Fig. 2.10, the required electronics are classified among the radiation-hardened one.

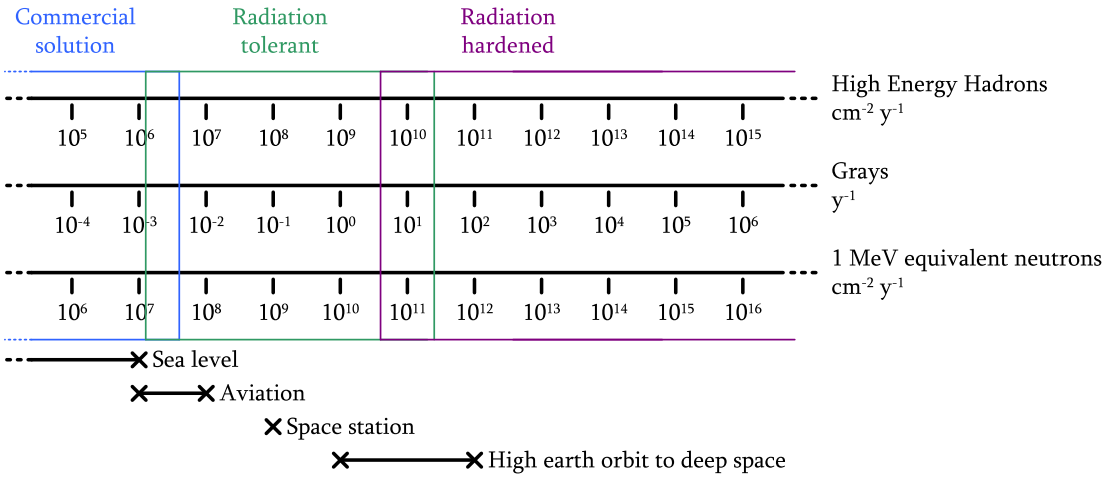

Figure 2.10: Radiation tolerance classes for electronic equipment.

Feature	**Value**
Overall input channels	16
Input dynamic range	$> 160\ dB$
Minimum target accuracy	10 %
Operative input range	$< 10\ pA$ — $1\ mA$
Achievable current resolution	$< 10\ pA$
Maximum input current	$5\ mA$
Current monitoring sampling rate	$100\ kHz$
Minimum response time	$10\ \mu s$
Target input detector	Ionisation Chamber
Maximum input source impedance	$5\ nF$
Front-End radiation TID tolerance	$> 3\ kGy$
Front-End radiation-induced fault tolerance	SEU resilient
Deviation of measurements due to TID	< 10 %
Electrical input protection	against ESD & Overload
Data communication medium	Single-mode optical fibre
Data communication redundancy	Double

Table 2.4: Key specifications of the upgraded BLM electronics, referred to the single Unit represented in Fig. 2.9.

Chapter 3

ASIC Converters

Electronic solutions for the digitisation of loss detector signals are presented, focusing on the current to frequency converter and the proposed $\Delta\Sigma$, which represents an innovative solution compared to state of the art in the considered context. The theory of their functioning principles is illustrated and the characteristics of the circuit integration are reported. Thus, it will be analysed the structure of the BLMASIC converter, which represents the key component of the novel BLM system. Details concerning its internal analogue and digital circuitry are discussed, providing an appropriate background for a better understanding of the project. This includes the techniques exploited to engineer a radiation-tolerant device.

3.1 Current Measurements

The key component of the BLM front-end electronics is represented by the ASIC converter. It includes the circuitry necessary to sample the currents generated by the IC detectors and yield digital data streams. By appropriate processing of the latter, an estimation of the currents is obtained. In fact, the converter represents the front-line interface between the physical quantity to be measured and the rest of the system, where the data transmission and required computations are achieved by employing exclusively digital devices.

During the development of the converter, the first inquiry concerned its digitisation architecture. Literature reports on a variety of possible circuit topologies to perform the measurement of currents [46]–[48]. Typical methods involve using a trans-impedance network, the output of which is monitored by a voltage ADC. For example, it can be composed by a precision shunt resistor and a possible active amplification (see Fig.s 3.1a and 3.1b). Differently, a canonical trans-impedance configuration based on an operational amplifier can be used (see Fig. 3.1c). On the other hand, the integrated ADC architecture must be chosen on the basis of the desired resolution and sampling rate. In particular, the latter must be high enough, to allow an effective observation of the input current variations, in agreement with the Nyquist-Shannon criterion.

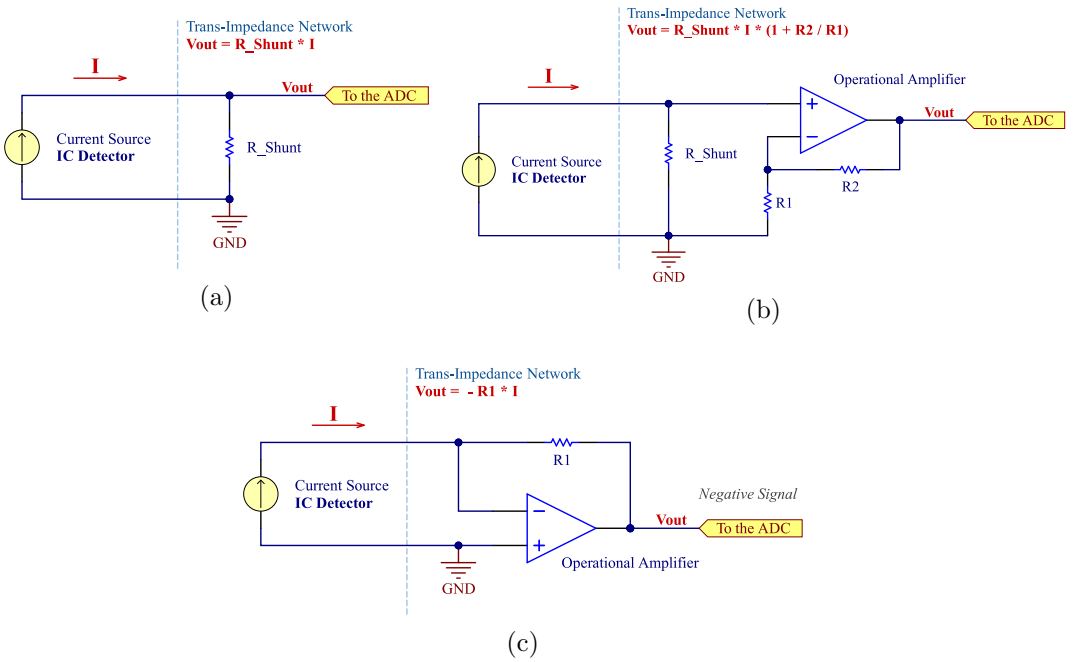

Figure 3.1: Simplified examples of trans-impedance networks.

However, since the detector currents in the BLM system are tracked in specified integration time windows, the conversion sampling rate is not strictly constrained to the bandwidth of signals, but accordingly to the time constants characterising the effects of the radiation losses. Indeed, the particle loss phenomena and the related detector currents may evolve in a much shorter timescale than the damages they cause by depositing energy into the magnets and the surrounding equipment. For instance, the IC detector time constant is about 300 ns [30], which is thirty times smaller than the sampling period stated in the specifications for the new system. Moreover, that is up to one billionth of the time required by a *slow* loss to pose a threat to the accelerator. Therefore, the converter architecture choice must be focused on the ability to carry out precise and continuous measurements of averaged input currents. Alternatively, this can be perceived as a measurement of the charge injected by the detector, modulo an integration time scale constant.

The presented architectures are not the most effective solutions for the IC-based BLM system. In fact, the full bandwidth current signal should be converted into digital, before carrying out an appropriate estimation of the injected charge. This requires a fast ADC, implying high power consumption and difficulties in the micro-electronic implementation. Instead, it is convenient to employ a conversion architecture, which performs the integration of input currents by its analogue circuitry. This enables the possibility of choosing the desired sampling rate within fairly loose margins, while ensuring that no input charge is lost.

3.2 Conversion Architectures

Two circuit architectures have been investigated and implemented during this work: the Current to Frequency Converter (CFC, employed in several CERN applications [35], [43], [49]–[51]) and the $\Delta\Sigma$ Converter [52], [53]. The working principle of such devices will be described in the following paragraphs.

3.2.1 The Current to Frequency Converter (CFC)

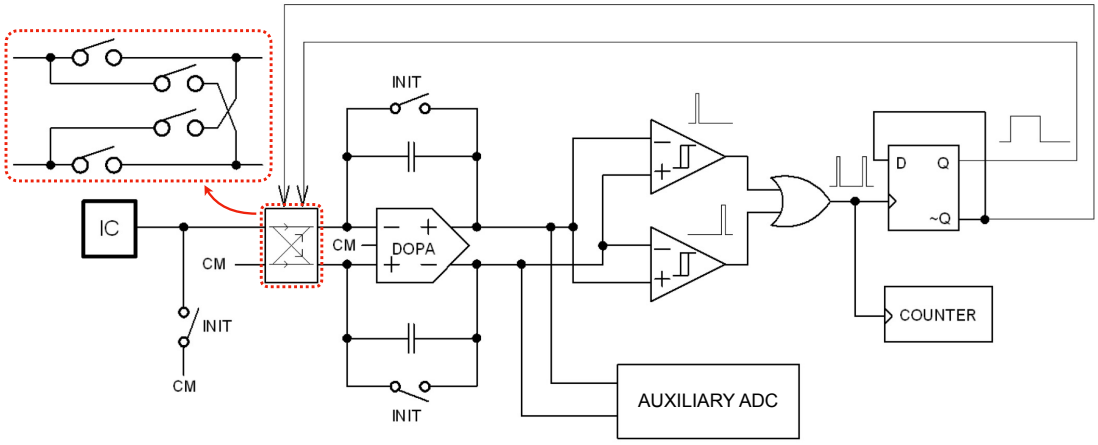

Figure 3.2: Simplified schematic of the CFC.

The CFC architecture is aimed to generate a periodic signal, the frequency of which is linearly related to the measuring input currents. The signal waveform is shaped in order to simplify the frequency tracking procedure. Fig. 3.2 shows a simplified schematic of the circuitry employed in the project, the described components of which are boldface denoted in the text. The **IC** block on the left represents the ionisation chamber detector, which can be ideally modelled as an independent current source. The **DOPA** is the input fully-differential amplifier, which is configured as a current integrator by the addition of two feedback capacitors. This is connected to the source through a changeover circuit, which provides the possibility of inverting the current polarity with respect to the integrator inputs. A couple of quantisers, a toggling flip-flop and a counter synchronous to it are shown on the right.

After the initialisation of the circuit, closing the **INIT** switches, any charge produced by the IC will be accumulated into the capacitors, producing a voltage swing at the amplifier outputs. Assuming that the changeover is injecting the current flowing towards the non-inverting terminal of the amplifier, the following equation holds

$$v_{out}(t) = \frac{2}{C} \int_{-\infty}^{t} i_{in}(\tau)\ \mathrm{d}\tau\ , \tag{3.1}$$

where $v_{out}(t)$ is the differential voltage across the positive to the negative output terminals of the integrator, $i_{in}(t)$ is the input current and C is the value of each of the

two capacitors. Factor 2 is due to the connection of the second changeover input to the output common mode reference voltage (**CM**). In fact, the same current flows through the capacitors in opposite directions, causing a doubled swing at the differential output. The **CM** can be considered as a virtual ground reference of the network and it is set to half of the supply voltages. In the studied circuit, these latter are set to V_{supply} and GND, exploiting a single powering bus.

The integrator output is later compared with two predefined thresholds, by using differential quantisers with hysteresis. These produce sharp pulses indicating whether the upper or the lower limits have been exceeded, respectively. Assuming that the **DOPA** is a Rail-to-Rail output amplifier, it is possible to maximise the swing of the integrator by setting the upper limit to a value close to the positive supply voltage and the lower one close to its opposite. With respect to a reference supply rail, the threshold values are set during the design, ensuring a certain margin for the hysteresis mechanism to function properly. The produced pulses are then combined with an OR gate and used to feed the clock of the flip-flop. This latter drives the changeover switches thanks to a feedback network, allowing to invert the input currents whenever a limit is reached.

To understand the dynamic behaviour of the system, it can be assumed that a continuous (DC) current is injected at the input of the circuit. According to the equation (3.1), the integrated current will cause a voltage ramp at the amplifier output. As soon as the voltage will reach one of the limits, the flip-flop is toggled and, consequently, the input polarity is reversed. This leads to the inversion of the ramp slope until the opposite limit is reached. For ideal circuits, such a mechanism is carried on iteratively, generating a triangular-shaped waveform at the output of the integrating amplifier. Examples of a real oscilloscope trace, produced by such a circuit, will be shown in Chapter 4 (cf. Fig. 4.14).

As a consequence of the considerations made above, the estimation

$$T_{rise} \approx T_{fall} = \frac{C\,V_{swing}}{2\,\overline{i_{in}}} \tag{3.2}$$

holds, being T_{rise} and T_{fall} the times required to increase the output from the lower to the upper thresholds and vice-versa, V_{swing} is the total voltage span between the thresholds and, finally, $\overline{i_{in}}$ represents the average current, injected into the circuit during such rising and falling periods. The waveform frequency is expressed by

$$f = \frac{1}{T_{rise} + T_{fall}}\,, \tag{3.3}$$

namely

$$f = \frac{\overline{i_{in}}}{C\,V_{swing}}\,, \tag{3.4}$$

saying that it is linearly proportional to the input current intensity.

Therefore, by knowing the circuit parameters and tracking the waveform frequency, it is possible to estimate a mean value of the currents injected during the observation time. If the thresholds are approximated to V_{supply} and GND, respectively, according to the considerations made above and using $V_{swing} = 2\,V_{supply}$, it holds

$$\overline{i_{in}} = 2\,f\,C\,V_{supply}\ . \tag{3.5}$$

The real device may actually present a slightly smaller V_{swing}, which will be described in detail in the implementation of Par. 3.3.5.

The employment of **COUNTER** allows the waveform frequency measurement. In fact, it is incremented by the quantiser pulses whenever one of the thresholds is crossed. The readout logic will query the counter values and perform the differences between two successive readings. With such an architecture, a good frequency estimation can be carried out whether the observation time is much larger than the waveform period. Furthermore, the circuit dimensioning can be optimised to guarantee an effective tracking of the input current variations. Indeed, the waveform oscillation frequencies should be much larger than the bandwidth of the input current signals.

A critical issue of such an architecture concerns the measurement of very weak currents. In fact, these are usually associated with long rising and falling times, with the consequent generation of very low frequency triangular waveforms. Despite the circuit can still operate correctly, as long as its intrinsic electronic noise is lower than the measuring current, a wider observation interval will be necessary to carry out the measurement with adequate precision. Furthermore, due to the gentle integrator slope, a long time is required before **COUNTER** increments and a non-zero value will be produced by the readout logic. Thus, it proves difficult to quickly ensure the correct operational status of the system.

To overcome the present limitations, the **AUXILIARY ADC** device can be integrated into the circuit, to gather samples of the integrator output. A real-time processing can after combine such information with the counter differences, to produce accurate results more quickly [54]. The ADC sampling frequency can be tuned accordingly to the minimum desired measurement rate.

3.2.2 The $\Delta\Sigma$ Converter

The $\Delta\Sigma$ converter shares a rather similar integrating architecture with the already presented CFC (Par. 3.2.1), however, it is characterised by a different operation theory. Looking at Fig. 3.3, it is possible to notice the addition of an internal reference current source, which can be connected to the integrator in parallel to the input. The injection is managed by a 2-way switch, which is driven by the feedback path, originating from the digital output bitstream. Unlike the CFC, the input currents are never reversed and a synchronous architecture is employed for the quantiser. Hence, the feedback control is time constrained by the system clock.

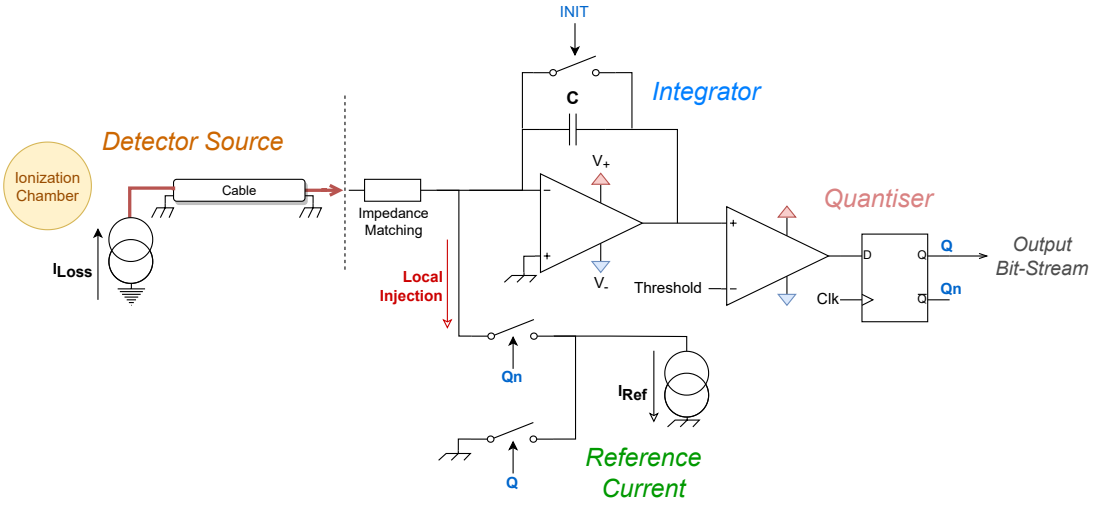

Figure 3.3: Simplified schematic of the $\Delta\Sigma$ modulator.

Assuming that a DC current is injected at the input and the internal source is disconnected, a voltage ramp is produced at the integrator output. Whether, in the event of a rising clock edge, the ramp has exceeded the quantiser threshold, this asserts its output **Q** synchronously. Thus, the feedback commands the injection of the internal source, the current of which is now subtracted from the input one due to its opposite direction. Consequently, the integrator capacitance is discharged and the ramp is inverted until it lies below the threshold (**Q** is de-asserted). The process is iterated similarly to what occurs for the CFC, despite the rising and falling ramp slopes are usually not symmetrical since unequal currents are integrated within the two phases. It can be proven that, established the steady state conditions, the charge drawn at the device input is equal to the one injected by the internal source in a certain observation time. Thus, the estimation of the averaged input current can be carried out by knowing: 1) the system clock period, 2) the integrating capacitance, 3) the observation time window, 4) the internal reference current and 5) the number of clock cycles during which the output bitstream was asserted. In Fig. 3.4, it is possible to see two examples of output bitstream patterns as a consequence of test input signals.

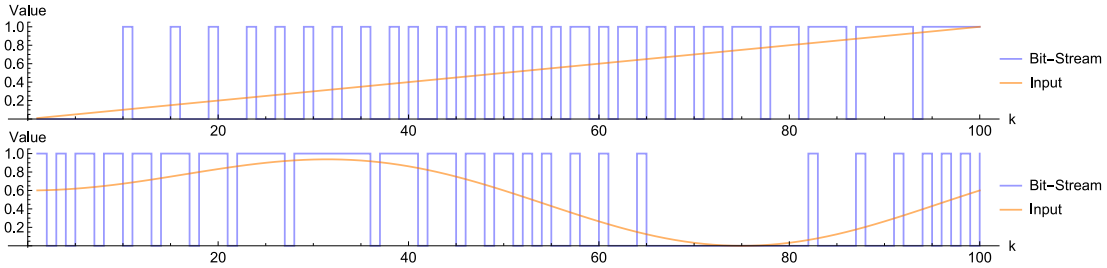

Figure 3.4: Two examples of $\Delta\Sigma$ output bitstreams superimposed to the input signals. The k axis refers to time samples in an arbitrary unit, while the signals are normalised.

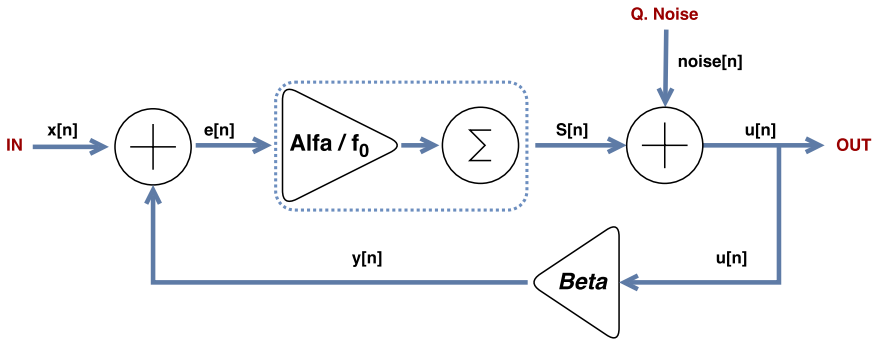

Figure 3.5: System model of a $\Delta\Sigma$ modulator.

A system-level model of the $\Delta\Sigma$ converter is reported in the Fig. 3.5, where the time continuous quantities have been replaced by their time sampled sequences: x_n is the input current, y_n is the internal current source, S_n is the integrator output and u_n is the converter output. The discrete time approximation for x_n and S_n is valid under the assumption that their bandwidth is much smaller than the frequency f_0, at which all the model signals are sampled. Such a hypothesis is generally fulfilled since the $\Delta\Sigma$ operates in compliance with the principle of oversampling. Furthermore, the quantisation introduced by the threshold comparator is modelled as a noise source $noise_n$, to be added to the integrator output. It actually represents the deviation between the integrator (S_n) and the whole system (u_n) outputs, which corresponds to the non-normalised bitstream. The following set of difference equations holds:

$$S_n = S_{n-1} + \frac{\alpha}{f_0}\, e_n \ , \tag{3.6}$$

$$e_n = x_n + \beta\, u_n \ , \tag{3.7}$$

$$u_n = S_n + noise_n \ . \tag{3.8}$$

Where the parameters α and β are proportional to the integrating capacitance value and the internal source current, respectively. Furthermore, note that the feedback gain β is dimensionless, while the integrator gain α has here the dimensions of a frequency. It is possible to obtain a single difference equation describing the system transfer function. In fact, by substituting the (3.7) and the (3.8) into the (3.6), one yields

$$\begin{aligned} u_n - noise_n &= u_{n-1} - noise_{n-1} + \frac{\alpha}{f_0}\,(x_n + \beta\, u_n) \\ &\Updownarrow \\ u_n\,(1 - \frac{\alpha\beta}{f_0}) - u_{n-1} &= \frac{\alpha}{f_0}\, x_n + noise_n - noise_{n-1} \ . \end{aligned} \tag{3.9}$$

The Eq. (3.9) can also be formulated in its z-Transform [55], [56] expression

$$u_n\,(1 - \frac{\alpha\beta}{f_0} - z^{-1}) = \frac{\alpha}{f_0}\, x_n + noise_n\,(1 - z^{-1}) \ . \tag{3.10}$$

The contributions of the input x_n and the quantiser $noise_n$ on the output u_n can be studied separately. In particular, one has

$$\begin{aligned} u_n^{input} &= \frac{\frac{\alpha}{f_0}}{1 - \frac{\alpha\beta}{f_0} - z^{-1}}\, x_n \implies \\ u_n^{input} &= \frac{\frac{\alpha}{f_0}\, z}{(1 - \frac{\alpha\beta}{f_0})\, z - 1}\, x_n \ , \end{aligned} \tag{3.11}$$

$$\begin{aligned} u_n^{noise} &= \frac{1 - z^{-1}}{1 - \frac{\alpha\beta}{f_0} - z^{-1}}\, noise_n \implies \\ u_n^{noise} &= \frac{z - 1}{(1 - \frac{\alpha\beta}{f_0})\, z - 1}\, noise_n \ . \end{aligned} \tag{3.12}$$

The equation (3.11) has a zero in the origin of the z plane and the (3.12) vanishes for $z = 1$, while both share the same real pole at $z_p = (1 - \frac{\alpha\beta}{f_0})^{-1}$. In order to guarantee the system stability, this latter must satisfy $-1 < z_p < 1$, constraining the parameters α and β to the conditions:

$$\alpha\beta < 0 \quad \vee \quad \alpha\beta > 2\, f_0 \ . \tag{3.13}$$

The first alternative is actually chosen for the BLM converter implementation. In fact, it yields negative feedback closed loop and ensures that the magnitude of the u_n^{input} transfer function is monotonically decreasing, with its maximum in the DC component. On the other hand, the u_n^{noise} transfer function has the feature of a high-pass filter. This allows a *shaping* procedure to be performed, which shifts the $\Delta\Sigma$ quantisation noise towards the high frequencies. In this way, the noise spectrum is separated from that of the measured signal and it can easily be suppressed by adopting successive filtering.

In Fig. 3.6 is reported an example of the magnitude and the phase plots versus the frequency for the transfer functions u^{input}/x and $u^{noise}/noise$. They have been obtained by applying the fundamental relation $z = e^{\frac{s}{f_0}}$ and varying s over the imaginary axis $s \to i2\pi f$. As system parameters, the following values have been chosen: $\alpha = 10^8$, $\beta = -1$ and $f_0 = 10^8$. In Fig. 3.7 the same features are shown as a function of z within its unit circle.

In conclusion, the described architecture exploits a coarse quantiser that, operating in over-sampling conditions, allows to carry out high-resolution measurements of comparatively slow varying signals. These qualitative considerations have to be quantified during the design process, in order to determine the suitable system parameters according to the requirements.

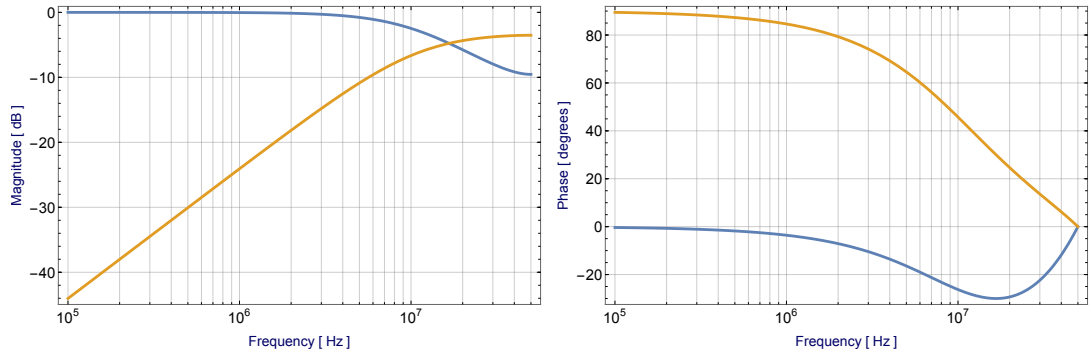

(a) Magnitude Plot (b) Phase Plot

Figure 3.6: Example of transfer functions u^{input}/x (in blue) and $u^{noise}/noise$ (in orange).

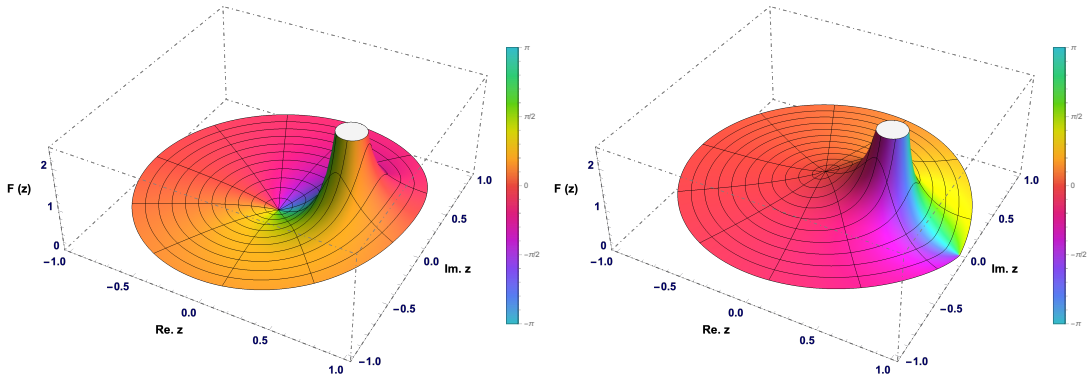

Figure 3.7: Example of transfer functions u^{input}/x (left) and $u^{noise}/noise$ (right) within the z unit circle.

3.3 Converter Integration

The BLMASIC, the implementation of the converter as an integrated circuit device, has been the result of the cooperation between the Beam Loss (*SY-BI-BL* [S4]) and the Micro Electronics (*EP-ESE-ME* [S15]) sections at CERN. During the design planning stages, it has been decided to proceed with the implementation of twin prototypes: one is based on the CFC and the other on the $\Delta\Sigma$ converter. A prototype of the former was already investigated in early implementations [57]–[60] employing a 250 nm CMOS technology and it was compared with the solutions based on discrete components [61]. Although the performance was promising, additional manufacturing iterations should have been required to solve some encountered issues, which led to delays. In the meanwhile, the newer 130 nm CMOS technology has reported good performance in similar applications [43] at CERN and it has become the most widely supported. Thus, it was decided to combine the valuable experience gained in the previous design with such a technology to develop a newer ASIC.

It was studied that either the proposed architectures could have represented a viable candidate for the goals and the selection of the optimal solution was postponed until after the characterisation of real components. Furthermore, the computer-aided

simulations do not permit an exhaustive description of the behaviour of circuits and specific assessments, like the radiation tolerances [62], [63], must be verified after the production of devices.

In the following paragraphs are reported the BLMASIC key specifications and the implementation details of its main architectural elements. Unless otherwise specified, the description will be focused on the CFC device, without loss of generality. Indeed, the two prototypes share the vast majority of their internal components.

3.3.1 Specifications

The converter technical specifications have been carefully tailored for the BLM purposes, accordingly to what was presented in Par. 2.6. Furthermore, large design margins were taken into account where feasible to ensure improved device reliability, allow compatibility with future system upgrades and pave the way for additional research-oriented applications. The key requirements are summarises in Tab. 3.1.

Feature	**Value**
Converter input channels	2
Measurement sampling rate	100 kHz
Input dynamic range	up to 180 dB
Full scale input current	1 mA
Maximum input current	5 mA
Achievable current resolution	< 10 pA
Tolerable internal leakage	< 5 pA
Maximum input source impedance	5 nF
System clock frequency	80 MHz
Logic high reliability techniques	Triple redundancy
Data transmission interface	CLPS E-Link compatible
Data transmission bitrate	80 $Mbps$
Data transmission encoding	8b/10b
Configuration interface	I^2C
High reliability techniques of interfaces	Double redundancy
Radiation TID tolerance	> 3 kGy
Deviation of measurements due to TID	< 5 %
Radiation-induced fault tolerance	SEU resilient
Protection of electrical inputs	against ESD & Overload
Manufacturing technology	TSMC CMOS 130 nm

Table 3.1: BLMASIC converter key specifications.

3.3.2 ASIC Structure

The BLMASIC constitutes a mixed-signal device, which includes the conversion circuitry, the digital readout, the control and the data transmission logic. It has been based on a 130 nm TSMC [S16] lithography process and manufacturing optimisations have been widely employed to cope with its use in a critical system. In particular:

- the best production plant for radiation tolerant transistors has been chosen on the basis of previous experiences;
- optimised layouts, has been adopted for the critical transistors, allowing a better radiation tolerance [64] and reducing the current leakage [65];
- independent power rails have been used for the analogue and digital sections;
- a silicon substrate (deep N-Well) has been placed underneath the digital circuit to improve the isolation of the very sensitive analogue components;
- the analogue quantisers have been equipped with triple redundancy to minimise the occurrence of spurious transitions;
- the control logics and configuration registers have been tripled and high-reliability voters have been included;
- the data transmission interfaces and I^2C configuration ports have been doubled in two redundant buses.

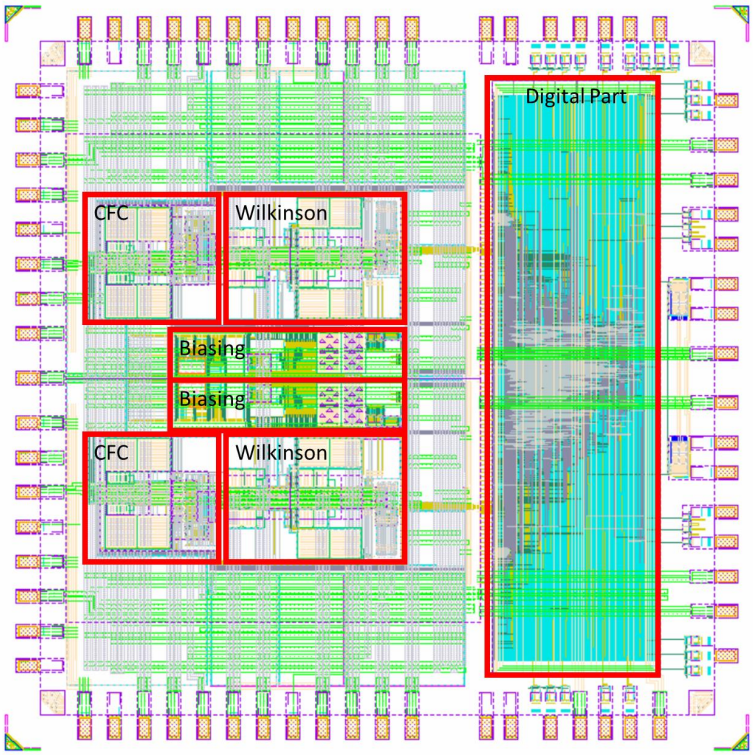

Figure 3.8: Internal layout of the BLMASIC-CFC.

A top view of the micro-electronics layout for the CFC device is shown in Fig. 3.8, where it is possible to distinguish the main integrated subsystems and their dedicated connection pads. The analogue section (left side) includes the two input stages, implementing the architecture presented in Par. 3.2.1, the transistor biasing blocks and the auxiliary ADCs (Wilkinson architecture [66]). The digital section is highlighted on the right. All remaining spare area has been exploited as decoupling capacitors, allowing a low impedance path for disturbances and the current spikes. The details of the modules will be described in the following paragraphs.

Both schematic and post-layout Montecarlo simulations of the merit factors have been undertaken by varying the process, voltage and temperature (PVT) parameters. It evaluated 11 possible meaningful corners, resulting from the combination of the MOS device speeds (Typical, Slow, Fast), supply voltages (1.15 V, 1.20 V) and temperatures (0 $°C$, 27 $°C$ and 60 $°C$). The model of the transistors employed in this process was the *BSIM3*, version 3.24.

3.3.3 Integrator Circuit

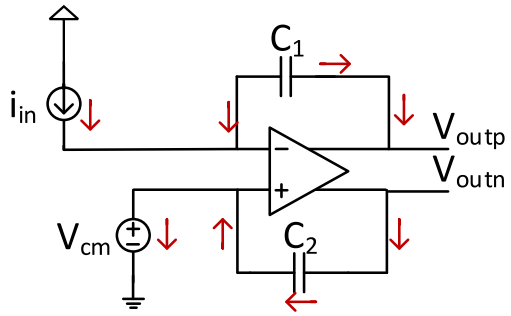

Figure 3.9: Example of current injection into the differential integrator. The intensities remain unchanged for the whole path.

The integrator circuit is realised through the use of a fully differential amplifier, around which is installed a feedback network made of two capacitors of 60 pF each. The implemented schematic is shown in simplified form in Fig. 3.2 and the associated behavioural equations have been described in the Par. 3.2.1.

A comparative study has been carried out about the benefits provided by the differential architecture with respect to the simpler single-ended one [47], [48]. In particular, due to the circuit symmetry, such a solution enables the two capacitors to contribute equally to the voltage swing when currents are injected (see Fig. 3.9). This leads to an improved linearity performance, to a doubling of the output signal and the theoretical integrator slew rate. The trans-impedance is also doubled, despite the waveform frequency remaining the same as the single-ended counterpart. Furthermore, considering both the thermal and the flicker noise contributions, it can be shown that these are mainly due to the input differential pair. Thus, doubling the output swing, the Signal to Noise Ratio (SNR) is also improved by a factor of 2. On the other hand, a major

drawback of employing a differential integrator lies in the operating characteristics of the IC detectors. Indeed, according to the device description of Par. 2.2, the IC is a single-ended current source by construction. Hence, it compels the introduction of the common-mode CM as a sink node for the secondary integrator input (see Fig. 3.2). This is actually an additional power rail, set to half of the supply voltages (600 mV), which must be able to draw the same current acquired by the converter. Furthermore, since any disturbance of CM is added up to the weak measured signals, very good stability and noise suppression must be taken into account when designing the overall front-end PCB. Finally, the advantages delivered by such a solution come at the cost of increased design complexity, a doubled capacitance area and higher current consumption. In fact, additional feedback circuitry is embedded into the amplifier to regulate the output common mode voltage.

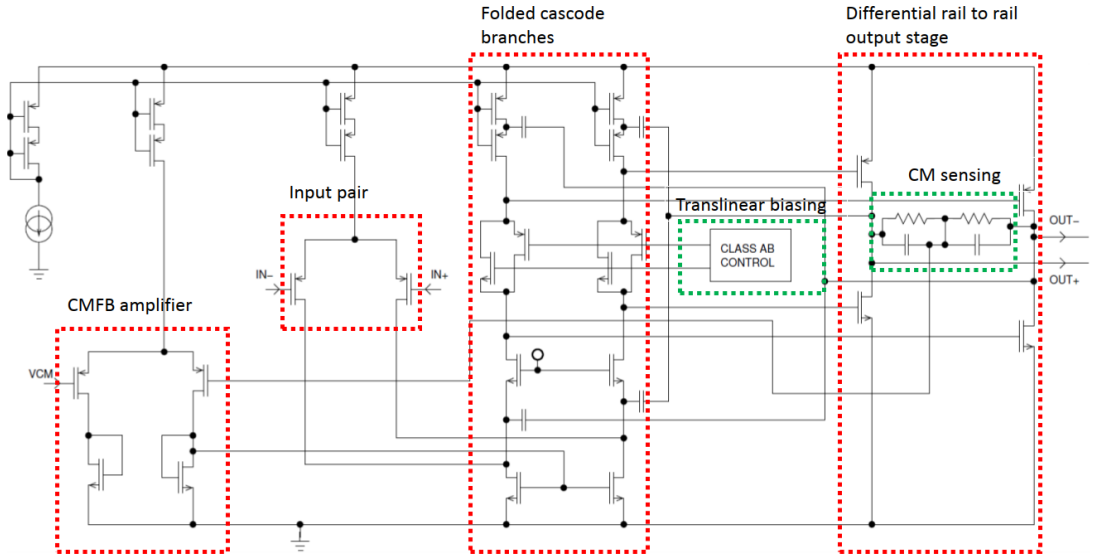

Figure 3.10: Transistor level schematic of the fully-differential amplifier.

The amplifier employed for the integrator is a two-stage CMOS, cascode compensated, with a rail-to-rail class AB output driver [67]. The first stage is configured as folded cascode with a P-MOS input pair. Since the second version of the BLMASIC, such a pair was made of thick-oxide FETs, to reduce the leakage and the input bias current, consequently. Instead, the output stage adopts a push-pull structure. This latter is biased by a trans-linear loop, in order to regulate its quiescent current. The design has been carried out according to the Hogervorst publication [68]. The transistor-level schematic is reported in Fig. 3.10, where the main modules are highlighted. Then, in Fig. 3.11 it is possible to observe the amplifier open-loop gain across the evaluated PVT corners (in various colours). These results are resumed in Tab. 3.2, where the performance of the amplifier controlling the output common mode (CMFB) is also reported. Furthermore, the same table includes the stability margins and the expected current consumptions.

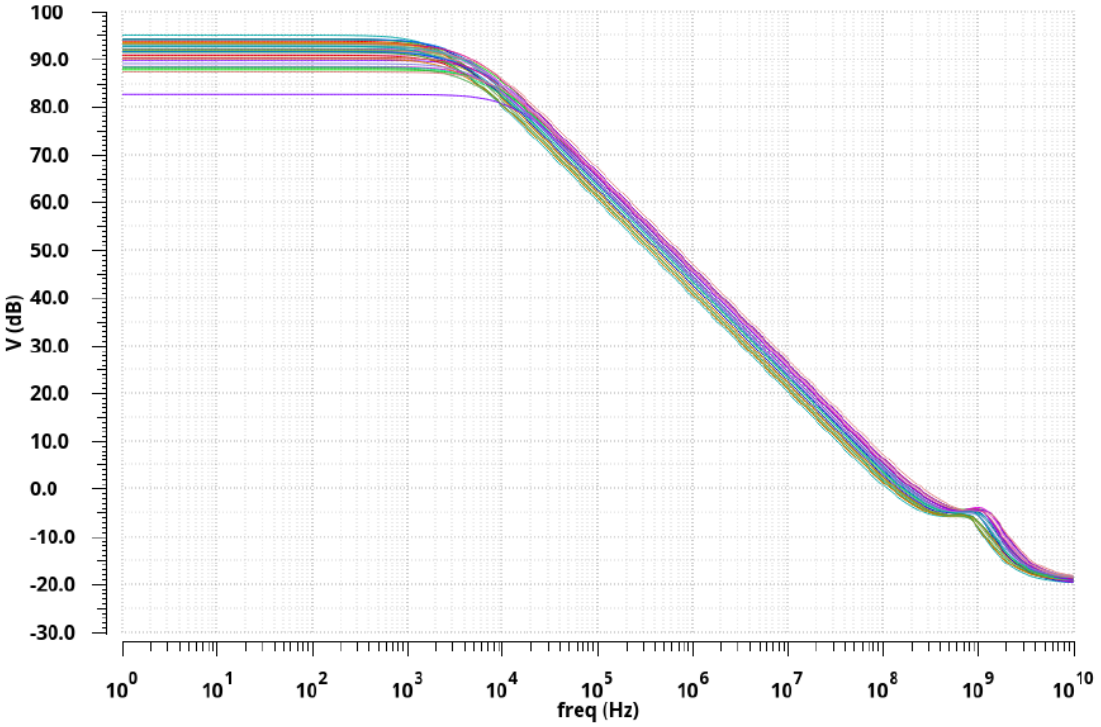

Figure 3.11: Simulated open loop transfer function of the amplifier, over different corners.

Parameter	Minimum	Typical	Maximum
Diff. DC Gain $[dB]$	91	95	97
Diff. GBW ($C_{in} = 350\ pF$) $[MHz]$	140	182	233
Diff. GBW ($C_{in} = 5\ nF$) $[MHz]$	133	152	212
Diff. Phase Margin $[deg]$	83	85	89.8
CMFB Amp. DC Gain $[dB]$	87	96	97
CMFB Amp. GBW $[MHz]$	22	24	25
CMFB Amp. Phase Margin $[deg]$	70	72	74
Offset Std. Deviation $[\mu V]$	...	654	...
Current Consumption $[mA]$	2.7	3.8	4.8

Table 3.2: CMOS fully-differential amplifier specifications.

3.3.4 Switches

Low leakage switches play a key role in such a sensitive device since they are widely employed in the chosen conversion architectures. They have been designed in complementary technology, which is obtained here by connecting a P and an N thick oxide FETs in parallel. The typical conduction state impedance has been set to 50 Ω.

The tasks carried out by the switches involve:

- input current reversing circuitry (CFC device only),
- internal source injection (ΔΣ device only),

- discharge of the integrator capacitors for the initialisation,
- input grounding during the reset.

3.3.5 Quantisers

A couple of differential quantisers are included in the circuit, to monitor the integrator amplifier outputs and command the corresponding feedback. The two components share the same implementation, but they are configured to produce a digital pulse whenever the upper and the lower thresholds are crossed, respectively. These digital signals are then combined by an OR gate.

In the CFC device, the quantisers are asynchronous and the OR gate output is used to clock the toggling flip-flop, controlling the input current direction. On the other hand, in the $\Delta\Sigma$, a version of quantisers equipped with a latch input is implemented. These are used to directly produce a clock synchronous bitstream output and the feedback.

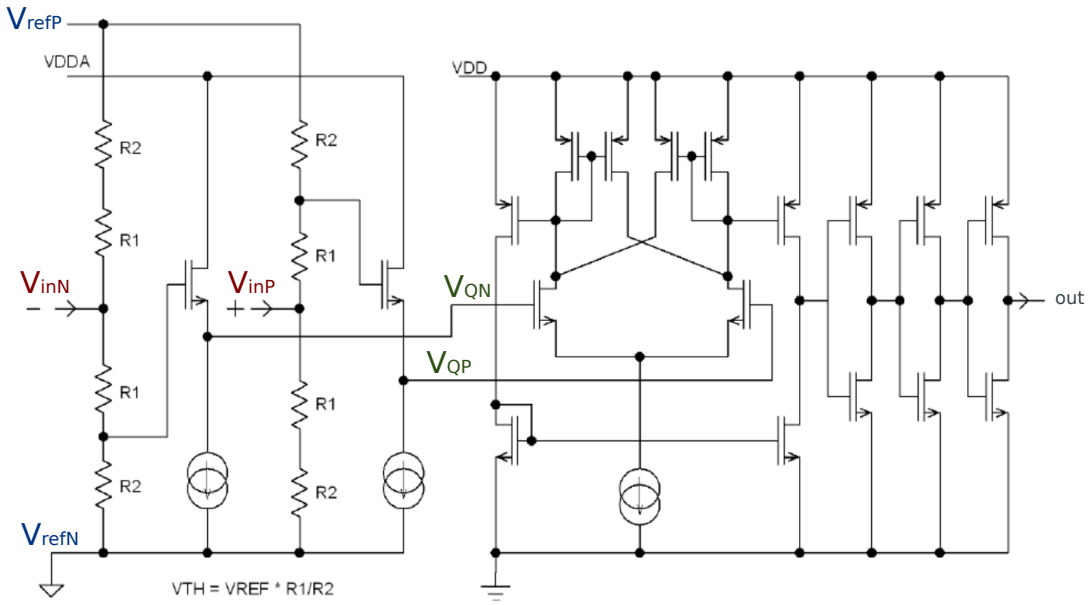

Figure 3.12: Transistor level schematic of the asynchronous quantiser implemented in the CFC.

Fig. 3.12 reports the transistor level schematic of a single asynchronous quantiser. The input stage on the left can be distinguished by the comparator core (right side). In fact, such sub-circuits have been separated into the analogue and the digital ASIC sides, respectively, to reduce the influence of digital switching spikes on the sensitive nets. The thresholds are set by tuning the input stage, which involves a differential divider composed of 8 resistors, sharing the values: $R1 = 8\ k\Omega$ and $R2 = 10\ k\Omega$. These feed two source follower buffers, which have been added to minimise the influence of the quantiser commutations on its own inputs (*Kick-Back* noise [69]). The resulting differential output V_Q is then moved to the comparator, which can detect the signal zero crossing.

Assuming the buffers gain equal to 1, the equation relating the V_Q voltage to the overall differential input V_{in} yields

$$V_Q = \frac{R1\ V_{ref} + R2\ V_{in}}{R1 + R2}\ , \tag{3.14}$$

with V_{ref} representing the differential reference voltage $V_{ref} = V_{refP} - V_{refN}$. In order to simplify the BLMASIC integration into the new front-end prototype, the V_{refP} reference has been tied to the analogue positive supply rail $VDDA = 1.2\ V$ and the V_{refN} is connected to its relative ground $GNDA$, thus $V_{ref} = 1.2\ V$. On the other hand, the complementary quantiser has the opposite connections to $V^c_{refP} = V_{refN}$ and $V^c_{refN} = V_{refP}$, so its V^c_{ref} is also inverted to $-1.2\ V$.

Substituting the resistor values into the equation (3.14), it is possible to deduce that the divider introduces an attenuation factor of 1.8 between the output V_Q and the input V_{in}. Furthermore, imposing the threshold condition $V_Q = 0$, or the complementary $V^c_Q = 0$, it is obtained

$$|V^{Thr.}_{in}| = \frac{R1}{R2}\ V_{ref} = 0.96\ V. \tag{3.15}$$

This value sets the typical oscillation range of the integrator output waveform. Indeed, if the voltage of this latter exceeds $V^{Thr.}_{in}$ or it drops below $-V^{Thr.}_{in}$, a pulse is produced by the quantiser and the slope will be soon reversed.

A critical situation concerns the case of small input currents injected into the converter, which cause a very slow integrator ramp. In fact, the output waveform recedes very slowly from the thresholds after its slope is reversed. This represents a vulnerable state for the converter since the signal noise could lead to an unexpected second threshold crossing, which brings the system to a deadlock state. Therefore, the device has been equipped with double protection against such situations. The first consists of the introduction of hysteresis of about 140 mV in the quantiser thresholding. The second is obtained by a dedicated logic, which has been interposed between the quantisers and the combining OR gate. It allows for the rejection of consecutive pulses produced by quantisers of the same polarity.

3.3.6 Calibration DAC

The BLMASIC converter is equipped with a calibration DAC, which injects DC currents that can be added to the measured ones. It provides an efficient tool to compensate for possible leakages, protect against deadlocks due to noisy inputs and calibrate the converter offset. Furthermore, it allows the continuous injection of a reference current into the device, which can be exploited to verify its effective functioning. This diagnostic feature has also been adopted in the operational BLM system [34], [35].

The current generator is implemented as a voltage DAC, which establishes a potential drop over an external resistor of 1 $G\,\Omega$ with a tolerance of maximum 1 %. Since the DAC shares with the integrator the same common mode reference voltage,

the resistor can be directly connected between its output and the converter input (cf. device pinout in Appendix D). The voltage can be configured by programming the dedicated registers (cf. Appendix C), allowing a variation between $-100\ mV$ to $128\ mV$, with respect to the common mode. In particular, an 8 bits register allows the tuning of the positive voltages, with a resolution of $0.5\ mV$, while an additional bit enables a fixed $-100\ mV$ offset. This results in the ability to inject fine-tuned currents, ranging between $-100\ pA$ to $128\ pA$. The theoretical $0.5\ fA$ DAC resolution is validated by the circuit Montecarlo simulations, which reported a maximum error of $0.23\ pA$ in the overall regulation range.

Finally, in the applications requiring larger offsets, it is possible to substitute the external resistor with one of smaller value. For example, a $100\ M\Omega$ resistor can be used to extend the adjustment range at $-1.00\ nA$ to $1.28\ nA$, with a tuning resolution of $5\ pA$. However, the maximum current delivered by the DAC must not exceed $256\ \mu A$.

3.3.7 Auxiliary ADC

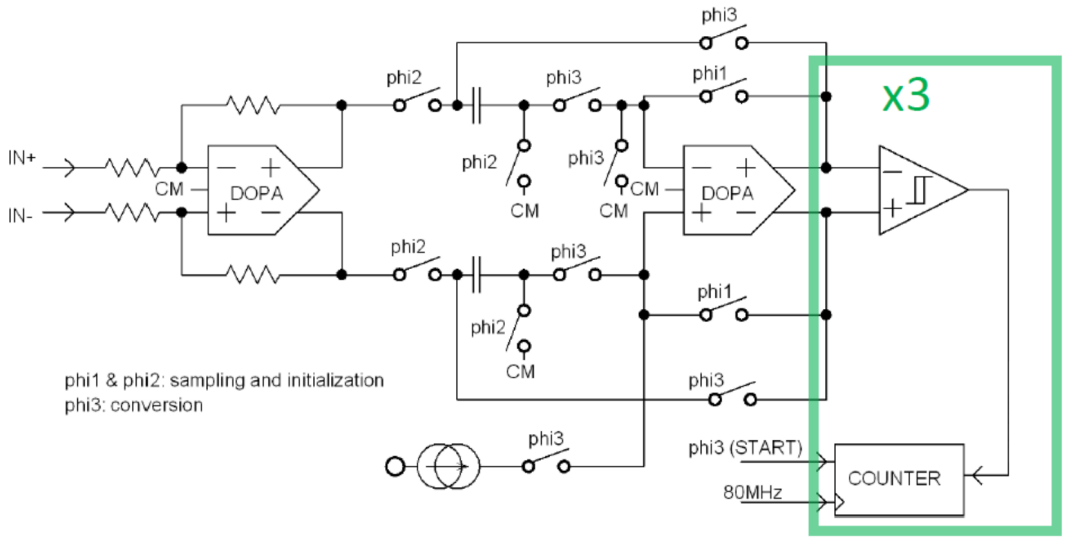

Figure 3.13: Schematic of the Wilkinson architecture ADC.

The *BLMASIC-CFC* is equipped with an additional ADC, which allows for faster and more accurate measurements (see Par. 3.2.1). Unlike the quantisers, which only detect that thresholds have been exceeded, the ADC samples the integrator output at intervals of $10\ \mu s$, providing a real-time measurement of its voltage.

The implementation is based on the Wilkinson architecture [66] shown in Fig. 3.13, which is composed of a unity gain differential buffer (on the left), a switched capacitor network (centre) a zero-crossing quantiser and an enabling counter (on the right). The buffer is used to isolate the sensitive integrating circuit from the rest of the ADC, which can eventually produce fast current spikes. The switched capacitor circuit deals with the signal sampling and it implements the comparison with an internal reference source.

In particular, it involves two working states:

- *Sampling Phase* (only **phi1** and **phi2** switches closed).
  The central differential amplifier is kept in a unity gain configuration and disconnected from the inputs. The capacitors are charged by the buffer and they sample the input voltage.
- *Integration Phase* (only **phi3** switches closed).
  The capacitors are connected to the differential amplifier, thus their combined voltage drop is presented at the output. The internal current source is meanwhile used to discharge them.

The *Sampling Phase* is programmed to last 1 μs, while the *Integration Phase* lasts 9 μs.

The quantiser and the counter are used to measure the time required for discharging the capacitors. In particular, synchronously to the closing of **phi3**, the counter is started at the beginning of the integration phase and it is stopped by the quantiser as soon as the output of the amplifier vanishes. Therefore, knowing: 1) the sampling capacitance (that is 50 pF), 2) the counter value at the end of the process, 3) the system clock frequency and 4) the internal source intensity, it is possible to estimate the input signal voltage.

Furthermore, the device is equipped with triple redundancy and a fail-safe mechanism that stops the counter after 8 μs, whatever the value produced by the quantiser. This guarantees a maximum value of the result, even in case of misbehaviours. The discharge current has a nominal value of 8 μA and it can be adjusted in $6-10$ μA, through the configuration of its dedicated 5 *bit* register (cf. Appendix C). The two differential amplifiers are based on the same design described in Par. 3.3.3 for the integrator, however, their bandwidth has been here reduced to 10 MHz. Indeed, this auxiliary ADC efficiently measures low currents, when the integrator waveform oscillation frequency is much smaller than the sampling one ($I_{input} \ll 10$ μA).

3.3.8 Digital Logics

The integrated logic has been developed employing the best consolidated methodologies for mission-critical applications [67] since high reliability is required for the accelerator machine protection. In particular, a circuit triplication has been adopted for the overall components, and a minimum spacing of 10 μm has been guaranteed for the homologues registers. The logic operational voltage is imposed by the technology to 1.2 V, which is the same used for most of the analogue circuitry. However, an isolated rail has been dedicated to the digital side, to avoid any injection of undesired disturbances into the sensitive converter components. In the following, the variety of different tasks performed by the logics are reported and the involved device pins (in Appendix D) are denoted boldface. They concern:

- the system reset and redundant clock management;
- the readout of the analogue circuitry;
- the redundant data transmission;
- the device configuration;
- the diagnostics.

Reset & Clock Management

The interfacing and propagation of resets and clocks are embedded into the digital area of the integrated circuit. The device initialisation is managed by a Finite State Machine (FSM), which is triggered by the external pin **RSTB** (active-low). De-asserting such a signal, the initialisation of the analogue circuit is launched and the full register map is restored to its default values. The **RSTB** is acquired asynchronously, but it is synchronised to the clock inside the device.

On the other hand, to drive the whole synchronous logic, an 80 MHz clock must be provided to the system. It is sourced by two redundant input clock interfaces, to which "ready" flags (**ARDY** and **BRDY** input pins) are associated. The flags are asserted by the clock source devices (the lpGBTs), whether their PLLs are stable and they are ready to accept new data. Hence, such flags are used by the BLMASIC to select the preferred clock between the two.

The input clock interface is compliant with the CERN Low Power Signalling (CLPS, included in the lpGBT documentation [S13]). This is a 100 Ω terminated differential lane, which is characterised by a common mode voltage of 600 mV and a differential amplitude between 200 mV and 800 mV.

Analogue Readout

This module represents the interface between the analogue and the digital signal domains. In fact, it acquires the flags produced by the quantisers, stores the information in dedicated registers and controls the internal feedback. Furthermore, it includes the CFC and the Wilkinson ADC counters.

The CFC readout involves two stop-less and unidirectional 10 *bit* counters, which accumulate the upper and the lower crossings of thresholds, respectively. Their registers have been tripled and a majority voter is employed to select the value, amongst the three, which has to be incremented and stored at the next clock edge. This allows a reliable and long-term operation even in the case of SEU, since the register containing the wrong value is always corrected by using the others. Furthermore, because the system is designed to reject consecutive increments triggered by the same quantiser, the values of upper and lower counters should not differ more than 1. However, additional precautions have been undertaken in the back-end firmware, allowing a correct operation even though this hypothesis is not valid.

The values of the CFC counters can be accessed in the device memory map and they are transmitted with a frequency of 100 kHz. Thus, the acquisition firmware must perform the differences between consecutive readings, in order to retrieve the number of positive and negative threshold crossing occurrences. The chosen number of bits guarantees that no more than a single counter overflow can occur within the 10 μs window. Furthermore, it is remarkable that the result of the differences must be truncated to 10 bit, in order to obtain a meaningful value even in case of overflow. Finally, the differences computed individually on upper and lower counters can be added together, to obtain the total ramp inversions during the latest 10 μs.

The counters of Wilkinson ADCs are also part of the register map and they are transmitted along with the former. Due to the component architecture, presented in Par. 3.3.7, the outcomes represent a direct measurement of the integrator output. Hence, unlike in the previous case, no further treatments are required except for eventual calibration.

Data Transmission

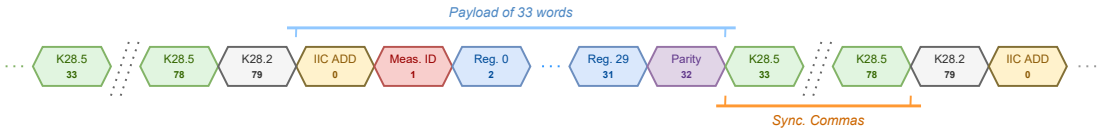

Figure 3.14: Full transmission frame structure.

The data transmission is achieved by point-to-point, system synchronous and redundant serial links, which operate at 80 $Mbps$. The development has been tailored to the features provided by the lpGBT device, which will be used as a transceiver chipset for the overall front-end. Specifically, the BLMASIC can be directly interfaced with its E-Links.

The transmission is enabled as soon as the device selects its working clock, and it will never stop unless both the clocks are disconnected or the "ready" flags are deasserted. Each payload (see Fig. 3.14) is propagated at the converter sampling period of 10 μs (100 kHz) and it is composed of 33 $byte$, with respect to the total 80 $byte$ transmitted within this time ($\approx$ 41 % of duty cycle). Furthermore, the bitstream is 8b/10b encoded [70], meaning that each byte of information is actually represented by using 10 bit, the most significant of which is transmitted first. Thus, the total bitrate is

$$BR_{ADC} = 80\ bytes \times 100\ kHz \times \frac{10}{8} = 10\ MBps = 80\ Mbps. \tag{3.16}$$

The payload starts right after the control symbol *K28.2*, provided by the standard. Instead, all the idle time is occupied by *K28.5* synchronisation commas.

While the chosen encoding allows a bit level error check, a parity vector byte has also been added to the payload, to provide a further frame integrity validation.

In particular, denoting by F_n the n_{th} byte of the payload F and with F_n^i its i_{th} bit, each parity vector bit P^i is set to obtain an even number of asserted i_{th} bits within all the bytes. Namely, it can be expressed as

$$P^i = \bigoplus_{n=0}^{N} F_n^i \quad (0 < i < 7) \,, \tag{3.17}$$

where $\oplus$ indicates the cumulative *XOR* operation and i is iterated over the payload bytes, excluding the parity one ($N = 31$). The meaningful content of the transmitted payload includes:

- F_0: the I^2C address of the device,
- F_1: the measurement ID, which is always incremented,
- F_2 to F_{31}: the entire BLMASIC memory space (30 registers),
- F_{32}: the parity vector byte.

The implemented redundancy provides the parallel transmission of such data over two independent serial links. This allows propagating the results towards two lpGBTs, thus improving the reliability of the front-end and ensuring a continuous transmission in case of unlocking of transceivers or failure. However, it should be pointed out that the BLMASIC operates with a unique clock and both the data links are synchronised with it, hence timing issues may arise on the receiving E-Link. The solution will be explained in detail in Chapter 6.

Device Configuration

Memory space of 30 registers, each of 1 *byte*, is integrated within the logic. It is used to store the overall device configuration and the results achieved by the conversion process. One of the capabilities provided to the user represents the possibility of finely adjusting the low-level analogue circuit parameters. Relevant examples are the biasing currents, the Wilkinson ADC range and the calibration DAC. Furthermore, it is possible to verify the clock selection status between the two, regulate the data transmission driver current and disable the 8b/10b data encoding. A reset of the analogue or digital domains can also be commanded. Nevertheless, for such critical features, triplicated bits must be programmed to validate the request. This avoids accidental triggering due to wrong register configurations or SEU. A detailed list of each register function is reported in Appendix C. The memory access can be performed by means of two redundant I^2C slave interfaces, which also represent the only available writing option. Each interface is intended to be connected to the corresponding lpGBT master and it is enabled according to the clock selection described in Par. 3.3.8. Both are compatible with bit rates up to 400 kHz.

Diagnostics

The device has been equipped with tools supporting the validation process. In comparison to the previous BLM front-end, which was assembled using discrete components, the integration of such a complex architecture inside an ASIC raised many concerns about the ability to detect possible faults. Thus, the designers have strived to allow the inspection of as many possible signals and internal components by the outer pins.

To effectively fulfil these goals, a differential analogue probe (pins **ATEST** and **BTEST**) and a single-ended digital one (pin **DTSTOUT**) have been installed. They are internally connected to a multiplexer, which can be configured by dedicated registers to monitor a wide range of internal nets (cf. the configuration examples in Appendix C). A frequent use case concerns the monitoring of the converter integrator output, to visualise the triangular waveform characteristics on the oscilloscope (cf. Fig. 4.14). At the same time, the digital multiplexer can be configured to probe the quantiser output, so a simultaneous study of the two observables can be achieved. The full list of selectable nodes is reported in Tab.s C.4, C.5 and C.6.

3.3.9 Package

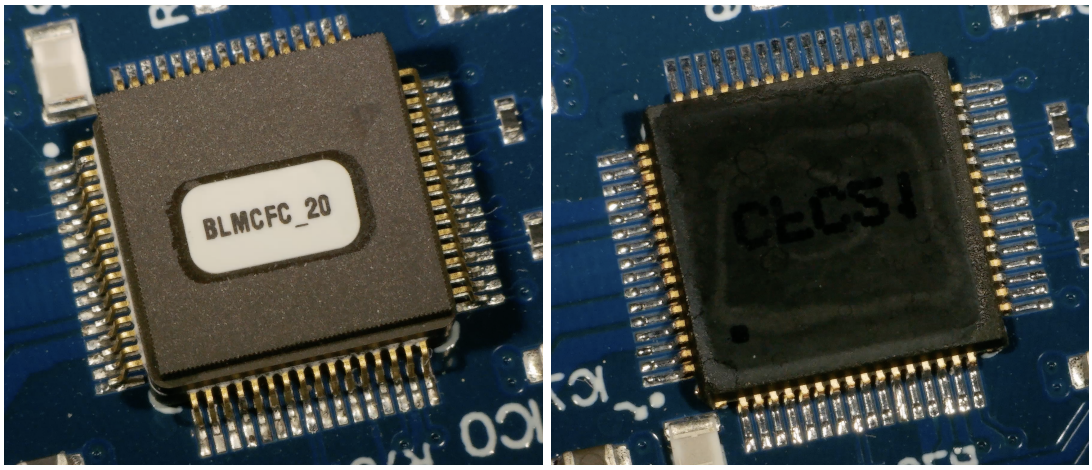

(a) Second version of BLMASIC encapsulated in the *Kyocera Corporation* ceramic package.

(b) Third version of BLMASIC encapsulated in the *SERMA Microelectronics* epoxy package.

Figure 3.15: Photos of the packaged and assembled BLMASIC converter.

The BLMASIC is housed inside a 64 pins quad flat package, with a 0.5 mm pitch. The first two prototype versions adopted a ceramic CQFP64 package (QC-064347-WZ), produced by *Kyocera Corporation* (in Fig. 3.15a). This was provided with a glued cap that could be removed for the direct optical microscope inspection and an effective X-Ray irradiation test. Instead, for the third and latest version (in Fig. 3.15b), a standard epoxy resin (type FP4470) filled package has been used, to get as close as possible to the conditions envisaged in the final production. This device was packaged by *SERMA Microelectronics* and the bond-wire inductance has been estimated to be

0.5 nH per connection. The supply and the common mode rails have been spread over multiple pins, in order to minimise parasitic effects and to provide low impedance paths for high-frequency signal components. The overall device pinout description is reported in Appendix D.

Chapter 4

ASIC Test-Bench

The architecture and the assembly of several test benches are presented. They aim to validate the main BLMASIC functionalities. Such activity involves the development of custom hardware equipment and the related firmware, to carry out targeted investigations. Furthermore, the parallelised software for the data analysis is described, underlining the challenges encountered and the opportunities of employment as a general purpose tool. The results of the basic features, the conversion behaviour, the achievable resolution, the X-Ray radiation tolerance and the operations of the digital interfaces are fully reported. It is also illustrated how these measurements revealed malfunctions in the device, leading to the release of two updated versions. Finally, promising outcomes obtained by the two conversion architectures are exhibited, affording an enlightened selection of the final device.

4.1 Objectives

These preliminary investigations aim to characterise the integrated BLMASIC converter (simply referred to as the ASIC) as an individual device, to validate its employment within the newer BLM front-end. They were focused on the critical requirements specified for the device, following a hierarchical analysis of the modules already described in Chapter 3. It must be pointed out that the raw data produced by the ASIC alone cannot fulfil the system level specifications, as described in Par. 2.6. This holds since the paired back-end processing is a critical step for achieving an accurate measurement of losses. Thus, lower-level studies will be carried out at this stage, postponing the complete performance validation to Chapter 8.

Even though the estimated device performance has been simulated using sophisticated tools, tests on the real component are essential to qualify its behaviour after manufacturing. There are physical effects not included in the simulations, depending on the fabrication run, and others that are intrinsically too complex to be easily implemented [62]. The versatile characterisation test-bench, built during this research, will also enable a direct comparison of the two architectures discussed in

Par.s 3.2.1 and 3.2.2, in order to decide which one choose for the final production. The activity included the design of a characterisation board, the selection of suitable laboratory instrumentation, the development of the acquisition and configuration firmware, as well as the software for subsequent batch data analysis.

4.2 Characterisation Board

The test-bench development began with the design of a printed circuit board (PCB), which hosts the ASIC and allows convenient access to its interfaces. Since this card is used for the overall preliminary validation, design efforts have been devoted to minimising the influences on the targeted study. Furthermore, the circuit has been organised in order to accommodate both the CFC and $\Delta\Sigma$ based prototypes.

In Fig. 4.1 is shown a 3D rendering of the card, highlighting its main sections. The lower part of the board has been dedicated to the fast digital interfaces, which can be accessed by employing 50 Ω SMA connectors. These include the two clocks, the serial data transmitters, the digital test probe and the "ready" signals, which must be employed according to the considerations done in Par. 3.3.8. Furthermore, the card provides access to the redundant I^2C interfaces, which can be linked to external programmers using the female RJ45 connector or the dedicated 100 *mil* headers. The measuring input channels are visible on the top and exploit two BNC connectors while, on the right, a LEMO plug is placed to power supply the card. Finally, many connections have been mirrored on alternative headers to simplify their access during the test and to improve setup flexibility. A relevant example is represented by the irradiation test sessions, when the embedded power supplier cannot be used.

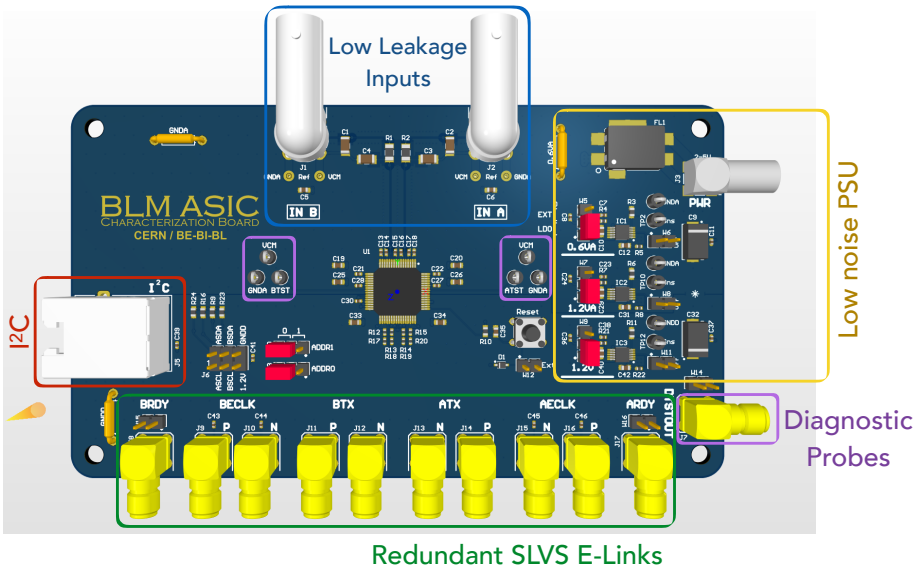

Figure 4.1: Rendering of the BLMASIC characterisation board.

4.2.1 Design of Schematics

The power supplies were the first modules to be designed (cf. bottom frame of Fig. 4.3), according to the ASIC requirements. In particular, these involve two rails operating at 1.2 V, for the analogue and digital circuitry, and a stable common-mode reference (VCM) at 0.6 V. Considering that the device is intended to measure very low currents, down to a few pA, a solution minimising the powering noise has been chosen. This is particularly relevant for the VCM reference, which directly influences the measurements due to the integrator implementation (see Par. 3.3.3). Independent rails and grounds have been exploited to effectively satisfy these needs. The voltages are controlled by a set of ultra-low noise LT3042 components, which are commercial linear regulators manufactured by Analog Devices [S17] and amongst the best solutions available for such a goal. In particular, it guarantees a 0.8 μV_{RMS} noise in the bandwidth of 10 Hz to 100 kHz. In addition, the Power Supply Rejection Ratio (PSRR) is higher than 90 dB down to 10 kHz and 79 dB at 1 MHz. The nominal output current of 200 mA is well inside the consumption expected by the ASIC. Furthermore, to perform the primary supply noise suppression and to allow the effective bypassing of the high-frequency disturbances towards the grounds, the following components have been integrated into the board:

- BNX025H01B EMI filter produced by Murata [S18], characterised by an insertion loss of 35 dB above 50 kHz and used for the main supply noise suppression,
- high-quality X7R 220 μF Tantalum capacitors as main supply bypass,
- SMD 0603 22 μF MLCC capacitors, close to the ASIC,
- low ESL SMD 0402 470 nF MLCC capacitors, placed underneath every device supply pins (cf. Fig. 4.2).

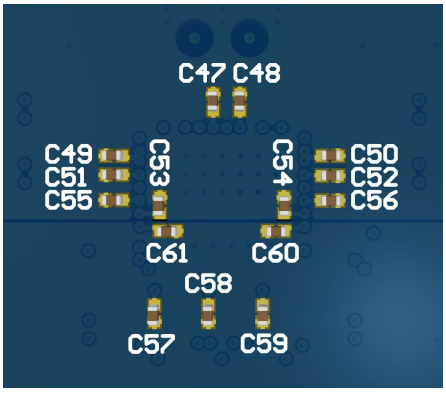

Figure 4.2: Detail of the bypass capacitors on the bottom layer.

Since the second version of the ASIC adopts an improved circuitry based on thick-oxide FETs, an additional rail of 2.5 V has been added (cf. Par 3.3.3). Hence, the board was modified to support this upgrade properly. Furthermore, an alternative dual supply

solution has also been investigated, which would involve a downshift of all the power buses of 0.6 V, hence promoting the VCM to an absolute ground reference. However, it has been excluded from the alternatives for the following reasons. First, the main argument of discussion was related to the possible increased leakage over the analogue transmissions, while a 0.6 V VCM is used. Indeed, such a reference polarises the core of the coaxial cables by design, which connect the front-end to the IC detectors. Thus, imposing 1 pA as the maximum acceptable leakage current for the cable, this should have an isolation resistance greater than 600 $G\Omega$, which is a challenging requirement. Furthermore, if the resistance has a lower value, undesired currents will be drawn from the converter in the inverse direction with respect to the measured signal and they can cause the system to deadlock (cf. Par. 3.3.6). Nevertheless, this last issue is reliably solved by the proper calibration of the internal DAC current (cf. Par. 3.3.6) and well-isolated cables are still required if $\mathcal{O}(pA)$ precisions want to be achieved. Secondly, a dual supply solution would require Rad-Hard negative voltage regulators, which would not exist yet during the installation of the final front-ends.

Regarding digital interfacing, differential pairs with a characteristic impedance of 100 Ω have been used for the clocks and the redundant data transmission. According to their implementation, the signals will be characterised by a differential amplitude of $\approx$ 400 mV and a common mode of $\approx$ 400 mV. Compatible specifications can be found in the CLPS standard (see the lpGBT documentation [S13]), which requires the same signal amplitude, but a common-mode voltage of 600 mV. In the test-bench measurements, the card is planned to be connected to an FPGA device, providing the correct clock for data acquisition. Such a class of components is usually equipped with LVDS interfaces only, thus a common-mode shifter is required between the ASIC interfaces and the FPGA. Furthermore, the cited commercial standard deals with higher amplitude signals (about $\approx$ 800 mV), but the employed FPGA supports weaker signals, so no incompatibilities are expected. Two resistive dividers have been used to shift the common-mode voltage right before the clock receiver (cf. top frame of Fig. 4.3). Their values were tuned to provide the fitting impedance termination to the line. Similar circuitry can also be used for receiving data on the FPGA side, if required. This is unnecessary for the employed test-bench since the serial data are DC balanced due to the 8b/10b encoding. Hence, capacitors can be employed to decouple the transmission lines between the FPGA and ASIC domains.

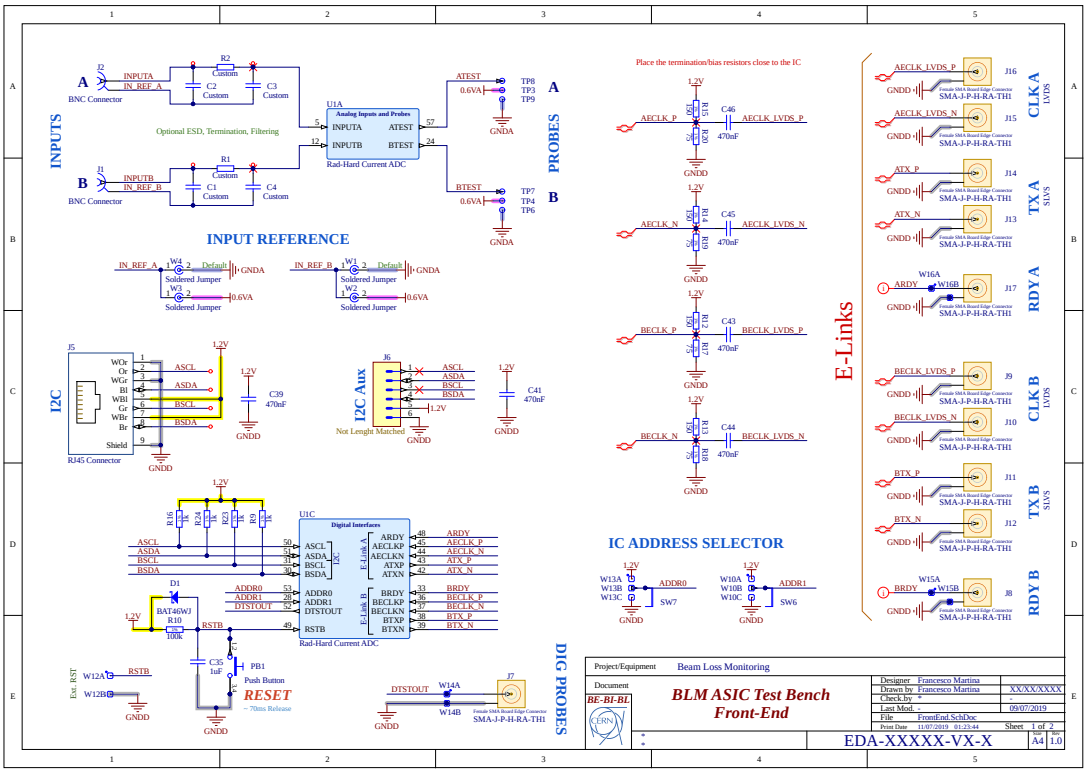

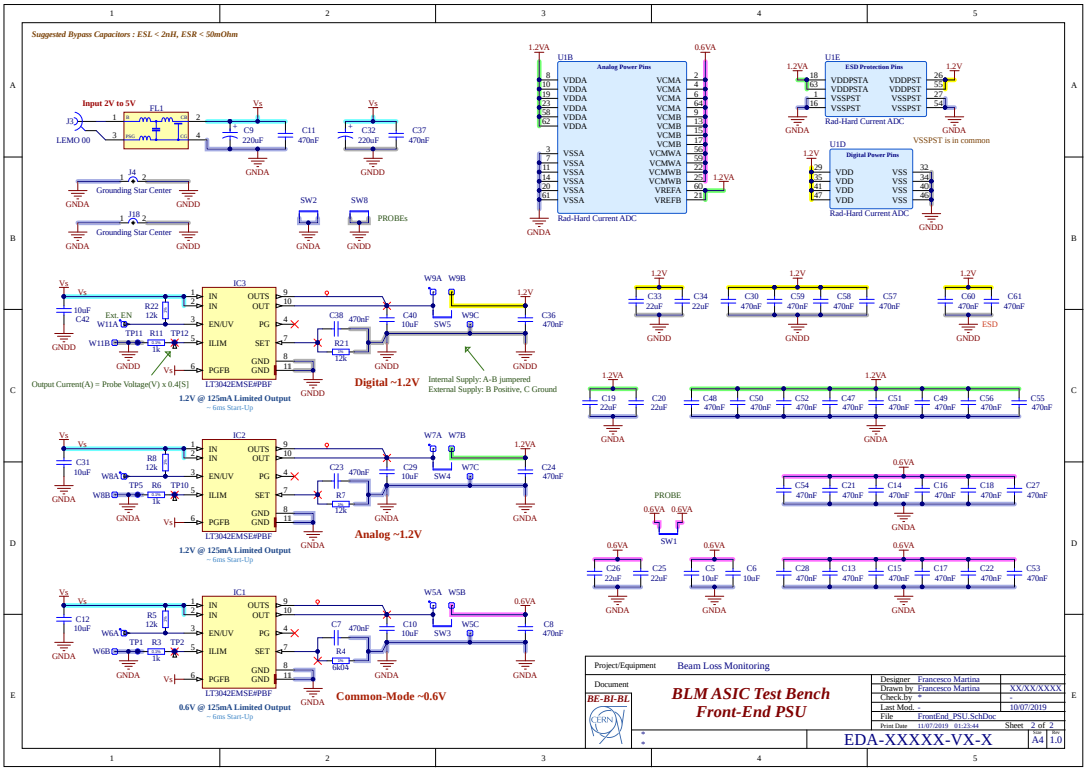

Figure 4.3: Characterisation board schematics.

4.2.2 Board Layout

A 6 layers FR4 stack-up has been chosen for the realisation of the PCB. Although the placement of components is not dense, this solution allows an optimal routing of the nets and a ruggedised power distribution. Following the criteria adopted for the power supply design, full-size power planes have been used to provide the 3 rails to the ASIC and the whole card has been subdivided into the digital and the analogue domains. This strategy allows easy reaching of all the powering pins of the ASIC, hence providing a low impedance path for the currents. The grounding has also been separated and a single star point has been placed close to the power input. This prevents the paths of noisy return currents from crossing and interfering with the sensitive nets. In Fig. 4.4, the top layer of the board is displayed and one can see the organisation of ground planes. The routing on all the other layers is carried out by adopting similar strategies.

Concerning the routing of the differential lanes, the geometry proposed by the factory was employed. They were located on the fifth layer from the top and implemented as striplines with 4 *mil* of track width and 6 *mil* gap. The length matching amongst the redundant pairs and the clocks was not mandatory since the ASIC only transmits data and the two redundant channels are independent. Nevertheless, a good symmetry has been deliberately kept to facilitate possible latency tests.

According to the converter design, the input signals could be directly connected to the device pins, since the embedded ESD protections should fulfil the specifications. However, extra protection composed of a π shaped filter has been envisaged. Since the component values were not strictly defined, empty SMD 1206 footprints were placed to be later populated.

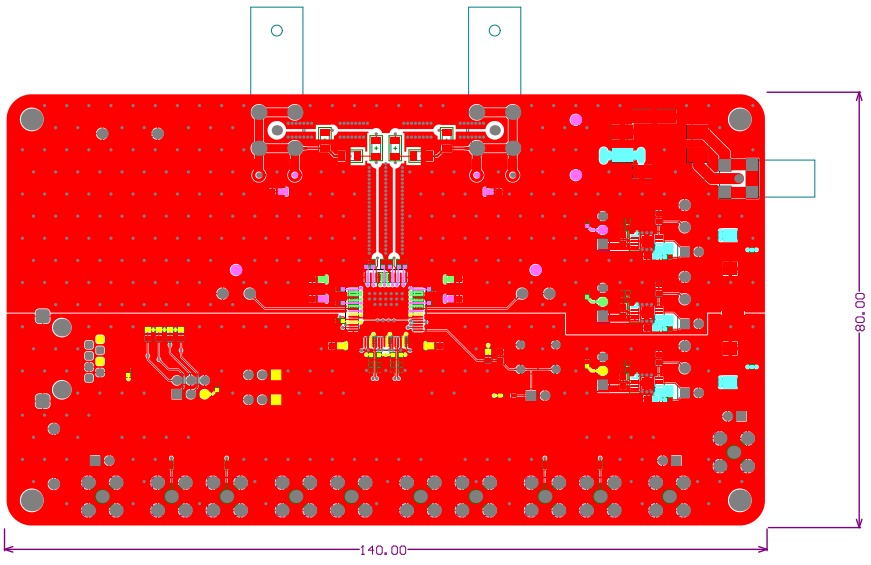

Figure 4.4: Top-layer of the characterisation board PCB.

During the PCB layout design, a great effort was undertaken concerning the input tracks routing, to keep the parasitic current leakage below a value of 100fA. To this end, the second layer from the top has been employed, seeking a compromise between the tracks shielding and the clearance from grounding planes and digital transmission lines. After a final rule check of the schematics and the layout, GERBER and Drill files were produced and sent to the Phoenix S.r.l. Company (Italy) [S19], which finalised the manufacturing of 20 bare boards.

4.3 Setup Architecture

The ASIC-focused validation has been carried out on three different test-benches, which were developed between 2019 and 2021. Although they involved similar components, the investigation of some characteristics, rather than others, required several iterations and updates. Furthermore, some processes have been gradually improved during the investigation, allowing easy management and greater flexibility. In particular, this concerns data acquisition and instrumentation control. A brief description of the various setup is reported below, while the actual methodology employed to carry out the validation will be presented later in Par. 4.4.

4.3.1 Initial Setup

Following the delivery of the first ASIC prototype, a preliminary test-bench was assembled to perform its initial power-up, verifying the functioning of its basic features. Its assembly scheme is shown in Fig. 4.5.

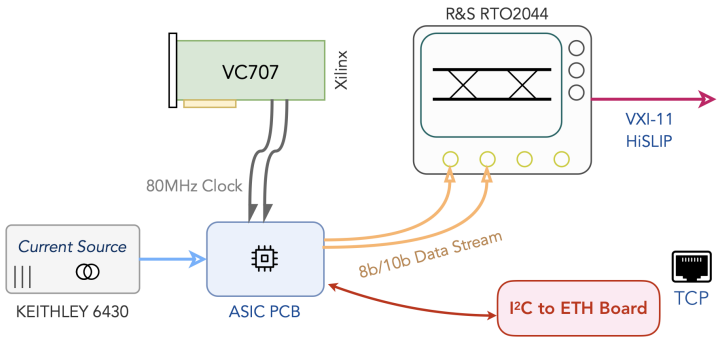

Figure 4.5: First test-bench architecture.

At the centre of the scheme, the ASIC characterisation board is visible, which was powered by an Agilent [S20] E3631A bench top supplier not shown in the figure. The system clock has been provided by an external FPGA, hosted on a Xilinx [S21] VC707 development board, which is programmed to generate low jitter 80 MHz reference on LVDS outputs. At the same time, the ASIC data stream was probed by a fast sampling oscilloscope (Rohde & Schwarz [S22] RTO2044), which allowed the signal integrity analysis and the decoding of the 8b/10b frames. Furthermore, the ASIC

analogue inputs were driven by a high accuracy source, injecting DC current references. A Keithley [S23] 6430 device has been chosen for this goal. It can be configured from the user panel and the GPIB interface, providing many automatisation features.

Finally, an external board has been programmed with custom firmware, to allow the I^2C interface access to the whole ASIC register map. Specifically, an MYiR [S24] Z-Turn embedded system has been chosen for this purpose. The card hosts a Xilinx Zynq 7020 SoC, a device which combines the performance of a Xilinx Series 7 FPGA logic with the flexibility of a dual-core ARM [S25] A9 processor, featuring numerous peripherals. To efficiently perform the I^2C programming in the test-bench, a TCP server has been included in the firmware. This allows a fast and reliable configuration of the ASIC through the network and extends the test-bench control to many concurrent users, even outside the CERN sites. Such a solution agrees with the concept of automatised and remote-controlled setup, which has been considered a helpful goal during all the development stages. In this case, the logic side of the Zynq device has not been used since the processor peripherals already included the I^2C and the speed requirements were not too demanding. The server firmware was based on the open source lwIP TCP network stack [71], [S26] and customised data packets have been used for the transmission. Finally, a set of configuration MATLAB [S27] scripts have been implemented on the user side to access the ASIC registers efficiently. Concerning the physical layer of the I^2C, a bi-directional voltage shifter has been assembled between the Z-Turn card and the device under test, in order to convert the 0 V - 3.3 V logic levels, available on the SoC, to the 0 V - 1.2 V ones used by the ASIC. To this goal, a Texas Instruments [S28] LSF0108 Bidirectional multi-voltage level translator has been used. Its maximum baud rate of 100 $Mbps$ is largely sufficient to satisfy the requirements. This voltage shifter has been used for all the test-benches of this characterisation. In Fig. 4.6, a part of the setup during the very first run is reported.

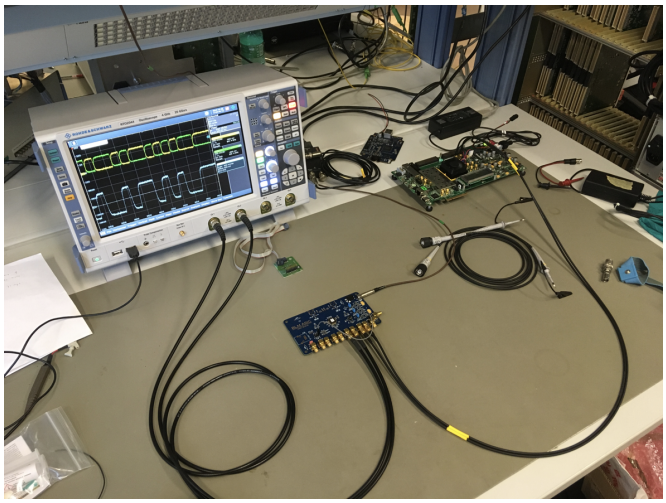

Figure 4.6: First test-bench assembled.

4.3.2 Converter Noise Characterisation Setup

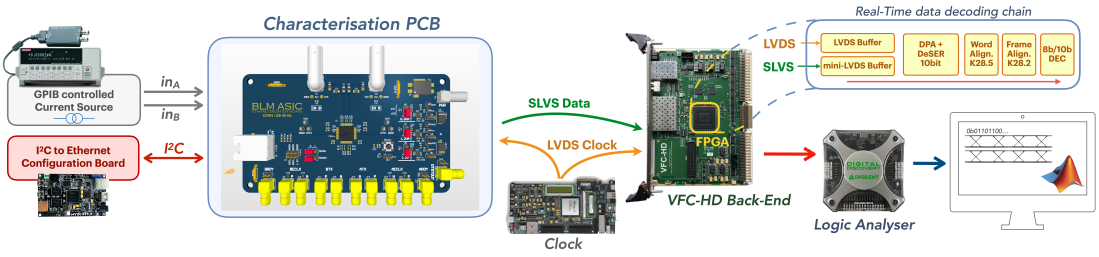

Figure 4.7: Architecture of the test-bench used for the converter noise characterisation.

By exploiting the first version of the test-bench, it was possible to acquire preliminary information about the basic device operations, the linearity of measurements above 1 nA and the input channels mismatch. However, its field of investigation was strongly confined to a narrow range of possible tests, because the oscilloscope buffer limited the acquisition time window to 4 ms. Since the characterisation of the device noise requires much longer capture time, an upgraded test-bench was developed and assembled [6]. This has provided far greater data storage capabilities.

In Fig. 4.7, the scheme of the upgraded setup is reported. Again, the characterisation PCB is connected to the testing current source and the I^2C configuration card, but a dedicated readout system now replaces the acquisition oscilloscope. It is based on the VFC-HD [45], [S14] carrier board (cf. Par. 2.5), the FPGA of which runs a firmware specifically developed to acquire the real-time data stream. As before, an external FPGA board sources the clock for the entire system. During these preliminary stages, the experience acquired with such a platform was fruitful for the overall project.

The VFC-HD performs the 8b/10b data stream decoding in real time and retrieves the measurement values from the related ASIC memory addresses. Unlike the previous case, it is no longer required to save the entire device memory, with significant benefits in terms of storage resource saving and faster subsequent processing. The data are moved to a personal computer through an external logic analyser (Digilent [S29] Digital Discovery), which is connected to the carrier board and controlled by a Python [S30] / MATLAB script. This saves long histories of acquired data into a file, allowing an efficient study of the conversion noise. Furthermore, other developed scripts can program the ASIC internal registers and inject the desired test signals by configuring the current source. This has been installed within the CERN network through a commercial Ethernet to GPIB adapter.

In order to verify the device performance under particular conditions, additional firmware was also developed. A relevant example is represented by digital modules that were able to identify glitches in the measurement data and faulty conditions. During the validation of the first device prototype, such modules were employed in conjunction with the probes (cf. Par. 3.3.8), to investigate the causes of an issue affecting its internal logic (changelog are reported in Par. 4.10). In case of a detected problem, the FPGA

could trigger the oscilloscope, while the same monitored the ASIC analogue probes. This feature proved useful to identify and correct rare erroneous conditions, which would otherwise have been very difficult to detect. Furthermore, it made it possible to pursue the circuit analogue performance characterisation, without waiting for the delivery of the successive ASIC version, which would have included the correction.

Each of the instruments involved in the setup has been directly connected to the CERN internal network, making it possible to run automated measurement sessions and to control the overall system remotely, even when the operators were teleworking. This feature has been extensively used during periods with COVID-19 restrictions.

4.3.3 Post-Irradiation Setup

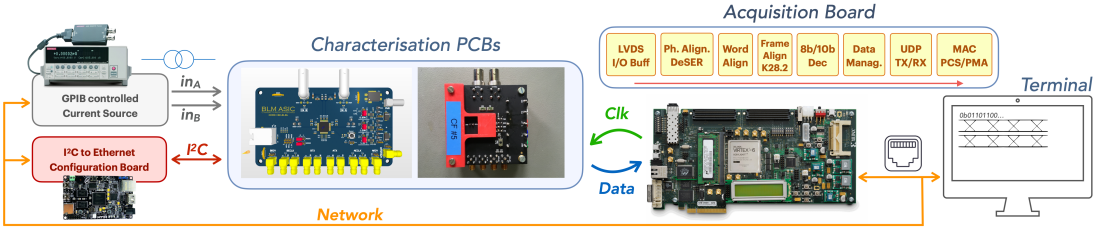

Figure 4.8: Architecture of the test-bench used for the characterisation after the X-Ray sessions.

The experience acquired on the ASIC validation activity during 2019 and 2020 has allowed standardising the investigation methodology [7] and optimising the employment of the instrumental resource. The complex setup described in Par. 4.3.2 has been simplified at this stage (see Fig. 4.8), integrating into a single readout board (Xilinx ML605) the tasks previously performed by several types of equipment. It embeds an FPGA device which has been programmed to: generate the ASIC clock, retrieve its memory content and achieve the data transmission to a personal computer, by using an Ethernet-based network. Hence, the commercial logic analyser is no longer needed and the VFC-HD has been temporarily shelved, to be later used in the validation of the complete BLM system (described in Chapter 8).

For operating in the HL-LHC environment, the BLMASIC must be radiation tolerant. Therefore, tests for the TID, the SEU and other effects of radiation (cf. Appendix A) must be performed, verifying the compliance with the specifications. In fact, 5 devices previously irradiated with X-Rays have been validated with the present setup. In particular, a conversion characteristic study has been carried out to investigate performance drifts amongst them. Finally, an alternative version of the characterisation PCB board has also been used during these tests, with simplified features compared to the mainstream one (cf. Par. 4.2).

4.4 Test-Bench Methodology

4.4.1 Basic Initial Tests

The validation of the very first prototype began with the basic electrical checks. The device was assembled on the characterisation PCB, according to the test-bench described in Par. 4.3.1. In order to prevent permanent damages, bench-top suppliers were used instead of the embedded one, since they would have allowed a more strict current limitation and monitoring during the first run. In particular, the limiters were set to less than twice the nominal expected currents and the consumptions were constantly monitored. Furthermore, the board has been examined using a thermal camera. This instrument helps to measure the device temperature, to detect and locate eventual excessive power dissipation spots or short circuits. This process was carried out either with or without the clock source.

4.4.2 Digital Interface Validation

The redundant serial data interfaces were tested individually using the oscilloscope and the equivalence of the two data streams was verified. Since those produce differential signals, many solutions have been examined to achieve effective probing. The preferred methodology involves the connection of the serial lane to a 100 Ω termination, monitoring the established voltage through a differential probe. The link between the characterisation board and the termination can be achieved by using two independent 50 Ω coaxial cables, used for the positive and the negative signals. An alternative method consists of the direct connection to different oscilloscope inputs if this latter provides internal 50 Ω terminations. Such a solution is usually simpler and cheaper since the differential probe itself could constitute a relevant fraction of the setup cost. However, a worse resolution will be achieved, because the difference between the signals will only be computed after the sampling and the digital conversion. Furthermore, it must be pointed out that the first solution does not influence the lane common mode, while the second may cause an undesired current drawn from this component. Such a current can lead to the transmission driver malfunctioning and should be considered when measuring the device consumption. This issue can only be removed by adding external decoupling capacitors or enabling the AC 50 Ω mode, if such an option is provided by the oscilloscope.

The proper horizontal and vertical factors are configured on the instrument and the trigger is set to 0 V (considering the differential probe use case). The eye diagram can represent a very effective tool for performing a complete signal integrity validation in the time domain. A double-edge trigger mode can be employed to get the goal or, preferably, the system clock can be used to trigger the oscilloscope externally.

Later, the study moved to a higher-level analysing of the transmitted digital data. Hence, the 8b/10b decoder has been enabled. It provided the fast visualisation of

the ASIC frame structure and the possibility of highlighting any error occurrences. Exploiting this feature has simplified the validation of the transmission logic and provided prompt results, before proceeding with the development of a custom readout system. Instead of the oscilloscope, a logic analyser can be used to perform the frame validation.

A similar methodology has been applied to check the I^2C integrity and transactions, connecting the oscilloscope in parallel to the bus. The protocol validation has been achieved both by the decoder embedded into the oscilloscope and the features included in the I^2C configuration board. Several bit rates in the range 10 kHz to 400 kHz have been tested. The logic analyser was also employed as an alternative diagnostic tool when necessary. Finally, resets and power cycles were iterated, verifying the reliability of the device initialisation and its prompt operability.

4.4.3 Analogue Circuit Characterisation

As described in Chapter 2, the main purposes of the BLMASIC and its hosting front-end consist in the acquisition of currents and the production of digital data, which the back-end will properly process. Thus, a preliminary study aims to extrapolate merit factors at the lowest level, compared to what generally happens in commercial ADCs characterisation. In particular, the described methodology involves:

1. the validation of the conversion architecture behaviour;
2. the converter noise characterisation;
3. a basic study of the converter linearity;
4. the analysis of post X-Ray irradiation drifts.

Task 1 is mainly achieved by injecting test currents into the ASIC and monitoring the integrator outputs. This is obtained by configuring the analogue multiplexer with an I^2C access, allowing the probing of the desired signals. Supposing correct operations, the waveforms described in Par. 3.2.1 should be qualitatively visualised on the oscilloscope and their characteristics must be dependent on the input current intensity (cf. Par. 3.2.1 and Par. 3.2.2). The initial test-bench architecture has been employed for this goal (cf. Par. 4.3.1).

The next objective concerned injecting a repertoire of predefined currents into the device and simultaneously acquiring its measurements. The procedure involves: 1) configuring the test current, 2) waiting for the output to stabilise and 3) record the ASIC memory content, by exploiting the oscilloscope data decoder (cf. Par. 4.3.1) or the dedicated readout systems (cf. Par. 4.3.2 and Par. 4.3.3). This is iterated for the various input currents of interest. The ability to store the entire device memory during the session has allowed us to perform batch processing directly on the raw data and to keep track of the device configuration. Such a method has the disadvantage of requiring larger storage space, but the additional acquired information may become convenient

when evaluating various signal processing methods. Therefore, once the optimal data treatment procedure is found, it will be implemented in the back-end cards to perform the computation in real time. In this study, one of the two redundant transmission links can be used at will since their equivalence was already validated. The linearity analysis can be performed with the described procedure, injecting a sweep of current into the device. Thus, conversion characteristic plots can be obtained after the processing. Due to the very high dynamic range specification (cf. Par. 3.3.1), logarithmic current sweeps have been used.

On the other hand, the converter noise study has been performed on the setup described in Par. 4.3.2 and it shares a similar methodology employed for the task 3. However, a single or few test currents are now employed. Here, the key point is represented by the very large amount of recorded data and the associated long measurement time (up to 10^3 s). In fact, the average over a big set of samples is required to evaluate the minimum achievable standard deviation. This represents the resolution limit of the converter, corresponding to the smallest detectable current variation or, alternatively, its intrinsic total noise. These tests can be repeated after the irradiation of the device, in order to investigate possible performance drifts. As it will be described in Par. 4.11, a preliminary linearity characterisation of the devices subjected to X-Rays has been performed, exploiting the test-bench of Par. 4.3.3. Additional irradiation campaigns, which more realistically emulate the final operational environment, will be carried out on the complete front-end and they are described in the Chapter 8.

4.5 Data Analysis

The research involved the employment of several data analysis tools to cope with the different types of investigations. The initial tests and the basic digital validations were performed thanks to the features embedded into the Section laboratory instrumentations. Instead, the successive analogue performance characterisation required the development of complex processing software, since a large amount of data had to be handled. The ASIC noise analysis represents a relevant example, performed using the test-bench described in Par. 4.3.2. There, more than 10^8 data frames have been acquired, including both the measurement registers and the device configuration. Hence, using a high-level MATLAB implementation, the time required to obtain the results of interest could become very long (several hours).

A multi-threaded C++ batch data analysis software has been implemented to efficiently perform the noise computation, allowing the operations to be optimised on multi-core computers. Such a solution has reduced the processing times from hours to less than ten minutes, allowing the results to be quickly verified for various configurations. In particular, the execution flow included the following steps:

1. the raw measurement samples are loaded from a file into the computer memory;
2. the samples are clustered in groups of a specified size, representing a certain time window;
3. the averaging is performed within each of the groups;
4. the standard deviation is computed amongst all averages;
5. the procedure is repeated for an increasing sweep of time windows.

This procedure is equivalent to processing the current measurements by a moving average filter and then decimating the data. Thus, the corresponding outcome error is computed by their standard deviation. The procedure is then iterated for increasing filter time windows, until a plateau of the error values is reached. The existence of a non-zero limit means that other noise sources exist in addition to the thermal one. Without loss of generality, subtracting the average contributions to the measurements, their processing can be expressed in terms of the continuous time model, described by

$$out_t = s_t \otimes w_t \ , \tag{4.1}$$

where out_t represents the processing result, obtained as the convolution $\otimes$ of the measurement signal s_t and the filter impulse response w_t. This latter is

$$w_t = \begin{cases} \frac{1}{T} & 0 \leq t \leq T \\ 0 & otherwise \end{cases} \ , \tag{4.2}$$

which can also be expressed in the frequency domain by its Fourier transform

$$W_f = \frac{Sin(\pi f T)}{\pi f T} = Sinc(fT) \ . \tag{4.3}$$

Considering the measurements and the results as broad stationary random processes with zero average, the equation relating the noise spectrum power densities Out_f and S_f respectively, reads

$$Out_f = S_f \, |W_f|^2 \ . \tag{4.4}$$

Exploiting the Wiener–Khinchin theorem [56] and the definition of the autocorrelation, it is possible to obtain the root mean square (rms) of the output process by its spectrum power density. In particular

$$out_{rms}^2 = \int_{-\infty}^{\infty} Out_f \ \mathrm{d}f = \int_{-\infty}^{\infty} S_f |Sinc(fT)|^2 \ \mathrm{d}f \ . \tag{4.5}$$

To focus the study on the noise contribution affecting the processed results, S_f should be set according to its expected input spectrum. Specifically, if it were only due to thermal effects, S_f would be a constant over the whole spectrum [55], [56].

Thus, substituting $S_f = S_0$ into (4.5) it results

$$out_{rms}^2 = \frac{S_0}{T} \, . \tag{4.6}$$

Therefore, the thermal noise contribution to the results is mitigated as the average times are increased until it is completely suppressed. However, the experimental data confirmed the existence of a non-zero plateau, thus further noise sources with different spectrums are present in the system. Indeed, in addition to the various types of electronic noise characterising the ASIC (Flicker, Shot, Generation / Recombination) [72], [73], significant contributions are due to the quantisation error, the specific analogue to digital conversion methodology and external effects.

The implementation of this analysis code has involved the dataset splitting parallelisation approach. The acquisition samples were clustered in many groups, the number of which depends on the selected averaging time window. Note that longer windows correspond to fewer groups and vice versa, since the total number of samples is invariant. Then, after the clustering, the averaging has to be performed within each group, producing a set of averaged outputs. Since this operation acts independently in each group, it can be distributed to different threads. The same approach has been employed to parallelise the standard deviation computation over the obtained averages. In particular, its numerator computation (the sum of squared variances) has been assigned to parallel threads and the division is performed at the end by a single one.

The challenge faced during this software development resided in the organisation of the computation jobs, which represent the set of operations and data references that the threads must handle. In order to optimise the heavy processing, the number of concurrent working threads should be kept equal to the available CPU logic cores for the whole calculation time. In modern computers, such a number (Q) can be the same as the physical cores or the double, whether the CPU is equipped with hyper-threading technology [74]. Following this methodology, the analysis time will be theoretically reduced by a factor Q, compared to the single thread implementation. Reaching this factor is actually not possible due to the overhead of both the operating system thread manager and the launcher of jobs, but it is possible to get close enough to it. Similar approaches have been exploited for the linearity study after the device irradiation, the results of which will be presented later in Par. 4.11.

4.6 Basic Electrical Results

Here, the measurements of the typical ASIC current consumptions are reported for the powering rails. No significant changes are expected over the different upgrades, consisting of 2 $\Delta\Sigma$ versions and 3 CFCs. The device was initialised with the default configuration and the clock was provided. The supply currents, measured for the BLMASIC-CFCv3 device, are:

- 16 mA on the main digital rail (1.2 V),
- 34 mA on the main analogue rail (1.2 V),
- 370 μA on the common-mode reference (0.6 V),
- 100 μA on the low leakage input stage rail (2.5 V).

Instead, the same for the BLMASIC-ΔΣv2 are:

- 20 mA on the main digital rail (1.2 V),
- 17 mA on the main analogue rail (1.2 V),
- 2.7 mA on the common-mode reference (0.6 V),
- < 100 μA on the low leakage input stage rail (reduced from 2.5 V to 1.7 V).

The additional bus introduced in version 2 (2.5 V) reported an issue which dramatically affected the current consumption and the conversion performance. In fact, the device drew an excessive current compared to what was expected by the simulations. Fig. 4.9 reports the measured currents versus the voltage supplying such a bus. The plot trend is exponential for voltages below 2 V and its growth fades for higher values. The initial exponential behaviour suggested an undesired junction polarisation. Indeed, after the design check, wrong connections were found in the cathode of input clamping diodes, which cause them to be directly polarised towards the 1.2 V bus. Such a problem forced the characterisation to be carried out with a reduced bus voltage of 1.7 V. It has been studied that this would not have influenced the measures. The third version was released with this issue solved and the rail has been supplied with the correct 2.5 V voltage.

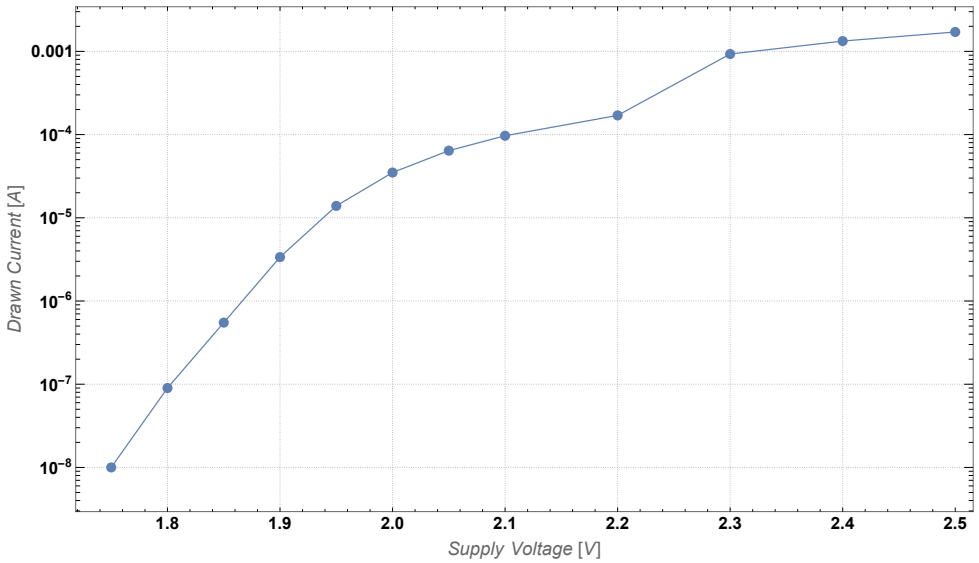

Figure 4.9: Actual current drawn from the 2.5 V rail in version 2 versus the supply voltage.

4.7 Digital Characterisation Results

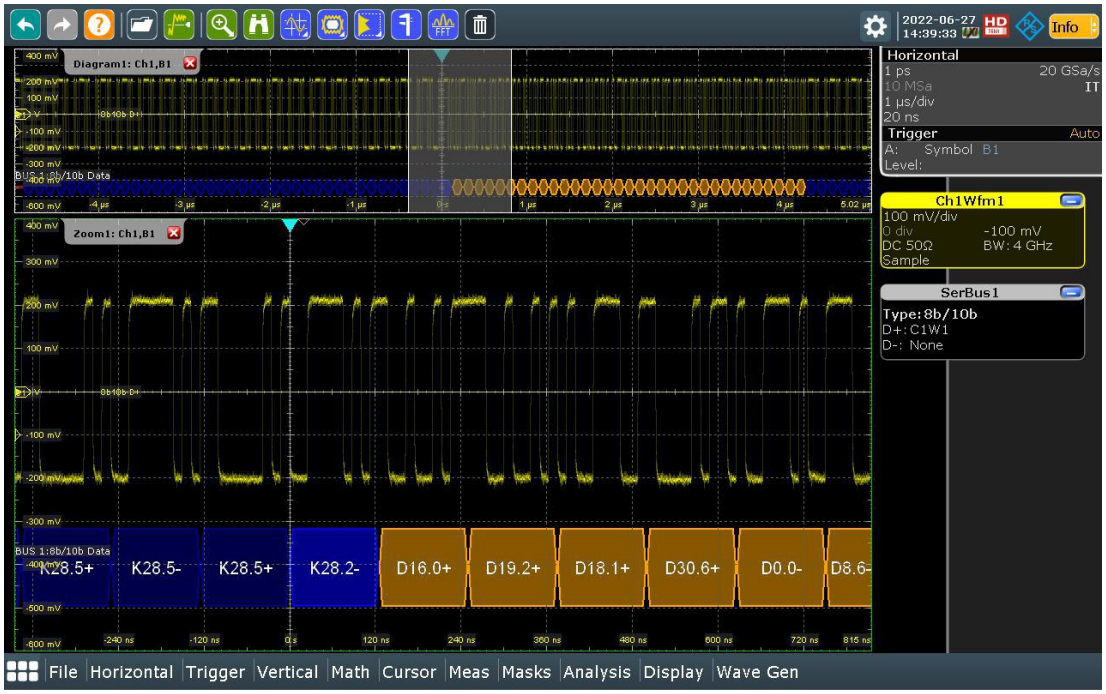

Figure 4.10: Visualisation of the differential data transmission, where the decoded words are shown in the bottom line.

During the initial test, the serial data interface has been connected to the oscilloscope (see Fig. 4.10) according to the differential probe methodology described in Par. 4.4. An overall $\sim 1.2\ \mu s$ acquisition window is shown, with the voltages reported in the vertical axis. A differential amplitude of $\approx 400\ mV$ has been measured, which agrees with the simulations carried out during the design. Furthermore, an eye diagram analysis was performed on such a data stream (results in Fig. 4.11), confirming the compliance with the specifications and good integrity of the signal.

Later, the high-level protocol decoding and checking were performed. The device passed these tests successfully and no data integrity issues or erratic behaviours have been observed. The screen in Fig. 4.10 also includes an example of the data decoding, in which the 8b/10b control symbols *K28.5* and *K28.2* (in blue) have been distinguished from the payload words (in yellow). The validation confirmed that the frame morphology agrees with the structure set during the design, presented in Par. 3.3.8. Furthermore, concurrent access by the I^2C interface has been employed to query the entire memory map and to perform a comparison with the received data. Their equivalence ensured the transmission data validity. These tests have been carried out for both the redundant links, driving the "ready" signals accordingly (cf. Par. 3.3.8).

Afterwards, the investigation turned on the current measurement registers. Concerning the first version of the CFC device, despite its main counters showing results compliant with the monitored integrator waveform, the Wilkinson ADC was affected by

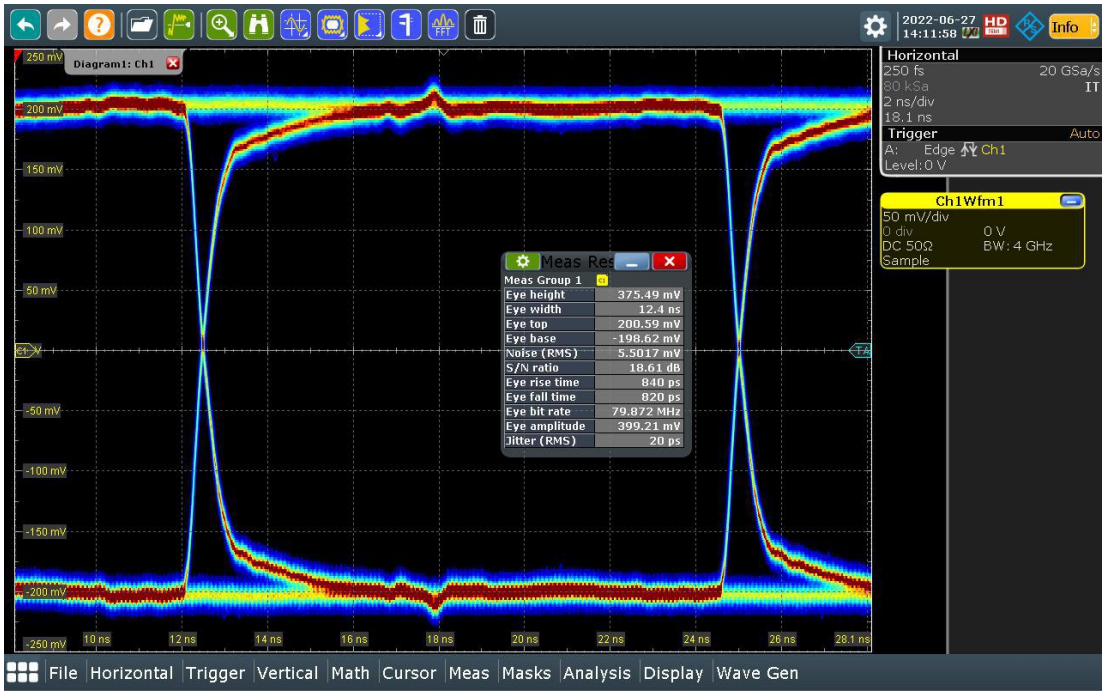

Figure 4.11: Eye diagram of the data transmission link at 80 *Mbps*.

an unexpected behaviour (cf. Fig. 4.12). Although the integrator output appears to approximate a symmetric triangular waveform properly, many glitches are present in the data. Hence, the probe multiplexers have been configured to monitor the integrator, to verify whether such effects are also detectable in the analogue signal. The investigation concluded that the glitches are not present in the measured analogue signal, so their origin should be in the auxiliary ADC implementation or in the device logic. Therefore, a specific acquisition firmware was implemented for such a goal. In particular, referring

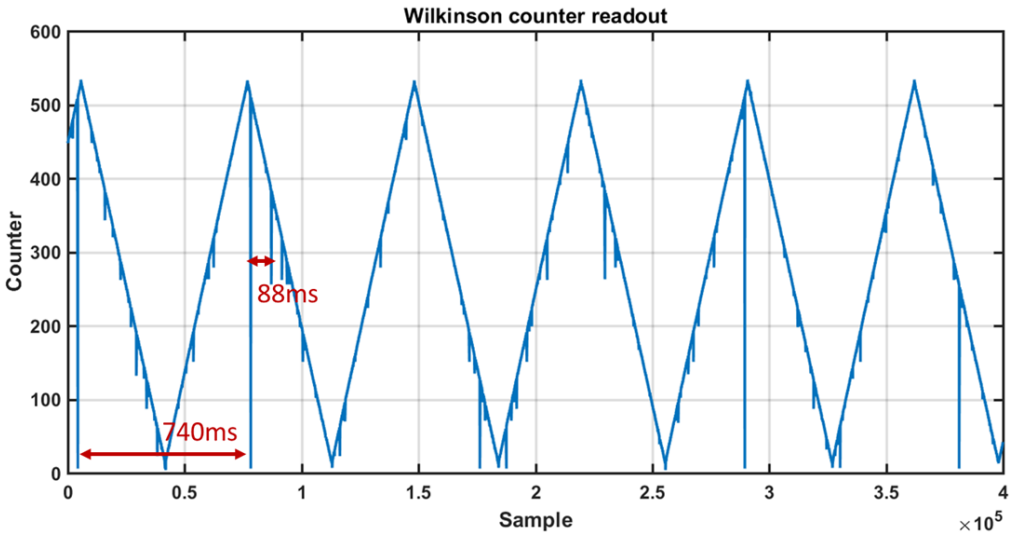

Figure 4.12: Glitches observed in the Wilkinson ADC data of BLMASIC-CFCv1. ADC readings are reported versus the acquired samples, employing a test current of 100 *pA*.

to the architecture presented in the Par. 3.3.7, a copy of the Wilkinson time counter was implemented within the FPGA and the **phi3** digital flag was probed out to control it. Thus, the internal device logic was bypassed and the successive test verified that a correct measurement could have been obtained in this way. So, the problem was isolated at the interface between the analogue and digital domains, where the counter is implemented. The study concluded that **phi3** was causing sporadic metastability phenomena [67], [75] in the internal counter, thus a synchronisation stage was included in version 2 of the device. The same test was repeated on the new device and did not report similar issues. Fig. 4.13 shows an example of correct Wilkinson data versus the time, for 1 nA injected current, while Fig. 4.14 refers to the probed integrator outputs in the same conditions.

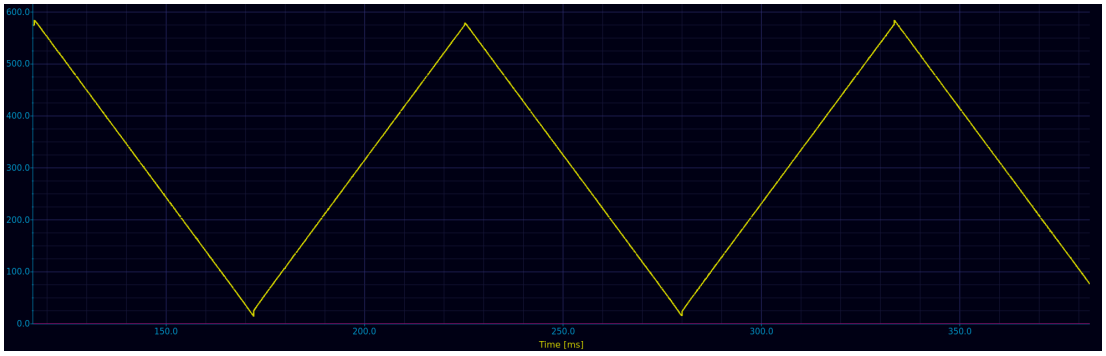

Figure 4.13: Wilkinson ADC data obtained by the BLMASIC-CFCv3 at 1 nA.



Figure 4.14: Integrator output waveform for the BLMASIC-CFCv3 at 1 nA. The histogram reports the distribution of the waveform frequencies.

Finally, when testing the device with different power supply sources, a sporadic internal FSM lock was observed if significantly fast ramps were applied to the digital power rail. In order to overcome this last issue, an upgrade of the device logic has been developed and released in the latest 3rd version.

4.8 Analogue Characterisation Results

4.8.1 Conversion Precision

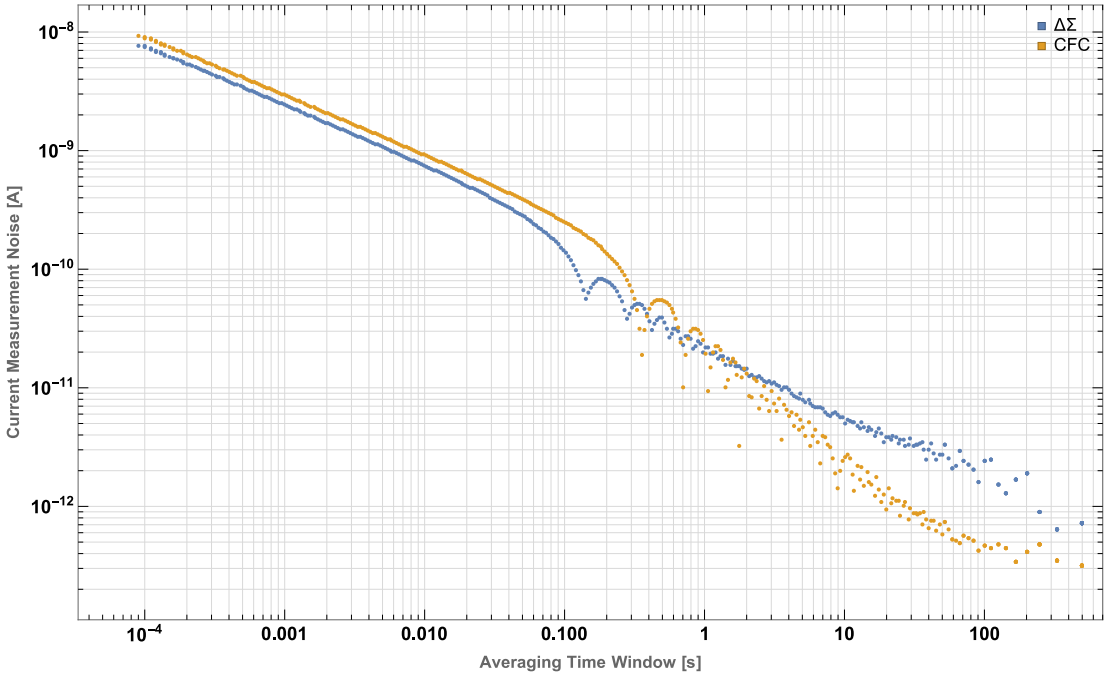

Figure 4.15: BLMASIC measurement noise for different averaging time windows, characterised at the test current of 1.5 pA.

The methodology described in Par. 4.4 and the optimised data analysis software (cf. Par. 4.5) provided the first estimation of the maximum resolution, also evaluating how long time is needed by the instrument to achieve a given precision. The results are reported in Fig. 4.15, where the total measurement noise has been plotted over a sweep of different averaging time windows. As expected, longer windows lead to smaller errors, since the converter noise spectrum is progressively phased out from the higher to lower frequency components. Because of the reduced number of clusters for times longer than 10 s, the estimations suffer from some fluctuations, which can be smoothed by acquiring a larger amount of data. Furthermore, it is possible to notice the gentle beginning of the plateau for the more extended integrations, as discussed in Par. 4.5.

In the BLMASIC specifications (cf. Par. 3.3.1), it was required to measure current signals below 10 pA. The presented characterisation proves that such precision is not limited by the overall converter noise in both the device architectures. In fact, it can be observed that an error below 1 pA can be achieved for averages longer than 200 s.

4.8.2 Parallelised Data Analysis

Regarding the efficiency of the parallelised computation, Fig. 4.16 shows the time required to perform the steps presented in Par. 4.5 with respect to the different average filter windows. Since the same number of threads and samples have been employed for each computed point, lower time values correspond to better exploitation of the computer resources, thus shorter overheads. As expected, the software runs optimally with medium-sized sample groups, with about 600 ms required for each iteration, which is typical for most cases. Instead, it suffers from performance degradation when the groups are very numerous (small time windows) or less but contain more samples (longer time windows). The first scenario is due to the overhead time caused by the manager of jobs. In other words, it is more expensive to organise the parallelisation than to perform the computation itself. The second concerns the operation on a few but huge sample sets.

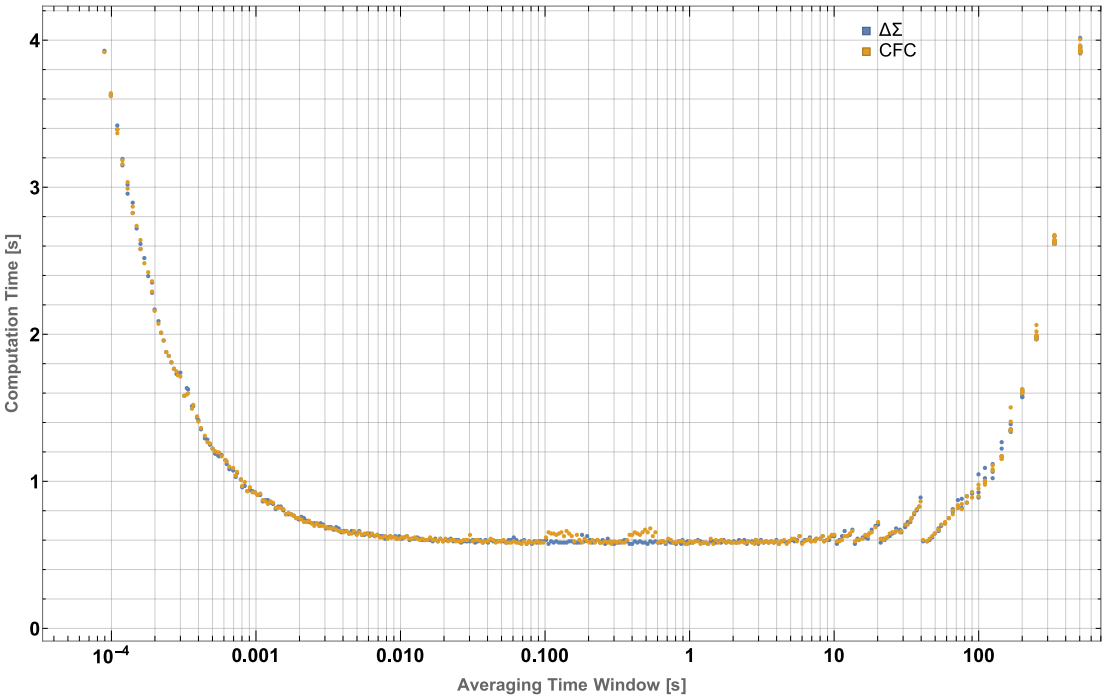

Figure 4.16: Computation time required by the measurement noise characterisation, plotted for the considered averaging time windows.

4.8.3 Maximum Achievable Resolution

An alternative type of measurement has also been carried out for the CFC device to check the stability of the asynchronous conversion loop over a long measurement time. The external current source has been disconnected and the calibration DAC has been set to 15 pA. The integrator outputs have been again monitored with the oscilloscope differential probe (at 1 MSs^{-1}) and data were acquired for about 4 days. The results about the waveform characteristics are reported in Fig. 4.17, showing a mean waveform

period of 6.809 s, with an error $\sigma = 33.14\ ms$ (thus $Jitter_{RMS} \approx 23.4\ ms$). The plotted histogram represents the expected Gaussian distribution of the period measurements over 32881 acquisitions.

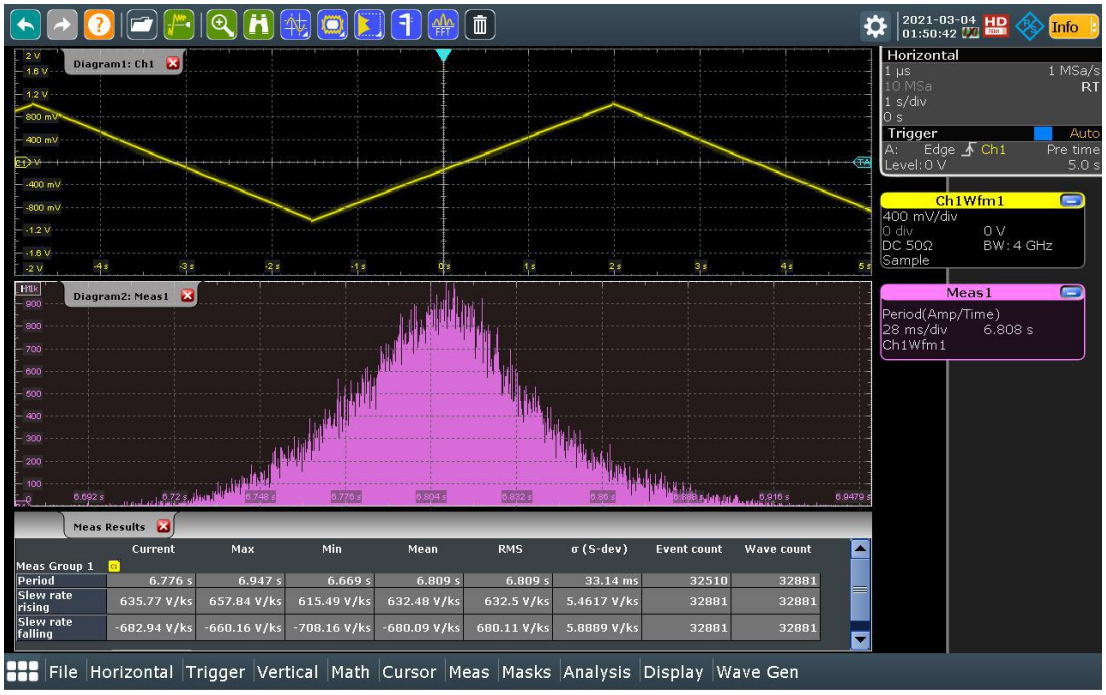

Figure 4.17: Integrator output waveform (top) of the BLMASIC-CFCv3 at 15 pA and distribution of the waveform periods throughout 4 days of measurements (bottom).

Complementary to the discussion of Par. 4.8.1, also these results provide valuable information about the maximum converter performance. In fact, since the period relative uncertainty is 0.487 %, the corresponding measurement absolute error is 73 fA, assuming that the calibration DAC is accurate. This value can be considered the resolution limit achievable by the device for currents of that magnitude. If the own device logic carries out the same measurement, it would likely be affected by a higher error because of the modulation noise introduced by the conversion architecture and the auxiliary ADC quantisation. The processing employed to minimise their effects is described in Chapter 5, and the final results are included in the overall system validation (Chapter 8).

4.8.4 Integrating Capacitance

The long duration acquisitions reported in Fig.s 4.14 and 4.17 have also been exploited to measure the BLMASIC-CFCv3 integrator capacitors. In particular, by using the equations presented in the Par. 3.2.1 and the quantiser threshold consideration of Par. 3.3.5, the relation

$$C = \frac{\overline{i_{in}}}{2\, f\, V_{thr}} \tag{4.7}$$

holds, for $V_{thr} = 0.96\ V$. Thus, averaging the results of the two measurements, it is obtained $C = 58.8\ pF$, which is within the 2% of the design value and well inside the requirements. The corresponding frequency to current ratio results

$$\frac{f}{i_{in}} = 8.86 \times 10^9\ Hz\ A^{-1}\ . \tag{4.8}$$

4.8.5 Integration Asymmetry

In Fig. 4.17 are also reported the ramp slew rates, measured during the rising and the falling times of the triangular waveform. Their averaged values result to be $SR_{rising} = 632.48\ mVs^{-1}$ and $SR_{falling} = -680.09\ mVs^{-1}$. They can be exploited to evaluate the integration asymmetry.

To this aim, it is necessary to calculate the T_{rise} and T_{fall} by the equation

$$SR_{rising} T_{rise} + SR_{falling}\, T_{fall} = 0\ , \tag{4.9}$$

which is true since the waveform is periodic. Being the sum $T_{period} = T_{rise} + T_{fall}$, we can obtain

$$T_{rise} = \frac{SR_{falling}\, T_{period}}{SR_{falling} - SR_{rising}}\ , \tag{4.10}$$

$$T_{fall} = -\frac{SR_{rising}\, T_{period}}{SR_{falling} - SR_{rising}}\ . \tag{4.11}$$

Furthermore, by using the equation (3.2) it is possible to obtain the currents flowing into the integrator capacitors during the rising and the falling ramps. The relations hold

$$\overline{i_{rise}} = \frac{C\, V_{swing}}{2\, T_{rise}}\ , \tag{4.12}$$

$$\overline{i_{fall}} = \frac{C\, V_{swing}}{2\, T_{fall}}\ . \tag{4.13}$$

Hence, the leakage current $\overline{i_{leakage}}$ is obtained by defining

$$\begin{cases} \overline{i_{rise}} = \overline{i_{input}} - \overline{i_{asym}} \\ \overline{i_{fall}} = \overline{i_{input}} + \overline{i_{asym}} \end{cases}\ , \tag{4.14}$$

thus

$$\overline{i_{asym}} = \frac{\overline{i_{fall}} - \overline{i_{rise}}}{2}\ . \tag{4.15}$$

Finally, substituting the measurements into the equations, it results: $T_{rise} = 3.53\ s$, $T_{fall} = 3.28\ s$, $\overline{i_{rise}} = 14.5\ pA$, $\overline{i_{fall}} = 15.6\ pA$ and $\overline{i_{asym}} = 0.55\ pA$. For a better accuracy, the integrating capacitance used in this calculation is the one measured by the $i_{in} = 15\ pA$ data only (see Fig. 4.17), that is $C_{15pA} = 53.2\ pF$.

4.8.6 Calibration DAC

This investigation focused on the measurement of the calibration DAC characteristic (cf. Par. 3.3.6). Its configuration register was set with a sweep of values, from 0 to 255, and the resulting CFC waveform frequencies were recorded. Thus, the converter itself was employed to validate the linearity of its embedded DAC. Only the positive currents have been tested, while the negative bias has been left deactivated during this study. Fig. 4.18 report the obtained results, where the currents have been computed by the frequencies using the Eq. 3.5 and the already measured integration capacitance.

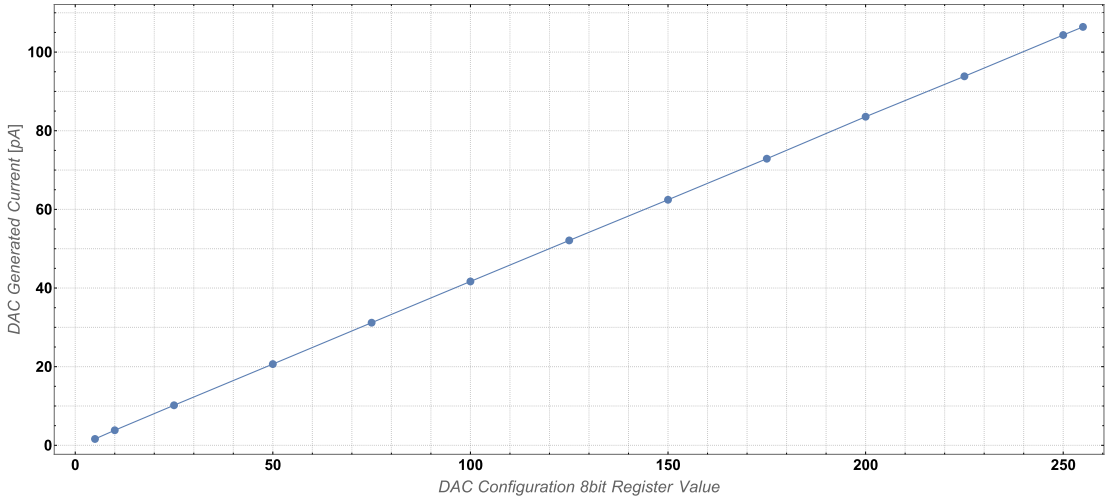

Figure 4.18: Calibration DAC characteristic.

A slightly lower gain is observable since the full-scale current should theoretically reach 128 pA, rather than the measured 106 pA. It is unlikely that such an effect is due to the DAC implementation, since a precision reference has been employed for its internal biasing. Probably, the integrator capacitance value or the integrator output swing have been mildly underestimated. However, it is possible to notice that good linearity is observable over the whole range.

4.8.7 Integrator Output Issue

Now, the accuracy of the triangular waveform (cf. Par. 3.2.1) has been taken under consideration. In particular, a persistent non-ideal transition has been noticed close to the current reversing peaks. The effect appears as a jump and can be observed both in Fig. 4.19 and the more detailed colour grade plot in Fig. 4.20. The rationale for the phenomenon is still under investigation and assumptions have been made about parasitic effects within the integrated circuit. The supply, the common mode and the input node voltages have been monitored during such events and no relevant spikes have been reported so far. Furthermore, it was not observed a marked correlation with the injected currents. On the other hand, it is proven that the height of jumps increases with the input source capacitance.

In summary, this issue seems not to cause dramatic degradation in the device conversion performance. Still, it is definitely a critical challenge to be addressed in the eventual production of a new version of the ASIC.

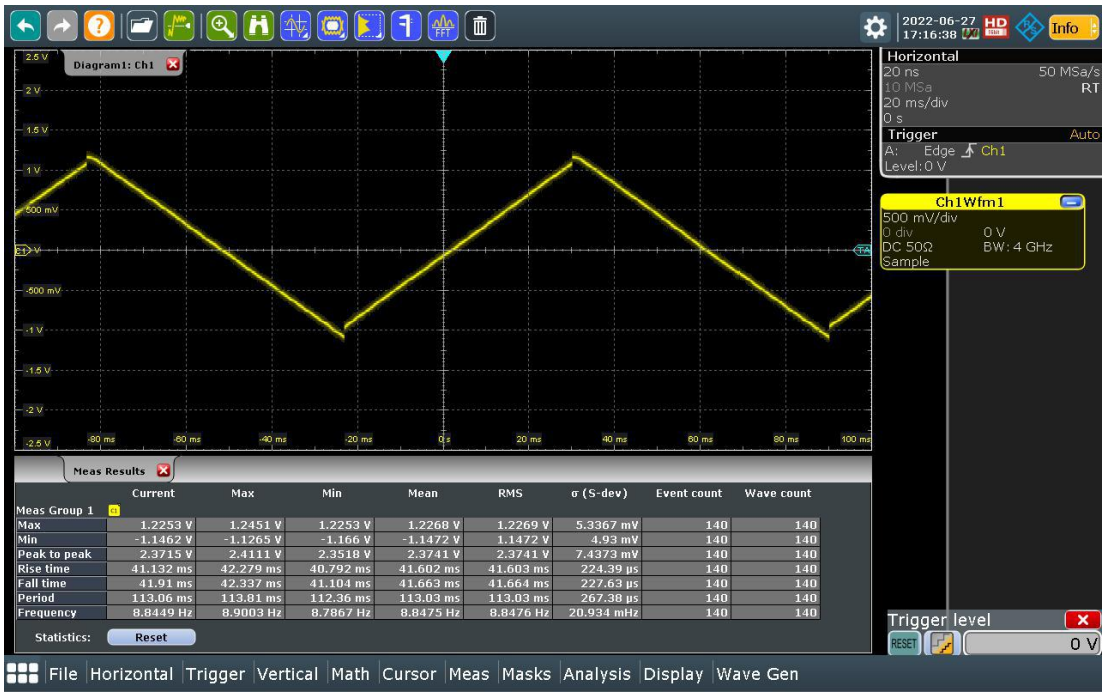

Figure 4.19: Jumping issues visible in the integrator output at the current reversing positions. An injected current of 1 nA is used and the source capacitance is 1 nF.

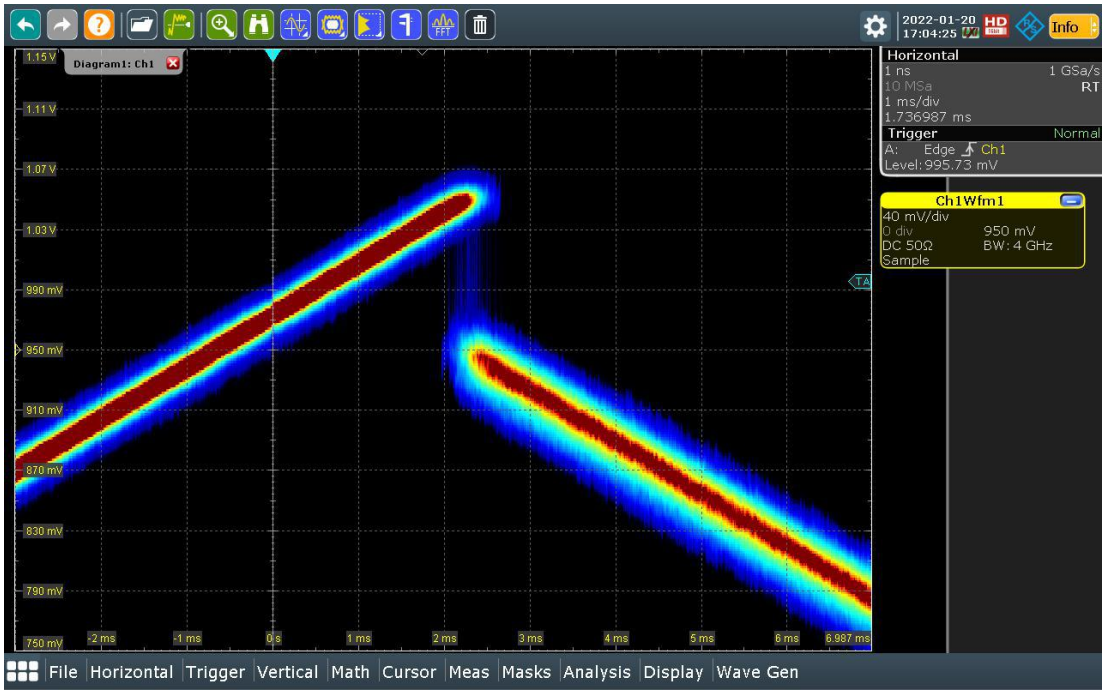

Figure 4.20: Colour grade detail of the integrator waveform jump.

4.9 Selection Verdict

The results of this preliminary characterisation have provided a better understanding of the two architecture capabilities within the weak current measurement context. Furthermore, although some minor design issues have appeared, the developed test-benches offered valuable tools to achieve their solution and release upgraded device versions readily. No significant criticalities have been found at this stage and both the CFC and the $\Delta\Sigma$ based converters have reported auspicious performance. However, following many discussions within the team, it has been decided to target the efforts on the CFC device, as it was also possible to notice from the amount of results collected.

The first reason concerns the better precision performance reported by such an architecture in the lower currents (cf. Fig. 4.15). Indeed, most of the installed BLM detectors of LHC operate within this range and great interest has arisen in enhancing the system performance in such a measurement window. Second, the permanent staff engineers of the Section possess excellent know-how with the CFC converter topologies, since these are implemented in the operational front-ends (cf. Par 2.3.1) [34], [35]. This can facilitate subsequent maintenance of the plant and greater awareness in case of possible failures.

In conclusion, although neither of the two solutions has been technically discarded, more confidence was given to the CFC, considering both the scientific outcome and the commission planning strategy for the HL-LHC upgrade.

4.10 Changelog of Versions

Here, a log of the device upgrades carried out after the first prototype production is reported. The term BLMASIC will be used hereafter only to identify the CFC-based device.

4.10.1 BLMASIC Version 2 (2020)

The characterisation of the first BLMASIC prototype showed promising results and no major issues had been encountered. However, the high current leakage of the input stage FETs drove out the measurements in the low current ranges from the desired accuracy. Thus, by substituting the FETs in input stage and the switches (cf Par.s 3.3.3 and 3.3.4) with ones employing thick oxide gate [67], [76], [77], the current leakage has been significantly reduced. Furthermore, an additional synchronisation stage has been inserted between the Wilkinson ADC quantiser and its dedicated counter to solve the metastability issue [67], [75]. This has removed the glitches affecting the data.

4.10.2 BLMASIC Version 3 (2021)

In this last release, the wrong connections of the input clamping diodes have been corrected (cf. Par. 4.6). Thus, it is now possible to operate the 2.5 V supply rail at its nominal voltage. Furthermore, the internal FSM lock issue, occurring with fast supply power-up ramps, has been solved. Finally, the ceramic package employed in the previous releases has been changed to the cheaper epoxy one (cf. Par. 3.3.9). In fact, the final hardware will most likely adopt such a solution.

4.11 X-Ray Irradiation Tests

DUTs	Flux [kGy/h]	Time [h]	TID [MGy]
CFC$_{\#}$1	44.3	23	1.02
CFC$_{\#}$2	15.4	52	0.80
CFC$_{\#}$3	44.3	23	1.02
CFC$_{\#}$4	44.3	23	1.02
CFC$_{\#}$5	15.4	65	1.00

Table 4.1: BLMASICv2-CFC X-Ray irradiation sessions.

In the first quarter of 2021, a preliminary study was performed to validate the device's radiation tolerances. It has allowed monitoring of any conversion performance degradation and validating of the main operational functionality of the digital interfaces. In particular, the conversion characteristic of a non-irradiated device was compared with five other units subjected to X-Rays. The irradiation sessions were undertaken at CERN and the details about the Flux and the TID are reported in Tab. 4.1.

The converter data analysis has involved similar steps to those presented in Par. 4.5 for the noise characterisation. However, the averaging time window has been fixed to 100 μs, corresponding to 10 samples, and the error computation has been iterated for a sweep of input currents. The sweep included 500 current intensities varying exponentially from 1 nA to 1 mA, while about 1000 samples have been acquired for

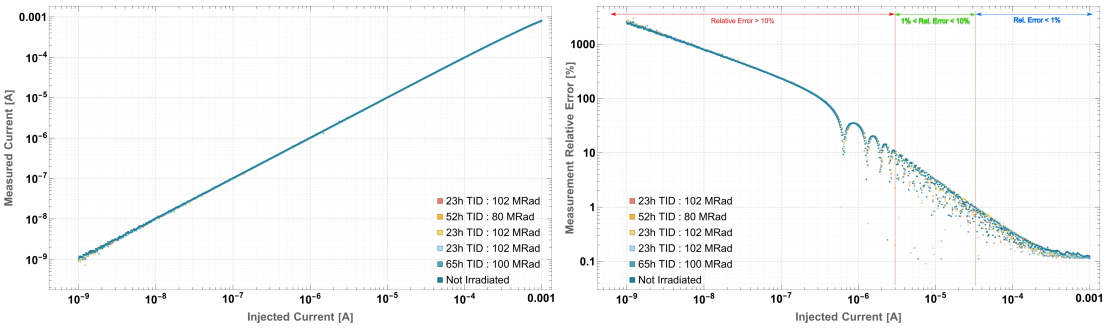

Figure 4.21: Conversion characteristics (left) and relative errors (right) by a logarithmic current sweep of 500 values from 1 nA to 1 mA. The averaging time window was set to 100 μs.

each iteration. The analysis results are shown in Fig. 4.21, where the left plot reports the averaged measurements and the right plot the corresponding errors versus the injected currents. The different colours in the legend specify which was the tested device.

No major performance drifts have been observed amongst the devices up to 1 MGy of TID and the digital interfaces operated correctly during the whole irradiation session. The results reveal errors smaller than 10 % for measurements of currents greater than 3 μA, while the same stays below the 1 % for currents above 35 μA. Although high errors can be observed for the currents below 3 μA, these achievements comply with the specifications, since the small average window used (100 μs) is only effective for the high current range. The ripples in the error plot are an effect due to the CFC counter embedded into the device. A detailed study of such a phenomenon and the signal processing needed to carry out precision measurements will be presented in Chapter 5.

Chapter 5

Signal Processing

The need for signal processing is introduced, focusing on the concepts of accuracy, precision and response speed required by the research objectives. After a summary of how the raw measurements are derived from the BLMASIC physical behaviour, it is described how the data of CFC counters are combined with the auxiliary ADC ones, to improve the precision. Thus, the Fourier spectrogram of the acquired data is shown, highlighting the features of the total converter noise. That is exploited in an original methodology for measurement error prediction. In parallel, the standard running sum data treatment is formalised in the realm of Signal Theory and novel techniques based on IIR filtering are introduced. To allow an effective comparison amongst these different algorithms, the notion of processing metrics, tailored for the BLM, is explained. Thus, the IIR filter tunings are set and the benefits brought to the measurement precision and the step response speed are demonstrated on the real data. Because of its sharp noise suppression, this method can be applied in more general contexts. Finally, instead of a typical batch data analysis, it is shown the convenience of prototyping the real-time filters by custom computer software.

5.1 Accuracy and Precision

Alongside developing new converters, improved signal processing methodologies have been tailored to cope with the requirements of the upgraded BLM system. In particular, the signal theory of the operating system has been formalised and merit factors have been introduced, to allow an effective comparison with the alternative solutions. In Par. 3.3.8, the way the BLMASIC logic handles current measurements was introduced and it was pointed out that they are propagated with a frequency of 100 kHz. The present research is based on considering such a stream of data as a signal sampled over time (m_k), which in steady-state conditions can be expressed as

$$m_k = m^{dc} + m_k^{noise} \,, \tag{5.1}$$

where: m_k is the actual sequence of values produced by the converter, m^{dc} is the continuous component of the measurements and m_k^{noise} is a zero average random variable, representing the total contribution of noise. Note that, in contrast to the real case, it has been assumed that the system has to monitor constant currents, the values of which correspond to m^{dc}, whether the converter accuracy is considered arbitrarily high. Both these hypotheses have been introduced to ease the discussion and they are sufficiently general, although they will subsequently be relaxed. In fact, despite the signal of interest can evolve quickly, its frequency spectrum is generally developed around the DC component (baseband signal) instead of the noise one. Furthermore, the implemented processing is intended to optimise the conversion precision, rather than the accuracy, which can be unattended at this stage.

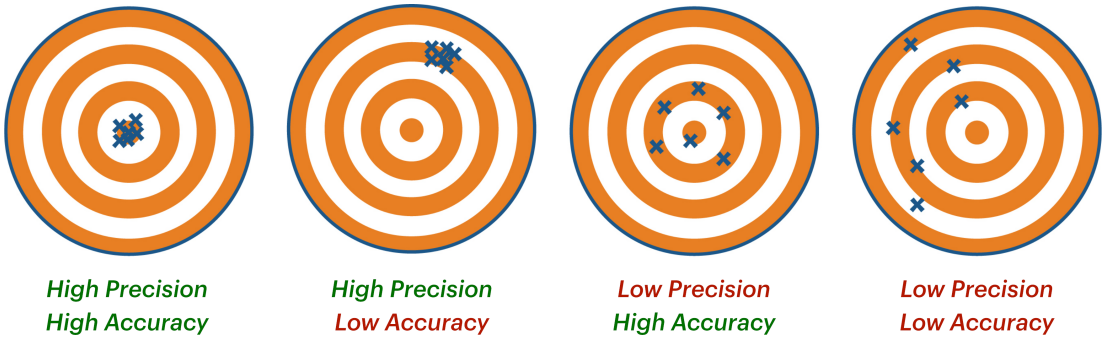

Figure 5.1: Graphic representation of precision and accuracy of a measuring instrument.

To better understand the processing aims, it is possible to refer to Fig. 5.1. An accurate instrument produces samples, the long-term average of which (m^{dc}) is very close to the real value, no matter how dispersed they are. Instead, the samples of a precise instrument are low dispersed, although their average may deviate far from the actual value. This can be translated into exhibiting a low root mean square (σ_m).

Ideally, the system should be both accurate and precise, but the former feature is initially more challenging to achieve. Indeed, the CFC converter does not include embedded absolute calibration electronics, hence the data are not inherently compared with an accurate reference. Nevertheless, the same objective can be effectively achieved by requiring the system to be precise and monotonous regarding the measuring quantity. In fact, under such conditions, it is possible to obtain a unique relationship linking the measurement signal with the real values. This could be obtained by performing the inference of data based on a model function, thus applying its inverse to data as the last step of the processing.

The BLMASIC architecture described in Chapter 3 and the preliminary validation of its capabilities (cf. Chapter 4) provided a detailed background of the how the analogue to digital conversion is performed. Here, the resulting data stream is analysed from a higher point of view. In particular, the following paragraphs describe how the measurement signal is composed and which filtering solutions can be adopted to minimise σ_m. Since real acquired data will be employed to extract the quantities of interest,

m_k^{noise} will eventually be considered a deterministic signal. Theoretically, the same considerations can be attained by studying it as a broad stationary random process with zero average. Finally, it must be underlined that the research is focused on improving the processing methods from a linear signal theory point of view. At this stage, the effects of quantisation errors can be neglected by employing a double precision floating-point computation. In Par. 7.8, it will be proved that such a numeric representation is also employable in the real back-end hardware. This further strengthens the validity of the presented methods in view of a possible role in future accelerators.

5.2 Measurement Composition

The measurement data stream includes the CFC counters and the auxiliary ADC values. The first report the number of times that the integrator output exceeded the set thresholds in a specific observation period, being directly related to the frequency of the produced triangular waveform (cf. Par. 3.2.1). Whereas the auxiliary ADC directly monitors the latter (cf. Par. 3.3.7), with a sampling frequency equivalent to the frame transmission of 100 kHz.

By design, it is established that the acquisition of weak currents produces very slow ramps at the integrator output, which the auxiliary ADC can efficiently track. However, as the current increases, the triangular waveform frequency exceeds the ADC limit of 50 kHz (by the Nyquist-Shannon criterion [56]), so the resulting samples will be affected by the aliasing. During the initial tests, the two data sources were processed separately, evaluating the measurements over current ranges chosen according to the handled source. The results confirmed that the CFC counters data are characterised by a good SNR from medium to high-intensity currents (above 10 nA), while the auxiliary ADC provides better precision in the lower range. On the contrary, this latter is not effective above $\approx$ 5.6 μA, since the aliasing effects degrade the measurements. Such results prompted the investigation of solutions able to combine the data sources, producing unique measurement values to be propagated to the subsequent treatment.

One of the studied proposals is based on the methodology already validated for the currently operational system [54]. It concerns the evaluation of the input current by measuring the charge acquired during a 10 μs period. This is obtained by combining the CFC counter measurements, the auxiliary ADC readings and the ramp slope information. In Fig. 5.2, three waveform examples are shown for different injected currents. Since the ADC sampling rate is equal to the overall data transmission one, the produced results s_k include the ADC voltage s_k^v, the number of waveform peaks (sum of positives and negatives) that occurred since the previous transmission s_k^{CFC} and the instantaneous slope polarity s_k^{slope} (rising, falling). According to the converter architecture (cf. Par. 3.2.1), the waveform visible in the figure results from the input current integration and the changeover circuit action. In particular, the integrator output equation (3.1) can be resolved for a single period, assuming that $i_{in}(\tau)$ can be replaced by its average

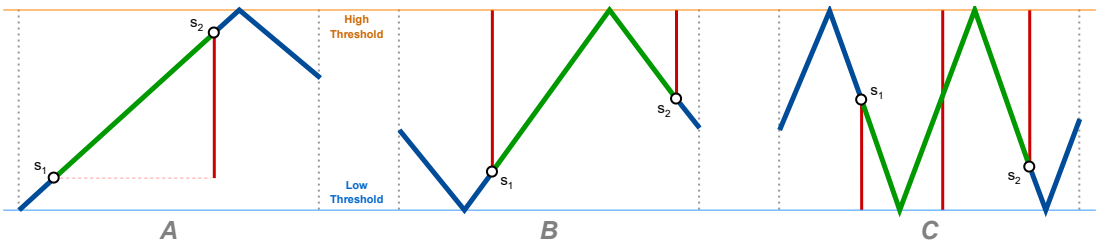

Figure 5.2: Example of the integrator waveform and its sampling by the auxiliary ADC for different currents (A is the lowest one and C the highest).

during the observation time $\overline{i_{in}}$. The relation holds

$$\begin{aligned} |v_{out}(t_2) - v_{out}(t_1)| &= \frac{2}{C} \int_{t_1}^{t_2} i_{in}(\tau)\ \mathrm{d}\tau = \\ &= \frac{2}{C}\ \overline{i_{in}}\ (t_2 - t_1) \implies \\ &= \frac{2}{C}\ \overline{i_{in}}\ T\ , \end{aligned} \tag{5.2}$$

where T represents the sampling period of 10 μs and C the integrator capacitors value (60 pF). As indicated on the left-hand side, since $\overline{i_{in}}$ is always positive by design, only the modulus of the voltage difference is meaningful. Furthermore, the expression requires the condition that $v_{out}(t)$ is monotonous within the same slope polarity. Hence, the average input current can be tracked by knowing the voltages of consecutive samples. Assuming $s_k^v = v_{out}(t_2)$ and $s_{k-1}^v = v_{out}(t_1)$, it is possible to estimate the current measurement $i_k^{in} = \overline{i_{in}}$ by the (5.2), being

$$i_k^{in} = \frac{C}{2\ T}\ |s_k^v - s_{k-1}^v|\ . \tag{5.3}$$

This solution can be directly applied whenever the latest two acquisitions (s_k and s_{k-1}) fall in the case **A** of Fig. 5.2. It can be observed that the length of the red line within the figure is actually equal to $|s_k^v - s_{k-1}^v|$, so it is proportional to the measured current i_k^{in}.

Similar considerations can be adopted for case **B**, which occurs when s_1 and s_2 are interspersed by a positive or a negative peak. Here $s_k^{CFC} = 1$ and the equation for the current (5.3) has to be modified by adding the two integral contributions independently, on the basis of the slope polarity informations s_k^{slope} and s_{k-1}^{slope}. For the example in figure, s_k^{slope} is falling and s_{k-1}^{slope} is rising, hence the modules become $v_{Thr}^H - s_{k-1}^v$ (*early*) and $v_{Thr}^H - s_k^v$ (*late*), which correspond to the lengths of the two red lines. The last case **C** is addressed by the same discussion, however $s_k^{CFC} > 1$, thus the contribution of full slopes (*centre*) must be added to the result.

By summarising these considerations, it is possible to obtain the equations valid both for cases **B** and **C**.

$$i_k^{early} = \begin{cases} \frac{C}{2\,T}\,(s_{k-1}^{v} - v_{Thr}^{L}) & with\ s_{k-1}^{slope}\ falling \\ \frac{C}{2\,T}\,(v_{Thr}^{H} - s_{k-1}^{v}) & with\ s_{k-1}^{slope}\ rising \end{cases}, \tag{5.4}$$

$$i_k^{late} = \begin{cases} \frac{C}{2\,T}\,(v_{Thr}^{H} - s_{k}^{v}) & with\ s_{k}^{slope}\ falling \\ \frac{C}{2\,T}\,(s_{k}^{v} - v_{Thr}^{L}) & with\ s_{k}^{slope}\ rising \end{cases}, \tag{5.5}$$

$$i_k^{centre} = \frac{C}{2\,T}(v_{Thr}^{H} - v_{Thr}^{L})\,(s_k^{CFC} - 1)\,, \tag{5.6}$$

with

$$i_k^{in} = i_k^{early} + i_k^{centre} + i_k^{late}\,. \tag{5.7}$$

Therefore, the processing firmware has to exploit the information received to detect the device working conditions and decide the equation to be used, between (5.3) and (5.7), for the estimation of the current. A prototype C++ code was developed to execute the presented method in real time and was widely adopted during the system validation.

The effectiveness of this method was also evaluated from a signal theory point of view. In fact, despite the reasoning carried out being very intuitive, it assumes the strong hypothesis of a noiseless signal. This is not valid in the actual case, since the difference $s_k^v - s_{k-1}^v$ behaves as a Finite Impulse Response (FIR) [55] high pass filter, emphasising the noise effect further. Thus, an additional condition has been imposed on the auxiliary ADC data to mitigate the drawbacks. In particular, new values s_k^v are neglected if they do not guarantee the monotony of the slope. It is possible to prove that this reduces the noise effects, but still allows the correct computation of the current.

Finally, alternative methods have been developed to exploit longer sequences of past data. For instance, linear regression can be computed over many acquired samples that belonged to the same slope. In fact, the resulting angular coefficient provides a direct estimation of the input current. This solution was used during the validation of the BLMASIC as part of the batch data analysis, and it has allowed us to carry out precise measurements in a relatively short time. However, its actual use in the upgraded system is still under discussion, since it involves a considerable implementation complexity within the target FPGA device.

5.3 Spectrum of Converter Measurements

The development of the processing architecture has begun by analysing the frequency spectrum of the signal m_k (cf. Par. 5.1) and its dependence on the injected current.

This allows to efficiently describe the converter behaviour over the whole measurement range and to carry out suitable tailoring of the required processing. In particular, the sequence m_k must be appropriately treated to produce an output r_k, which includes the same information of interest, but with an error $\sigma_r < \sigma_m$. Of course, this leads to some adverse implications, described later, which can anyway be properly controlled.

The spectrum analysis was carried out exploiting the system validation test-bench, the details of which will be described in Chapter 8. It involved the following steps:

1. a list of 5×10^3 test currents is calculated by an exponential sweep from 10 nA to 1 mA;
2. a test current is configured into the reference source;
3. a large number of contiguous samples (10^4) are acquired from the digital stream;
4. the Discrete Time Fourier Transform (DTFT) is computed over the samples;
5. the procedure is iterated for all the test currents.

Here, the DTFT is defined by

$$M_k = \frac{1}{N} \sum_{j=0}^{N-1} m_j \; e^{-i2\pi \frac{j \times k}{N}} \; , \tag{5.8}$$

being N the number of samples (10^4), m_j the input data and M_k the transformed complex variable. In particular, M_0 represents the DC component of the measurement (equal to m^{dc}) while, for $k > 0$, the values M_k represent all the remaining spectral footprint. In general, these latter will include both the total contribution of noise and the variable input components. Since DC currents are injected during the test, the main content is provided by M_0. The other contributions are due to m_k^{noise} and their magnitudes $|M_k|$ have been employed for the evaluation of σ_m, while the phase of M_k is not relevant in this analysis.

The measurement spectrum for a single test current of 100 nA is shown in Fig. 5.3. The yellow plot is obtained by exploiting s_k^{CFC} data only, while the blue one also includes the auxiliary ADC data s_k^v, which were combined by the method described in Par. 5.2. It is possible to notice that the DC component equals the input current, while the large repertoire of disturbances presents considerable intensities. Furthermore, the benefit provided by the data combination reducing the intensities of $|M_k|$ by approximately a factor of 10 is manifest.

Collecting the same kind of results for the entire sweep makes it possible to plot a spectrogram, which summarises the correspondence between the input currents and the measurement signal footprints. In this case as well, the difference between s_k^{CFC} samples (in Fig. 5.4) and combined data (in Fig. 5.5) were evaluated. The DC component has been removed from the plots to highlight the variable contributions. It is possible to observe the frequency modulation behaviour of the CFC architecture, characterised by dominant components easily distinguishable up to a few μA. They reach the maximum

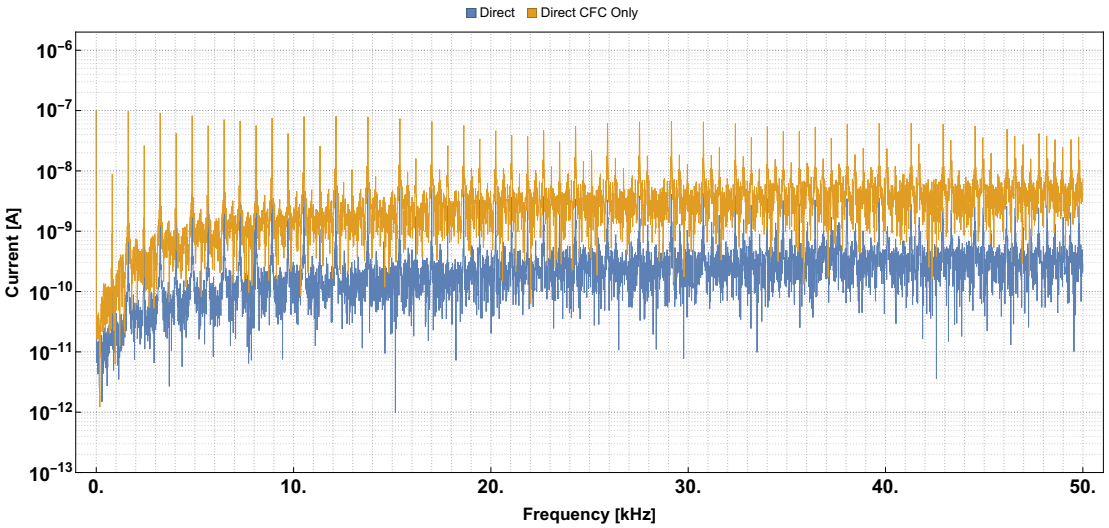

Figure 5.3: Compared measurement signal spectrums with an input current of 100 nA.

detectable frequency of 50 kHz, before symmetrically reverting due to the aliasing. Studying the CFC waveform frequency respect to the current (cf. Eq. (3.4)), it can be envisaged that it is lower than 50 kHz for currents below ≈ 5.6 μA. This means that a single threshold crossing occurred at most between two transmitted samples s_k and the s_k^{CFC} data will be composed by 1 or 0 only. Thus, below 5.6 μA, the spectrums will match those of periodic pulse signals, until a constant s_k^{CFC} sequence of 1 is transmitted at that limit. For higher currents, the process is repeated with the addition of a bias, which contributes to the DC. Moving towards the right of the plots, the dominants become not distinguishable due to the fast change of the exponential sweep and the introduction of secondary noise effects.

The combined data spectrogram (Fig. 5.5) is more difficult to describe, since it blends characteristics of the CFC data, the auxiliary ADC and the computation described in Par. 5.2. However, similar features can still be identified, together with an overall reduction of the background noise. Furthermore, the presence of quiet bands above 100 μA is more evident here, which are likewise related to the combined effects of the converter waveform frequency and the aliasing. In particular, they appear close to the currents making the frequency of waveform peaks a multiple of the data transmission rate. In fact, according to the internal analogue readout circuitry (cf. Par. 3.3.8), the sequence s_k will be constant, contributing exclusively to the DC component and cutting down the remaining noise. This happens because the triangular waveform is synchronous to the sample transmission, hence the quantities s_k^{CFC}, s_k^v and s_k^{slope} will remain constant in k, as well as the resulting measurements m_k. Such a behaviour explains the presence of valleys both in the noise characterisation results (cf. Par. 4.8) and the post X-Ray validation plots (cf. Par. 4.11).

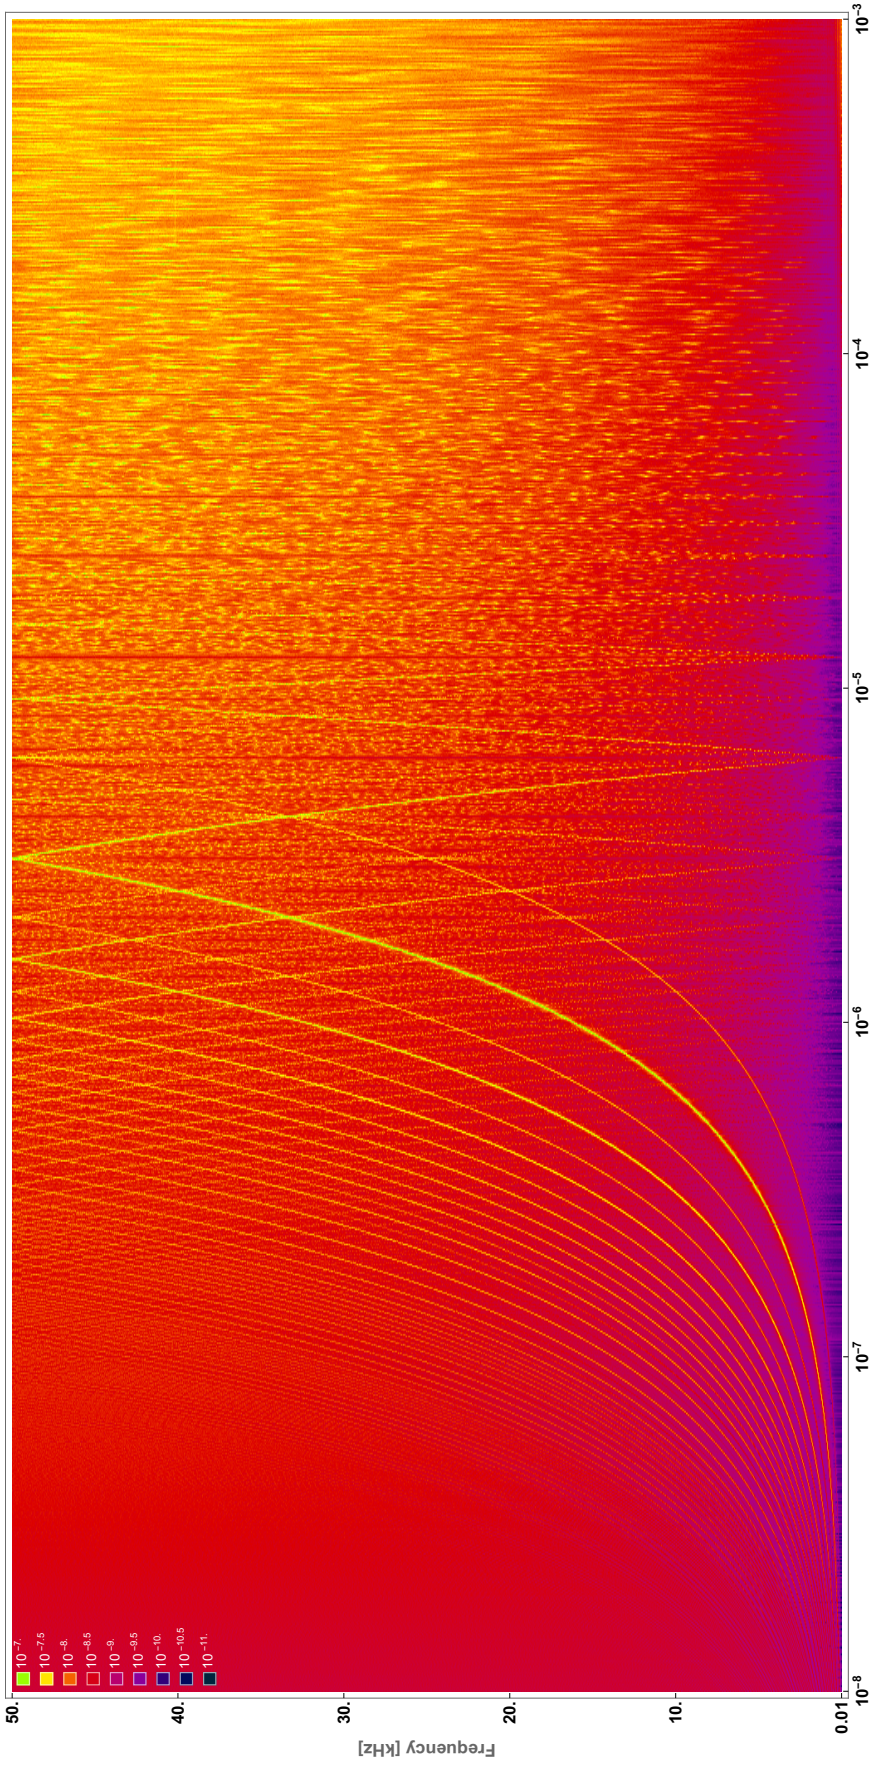

Figure 5.4: Spectrogram of converter measurements obtained by s_k^{CFC} data only versus the injected currents. The colour scale represents the intensity of current frequency components in $[A]$.

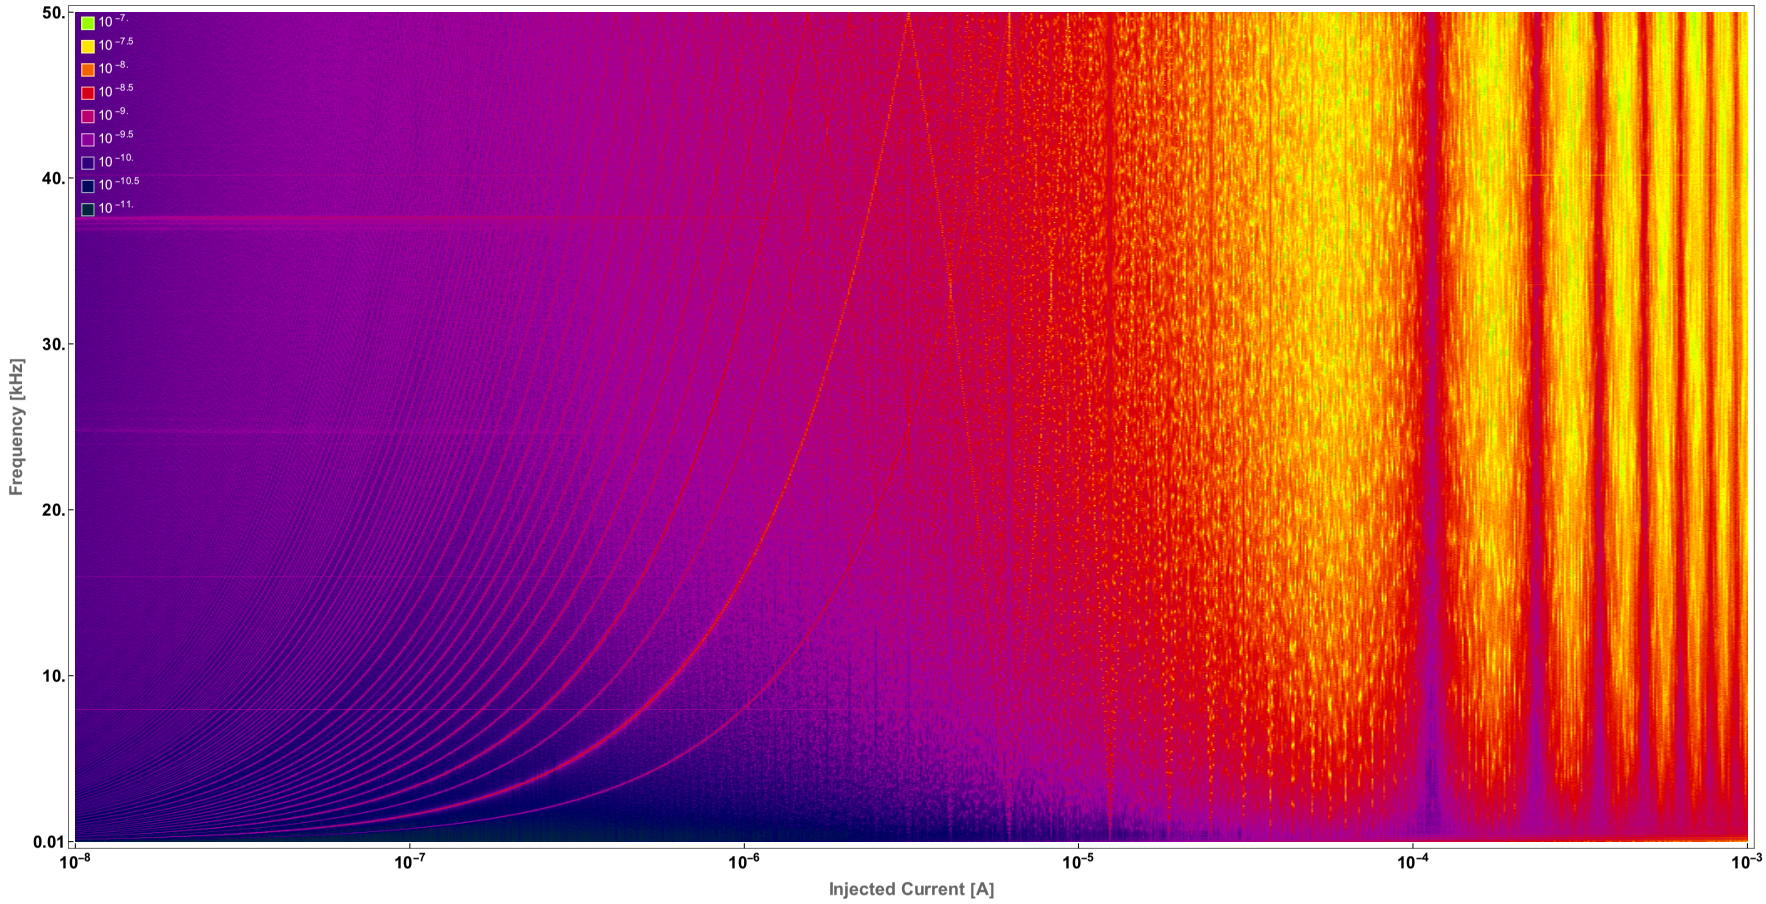

Figure 5.5: Spectrogram of converter measurements obtained by combined s_k data versus the injected currents. The colour scale represents the intensity of current frequency components in $[A]$.

5.4 Noise Considerations

The total noise contribution can be evaluated on the basis of the above spectrograms or, alternatively, by performing the standard deviation over the time domain samples. In fact according to DTFT defined in (5.8), the Plancherel-Parseval theorem states

$$\sum_{k=0}^{N-1} |M_k|^2 = \frac{1}{N} \sum_{j=0}^{N-1} |m_j|^2 \ . \tag{5.9}$$

Furthermore, if the DC component M_0 is deliberately removed, not contributing to the noise, it is possible to obtain

$$\sum_{k=1}^{N-1} |M_k|^2 = \frac{1}{N} \sum_{j=0}^{N-1} \left| m_j - m^{dc} \right|^2 = \sigma_m^2 \ . \tag{5.10}$$

This confirms that the study of the measurement error can be carried out equivalently in both domains. Nevertheless, the frequency one provides more effective information to understand where to intervene for its reduction. In fact, the analysis of the spectrogram allows for the identification of the frequency bands which contain the data of interest and the others that can be attenuated by suitable filtering. For example, in the case studied in Par. 5.3, a DC current was injected at the input, so an ideal constant measurement should also be expected in the post-processed r_k sequence. Indeed, all the other frequency components are disturbances, which increase the value of σ_m. It is therefore evident that a low-pass filter must be adopted and its cut-off frequency should be set as low as possible, to maximise the measurement precision. Ideally, if its frequency response $H(f)$ can be pushed to the limit $H(f) = \delta(0)$, the noise power σ_m^2 would be completely nullified.

The considerations made describe how the precision of a system can be ideally taken to infinity. However, such a target is neither feasible nor useful in reality, as such a filter would produce a constant signal for an infinite time, which is nonsense. In fact, DC signals are only abstract representations, which do not actually exist in nature, nor can they represent any scientific interest. Therefore, a different approach must be adopted to select the suitable $H(f)$. Primarily, weakening the assumption that the meaningful signals belong to the DC component is necessary. Indeed, it is clear that the beam loss can even vary quickly (cf. Par. 2.1) and the response speed of the BLM system should be kept as high as possible, to protect the accelerator effectively. However, considering the baseband nature of the signal and the low-pass effect introduced by the cabling, the spectrum of interest should be mainly contained below 25 kHz. On the contrary, one can observe in the plots that higher noise components affect the upper half of the frequency range. Thus, at least qualitatively, a separation between the bands can be identified.

Therefore, the research of a suitable processing method is reduced to the choice of a low-pass filter model $H(f)$ and the tuning of its parameters. In particular, it will be necessary to identify the suitable trade-off between the response speed and the target measurement precision. In fact, a filter with a lower cutting frequency can attenuate the noise of measurements over a wider band, thus reducing σ_m, but it will respond less rapidly to an abrupt current change.

Finally, to rate the processing performance or compare different techniques over the whole input range, the trend of σ_m has to be computed as a function of the measured current. In Fig. 5.6, an example obtained by raw converter data is shown without applying any filtering. The s_k^{CFC} measurements (*Direct CFC Only*) have been compared with those obtained using the combination method described in Par. 5.2 (*Direct*). The results have been obtained by integrating the noise components over the entire spectrum of Fig.s 5.4 and 5.5, in agreement with the equation (5.10). It can be noticed that the combined data are affected by less noise for low to medium currents (up to $\approx 10\ \mu A$), so the measurements will be more precise. However, significant noise peaks are reported above, indicating that the method may be improved for operating with high currents. This is due to ineffective sampling of the auxiliary ADC in this range, which had been envisaged by design. In fact, an analysis of its raw data showed that almost constant output values are produced at the middle of the scale. Presumably, this happens because the triangular waveform is inherently averaged, having an oscillation frequency much higher than the ADC bandwidth. A possible solution to mitigate the peaks may involve weighing the two data sources according to the input trend, or more advanced estimation techniques [78], [79]. The maximum reported noise is 3.08 μA for pure CFC data and 6.02 μA for combined one. The first result can be considered the threshold of maximum precision achievable with raw BLMASIC data, without employing further filtering processing.

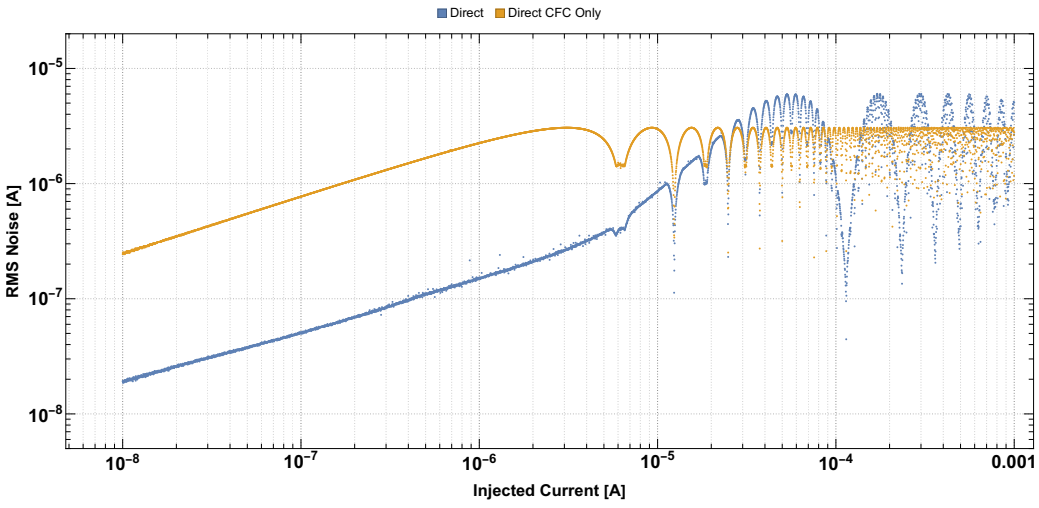

Figure 5.6: RMS noise of the converter measurements. Raw CFC counter data are compared with CFC / ADC combined ones.

5.5 Processing Techniques

The activity involved the mathematical modelling of the currently adopted processing and the development of alternative solutions capable of improving the system response speed or, alternatively, the precision of measurements . Equivalent time-continuous models are initially presented for an easier understanding. In particular, since linear systems are investigated, the processed measurements $r(t)$ will be the result of the convolution between the raw data signal $m(t)$ and the filter impulse response $p(t)$ [56]. After, the investigation will be extended to the discrete-time form, assuming the sampling frequency f_s is equal to the BLMASIC data transmission one (100 kHz). Both the time and the frequency behaviours will be analysed, to describe in detail how the processing techniques will affect the measurements. Finally, the effects of data quantisation will be neglected, according to the considerations made at the end of Par. 5.1.

5.5.1 Moving Average Filters

The computation employed in the operational system was described in Par. 2.3.2 and involves a set of running sums that accumulate raw samples over different windows, producing corresponding integrations. From a signal theory point of view, those can be studied as a family of moving average filters, barring multiplicative constants equal to the total of the added samples. In particular, 11 different integration windows have been configured over widths from 2 to about 2×10^6 samples (cf. Tab. 2.3), actually generating a bank of low-pass filters, independently operating on different bandwidths. The shortest running sum **RS1** (time window of 80 μs) can track fast losses at the expense of higher errors, while the longest **RS11** has a very slow ramp of ≈ 84 s, but providing high precision results. Due to the limited resources, the decimation of the outputs was widely employed in the old system. However, this is not planned to be required for the renewed electronics, which will adopt a higher-performance FPGA device (cf. Chapter 7). Hence, the input and output data rates will considered equal.

The model of a moving average filter, hereafter referred to as AVG, can be introduced with its equivalent time continuous impulse response $avg(t)$, defined as

$$avg(t) = \begin{cases} \frac{1}{W} & 0 \leq t \leq W \\ 0 & otherwise \end{cases}, \tag{5.11}$$

with W representing the integration time window. Furthermore, it is useful to express its response $avg_{step}(t)$ to a step, an abrupt change in the measuring currents, by the convolution $avg(t) \otimes step(t)$, where

$$step(t) = \begin{cases} 1 & t \geq 0 \\ 0 & t < 0 \end{cases}, \tag{5.12}$$

namely

$$avg_{step}(t) = avg(t) \otimes step(t) = \begin{cases} 0 & t < 0 \\ \frac{t}{W} & 0 \leq t \leq W \\ 1 & t > W \end{cases} . \tag{5.13}$$

In Fig. 5.7 is reported an example of the presented expressions, which were carried out by an AVG filter of window $W = 1\ s$.

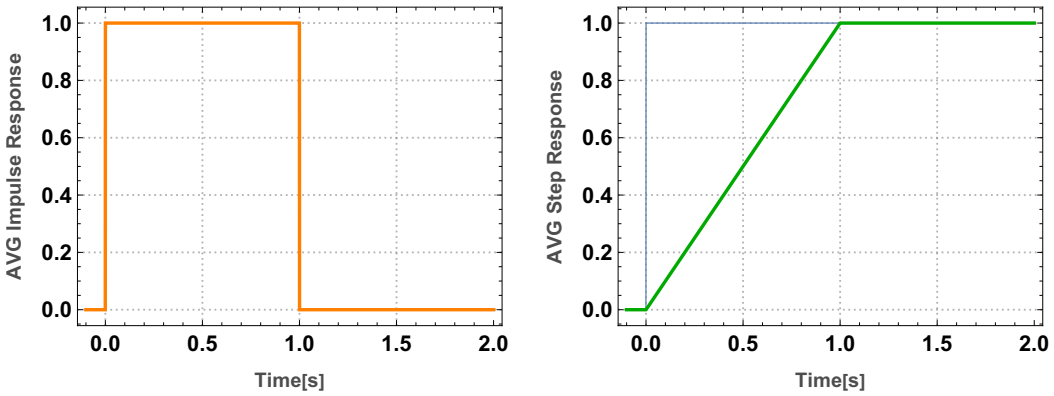

Figure 5.7: Impulse (left) and step (right) responses of a moving average filter with $W = 1\ s$.

From the equation (5.11) it can also be derived its frequency response $AVG(f)$, applying the continuous Fourier Transform. It results

$$AVG(f) = W\ Sinc\ (Wf)\ , \tag{5.14}$$

which vanishes for $f = k\ W^{-1}$ and decreases as f^{-1}, corresponding to $-20\ dB/decade$. The knowledge of this expression provides valuable information about measurement processing. Indeed, looking at the example plot of Fig. 5.8, it is easy to distinguish which frequency bands will be attenuated and which will remain untouched. The presence of notches is evident above 1 Hz and their distribution depends only on W, as stated before. It was probably noticed that such a pattern is recurrent even in raw converter data. In fact, the accumulation performed by the CFC counters during the period of 10 μs has the same effect on the signal spectrum.

Finally, the AVG filter can also be studied in its discrete-time counterpart, which allows modelling more accurately the real implementation. It is actually an FIR [55] filter, the coefficients of which are all equal to the inverse of the number of averaged samples. In fact, the discrete impulse response avg_k will be very similar to its continuous version (cf. equation (5.11)), but the time is now substituted by the index k, obtaining

$$avg_k = \begin{cases} \frac{1}{W} & 0 \leq k \leq W \\ 0 & otherwise \end{cases} , \tag{5.15}$$

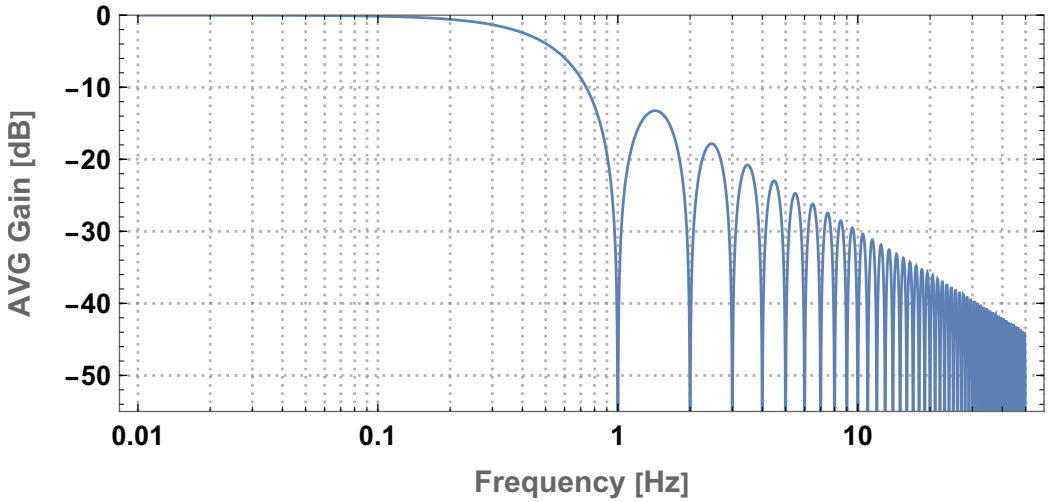

Figure 5.8: Frequency response $AVG(f)$ of a moving average filter with a time window of 1 s.

with W being here a natural number. The corresponding frequency response AVG_j can be computed by using the already described DTFT definition in equation (5.8). However, the coefficient in front of the sum can now be set to one since avg_k is already normalised. The resulting AVG_j Transform can represent a useful tool to estimate the post-processing data spectrum, which is

$$R_j = AVG_j \times M_j \ , \tag{5.16}$$

where M_j is the direct measurement spectrum acquired in Par. 5.3. Once R_j is obtained, the resulting error σ_r can be computed by the equation (5.10). This makes it possible to estimate the measurement precision of the real system, which would be difficult to achieve otherwise. In fact, it must be pointed out that the typical decreasing of σ_r as $1/\sqrt{W}$ is only valid for white spectrum noises (cf. Par. 4.5), which may be far from the real case. Furthermore, due to the complexity of the conversion architecture, it is not easy to formulate a complete theoretical model estimating the system's precision.

5.5.2 Infinite Impulse Response Filters

After carrying out an exhaustive study of the currently operational processing, novel methods have been investigated, to improve the characteristics of the upgraded BLM system. The analysis has focused on Infinite Impulse Response (IIR) [55] filters, because they provide suitable spectrum features while requiring limited resources. In fact, although the computational logics needed for a single channel are not overwhelming, the back-end must handle the data coming from 2 different front-end cards, for a total of 16 channels. Furthermore, each of them will feed the filter bank to track the various loss events, thus the resources economy must be considered.

Initially, first-order filters have been compared with the AVG, reporting a slightly inferior ability to cut the measurement noise. In fact, they are characterised by

an off-band decay of $-20\ dB/decade$, which is the same as the AVG filter, but the lack of notches allows more disturbances to reach the results (cf. Par. 5.5.1). Nevertheless, they remain a valid alternative to the previous method regarding resources and flexibility. Indeed, unlike the AVG, their computational requirements (number of stored samples and basic operations) do not change with the cutting frequency, instead, it can be easily tuned by very few parameters, even during the operations. Furthermore, the filter bank implementation is dramatically simplified since the same logic can be replicated for all the required bands.

To merge the benefits of the IIR implementation with an effective noise suppression, second order architectures have been analysed. These are described by the difference equation [55]

$$a_0\ y_k + a_1\ y_{k-1} + a_2\ y_{k-2} = b_0\ x_k + b_1\ x_{k-1} + b_2\ x_{k-2}\ , \tag{5.17}$$

with y_k representing the filter output and x_k its input. In the context of the processing system, $x_k = m_k$, while $r_k = y_k$. Furthermore, it can be implemented by using the scheme of Fig. 5.9, which is commonly known as *Direct Form II* [55].

The equation (5.17) can also be expressed in z Transform, to study the stability and the frequency behaviour. The transfer function is

$$G_d(z) = \frac{Y_d(z)}{X_d(z)} = \frac{b_0 + b_1\ z^{-1} + b_2\ z^{-2}}{a_0 + a_1\ z^{-1} + a_2\ z^{-2}}\ , \tag{5.18}$$

which can be alternatively written as

$$G_d(z) = \frac{b_0\ z^2 + b_1\ z + b_2}{a_0\ z^2 + a_1\ z + a_2}\ , \tag{5.19}$$

where X_d and Y_d are the input and output signal Transforms, respectively. In order to achieve the desired frequency response, the equation coefficients must be properly configured. Literature includes a large variety of employable methods, optimising different target characteristics [55], [80]. On the other hand, it has been preferred to

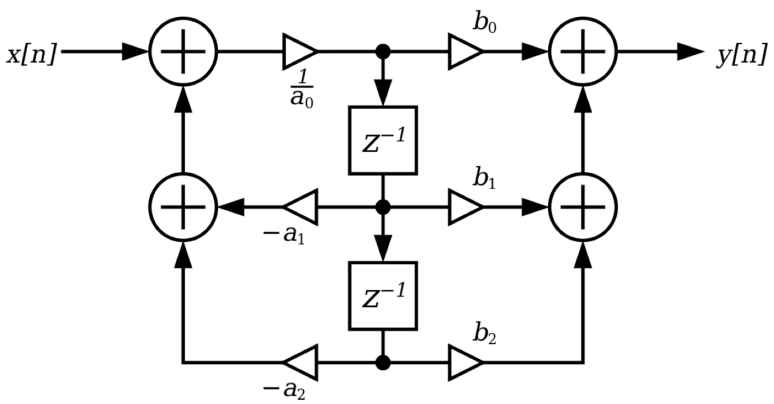

Figure 5.9: Implementation structure of the second order IIR filter (*Direct Form II*).

start with the time continuous model, providing a more effective overview of the features involved. Initially, it has been considered the second order low-pass filter [56], the response of which is described by the expression

$$G_c(s) = \frac{Y_c(s)}{X_c(s)} = \frac{w_0^2}{s^2 + 2\zeta w_0\ s + w_0^2}\ , \tag{5.20}$$

being G_c the system transfer function, s the Laplace variable, w_0 the natural resonance pulsation and ζ the damping factor. For convenience, w_0 can be substituted by the corresponding frequency $f_0 = w_0/(2\pi)$, obtaining

$$G_c(s) = \frac{4\pi^2 f_0^2}{s^2 + 4\pi\zeta f_0\ s + 4\pi^2 f_0^2}\ . \tag{5.21}$$

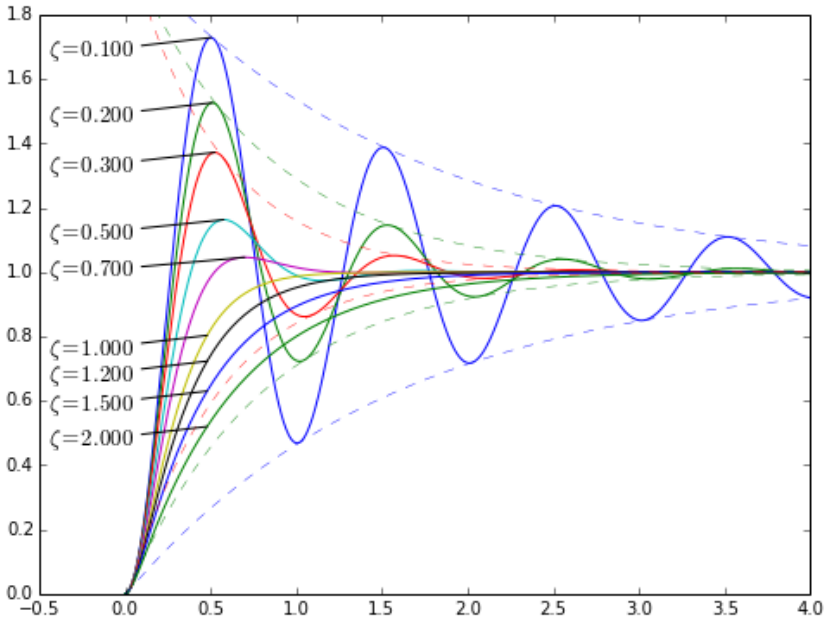

Figure 5.10: Step response of a second order system with $f_0 = 1\ Hz$ at various ζ dampings.

The frequency f_0 defines the boundary between the pass and stop bands. The information signal desired to propagate at the output should only occupy the first one. Instead, the latter should only include the noise to be suppressed. According to the consideration made in Par.s 5.4 and 5.5.1, tuning this parameter results from a compromise between an appropriate response speed and the target precision. The coefficient ζ represents a further degree of freedom in the filter design. In fact, it directly influences the frequency behaviour near the boundary and, consequently, the system step response. It can be observed that if $\zeta = 0$, $G_c(s)$ has pure imaginary poles, thus it behaves like an ideal oscillator at frequency f_0, which is useless for the goals. If $0 < \zeta < 1$, the poles are complex with a strictly negative real component, which translates into a damped resonance. Instead, the poles are pure reals for $\zeta > 1$, meaning that the system is not resonant. The output behaviour for different ζ under

a unit step solicitation is reported in Fig. 5.10. By using the inverse Laplace Transform, it is also possible to relate ζ to the percentage overshoot $OS_{\%}$, resulting

$$OS_{\%} = 100 \times e^{-\frac{\pi\zeta}{\sqrt{1-\zeta^2}}} \ . \tag{5.22}$$

It has been observed that the case $\zeta = 1$ represents the best compromise for the targeted filter, which implies a double real pole. In fact, this solution removes the overshoot experienced in the resonant case, which may lead to overestimated (or underestimated) measurements, but it still maximises the ramp slope. Under such a condition, the time step response of the system is expressed for $t > 0$ as

$$g_{step}(t) = 4\ \pi^2 f_0^2\ t\ e^{-2\pi f_0 t}\ , \tag{5.23}$$

an example of which is shown in Fig. 5.11 together with the impulse response. Furthermore, the corresponding magnitude frequency response can be found in Fig. 5.12.

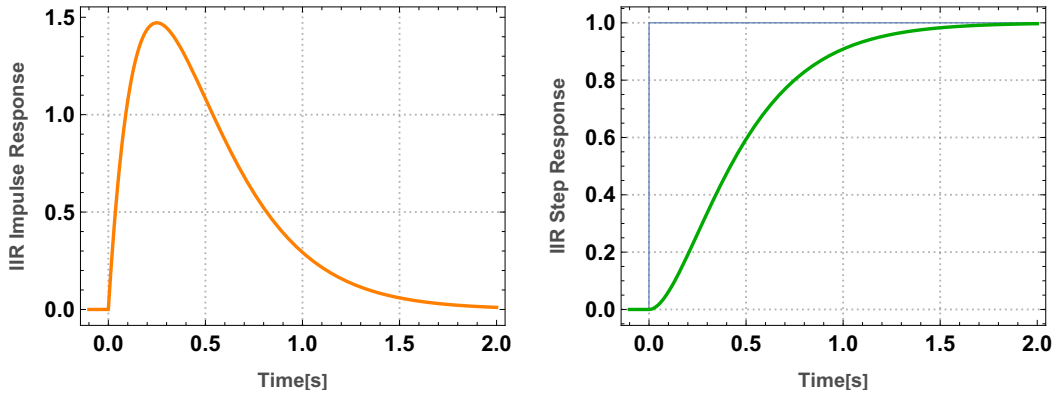

Figure 5.11: Impulse (left) and step (right) responses of second order low-pass filter, with $f_0 = 1\ Hz$ and $\zeta = 1$.

It is necessary to assess the relation between the parameters of such a continuous model and the discrete-time implementation coefficients. For this goal, the bilinear transform (called the Tustin's method [55]) can be employed, which remaps the Laplace s plane into the z one, by

$$z = e^{sT} \approx \frac{1 + \frac{T}{2}s}{1 - \frac{T}{2}s}\ , \tag{5.24}$$

or

$$s \approx \frac{2}{T}\frac{z-1}{z+1}\ , \tag{5.25}$$

where $T = f_s^{-1}$ and the approximation in the right hand sides hold for $|s| << 2\,f_s$, or $z = e^{i\theta}$ and $|\theta| << \pi$. For $\Re(s) > 0$ one obtains $|z| > 1$, while the opposite holds for $\Re(s) < 0$. This guarantees that the discrete-time counterpart of a stable system is also stable. Finally, the imaginary axis of s is transformed into the unit circumference.

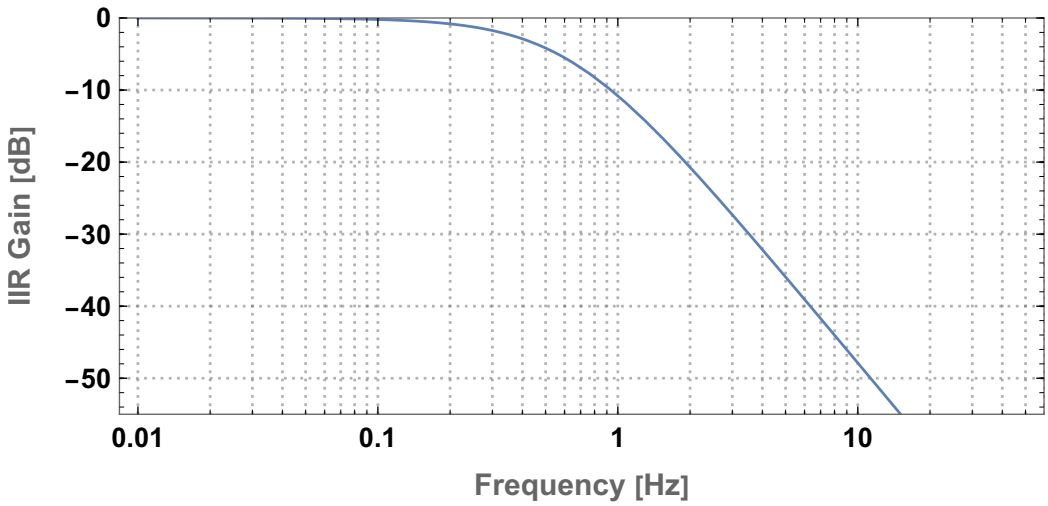

Figure 5.12: Frequency response of a second order low-pass filter, with $f_0 = 1\ Hz$ and $\zeta = 1$.

The presented method allows a fast transposition of the continuous time transfer function $G_c(s)$ into the discrete one $G_d(z)$, providing a useful tool for developing IIR filters. In fact, rational expressions in the s variable belong to the same class in the z domain. In turn, they can be implemented by exclusively employing arithmetic operations and delays. However, the effects introduced by the approximation must be taken into account during the design. In particular, if $G_c(s)$ includes a certain number of zeros and poles, they will also characterise $G_d(z)$, but the corresponding frequencies in z may, in general, be slightly different.

Such an issue can be solved by applying a warping function at significant frequencies for the filter characteristics, such as f_0, before proceeding with the domain transformation. In fact, by using the Eq.s (5.24) and (5.25) and imposing the local equality of the two transfer functions in f_0, it holds

$$\begin{aligned} G_d\left(e^{i2\pi f_0 T}\right) = G_c\left(\frac{2}{T}\frac{z-1}{z+1}\right)_{z=e^{i2\pi f_0 T}} = \\ G_c\left(\frac{2}{T}\frac{e^{i2\pi f_0 T}-1}{e^{i2\pi f_0 T}+1}\right) = G_c\left(i\ \frac{2}{T}\ Tan\left(\pi f_0 T\right)\right)\ . \end{aligned} \tag{5.26}$$

This relation states that the behaviour of the discrete system G_d at frequency f_0 is equal to the time-continuous one G_c at

$$f_0^* = \frac{1}{\pi T}\ Tan\left(\pi f_0 T\right) = \frac{f_s}{\pi}\ Tan\left(\pi\ \frac{f_0}{f_s}\right)\ , \tag{5.27}$$

meaning that the desired frequency f_0 must be warped to f_0^* before substituting it in the equation (5.21). Only after such a correction, the Eq. 5.25 can be applied to obtain

the discrete system transfer function

$$G_d(z) = \frac{(1+z)^2 \; Sin\left(\pi \; \frac{f_0}{f_s}\right)^2}{\left[1+\zeta \; Sin\left(2\pi \frac{f_0}{f_s}\right)\right] \; z^2 - 2 \; Cos\left(2\pi \frac{f_0}{f_s}\right) \; z - \zeta \; Sin\left(2\pi \frac{f_0}{f_s}\right) + 1} \; ,$$

the coefficient of which are reported in Tab. 5.1.

Coefficient	**Value**
a_0	$1+\zeta \; Sin\left(2\pi \; \frac{f_0}{f_s}\right)$
a_1	$-\; 2 \; Cos\left(2\pi \; \frac{f_0}{f_s}\right)$
a_2	$1-\zeta \; Sin\left(2\pi \; \frac{f_0}{f_s}\right)$
b_0	$Sin\left(\pi \; \frac{f_0}{f_s}\right)^2$
b_1	$2 \; Sin\left(\pi \; \frac{f_0}{f_s}\right)^2$
b_2	$Sin\left(\pi \; \frac{f_0}{f_s}\right)^2$

Table 5.1: Second order IIR filter coefficients.

If the damping factor ζ is set to 1, the system has a double zero for $z = -1$ and a double pole depending on the chosen cut frequency. These features can be identified in the plot of Fig. 5.13, showing an example of the function $G_d(z)$ computed within the z circle.

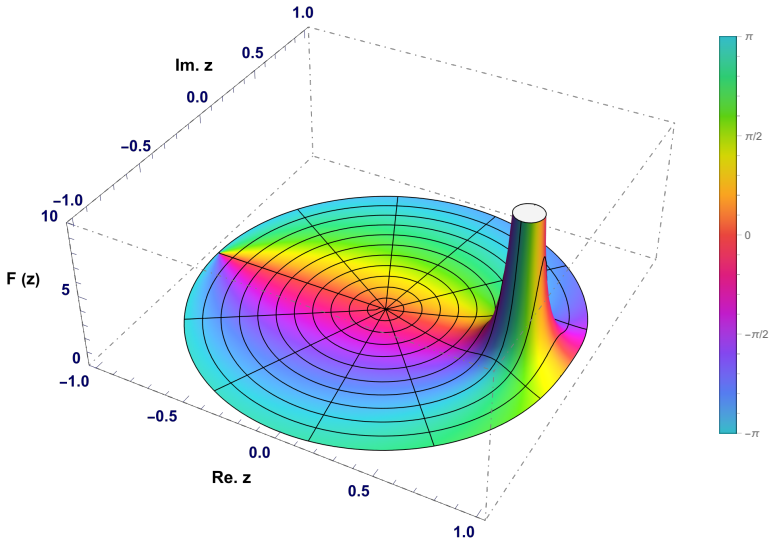

Figure 5.13: Example of the filter transfer function $G_d(z)$ within the z unit circle. The parameters are set to: $f_s = 100 \; kHz$, $f_0 = 5 \; kHz$ and $\zeta = 1$.

5.6 Processing Metrics

The described low-pass filter solutions may constitute valid candidates for future employment in the back-end firmware. However, appropriate metrics must be introduced at this stage, to allow a fair comparison of their performance. Indeed, both the precision of results and the step characteristics of the instrument are equally important when evaluating the processing effectiveness (cf. introduction of Par. 5.5.1). The precision can be estimated by filtering the real converter data e computing σ_r or, alternatively, by exploiting the measurement spectrogram and the filter frequency response (cf. Par. 5.3). On the other hand, the step characteristics are settled at the design stage. Therefore, equivalence metrics have been introduced to assess the suitable correspondences between the parameters, such as the AVG window time width W and the IIR cutting frequency f_0.

5.6.1 Equal Loss Charge

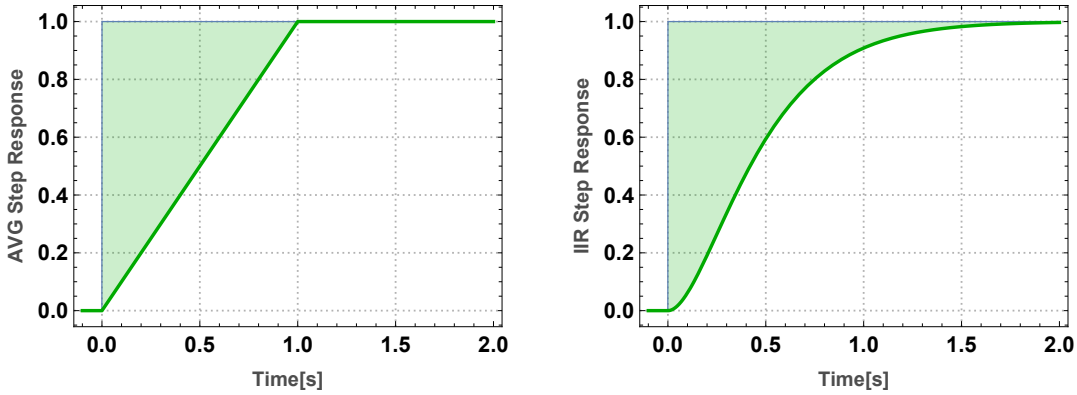

Figure 5.14: Areas corresponding to the lost charges for AVG ($W = 1\ s$) and the second order IIR ($\zeta = 1$ and $f_0 = 1\ Hz$).

Observing the step response of AVG and IIR filters (see Fig.s 5.7 and 5.11), it is possible to notice a delay introduced by the system subsequently to the ideal step. This is due to the filter causality and it can be adjusted by modifying its parameters, although never completely cancelled. A region can be identified between the ideal and the real step response curves (see Fig. 5.14), the area of which actually corresponds to a certain input charge lost during the measurement process. Such a quantity represents a valid comparison factor between the different solutions. In fact, it would be desirable to evaluate the noise suppression performance of filters, knowing that they will cause the same loss of charge for a given input signal. By using the AVG step response in Eq. (5.13) and the inverse Laplace Transform of $G(s)/s$, by the Eq. (5.21)), it is possible to estimate the charge values, which are reported in Tab. 5.2.

Filter	Lost Charge
AVG	$\frac{W}{2}$
IIR	$\frac{\zeta}{\pi\ f_0}$

Table 5.2: Charge lost due to an ideal unit current step filtering.

Finally, such a feature can be employed to relate the parameters according to the expression

$$f_0 = \frac{2\ \zeta}{\pi\ W}\ , \tag{5.28}$$

which allows to tune the step response of filters like in Fig. 5.15. In fact, the area highlighted in green is equal to the one in red.

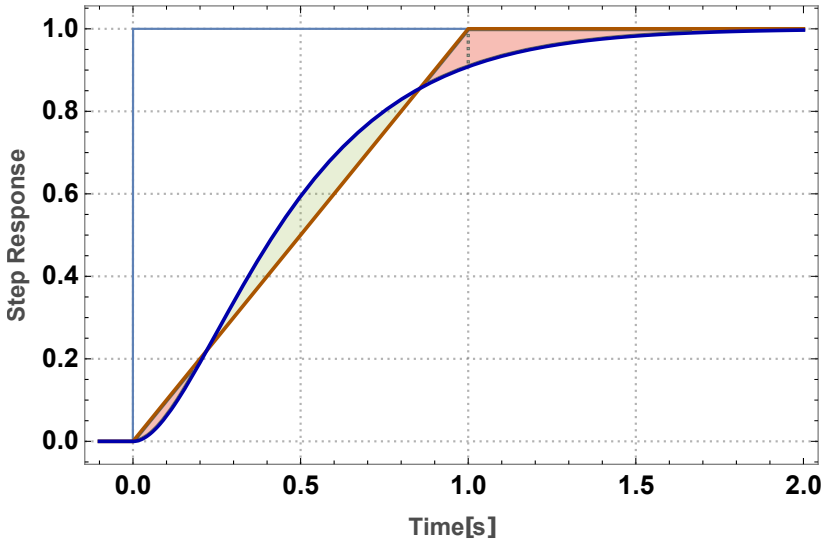

Figure 5.15: Responses of filters tuned with the loss charge equalisation criterion. The parameters are set to: $W = 1\ s$, $\zeta = 1$ and $f_0 = 637\ mHz$.

5.6.2 Equal Rise Time

This other approach focuses on equalising the rise time of the step responses. Usually, such a factor is defined as the time required by the system output to swing from 20 % to 80 % of the whole step [47] (other definitions use the thresholds 10 % and 90 %), providing an average estimation of how fast is the transient. However, for the BLM measurement goals, it would be more appropriate to identify a latency required by the data to reach the steady state condition since the loss step occurred. Thus, the first threshold has been theoretically decreased to 0 %. Furthermore, to be sure that the IIR long tail would not overextend the transition up to the steady state, an upper threshold of 95 % has been set.

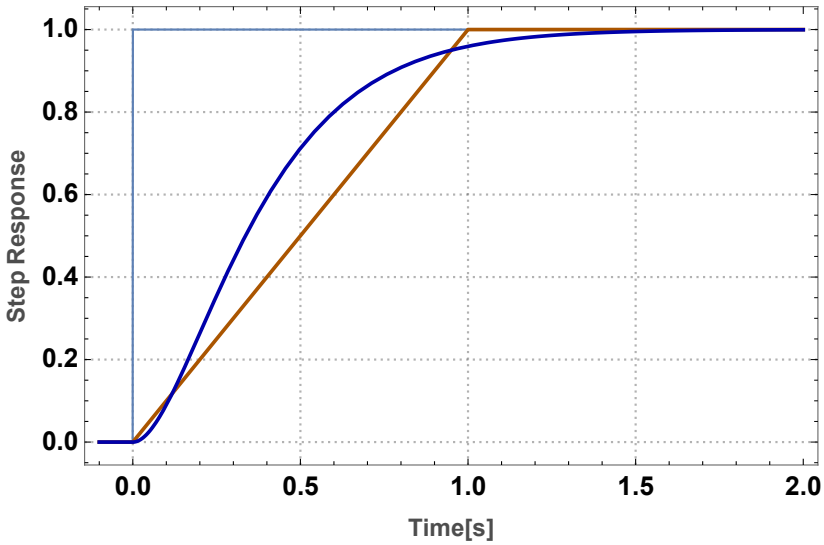

Figure 5.16: Responses of filters tuned with the rise time equalisation criterion. The parameters are set to: $W = 1\ s$, $\zeta = 1$ and $f_0 = 795\ mHz$.

The equalisation is performed by computing the rise times from the same equations employed in Par. 5.6.1. However, in this case, a numerical solver was exploited to handle the transcendent equation describing the IIR. Furthermore, the parameter ζ was fixed to 1, to reduce the computation complexity. It has been proved that the relation between the parameters is here defined by

$$f_0 \approx \frac{0.795}{W} \ . \tag{5.29}$$

In Fig. 5.16 is visible an example of the responses obtained after such a tuning. It is noticeable that the IIR curve is well above the AVG one during most of the transient, proving to be very effective in tracking fast-changing currents. In Par. 8, its noise rejection performance will be presented, which proved suitable for the goals.

5.7 Comparison of Methods

Once equivalence criteria amongst the various solutions have been defined, it is possible to estimate their output σ_r performance (cf. Par. 5.4). Initially, the frequency response of the discrete-time filters has been calculated and plotted in Fig. 5.17. This can be obtained by carrying out the DTFT over their impulse responses or by evaluating the module of the z domain transfer function all around the unit circumference, that is $z = e^{i2\pi\frac{f}{fs}}$. Despite the AVG attenuating more than the IIR close to the cutting frequency, thanks to its notches (see area highlighted in red), the steeper IIR stop-band slopes allow a better rejection of the high-frequency components. This region, highlighted in green and yellow, is wider than the red one, because a logarithmic scale is used.

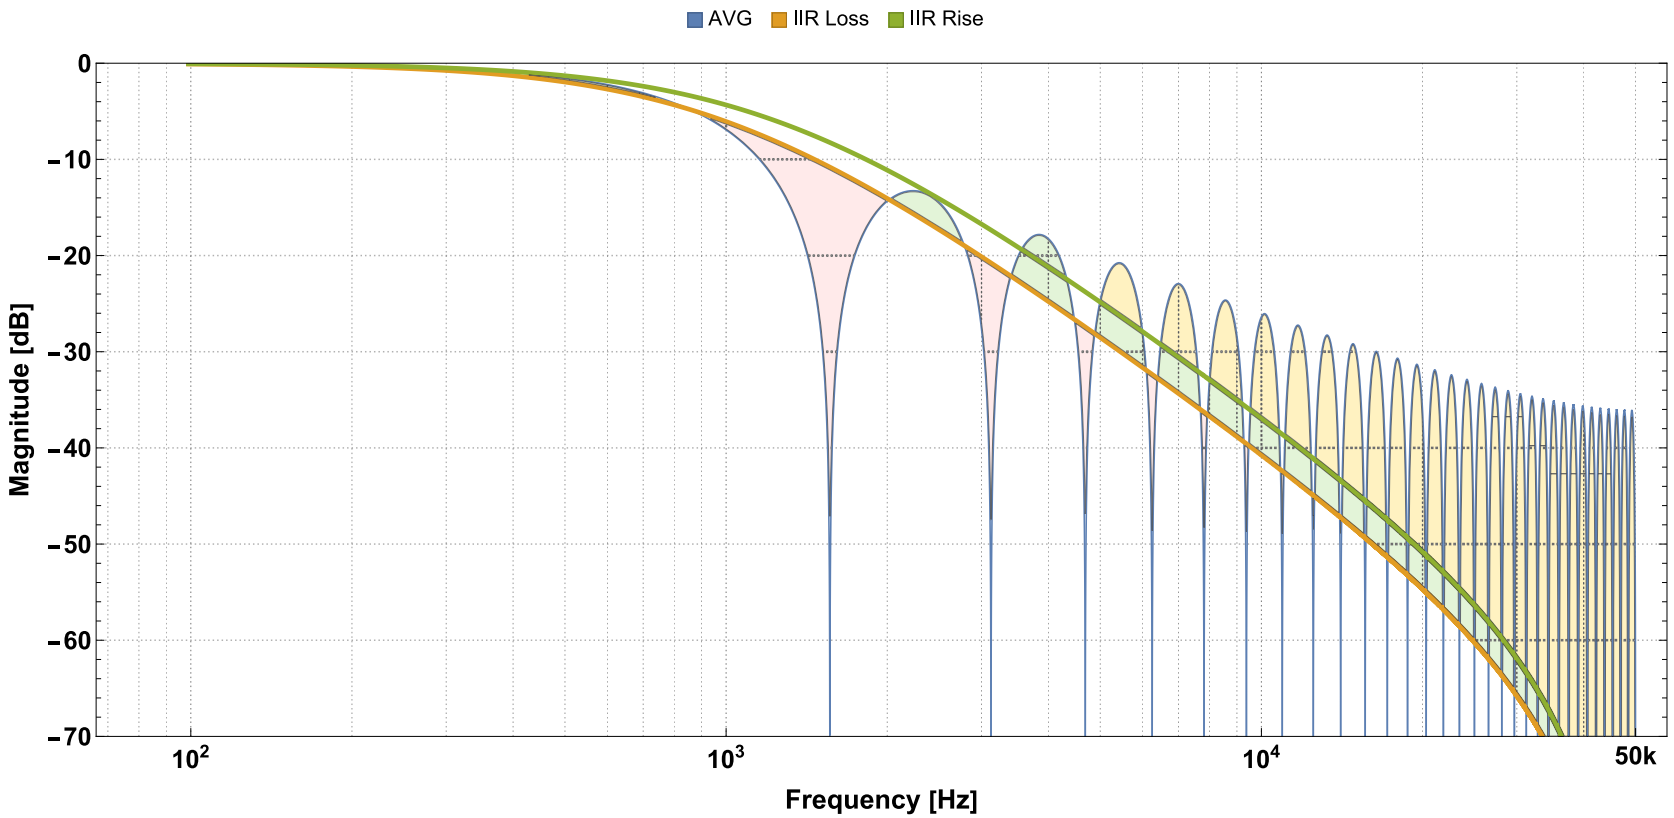

Figure 5.17: Comparison of the frequency responses of discrete-time filters, tuned applying the equalisation criteria. The parameters are set to: $W = 1\ s$, $f_0^{Loss} = 637\ mHz$, $f_0^{Rise} = 795\ mHz$.

Furthermore, as discussed in Par. 5.4, most of the measurement data noise is located at high frequency, hence better precision (lower σ_r) is expected employing the IIR. In any case, the test with real data is recommended to confirm such an assumption, due to the high variability of the noise spectrum as a function of the current (see. Fig. 5.5).

To facilitate the comparison of the upgraded instrument precision with the previous one, it has been decided to tune the filters following the original set of running sums (cf. Par. 2.3.2). Furthermore, the signal nomenclature RS_X is still used here for identifying the results having the same step response speed. For instance, the former **RS5** running sum produces a ramp of about 10 ms, which is exactly the same as the one generated by the AVG **RS5**. In such a way, the new result precision can be directly compared with the former, guaranteeing that the system will require about the same time to reach the steady state. However, it must be recalled that the sampling rate is now 4 times faster (100 kHz) than the old one (25 kHz), so the number of samples involved in the integrations is quadrupled. Hence, if the legacy thresholding module is desired to be migrated into the new system, the integration windows should be adapted accordingly. The parameters employed in the new bank of filters agree with the equalisation metrics, described in Par. 5.6 and they are reported in details in Tab. 5.3.

Signal	**Legacy Samples** ($40\mu s$)	**AVG Filter**		**IIR Filter (Loss)** f_0 [Hz]	**IIR Filter (Rise)** f_0 [Hz]
		Samples ($10\mu s$)	**Width [s]**		
Direct			10μ	Not Used	Not Used
RS0	1	4	40μ	15.92k	19.87k
RS1	2	8	80μ	7.958k	9.934k
RS2	8	32	320μ	1.989k	2.484k
RS3	16	64	640μ	994.7	1.242k
RS4	64	256	2.56m	248.7	310.4
RS5	256	1024	10.24m	62.17	77.61
RS6	2048	8192	81.92m	7.771	9.702
RS7	8192	32768	327.68m	1.943	2.425
RS8	32768	131072	1.31072	485.7m	606.3m
RS9	131072	524288	5.24288	121.4m	151.6m
RS10	524288	2097152	20.97152	30.36m	37.90m
RS11	2097152	8388608	83.88608	7.589m	9.474m

Table 5.3: Filter bank parameters.

A first analysis was carried out by comparing the spectrum of raw converter data M_j with the AVG and IIR processing result R_j. A DC current of 100 nA was injected at the input and the **RS3** signal has been employed for this test (see Fig. 5.18). It is possible to recognise the typical aisle pattern of the AVG filter, which reduces the noise baseline of a factor > 10 with respect to the direct data. A further improvement is achieved by the IIR filter, the spectrum of which lies below the others until a suppression of more than 3 orders of magnitude is attained.

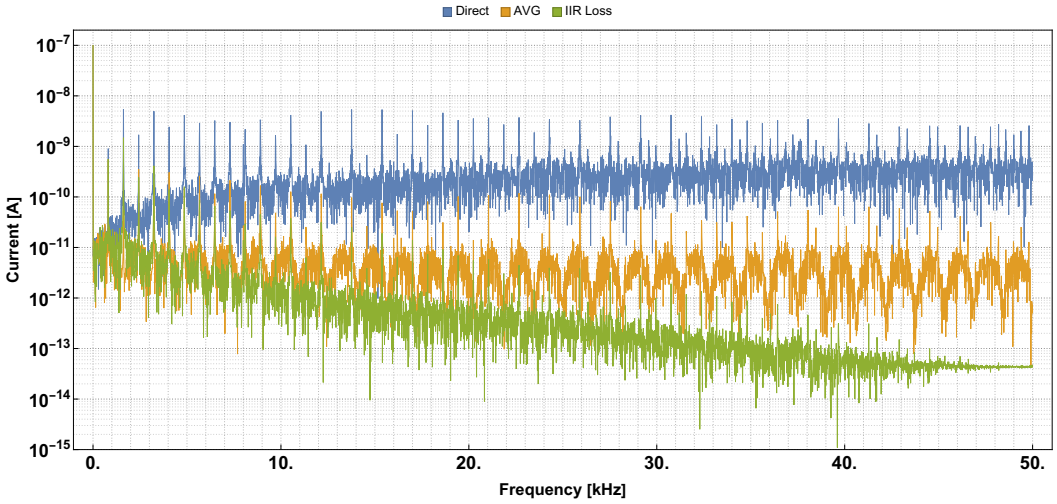

Figure 5.18: Comparison of raw data spectrum versus the one processed by the filters.

Following the same procedure employed for the converter noise analysis (cf. Par. 5.3), the spectrum has been calculated for a sweep of input currents, varying between 10 nA and 1 mA. The spectrograms shown in Fig. 5.19 refer to the different processing methods, which will be employed throughout the research (AVG, IIR equal *Loss* and equal *Rise Time*). Again, the noise suppression capabilities of the IIR filters are evident in comparison to the AVG implementation, especially for the frequency components close to the limit of 50 kHz. To provide a global characterisation of the precision, the RMS noise σ_r has been calculated for the 3 methods, following the procedures described in Par. 5.4. Its variation in function of the input currents is reported in Fig. 5.20, which manifests how the processing improves the measurement precision. Furthermore, the same plot displays the first result demonstrating the advantage of the IIR filtering over a classic AVG. In fact, the average noise level is significantly reduced, while the worst cases report equivalent results. Only with a current of 100 nA, the noise level from the IIR processing showed slightly higher. This probably occurs because a narrow frequency band solicitation was present in the raw data and tuned precisely to an AVG notch. However, this is an isolated occurrence and the IIR still reported a noise level only 1 nA higher. The performance of the rise time equalised filter has not been reported in the same picture to facilitate readability.

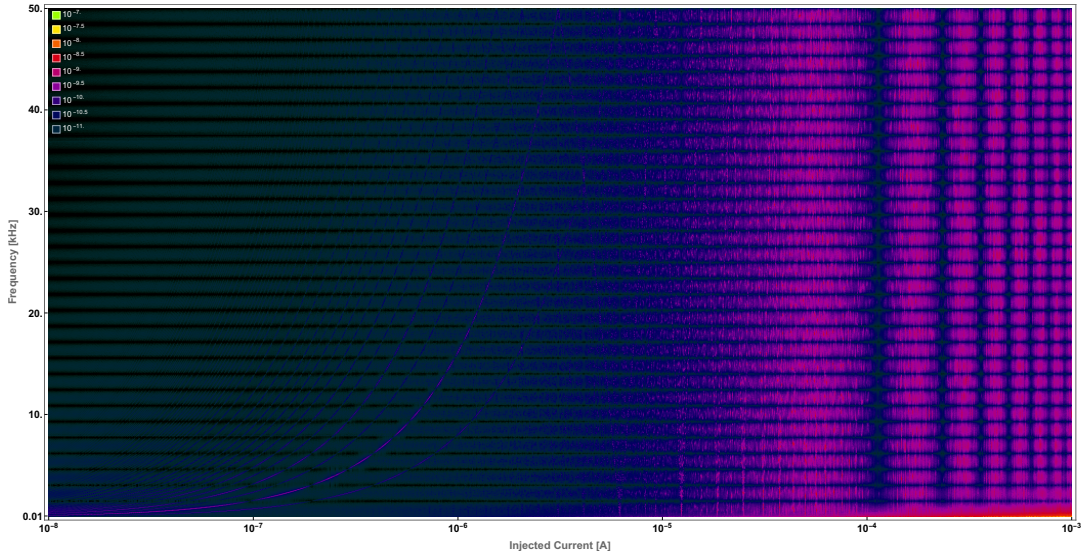

(a) AVG filter **RS3** ($W = 640\ \mu s$).

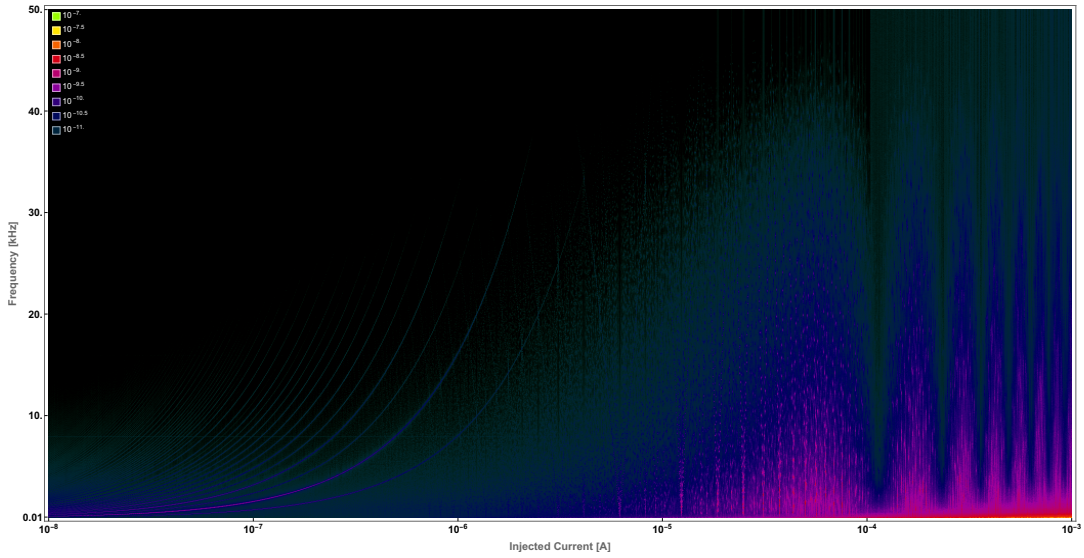

(b) Equalised loss charge IIR filter **RS3** ($f_0 \approx 1\ kHz$).

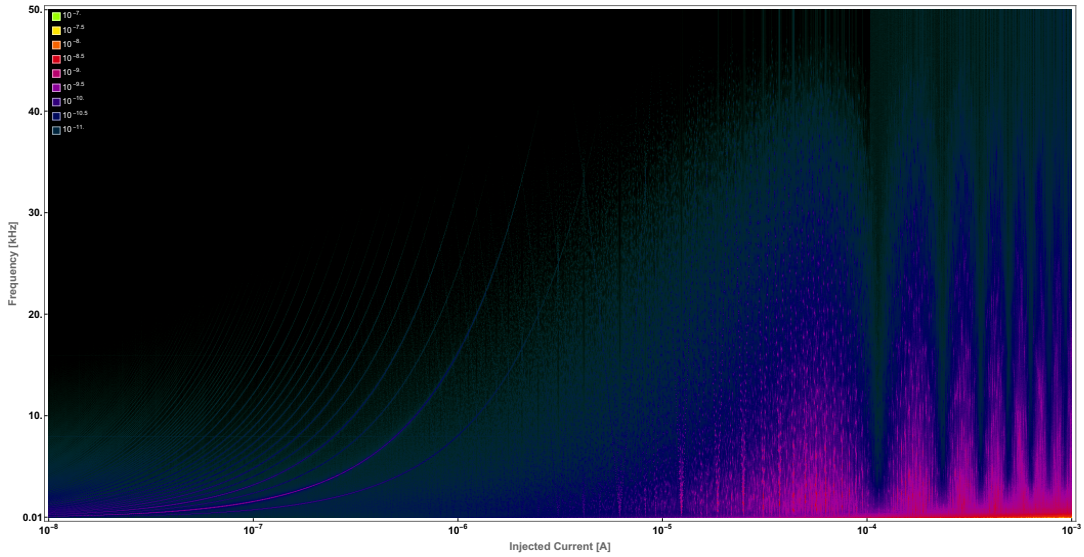

(c) Equalised rise time IIR filter **RS3** ($f_0 \approx 1.24\ kHz$).

Figure 5.19: Spectrogram of measurements processed with the different filters.

However, it has also been analysed during this study. In particular, the comparison between the two IIR filters can be observed in Fig. 5.21. As expected, the noise from the loss charge equalised filter is slightly lower than the other since its cutting frequency is smaller. However, their performance remains more or less similar and the gap is further reduced with increasing currents.

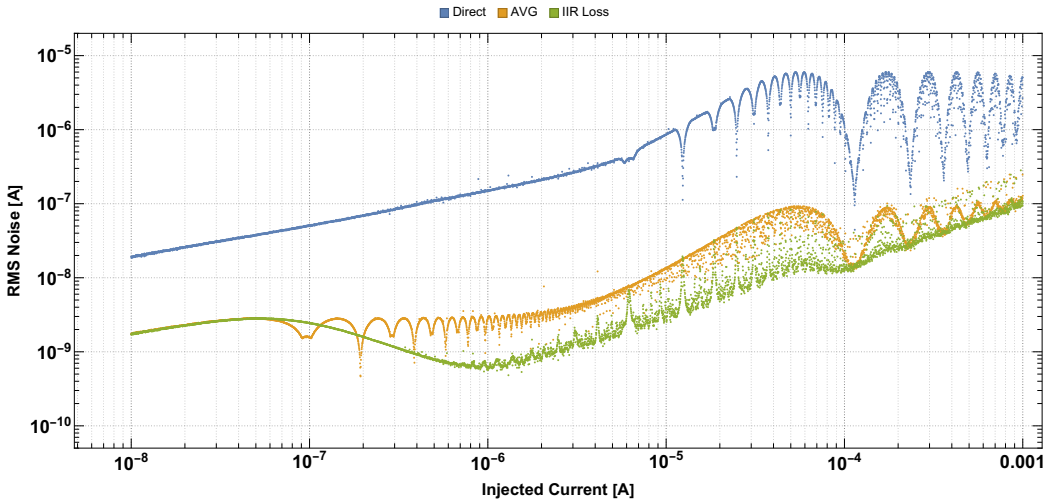

Figure 5.20: Precision σ_r comparison amongst raw versus **RS3** processed measurements.

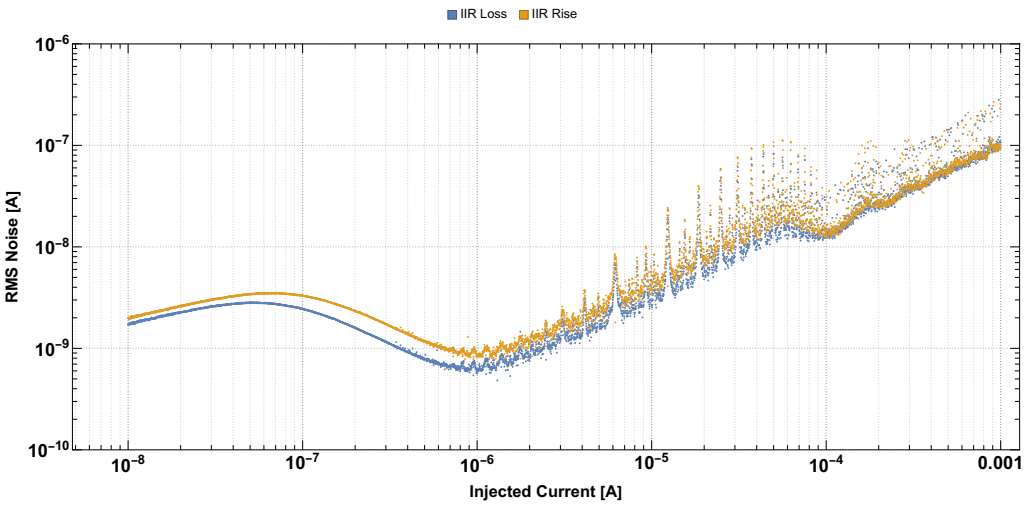

Figure 5.21: Precision σ_r comparison between the two **RS3** IIR tuning criteria.

5.8 Prototype Implementation

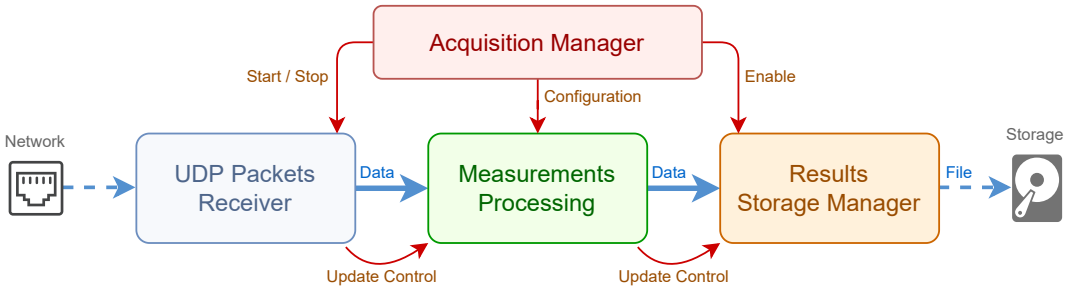

Figure 5.22: Architecture of the computer processing software.

In order to validate the performance of the described filters, custom software has been developed, which handles the real-time data stream, performs the required processing and stores the measurements in a file. Afterwards, these can be loaded into data analysis software to extrapolate the characteristics of interest. The whole software is written in C++ and based on the *Qt* [S31] framework, which has been chosen for its multi-threaded programming flexibility and computation performance. Fig. 5.22 reports the overall architecture, every block of which is executed in a dedicated thread.

The data chain source is represented by the UDP receiver, which acquires new network packets, produced by the back-end board. These latter include the raw converter measurements and further diagnostic reports, which are all managed by a specific transmission firmware, the details of which will be presented in Par. 7.10. As soon as a new packet is received, the software stores it in a newly allocated memory space, the pointer of which is then propagated to the measurement processing thread. Such a mechanism is handled by using the central Qt messaging system (signals / slots) [S32], which ensures that the memory space is only read after the successful data production. Hence, the processing thread takes care of freeing the memory once the work on the new sample is completed. All the packet reception and propagation throughout the chain have been optimised to allow an uninterrupted stream. Profiling of the compiled software showed that it could successfully cope with the ≈ 55 *Mbps* data stream produced.

The processing thread represents the heart of the application, because it runs in real-time the filtering techniques presented in this chapter. In particular, C++ classes have been developed to implement the AVG with configurable windows and IIR filters of any desired order and coefficients. Both adopt an auxiliary custom object that realises a FIFO buffer, optimised for signal processing. The code implementing the AVG class is reported following.

```cpp
#include "buffer.h"

class Moving_Average_Filter{
public:
    Moving_Average_Filter(int window_width_arg);
    ~Moving_Average_Filter();
    double new_sample(double input);

private:
    int window_width;
    Buffer* filter_buffer;
    double running_output;
};

//Constructor
Moving_Average_Filter::Moving_Average_Filter(int window_width_arg) :
    window_width(window_width_arg)
{
    //allocate the filter buffer
    filter_buffer = new Buffer(window_width);
    filter_buffer->fill_zero();
    running_output = 0;
}

//Destructor
Moving_Average_Filter::~Moving_Average_Filter(){
    delete filter_buffer;
}

//Feed a new sample and produce the result
double Moving_Average_Filter::new_sample(double input){

    //get the oldest sample
    double oldest_sample;
    filter_buffer->getPastSample(oldest_sample, -window_width);

    //add the newest sample and subtract the oldest
    running_output += input - oldest_sample;

    //push the newest sample into the buffer
    filter_buffer->fpush(input);

    //produce the output
    return running_output / window_width;
}
```

The class constructor initialises and resets the buffer, specifying the required averaging window width W. Thus, the received sample is pushed into the *new_sample* method, which is dedicated to performing the actual computation and returning the processing output. During the initial ramp, *running_output* accumulates the new sample values, while these latter also feed the buffer until it is full. Hence, at every iteration, the upcoming sample is added to *running_output*, while the oldest is subtracted from the total. In this way, the system keeps track of the sum of the latest W samples, from which it obtains the averaged output. Finally, the class destructor frees the dynamically allocated memory space, which only involves the sample buffer.

The IIR filter class shares methods similar to those presented for the AVG and its implementation code is shown here.

```
#include "buffer.h"

class IIR_Filter{
public:
    IIR_Filter(int filter_order_arg, double* num_coeff_arg, double*
    den_coeff_arg);
    ~IIR_Filter();
    double new_sample(double input);

private:
    int filter_order;
    double *num_coeff;
    double *den_coeff;
    Buffer* filter_buffer;
};

//Constructor
IIR_Filter::IIR_Filter(int filter_order_arg, double* num_coeff_arg, double*
    den_coeff_arg) :
    filter_order(filter_order_arg){

    //allocate the coefficient array
    num_coeff = new double[filter_order + 1];
    den_coeff = new double[filter_order + 1];

    //copy the coefficients
    for(int i = 0; i <= filter_order; i++){
        num_coeff[i] = num_coeff_arg[i];
        den_coeff[i] = den_coeff_arg[i];
    }

    //allocate the filter buffer
    filter_buffer = new Buffer(filter_order);
    filter_buffer->fill_zero();
}

//Destructor
IIR_Filter::~IIR_Filter(){
    delete[] num_coeff;
    delete[] den_coeff;
    delete filter_buffer;
}

//Feed a new sample and produce the result
double IIR_Filter::new_sample(double input){

    //network nets
    double feed_forward = 0;
    double feed_back = 0;
    double past_values;

    //nets computation
    for(int i = 1; i <= filter_order; i++){
        filter_buffer->getPastSample(past_values, -i);
        feed_forward += past_values * num_coeff[i];
```

```
54          feed_back += -past_values * den_coeff[i];
55      }
56
57      //push the center signal into the filter buffer
58      double center_signal = (input + feed_back) / den_coeff[0];
59      filter_buffer->fpush(center_signal);
60
61      //output sample computation
62      double out = (num_coeff[0] * center_signal + feed_forward);
63
64      return out;
65  }
```

In this case, the class constructor is used to allocate the vectors of coefficients and configure their values according to the passed arguments. Thus, even in this case, the *new_sample* method is called for every new received sample and its returning value represents the filter output. The actual computation reproduces the Direct Form II structure, shown in Fig. 5.9. In particular, the buffer is used to implement the central delays (z^{-1} blocks) and its stored values are employed to compute both the feed-forward and feed-back contributions.

The latest thread, visible on the right of Fig. 5.22, is used to store the results in a CSV file. They include the timestamps, the current values and other diagnostic information. Wolfram Mathematica [S33] has been subsequently used to carry out the batch analysis of the spectrum and the noise characteristics. The same software has been widely employed during this research, to perform further symbolic computation and to produce most of the presented plots.

Finally, the experience gained in the course of this activity made it possible to compare the advantages of developing real-time processing prototypes versus performing a complete batch data analysis, as it was done during the preliminary characterisations of the BLMASIC. (cf. Par. 4.8). In particular, it has been observed that the development time required to implement a performant batch analysis software is comparable to that employed for a real-time processing one. The former has the disadvantage of requiring much more computation memory and storage space. On the other hand, real-time processing should only be employed after the basic device features have been successfully validated. In fact, such a higher-level approach could eventually mask some issues affecting the raw data since these latter are never saved in a file. In general, the analysis methodology was optimised according to each project step but, advancing towards the final stages, such a latest strategy is strongly suggested. In fact, it allows us to face the challenges pertaining to the definitive processing ones and to inspect the system in a broader view. Furthermore, the time required for software programming is more rewarded by the flexibility and reusability of the code. Indeed, minor modifications would allow to implement it into an embedded system or to perform a high-level FPGA synthesis (HLS) [S34], [81], [S35], which provides the possibility of testing it directly into the commissioning hardware.

5.9 Measurement Error Forecasting

The described implementation of filters enables an effective characterisation of their performance over real data streams. In fact, computing the standard deviation over a large set of output samples yields the measurement error σ_r for specific injected currents m^{dc}. Such a procedure will actually be adopted during the characterisation of the digitiser (cf. Chapter 8).

However, once known the measurement noise spectrogram $M_k(m^{dc})$ (cf. Par. 5.3), it is alternatively possible to carry out an accurate estimation of the output error $\sigma_r(m^{dc})$, without acquiring any further data from the test-bench. As a first step, a desired linear low-pass and unity-gain filter is chosen with transfer function $G_d(z)$. The same can be expressed in the DTFT counterpart H_k, by applying the transformation

$$H_k = G_d(z)\Big|_{z=e^{i2\pi\frac{k}{N}}} , \tag{5.30}$$

where N is the number of samples acquired for every injected current. Thus, the spectrum at the filter output can be obtained as

$$R_k(m^{dc}) = H_k \times M_k(m^{dc}) , \tag{5.31}$$

where $R_0 = M_0 = m^{dc}$, while for $k > 0$ it represents the noise spectrum. So, rejecting the continuous component ($k = 0$), like in equation (5.10) the $\sigma_r(m^{dc})$ can be calculated in the form

$$\sigma_r(m^{dc}) = \sqrt{\sum_{k=1}^{N-1} |R_k(m^{dc})|^2} . \tag{5.32}$$

This methodology provides a direct estimation of the measurement error after the processing, without redoing the R_k measurements. Hence, a wide range of filters and tunings can be evaluated to identify the optimal ones, before proceeding with their actual implementation. To this aim, an automatised procedure can be performed by a computer.

Chapter 6

Front-End Development

The core of the upgraded BLM electronics is described here, starting with the technical specifications and the original architecture of the board. The implementation of radiation hardening techniques is illustrated. Afterwards, the circuit schematics will be presented, detailing the structure of each building block. The solutions adopted for the redundancy of the data path, the board control and the improved diagnostic features are notable. Finally, the PCB layout design is exhibited, highlighting the methodologies deployed to optimise performance and reliability. In the latest section, a template of the board configuration is reported, providing a quick start to the operations.

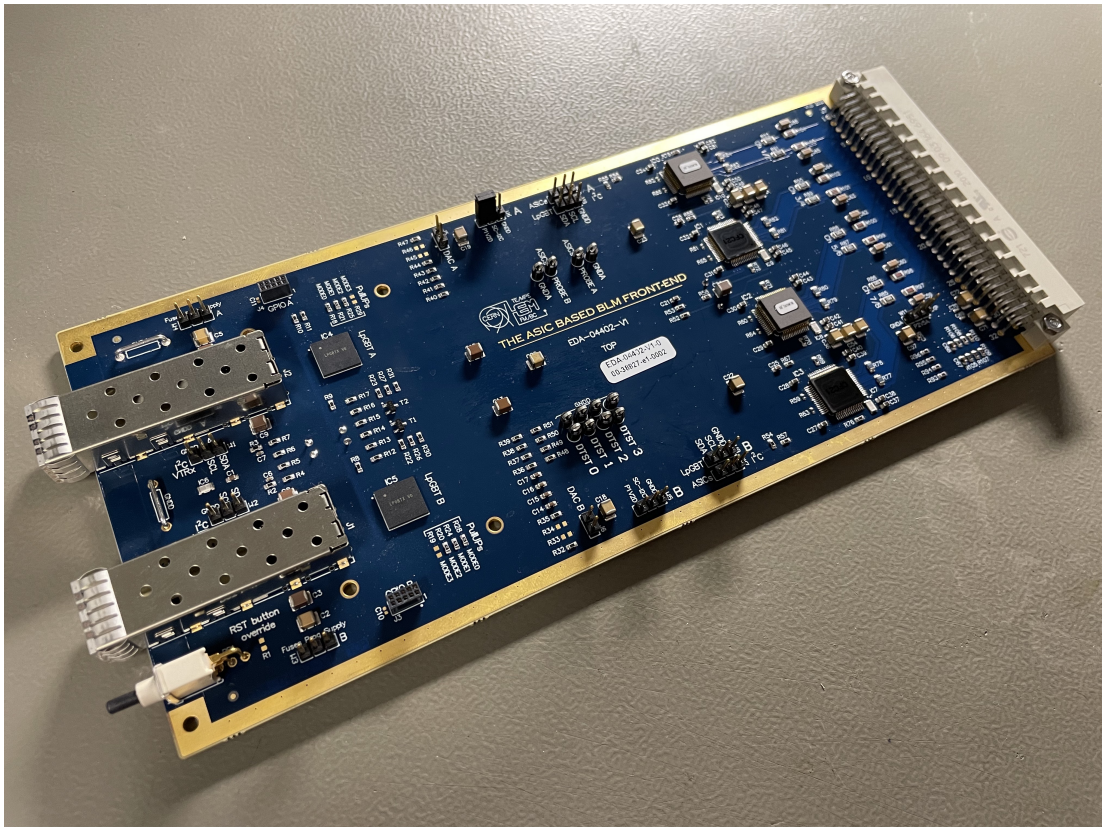

Figure 6.1: The ASIC based BLM Front-End.

6.1 Board Introduction

The new front-end development has represented one of the most challenging activities of this research. Great care has been devoted to each component and a constant search for quality details accompanied the entire project. The objectives were not limited to a feasibility study of the new concepts, but to the actual release of a product, representing a milestone in the novel BLM electronics for future accelerators.

The results achieved with the BLMASIC characterisation (cf. Chapter 4) made it possible to familiarise oneself with its performance and improvement aspects. Thus, the next development step has focused on the card that hosts such converters (the BLEIC) and provides the features necessary for successful integration within the system. The details concerning the board architecture and the involved circuitry are reported following.

6.2 Technical Specifications

Feature	**Value**
Input channels	8
Operative input range	$< 10\ pA$ — $1\ mA$
Maximum input current	$5\ mA$
Local configuration bus	I^2C
External communication interface	5.12 *Gbps* Versatile Link [S13], [S36]
Communication medium	Single-mode optical fibre
Optical wavelengths	$1270\ nm$ — $1330\ nm$
Digital data high-reliability techniques	Fully double redundancy
Front-End radiation TID tolerance	$> 3\ kGy$
Front-End radiation induced fault tolerance	SEU resilient
Electrical input protection	against ESD & Overload
General diagnostics	Temperature, Voltages and Optical
Advanced diagnostics	BLMASIC internal probes
Maximum power consumption	$5\ W$
Board PCB size	$220\ mm\ \times 100\ mm$
Board chassis frame	Backward compatible
PCB material	Panasonic MEGTRON 6

Table 6.1: Front-End board key specifications.

The technical specifications agreed upon during the preliminary design stages for the front-end board are reported in Tab. 6.1. It is clear that the BLMASIC conversion characteristics are intrinsically part of these (cf. Tab. 3.1) however, in accordance with the goals introduced in Par. 2.4, higher level features are presented in this context.

The core targets have been tailored based on of the expertise acquired with the system currently in operation (cf Par. 2.3.1) and improvements have been introduced to overcome the bottlenecks affecting the former solution. They were the result of many discussions and reviews carried out within the Section, to fulfil the various demands and allow a smooth upgrade of the electronics. The presented features, in conjunction with the back-end processing ones, will allow achieving the final specifications, demanded for the overall upgraded system (cf. Par. 2.6).

6.3 Board Architecture

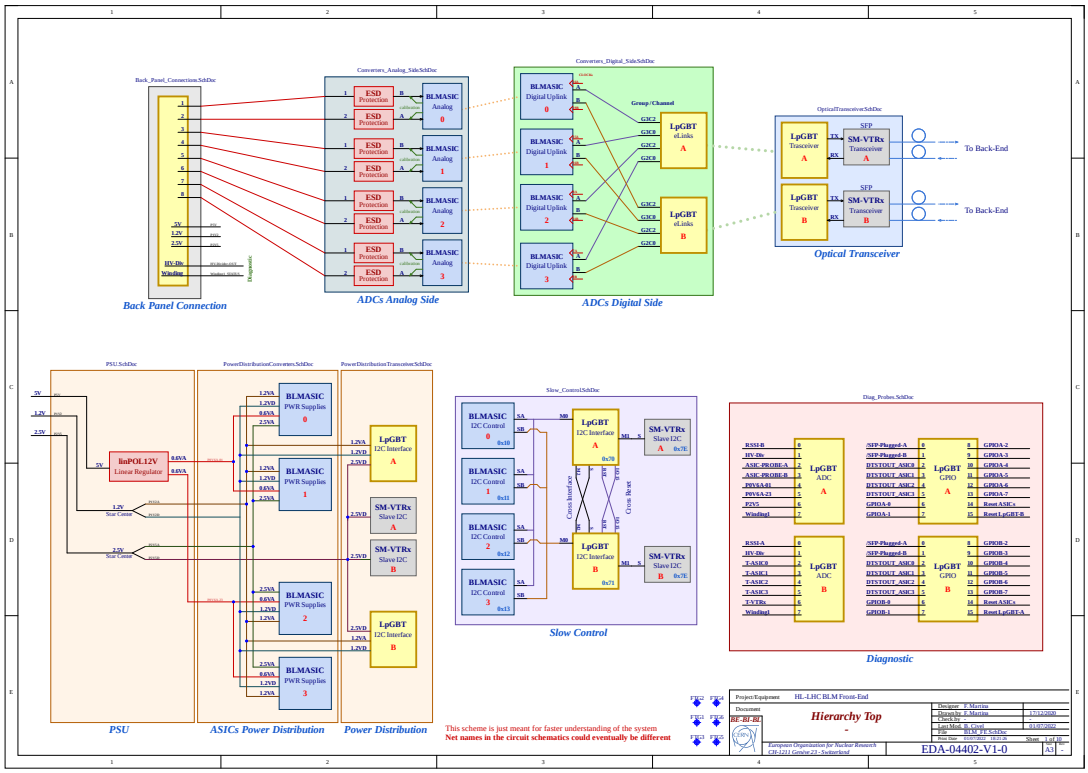

Figure 6.2: Front-End board architecture overview, highlighting the main operational blocks.

The title *"ASIC Based BLM Front-End"*, assigned to this project, embodies the spirit behind its concept and realisation. Indeed, by employing CERN-developed ASICs exclusively, such a card incorporates all the features necessary to convert the beam loss currents, produced by the IC detectors, into properly coded optical data streams. Furthermore, it provides a range of diagnostic functionalities, guarantees reliable operations within the harsh environment of the HL-LHC tunnel and establishes two-way redundant communication with the electronics located at the surface.

A hierarchical methodology has been employed during the development, subdividing the schematics into the various tasks undertaken by the circuit. An overview of the architecture is shown in Fig. 6.2, which will be analysed in the following paragraphs.

Each block in the diagram concerns a specific task and summarises the content of the unique schematic page to which it is referred. The upper row blocks include the beam loss signal path, which propagates from left to right. The currents are acquired by the back panel connector, converted into digital data by the BLMASICs, packetised by the *lpGBT* [44], [S13] devices and transmitted through optical transceivers (the *CWDM-SM-VTRx* [82]). Instead, following the same direction in the bottom row, the 3 powering supply and distribution blocks are visible, together with the slow control connections and the diagnostic probes.

Many of the complex component functions and interfaces have been divided into different graphical units to better organise the schematics. These can also be spread over different sheets, although they refer to a single hardware device. Finally, to facilitate the schematic referencing throughout the text, the nomenclature *sch.x.yz* will be employed, being: x the schematic page (from *1* to *10*), y the sheet row coordinates (from A to E) and z the column coordinates (from *1* to *5*). Parentheses are used to identify the spans between columns or rows.

Both the schematics and the layout development have been carried out using Altium Designer [S37].

6.3.1 Component Selection

The radiation tolerance requirements certainly represented the most demanding design challenge. In fact, the full employment of commercial components would have greatly simplified the board development and expanded the portfolio of solutions. During the early development stages, a survey about the effects of radiation on electronic components allowed the formulation of a repertoire of guidelines to be respected. In particular, it was assessed that active components are generally more vulnerable to radiation than passive ones. This is due to the wider range of damages that ionising radiation could cause to semiconductors and micro-electronic structures (cf. Appendix A) [83]–[87]. Furthermore, the complexity of the integrated circuits entails an higher probability of failure. Single Event Effects (SEE, cf. Appendix A) and drift of characteristics can be difficult to detect and they could eventually lead to serious system malfunctions. Therefore, to guarantee maximum radiation tolerance performance, it has been decided to completely avoid the employment of active commercial components. Hence, the instrument will entrust all the complex tasks to ASICs, specifically designed at CERN for radiation-tolerant applications.

Due to their extensive use, the passive components have also been subjected to a meticulous selection process. Indeed, although there are less restrictive rules about their employment in radiation environments, they significantly affect the long-term reliability of the card. Thus, high-quality and durable components, specifically the capacitors, have been chosen for the whole design, preferring those compliant with industrial-grade or automotive standards (AEC-Q200 [S38]). Furthermore, the demand

for accurate characterisation procedures, including reverse engineering, the study of suitable materials, radiation transport modelling and the validation of the radiation tolerance properties, is always time consuming and costly [88]. For all such reasons, it was preferred to employ well-tested X7R / C0G based ceramic capacitors (MLCC) [S39], [S40], which can withstand such harsh environments and they benefit from excellent temperature stability and low parasitic effects.

Finally, efforts have been undertaken to limit the use of plastics as much as possible, because they tend to deteriorate and become brittle after prolonged exposure to radiation. This was not possible for the back-panel connector (PBT material), which must comply with the existing installation and the front reset button. However, alternative solutions have been discussed to overcome the issue in stand-alone applications involving very high TID. The board is also equipped with auxiliary headers and probes, using PBT, LCP or Nylon based insulation. These are assembled into the first prototypes, to access a variety of diagnostic tools required during the characterisation in the laboratory. However, they will not be present in the commissioning version.

6.4 Circuit Design

The following paragraphs include the description of the several circuit modules in the same order as reported in the board schematics (cf. Appendix E). For easy visualisation, the shown schematic particulars might result narrower than in the original sheets, but the content is unmodified.

6.4.1 Back-Panel Connection

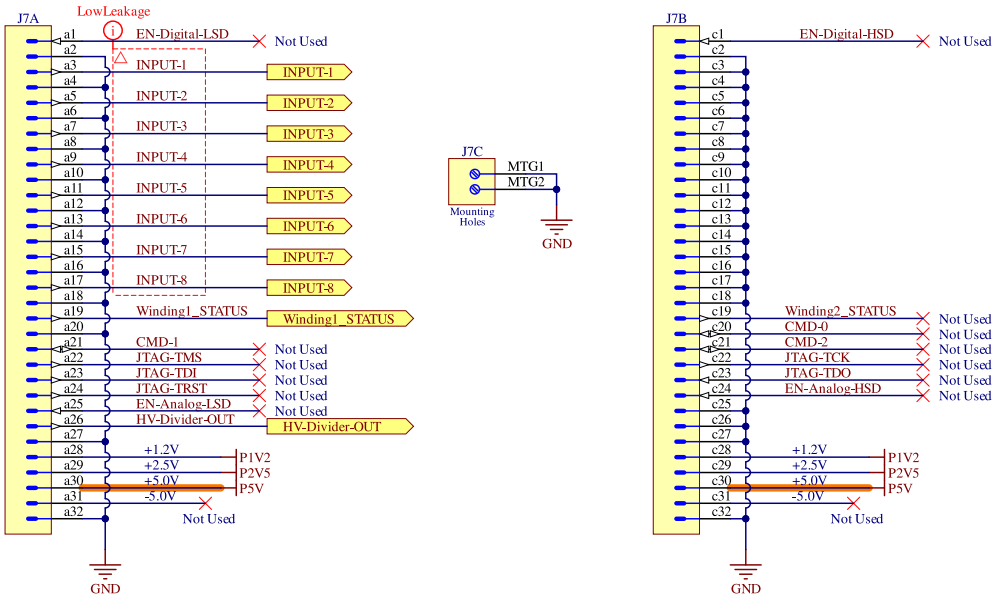

Figure 6.3: Diagram of the back-panel connections. Detail from *sch.2.C2* and *sch.2.C4*.

The front-end developed in this work has quite a different architecture in comparison with the operational one, described in Par. 2.3.1. However, they share the same physical geometry and connection to the hosting crate. In fact, it has been decided to employ the same back-panel connector (*DIN 41612*) and a compatible pinout, allowing a smooth transition amongst the systems. All the electrical connections between the card and the rest of the plant are carried out by means of this interface, the signals of which are reported in the schematic detail of Fig. 6.3.

First, the crate provides the power supply rails to the board, including the 1.2 V, 2.5 V and 5.0 V buses. In particular, the 1.2 V is not standard with respect to the operational card (internal ref. EDA-00593), which instead requires 1.5 V, representing the most significant compatibility mismatch between the two versions. However, a solution involving different crate power suppliers has been studied, which can be easily exchanged to meet the two requirements. Further details about the board powering will be described in Par. 6.4.9.

On the left of Fig. 6.3, it is possible to notice the 8 measurement input channels, which come from the IC detectors and the high voltage monitoring signal. This latter is a 60 dB attenuated replica of the IC bias source, which the system must continuously gauge to validate the loss detection. Excluding the ground connections, the remaining pins were left unused in the previous system and discussions are still ongoing about their final employment. But, in agreement with the crate design, it has been decided to dedicate one of them for measuring the transformer winding voltage (*Winding1_STATUS* pin). This represents the common source for all the power regulators and is integrated within the Input Power Unit module (IPU, internal ref. EDA-03918). Finally, a future upgrade will likely employ the additional pins to enable direct configuration access through the crate. This will be discussed later in Par. 6.4.5.

6.4.2 Converter Inputs

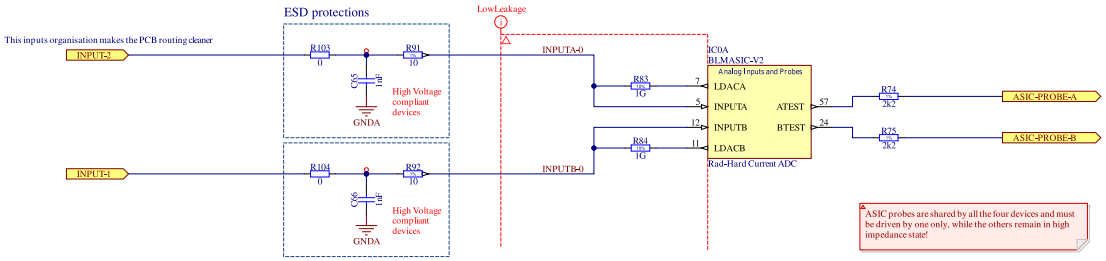

Figure 6.4: Input ESD protections and calibration DAC connection. Detail from *sch.3.A(1-4)*.

The connection between the input channels and the BLMASICs is performed according to the schematic in Fig. 6.4, which reports the analogue section of one of the total four converters. The electrostatic discharge (ESD) protections integrated within the device have been selected as the sturdiest provided by the foundry. However, because of commercial reasons, they have only been validated with the human body model

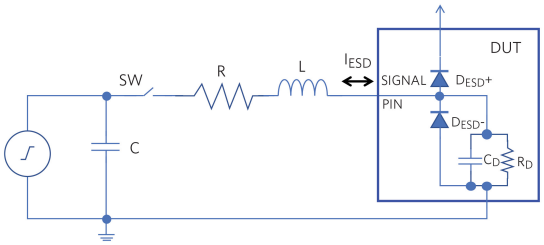

Figure 6.5: Electrostatic discharge model circuit.

(HBM, JS-001-2012 [S41] and MIL-STD-883H [S42]), which concerns a finger touching the device pins. This situation can be emulated using the schematic of Fig. 6.5, which also includes a typical ESD guard embedded into the integrated circuit. The source on the left has high internal resistance ($\approx 1\ M\Omega$) and it charges the body capacitor while the switch is open, meaning that the finger is not touching the device. Then, the source is disabled and the capacitor charge is transferred to the circuit by closing the switch (finger is touching). Suddenly a discharge current I_{ESD} flows into the device, producing a fast release of energy from the internal components D_{ESD} and R_D, due to the Joule effect. The protection is designed to dissipate the generated heat properly and to withstand the I_{ESD} transient without any damage. In the context of the HBM, the involved values of components are:

- $V_{ESD} \geq 2\ kV$,
- $C = 100\ pF$,
- $R = 1.5\ k\Omega$,
- $L = 750\ nH$.

Similar behaviour can also occur in the BLM installation if the IC is biased before being connected to the front-end. This happens because the stabilisation capacitor of the chamber will have already charged up to the bias voltage $V_{IC} \approx 1.5\ kV$, but the inherent capacity of plates could still be discharged (for $C_{IC} \approx 500\ pF$). In the worst case, this leads to a dissipation of energy about three times larger than the one evaluated for the HBM, so it could exceed the ESD protection margins. Thus, an additional safe circuit has been included in the board to improve the system reliability. To avoid the use of diodes, due to the radiation issues already discussed in Par. 6.3.1, a solution employing only passive components has been adopted. This is highlighted on the left of the schematic in Fig. 6.4. The exact simulation of the internal ASIC protection is not feasible, because the solution proposed by the foundry is subject to confidentiality. However, several SPICE models have been developed employing diodes that might be similar to the embedded ones. In particular, referring to a single channel (INPUT-2), it has been concluded that the capacitors C65 and the resistor R91 perform an effective transient suppression, dampening the power peak inside the ASIC. A larger capacitor provides more effective protection, but it cannot be increased too much since it reduces

the system response speed. Furthermore, the resistor value must be increased along with the capacitance, to avoid instability in the converter integrator.

The input channels were partitioned on the 4 BLMASICs according to Tab. 6.2. The mapping between the device inputs and the 8 channels might appear counterintuitive, but such correspondences greatly improve the routing quality. This does not affect the ease of access to results, because the subsequent firmware will remap them.

Input Channel	**BLMASIC Designator**	**BLMASIC Channel**
Ch.1	IC0	CH1 *(B)*
Ch.2	IC0	CH0 *(A)*
Ch.3	IC1	CH1 *(B)*
Ch.4	IC1	CH0 *(A)*
Ch.5	IC2	CH1 *(B)*
Ch.6	IC2	CH0 *(A)*
Ch.7	IC3	CH1 *(B)*
Ch.8	IC3	CH0 *(A)*

Table 6.2: Mapping of input channels over the 4 BLMASICs.

The schematic also shows how the internal calibration DAC is connected to the input, producing a current that is added to the detector one. According to the DAC operations, presented in Par. 3.3.6, this solution allows the injection of bias currents within the range $-100\ pA$ to $128\ pA$. Finally, the device analogue probes (cf. Par. 3.3.8) have been connected to the rest of the circuit by the resistors showed on the right side of the picture. This diagnostic feature will be described in Par. 6.4.8.

6.4.3 Measurement Data Path

To reliably transfer the data streams produced by the BLMASICs towards the processing back-end, they need to be collected and properly encoded. Such a task has been entrusted to the lpGBT device, presented in Par. 2.4. It allows for the implementation of a multi-purpose optical link, which fulfils a variety of features required in high-energy Physics experiments. Referring to Fig. 6.2, it is possible to notice how the redundancy is implemented by connecting the BLMASIC lanes to two independent lpGBTs. In fact, the beam loss data are propagated alongside the devices, which provide all the procedures required for their successive transmission. In ordinary conditions, the back-end can compare the two links and perform a low-level data validation for enhanced reliability. On the other hand, in the case of an optical link failure or a momentary unlock, the data can still reach the back-end, which could flag a warning but continue its accelerator protection. This improves the availability of the entire system, reducing the frequency of dumps due to a link loss.

The detail schematics of the data and clock connection amongst the 4 BLMASICs and the 2 lpGBT devices is reported in Fig.s 6.6 and 6.7. According to the BLMASIC design (cf. Par. 3.3.8), the data related pins can be directly linked from the source to the destination using a differential lane. No additional components are required, since 100 Ω terminations are embedded into the lpGBT and the BLMASIC uplink interfaces are compatible with the E-Links. Whereas, the 80 MHz clocks driving the BLMASICs require external terminations (visible in Fig. 6.7). Finally, the clock source ports differ between the two lpGBT devices due to routing optimisation.

The receiving E-Links include a variety of possible configurations, which have been investigated to select the best suitable for the application. In fact, the lpGBT provides 7 E-Link groups of 4 interfaces each (for a total of 28 E-Links), but only 4 of these are required for the purposes. This led to great flexibility in the development and allowed the choice of the solution providing the best nets routing. The devices have been assigned to the E-Links according to Tab. 6.3.

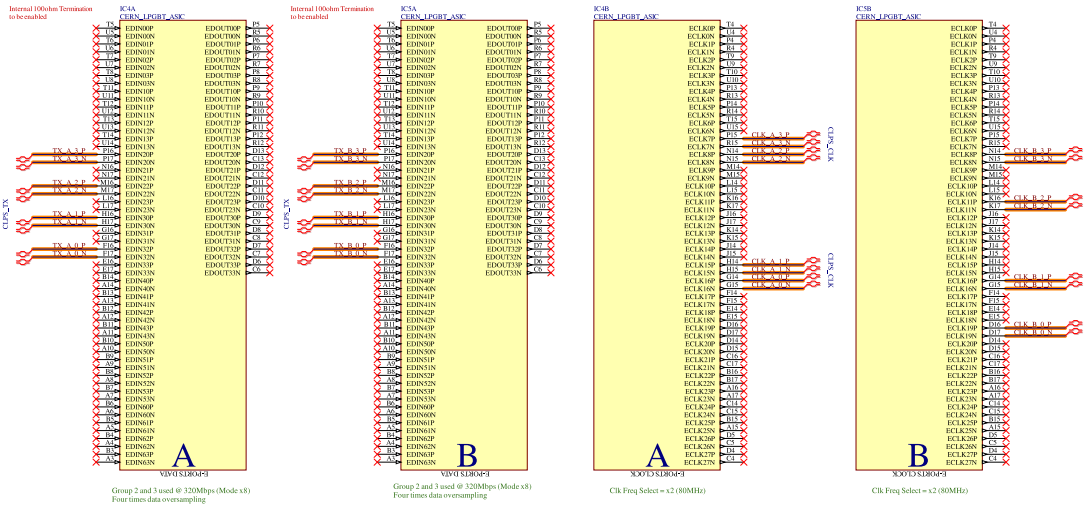

Figure 6.6: lpGBT connection of E-Links. Detail from *sch.4.(A-C)(1-5)*.

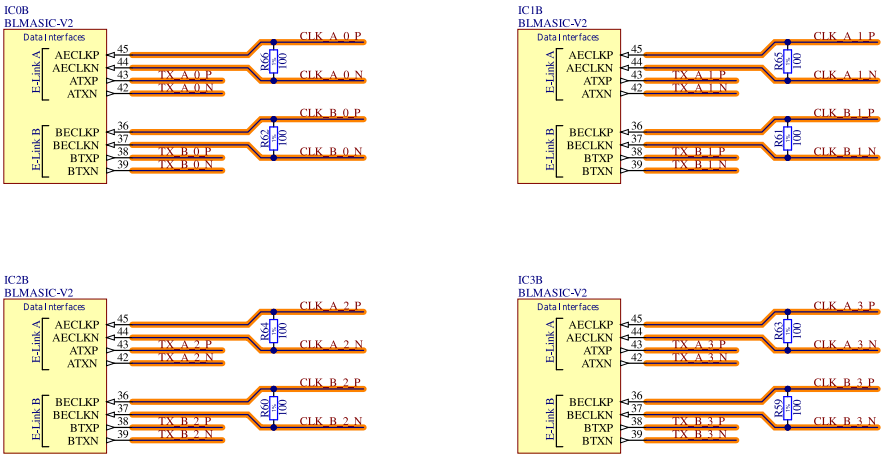

Figure 6.7: Connection of BLMASIC data lanes and clocks. Detail from *sch.4.(D-E)(1-3)*.

BLMASIC Designator	E-Link Group	E-Link Channel
IC0	3	2
IC1	3	0
IC2	2	2
IC3	2	0

Table 6.3: BLMASICs to E-Links assignment. Equal for both the lpGBT **A** and **B**.

Each E-Link logic can be configured to sample the input data at a bit rate between 160 *Mbps* and 640 *Mbps*. However, the full band allocated to the group cannot exceed 640 *Mbps*. Thus, the 4 interfaces have to be set accordingly. Since the BLMASIC rate is only 80 *Mbps*, data will be oversampled in all the possible configurations. Specifically, it has been decided to adopt the compromise solution of 320 *Mbps*, which provides a 4 times bit oversampling. This will enable efficient edge detection for the successful alignment of the stream phase. The lpGBT cannot perform an automatic phase tracking of the 80 *Mbps* bitstream. Thus, such a procedure must be carried out by the dedicated back-end firmware module, the detail of which will be described in Par. 7.5. The necessity of fulfilling this function is reinforced by the fact that the clock-to-data phase relation cannot be fixed during the operations, despite having adopted a routing length matching over the 8 redundant lanes and clocks. Indeed, even assuming that the data are synchronised to the clock for a single receiving lpGBT device, this does not ensure the same condition for the second one. For instance, when the **B** transceiver is used, due to the failure of link **A** (see Fig. 6.8), the BLMASIC selects the lpGBT **B** clock by checking its asserted "ready" flag. So, the produced data have a fixed phase relation with clock **B**, dependent only on the logic latencies and the PCB propagation delays. On the other hand, if both the links are working, the BLMASICs can indistinctly select one of the two clocks. Thus, the data received by the non-selected lpGBT will have an unknown phase relation. However, the back-end can still recover the data, assuming that the lpGBTs employ synchronous clocks. This is valid for the new system since the back-end (VFC-HD) transmits the same clock over the 2 redundant optical downlinks.

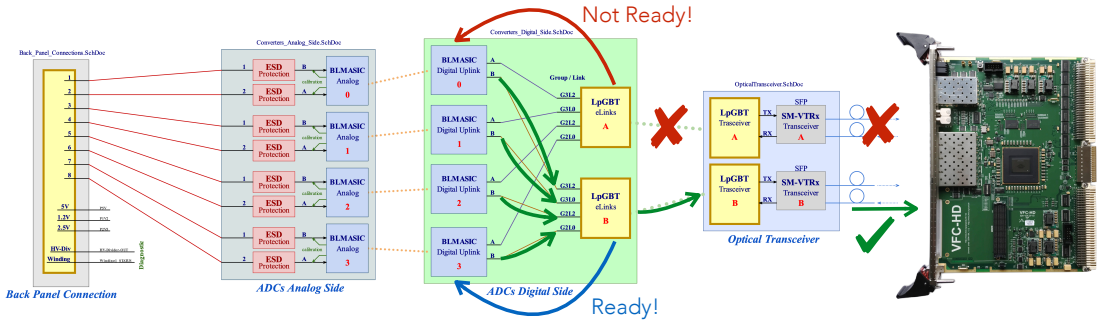

Figure 6.8: Measurement path redundancy in case of link A failure.

6.4.4 Fast Lanes

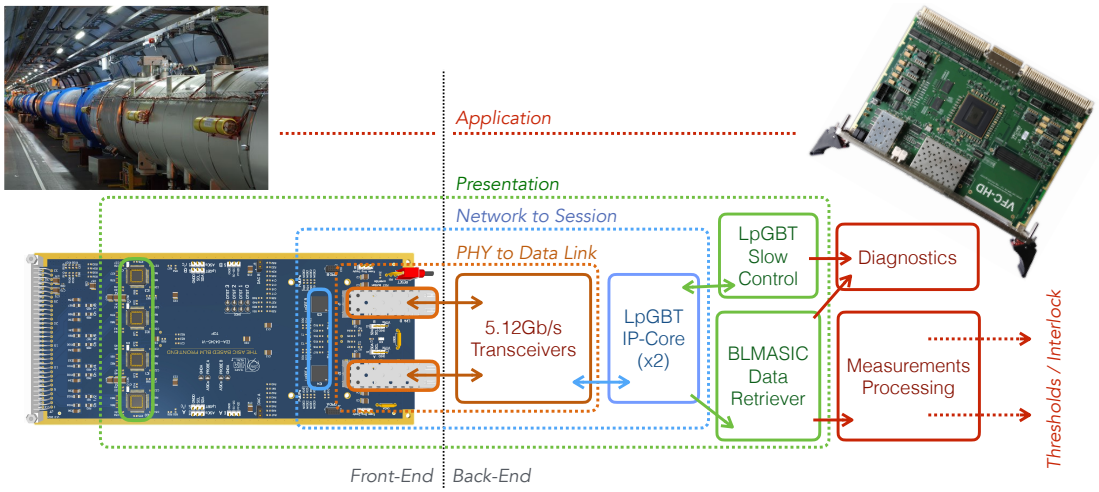

Figure 6.9: Data connection layers according to the OSI model.

The data transfer between the digitiser and the readout boards is achieved by a redundant optical link, operating at the bitrate of 5.12 *Gbps*. The different communication layers implemented within the system are reported in Fig. 6.9, providing an analogy with those defined in the Open Systems Interconnection model (OSI model) [89]. The physical medium and the optoelectronic modulators represent the lowest layer, which includes the single-mode optical fibres and the transceivers. A pair of CWDM-SM-VTRx [S36], [82] devices were employed for this goal, these being specifically designed at CERN for $\mathcal{O}(km)$ data transmissions within radiation environments. The transceivers, hereafter referred to as the VTRx, can operate with a bitrate up to 10.24 *Gbps* with a Bit Error Ratio of 10^{-12}, which is twice as faster as required in the application. Furthermore, they are characterised by a typical transmission power of 2 dBm, while the receiver sensitivity and the saturation threshold are -15.4 dBm and 0.5 dBm, respectively. These specifications allow an effective pairing with COTS single-mode transceivers. In fact, since radiation tolerance is not demanded in the back-end electronics, standard networking SFP modules will be employed.

The detailed integration of the VTRx inside the board circuit is reported in Fig. 6.10. Unlike the previous VTRx versions, the module used here has been commissioned for this specific BLM front-end upgrade and it is packaged within a standard SFP case (the cages are shown in Fig. 6.11). This provides a more sturdy mechanical support and more accessible system maintenance. The pinout is similar to the standard one, but the VTRx requires a 2.5 V rail for its operations instead of the commercial value of 3.3 V. Unfortunately, except for the radiation issues, this is the only restriction on the possible employment of a commercial SFP module within this board, but future upgrades could enable this option for possible uses outside the accelerator tunnel.

The module is interfaced to the lpGBT by a 100 *Ohm* differential lane and the terminations are embedded into the devices. The I^2C connection provides access to the laser driver (TOSA) configuration and the dedicated **RSSI** signal allows measurement of the received optical power. Finally, the module documentation reported that the

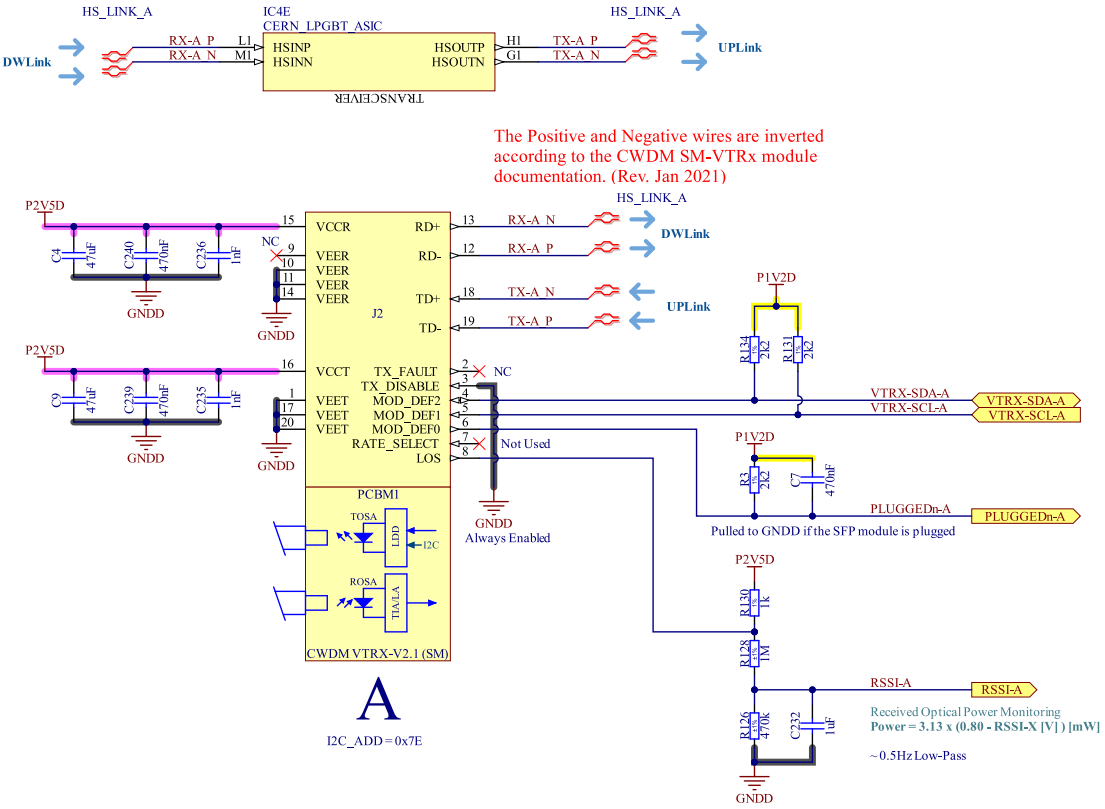

Figure 6.10: Optical transceiver socket connections. Detail from *sch.5.(A-C)(1-3)*.

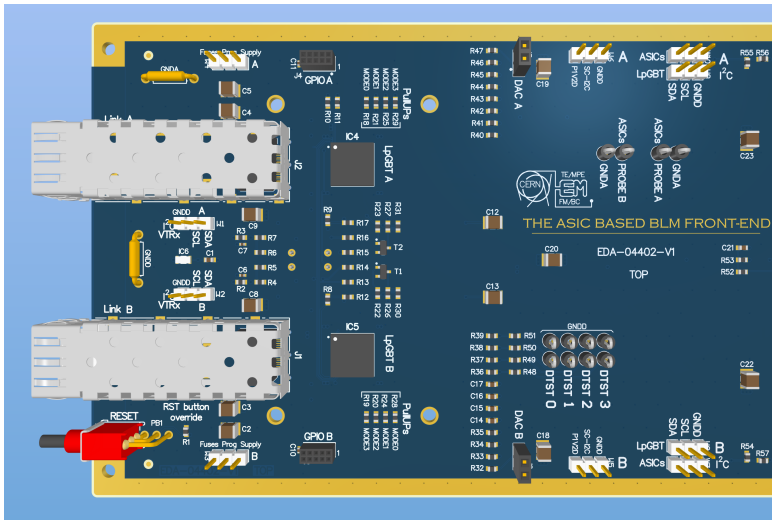

Figure 6.11: Transceiver area. The lpGBTs and the SFP cages are visible on the left.

SFP pin **MOD_DEF0** is tied to the ground. Hence, the diagnostic **PLUGGEDn** signal has been introduced, which is pulled up by a 2.2 $k\Omega$ resistor when the SFP cage is free, while it will be de-asserted whenever the module is inserted correctly.

6.4.5 Board Control and Redundancy

This circuit section involves the repertoire of features necessary to configure the card and access the internal memory of the various devices. In fact, all the slow rate interfaces, which are not necessarily operating on the BLM measurements, are part of this set of tools.

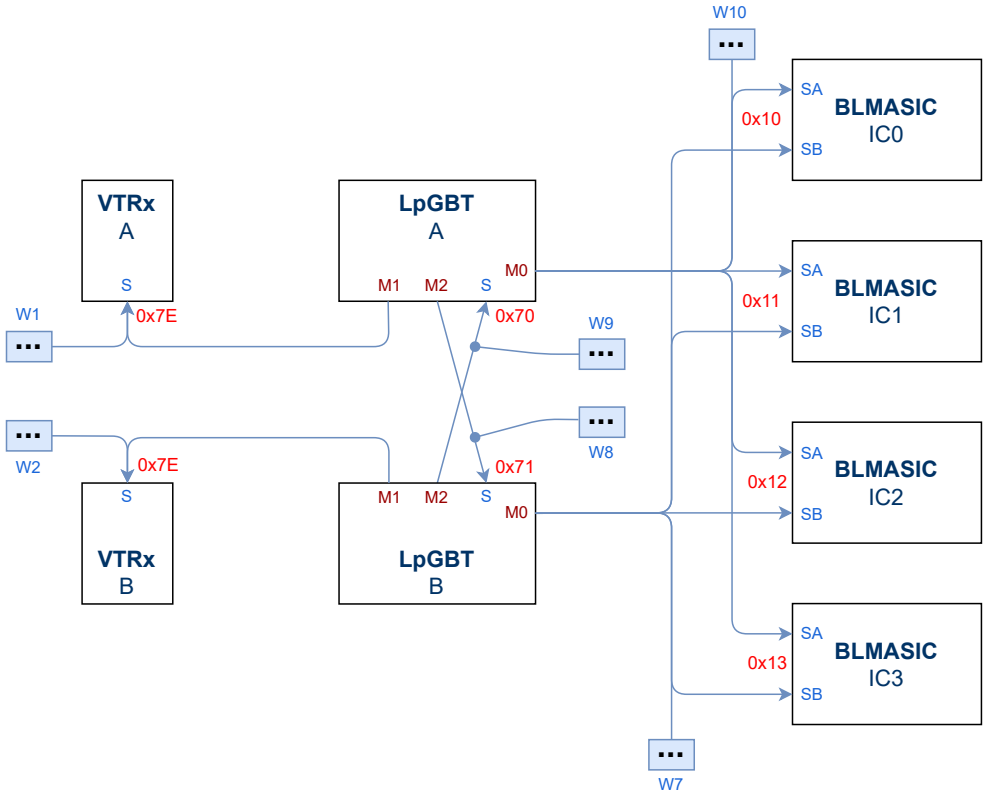

Figure 6.12: I^2C interface tree throughout the front-end, including the device addresses and the external header designators.

As it was already presented for the BLMASIC test-bench (cf. Chapter 4), the main bus exploited for these purposes is the I^2C, which can set a flexible connection amongst the 4 BLMASICs, the 2 lpGBTs and the 2 VTRx transceivers. In particular, all the devices are equipped with a slave I^2C port and the lpGBTs also include 3 masters. Thus, the connection tree is developed according to the architecture shown in Fig. 6.12 and implemented in the hardware as the schematic of Fig. 6.13 reports.

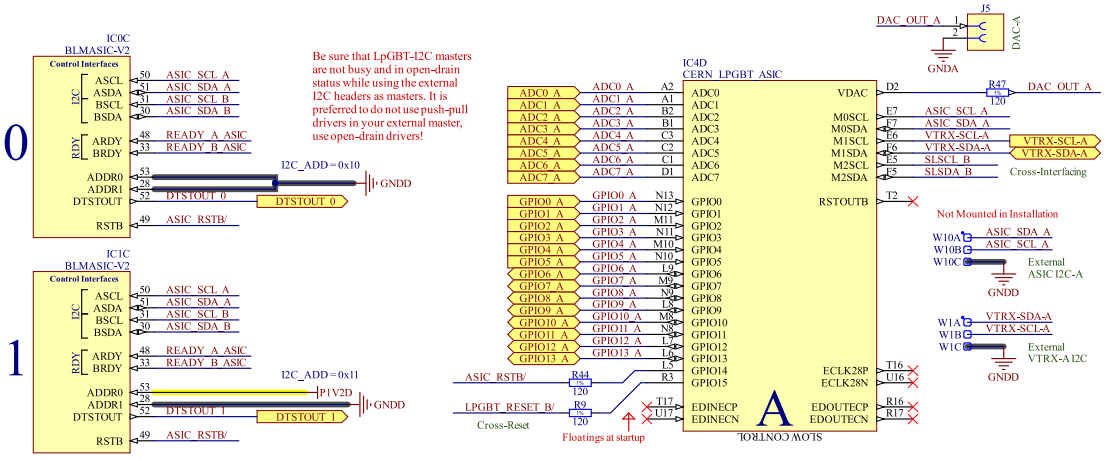

Figure 6.13: Connection example of the slow control interfaces, reporting the lpGBT **A** (on the right) together with the BLMASICs **IC0** and **IC1** (left). Detail from *sch.6.(A-B)(1-3)*.

The connection tree also shows the device designators, including the interface accessing headers (visible in Fig. 6.11), in order to better identify them on the schematics. The master port **M0** is connected to a bus including 4 BLMASICs, the addresses of which are properly configured by **ADDR0** and **ADDR1** pins, as reported in the figures. Instead, the master port **M1** is used to access the VTRx module.

A meticulous study was carried out to guarantee that all the devices could be accessed even in case of a single link failure. If at least one optical link is ready, the back-end firmware can query the connected lpGBT by the downlink control (named IC). This enables the launch of I^2C transactions throughout the card using the master interfaces. In addition, a cross-connection between the lpGBTs has been included to allow a prompt investigation of the problem affecting the faulty link. This is achieved by connecting the master **M2** to their slave **S** ports. In fact, such an architecture provides access to the disconnected device and its related VTRx by using the operative one. The mechanism is shown in Fig. 6.14, highlighting the benefits provided.

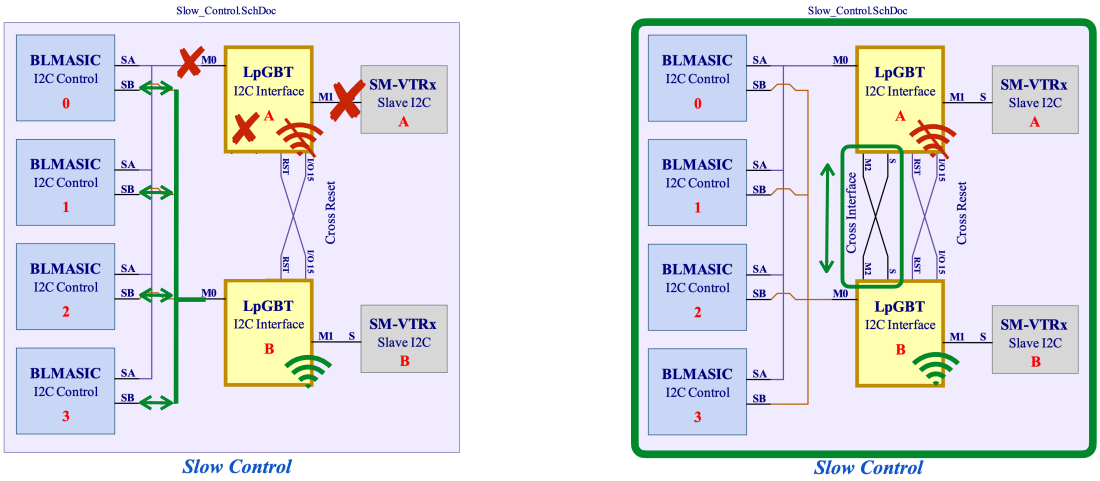

Figure 6.14: I^2C Access in case of **A** link failure. The right example shows how the introduction of the cross interface allows the connection of the whole components.

The latest discussions have paved the way for enabling the I^2C control of the card by its back-panel connector. In fact, some of the connector pins are actually not used, like the JTAG (cf. Par. 6.4.1), since the corresponding features are not available in the employed devices. Thus, for the next upgrade, it has been planned to replace the JTAG with I^2C interfaces, to directly query the front-end from the crate. Concerning the connection tree, several options are available depending on the number of additional masters. A suitable solution would involve the connection of such external ports to the slave interfaces of lpGBTs. If required, those buses can also be tied together since the devices have different addresses. The only strict requirements concern the compliance with the 1.2 V front-end digital signals and the avoidance (or management) of conflicts, which may occur among the masters.

6.4.6 Hardwired Configuration

The lpGBT requires a fixed setting for some of its pins before the start-up (cf. Fig. 6.15). They concern the I^2C device address, the operating mode and the booting configuration memory. In particular, the transceiver mode has been adopted here, because bi-directional communication is required. The lower speed of 5.12 *Gbps* was chosen and paired with the most reliable Forward Error Correction (FEC) [90] method available: the *FEC12*. This solution enables the back-end with the capability of performing corrections of up to 12 consecutive wrong bits, representing an important feature in such a critical application.

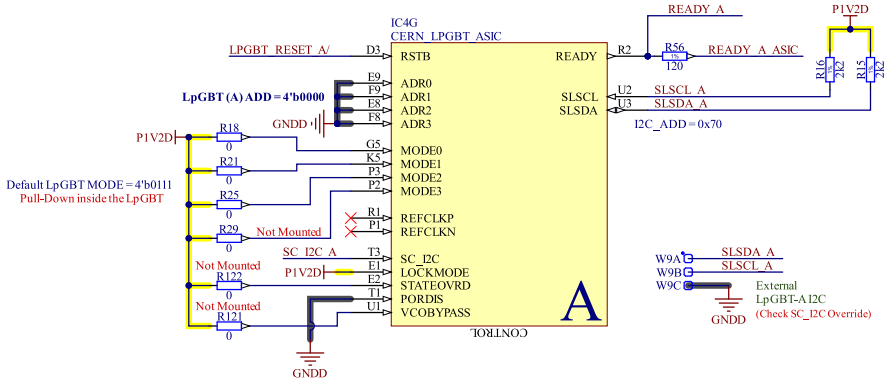

Figure 6.15: Hardwired lpGBT configuration, reset and ready flag. Detail from *sch.6.(C-D)3*.

6.4.7 Reset

The reset controls of the devices are all active low and they are driven by a capacitor and resistor low-pass circuit (cf. Fig. 6.16). Their values have been tuned according to the system power supply, to allow the reset release only once the rail voltages are stabilised. Furthermore, the large capacitors of 47 μF guarantee the suppression of eventual spikes on the signal and the reset button debouncing. Whenever it is required, this latter can be used to force a manual reset of the card.

The reset signals can also be driven directly by the lpGBT GPIOs, allowing for a partial or a complete board restart commanded by the back-end. In particular, both the lpGBT **A** and **B** can trigger the reset of the BLMASICs and each other. The GPIOs are controlled by the dedicated tripled register, which can be programmed either by the I^2C or the optical IC (downlink) interfaces, and they are reliable in the case of SEU. It is suggested to drive the reset of the devices common to **A** and **B** either by forcing the signals to low or leaving them floating, to avoid conflicts between the two lpGBTs. Anyway, the 120 Ω resistors will ensure that the current drawn, or supplied, by the controlling GPIO pins will be limited to the device specifications.

The typical engaging and releasing timings are reported in Fig. 6.16. If the application requires fast release times, it is possible to force the reset signal to a high state using the GPIOs or install a recycling Schottky diode, as illustrated. However, a radiation-tolerant component must be employed for the latest solution.

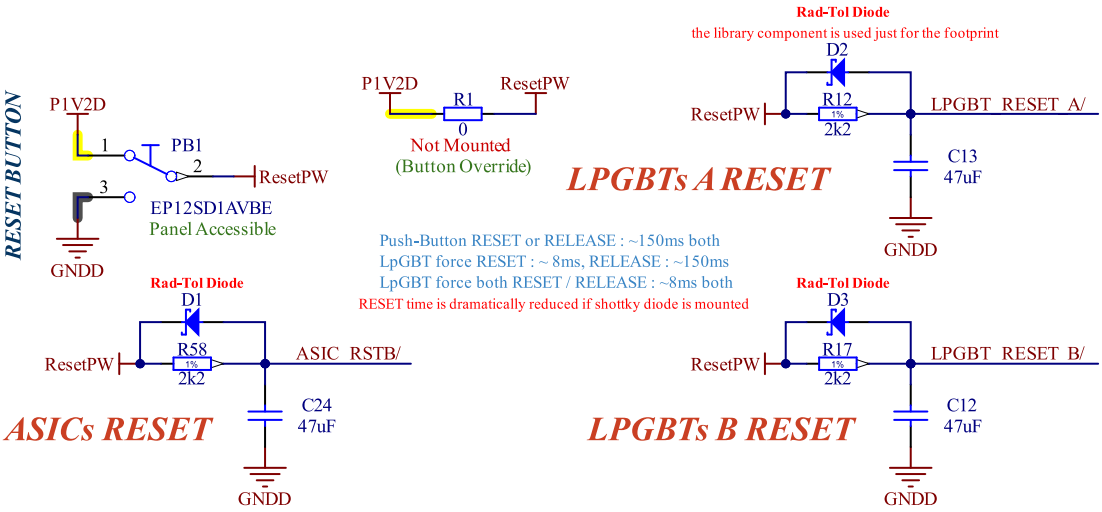

Figure 6.16: Device reset circuitry, including the typical engaging and releasing times.

6.4.8 Diagnostics

The new diagnostic features represent a crucial factor in this electronic upgrade. In fact, the previous system was limited due to the single direction data link used for the communication towards the back-end board (cf. Par. 2.3.1). Instead, the new architecture allows to query the system status whenever needed and set configurations for advanced analysis. These enable a comprehensive screening of the card behaviour, even while the human intervention in the tunnel is restricted. In addition to accessing device registers, through the I^2C interfaces (cf. Par. 6.4.5), the signal monitoring capabilities of the lpGBTs have been widely exploited in this circuit. In particular, the device includes 16 GPIO pins and a multi-channel ADC, characterised by 10 *bit* of resolution, $\approx 1\ V$ of full scale and a conversion time $< 1\ \mu s$. The following list summarises the monitoring features that have been implemented by using these peripherals.

The reported attenuation factors are due to the total contribution of embedded and external dividers (if present), which are introduced between the measured signal and its ADC input.

- Power Rails:
  - **P1V2D** (1.2 V): digital bus monitoring. $Att. \approx 1 : 2.38$,
  - **P1V2A** (1.2 V): analogue bus monitoring. $Att. \approx 1 : 2.38$,
  - **P2V5 A/D** (2.5 V): bus monitoring. $Att. \approx 1 : 5$,
  - **P0V6-01** (0.61 V): **CM** of IC0-1 BLMASICs. $Att. \approx 1 : 1$,
  - **P0V6-23** (0.61 V): **CM** of IC2-3 BLMASICs. $Att. \approx 1 : 1$,
  - **Winding1_STATUS**: IPU transformer winding. $Att. \approx 1 : 20$.
- Temperature Sensors (for a total of 7):
  - Proximity to BLMASICs IC0-3,
  - Proximity to VTRx,
  - Internal to lpGBT A/B.
- Ionisation Chamber Detectors:
  - High-Voltage bias (1.5 kV). $Att. \approx 1 : 2000$.
- High-Speed Lanes:
  - Received optical power (RSSI) on VTRs A/B,
  - SFP VTRx plugged modules,
  - lpGBT Eye Opening Monitor Threshold DAC (EOM).
- Advanced Features:
  - Internal BLMASIC analogue probes (one at a time),
  - Internal BLMASIC digital probes.
- Auxiliary Features:
  - 16 free of use GPIOs (on the external header),
  - 2 channels 12 *bit* voltage DAC.

Due to a large number of monitored quantities and the limited input channels, they have been spread over the two lpGBTs. Then, it is required to query both devices to collect comprehensive diagnostic data. Nevertheless, such a solution is not limiting, even in case of a link failure, thanks to the cross I^2C access presented before in Par. 6.4.5. However, the critical signals, like the detector bias voltage and the transformer winding

voltage, have been duplicated and can be directly read from both the lpGBTs. Furthermore, many signal probes include a filtering capacitor, which adds a low-pass pole to reduce the high-frequency noise components. Details about the chosen cut frequencies can be found in the schematics. Finally, all the resistors have been tuned considering the source impedances, even those external to the board (see the example of Fig. 6.17).

The proximity temperature sensors are based on PT1000 thermistors, which have a resistance of 1000 Ω at $0°C$ and a thermal coefficient of 3.85 Ω K^{-1} (see Fig. 6.18). They can be directly connected to the lpGBT inputs since adjustable current sources can be enabled in parallel to the ADCs. Finally, the BLMASIC probes have also been mirrored on dedicated headers, to allow their monitoring by using external instruments during the validation campaign. They are visible at the centre right of Fig. 6.11.

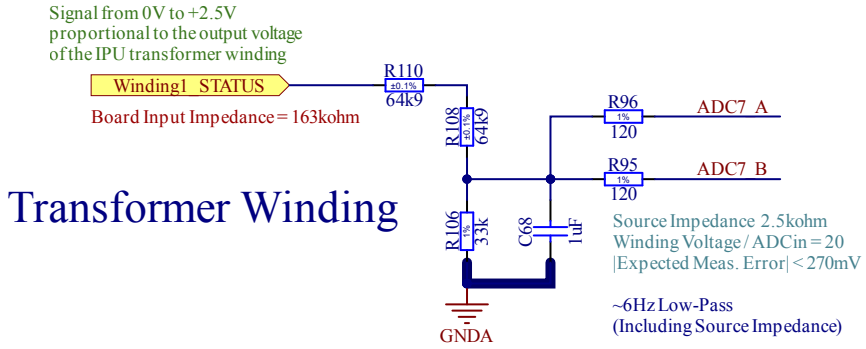

Figure 6.17: Divider used for the transformer winding monitoring. Detail from *sch.7.E(1-2)*.

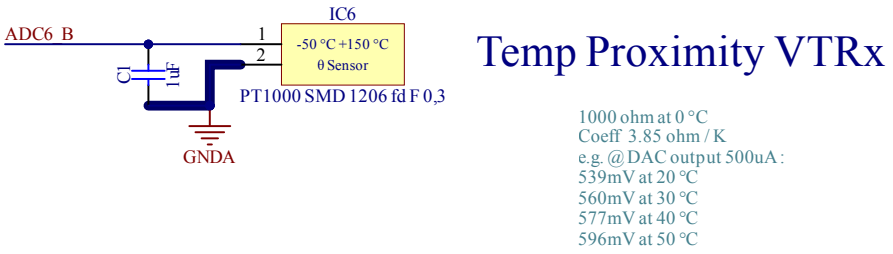

Figure 6.18: Temperature sensor detail from *sch.7.E(2-3)*.

6.4.9 Power Distribution

Reliable power distribution is mandatory in such a sensitive instrument since high-speed components are electrically paired to pico-ammeters. Thus, great care was exercised to minimise the impedance of rails and keep the digital supplying paths isolated from the analogue ones. Since the main regulators are placed outside the card, in the PSU unit (internal ref. EDA-03919), the goals have been achieved by optimising the routing and adopting an effective bypass of disturbances towards the ground.

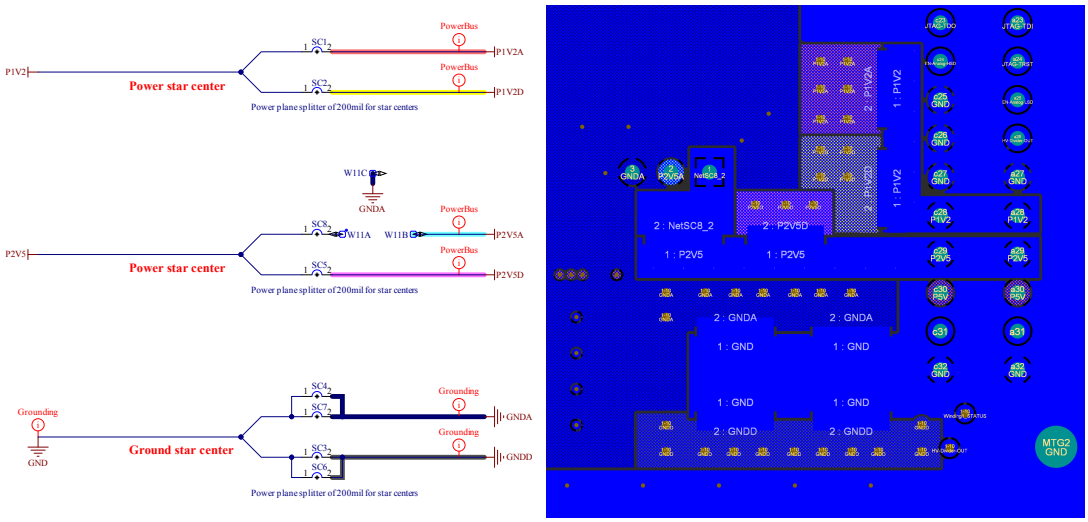

Figure 6.19: Schematics of powering rail star centres and relative layout (right). Detail from *sch.10.(A-C)3* (left).

The isolation of the different rails and grounds was obtained using dedicated copper planes in the PCB, which bring power to the various devices. The connection of those which share the same voltage, thus the same source, is performed in the so-called *star centres* (see Fig. 6.19). These are the only locations where analogue rails can join the digital ones and are placed very close to the back panel connector (see Fig. 6.19 at right). Such a solution prevents supply pins of sensitive components from being located in the flow path of the currents disturbed by the digital switching circuitry. Furthermore, a hierarchical strategy of bypass capacitors has been adopted for the whole board. Ceramics of 100 μF, with larger packages (Imperial SMD size 1210), have been spread symmetrically in the area surrounding the components, to stabilise the powering planes. While, smaller SMD 0402 or 0201 ones, of values 470 nF and 220 nF respectively, have been placed underneath the supply pins (on the bottom layer), for a near field bypassing (see Fig. 6.20).

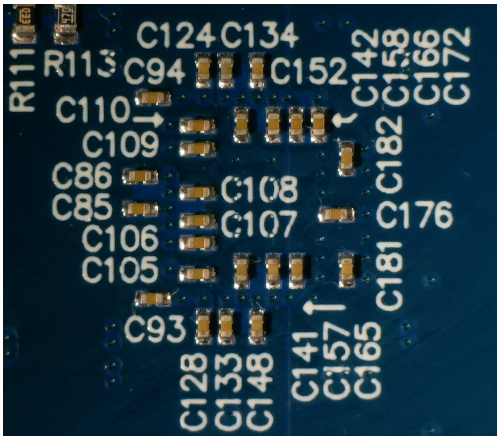

Figure 6.20: BLMASIC bypass capacitors, placed on the bottom layer.

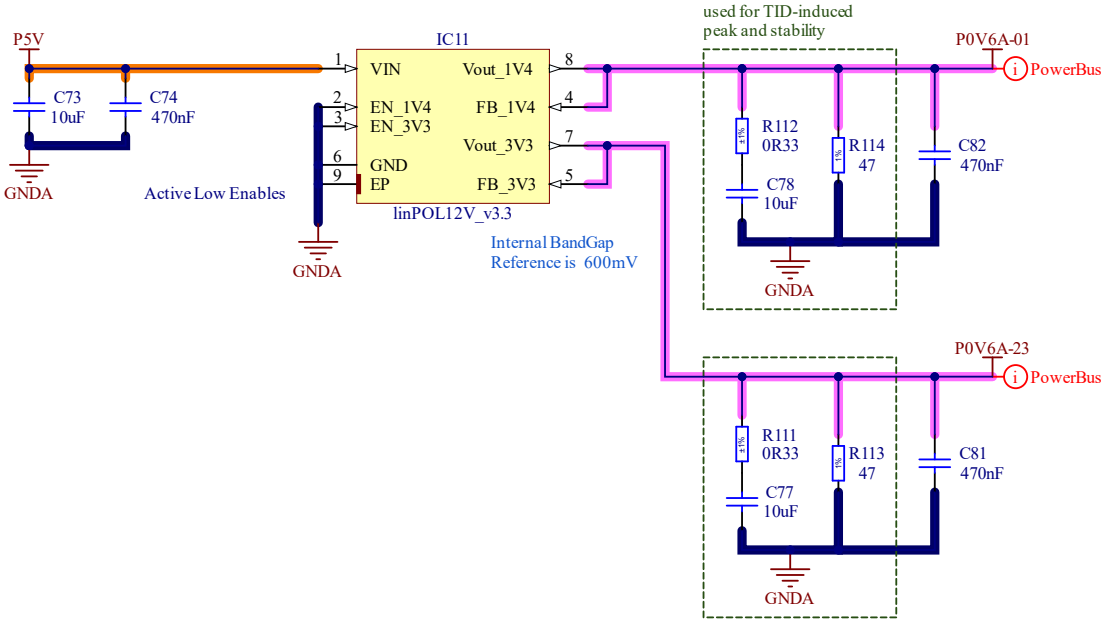

Figure 6.21: linPOL12V Linear regulator for the P0V6 rails. Detail from *sch.10.(D-E)(2-3)*.

This solution minimises the current path, thus reducing the track inductance. Such small packages have been chosen for their reduced parasitic components and to enable an effective placement in the dense circuit zones. An investigation of the manufactured capacitors assessed that the selected values provide optimal impedance characteristics for this application.

The card houses only a linear regulator, which generates the very sensitive BLMASIC common mode reference (CM) at approximately 620 mV. It is represented by the radiation-hardened *linPOL12V* ASIC [S43], which has been supplied at 5 V and configured according to the schematic in Fig. 6.21. The resistors R113 and R114 ensure that the device sources a minimum operational current in all possible conditions. This feature is crucial since the P0V6 rails act as virtual grounds for the BLMASICs. Thus, the measuring input currents are actually subtracted from the preloading value of $\approx 13\ mA$.

For the various components, the detailed power connections have been designed according to their datasheets and they are reported in Fig.s 6.22, 6.23 and 6.10. The BLMASIC pins **VREFA** and **VREFB**, shown in Fig. 6.22, require a special care. Indeed, these regulate the CFC quantiser thresholds and consequently its measurement gain, according to the description of Par. 3.3.5. In this prototype, such pins have been tied to the analogue 1.2 V supply (P1V2A), to simplify the system in view of its first validation. However, this means that the measurement results are subjected to supply voltage fluctuations. In other words, the card must be calibrated every time the PSU module is replaced or in the case of major temperature drifts. This may eventually be impractical for the front-end commissioning. Thus, a stable voltage reference should be available in the circuit. For instance, a second linPOL12V device can be installed.

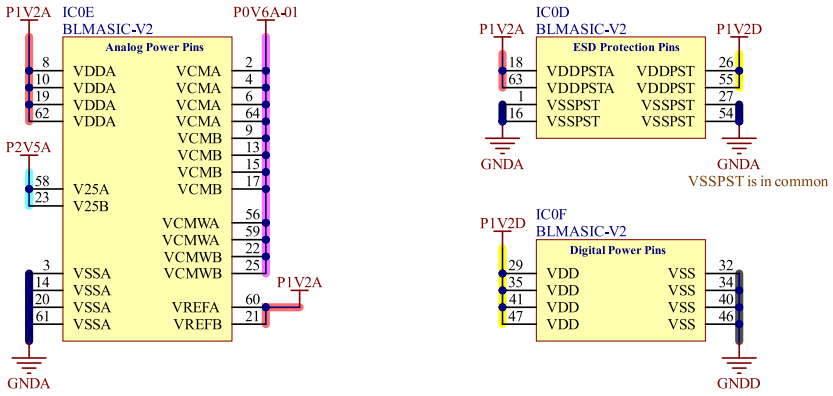

Figure 6.22: BLMASIC (IC0) power supply schematic. Detail from *sch.8.(A-B)(1-2)*.

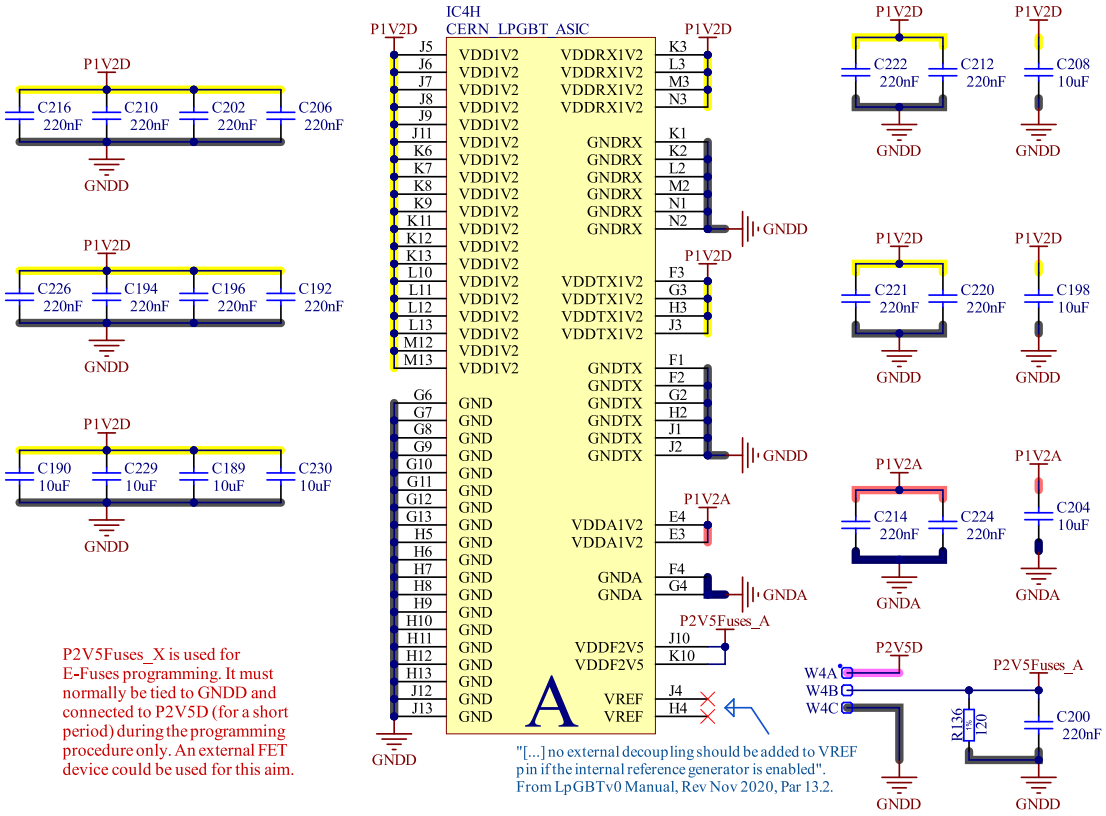

Figure 6.23: lpGBT (A) power supply. Detail from *sch.9.(A-E)(1-3)*.

6.5 Layout Design

6.5.1 Board Geometry

The board shape and size have been strictly set according to the requirements of its housing cassette, which is the same used for the older front-end. Furthermore, it also shares with this latter the same back-panel connector position, allowing in principle the compatibility with the current crates, as long as the 1.5 V power rail is modified to 1.2 V (cf. Par. 6.4). The entire board measures 220 $mm \times 100$ mm, and its maximum thickness is < 15 mm, when populated by components.

6.5.2 PCB Stack-Up

Many design iterations were undertaken to find the best solution for the strategy of vias, the number of layers and the PCB materials. These have been evaluated based on the routing quality, reliability, signal integrity requirements and cost.

The first layout concept involved employing only through hole vias, to manufacture the board using the most common and validated technology. However, it was assessed that the drilling of tiny holes (0.15 mm) should have been required below the highly dense package of the lpGBT. The company in charge of the board production has stated that this is feasible for them. However, according to the IPC-6012D/DS standard for the PCB design [91], such a solution would not be sufficiently reliable for critical applications. In fact, the hole diameters are too small compared to the board total thickness of ≈ 1.6 mm, with a ratio smaller than $1:8$, which is the recommended one.

Thus, different solutions were investigated, in particular those based on laser drilled via technologies, which represent the state of the art for high-density circuits. They allow the manufacturing of blind and buried vias, which can connect different layers without going through the entire board thickness. According to the proposals, two options for a 10 layers stack-up have been analysed. The first involved blind vias between the layers 1 to 2 and 9 to 10, while the 2 to 9 connection is drilled (named *1+N+1*). Instead, the second option involves a more flexible combination of blind and buried vias, in which the drilled hole is only necessary between layers 4 to 7 (*3+N+3*). In conclusion, the latter option was selected for the prototype since it allows a more effective routing of the E-Links (cf. Par. 6.4.3). The manufacturing technology set the hole diameter of vias to 4 mil, while the pad is 10 mil wide. On the other hand, the internal drilled holes present a diameter of 8 mil and a pad 16 mil wide.

Finally, a high-speed and low-loss certified material has been selected as the best candidate for the project. It is produced by Panasonic with the name *MEGTRON 6* [92] and involves the Laminate type R-5775G and the Prepreg R-5670G. These are advanced materials, designed for network equipment, mainframes, integrated circuit testers and high-frequency measuring instruments. Their main features concern: low dielectric constant and dissipation factors, low transmission loss and high heat stability (decomposition at $Td = 410^{\circ}\ C$). The selected stack-up specifications and the via technology are shown in Fig. 6.24, including the thickness of copper and insulation layers and the electric constants.



#	Name	Material	Type	Weight	Thickness	Dk	Df
	Top Overlay		Overlay				
	Top Solder	Solder Resist	Solder Mask		0.01001mm	3.5	
1	Top Layer		Signal	1/2oz	0.0175mm		
	Dielectric Top	R-5670G	Prepreg		0.076mm	3.47	0.002
2	L2	CF-004	Signal	1/2oz	0.0175mm		
	Dielectric 23	R-5670G	Prepreg		0.076mm	3.47	0.002
3	L3_pln	CF-004	Plane	1/2oz	0.0175mm		
	Dielectric 34	R-5670G	Prepreg		0.076mm	3.47	0.002
4	L4	CF-004	Signal	1oz	0.035mm		
	Dielectric 5	R-5670G	Prepreg		0.076mm	3.47	0.002
	Dielectric 45	R-5670G	Prepreg		0.12499mm	3.71	0.002
5	L5_pln	CF-004	Plane	1oz	0.035mm		
	Dielectric 56	R-5775G	Core		0.50038mm	3.71	0.002
6	L6_pln	CF-004	Plane	1oz	0.035mm		
	Dielectric 1	R-5670G	Prepreg		0.12499mm	3.71	0.002
	Dielectric 6	R-5670G	Prepreg		0.076mm	3.47	0.002
7	L7	CF-004	Signal	1oz	0.035mm		
	Dielectric 2	R-5670G	Prepreg		0.076mm	3.47	0.002
8	L8_pln	CF-004	Plane	1/2oz	0.0175mm		
	Dielectric 3	R-5670G	Prepreg		0.076mm	3.47	0.002
9	L9	CF-004	Signal	1/2oz	0.0175mm		
	Dielectric 4	R-5670G	Prepreg		0.076mm	3.47	0.002
10	Bottom Layer		Signal	1/2oz	0.0175mm		
	Bottom Solder	Solder Resist	Solder Mask		0.01001mm	3.5	
	Bottom Overlay		Overlay				

Figure 6.24: PCB stack-up specifications and arrangement of vias.

6.5.3 Component Placement

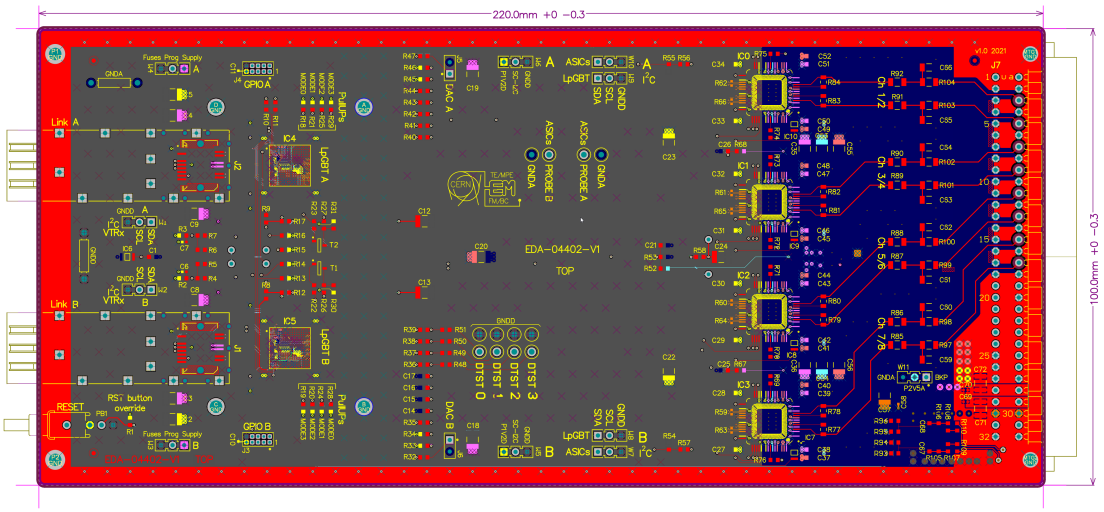

Figure 6.25: Top layer of the front-end PCB.

The component placement and orientation have been carried out to facilitate the connection of the nets sensitive to integrity or leakage issues. For example, they include the high-speed lanes, the E-Links and the analogue input channels. In addition, solutions for effective routing and the separation of power supplies have been adopted. In particular, the BLMASICs are placed close to the back panel connector, which provides the detector currents. Instead, the lpGBTs are situated on the opposite side. Thus, the routing involves very short connections towards the VTRx transceivers, the cages of which are exposed on the front panel. The placement can be viewed in Fig. 6.25, in which the different ground planes are highlighted with colours and the input channel tracks can be easily distinguished on the right. The grey plane represents the digital ground **GNDD**, the blue one is the analogue ground **GNDA**, while the surrounding red is directly connected to the chassis **GND**. They are all tied together but in a unique star centre placed on the bottom layer, as described in Par. 6.4.9.

After establishing the position of the back panel connector, the lpGBTs, the VTRx transceivers and the BLMASICs, the work was focused on the arrangement of bypass capacitors. In particular, those needed in the proximity of components were installed first (cf. Fig. 6.20). Then, the temperature sensors were distributed on the card to get the most accurate monitoring of hot spots. Finally, the placement was concluded by arranging the diagnostic probes to be easily accessible and reserving convenient spaces for the rest of the components. Both the board faces have been populated, but the protrusions of devices have been limited to $< 2\ mm$ in the bottom layer, to avoid undesired contact with the containing metal box.

LAYER IDENTIFICATION		LINE (mm) WIDTH	LINE (mm) WIDTH/GAP	FINISHED COPPER THICKNESS	CONTROLLED IMPEDANCE +/−10% Single	CONTROLLED IMPEDANCE +/−10% Differential	DIELECTRIC THICKNESS *=CRITICAL	Er
TOP LAYER	SIGNAL LAYER	–	3.5/6.5	17.5 um	–	100 Ohms	* 0.076 mm	3.47
LAYER 2	SIG/PLN LAYER	–	3.0/8.0	17.5 um	–	100 Ohms	* 0.076 mm	3.47
LAYER 3	GND PLANE			17.5 um			* 0.076 mm	3.47
LAYER 4	SIGNAL LAYER	–	3.5/8.0	35 um	–	100 Ohms	* 0.20 mm	~3.6
LAYER 5	POWER PLANE			35 um			0.50 mm	3.71
LAYER 6	POWER PLANE			35 um			* 0.20 mm	~3.6
LAYER 7	SIGNAL LAYER	–	3.5/8.0	35 um	–	100 Ohms	* 0.076 mm	3.47
LAYER 8	GND PLANE	–		17.5 um			* 0.076 mm	3.47
LAYER 9	SIG/PLN LAYER	–	3.0/8.0	17.5 um	–	100 Ohms	* 0.076 mm	3.47
BOTTOM LAYER	SIGNAL LAYER	–	3.5/6.5	17.5 um	–	100 Ohms		
FINISHED BOARD THICKNESS +/−10% :							1.6 mm	

Differential lines width / gap will be tuned during the pre−manufacturing

Figure 6.26: Geometries of 100 Ω differential transmission lanes.

6.5.4 Routing

The development of the board layout began with the general definition of the design rules and the assignment of purposes for the various layers. In particular, it has been decided which layer must support the net routing or be used for the power distribution (cf. layout layers in Appendix F). This has been investigated alongside the definition of the reference planes for differential transmissions. A stripline typology has been adopted for the latter, to reduce the electromagnetic emissions and minimise the interferences on the measurements. Their routing was performed on layers 2 and 4, employing the 1 (Top), 3 and 5 as **GNDD** references. The adopted geometries have been discussed with the manufacturer, to achieve a precise 100 Ω characteristic impedance and they are reported in Tab. 6.26. Great care was taken to prevent other tracks or components from breaking the return path of the transmission line currents. In fact, by comparing the layers 1 to 5 (in Appendix F) it is possible to notice that there are no obstacles placed both above and below the striplines.

The routing procedure prioritised the sensitive nets, as it was pursued during the component placement. Initially, the E-Links and the high-speed lanes were routed, adopting the length matching techniques. Specifically, this was carried out inside the differential pair and amongst the pairs. The first ensures that the mismatch between the two polarities is < 10 *mil* (0.254 *mm*), which is recommended for proper signal integrity. On the other hand, the second facilitates the phase tuning of E-Links and the circuit debug, although it is not strictly required for the correct operations of the system. A maximum length difference of 50 *mil* (1.27 *mm*) was accepted in this case. The resulting traces of the E-Links **B** and the VTRx high speed lanes are visible in Fig. 6.27, while the E-Links **A** can be found in Appendix F. All the data lanes from the BLMASICs to the lpGBTs have been equalised to ≈ 121.9 *mm*, while the clock ones are long ≈ 124.6 *mm*. The high-speed striplines connecting the lpGBT to the VTRx measure 13.2 *mm* and 13.9 *mm* for the data receiver and transmitter, respectively.

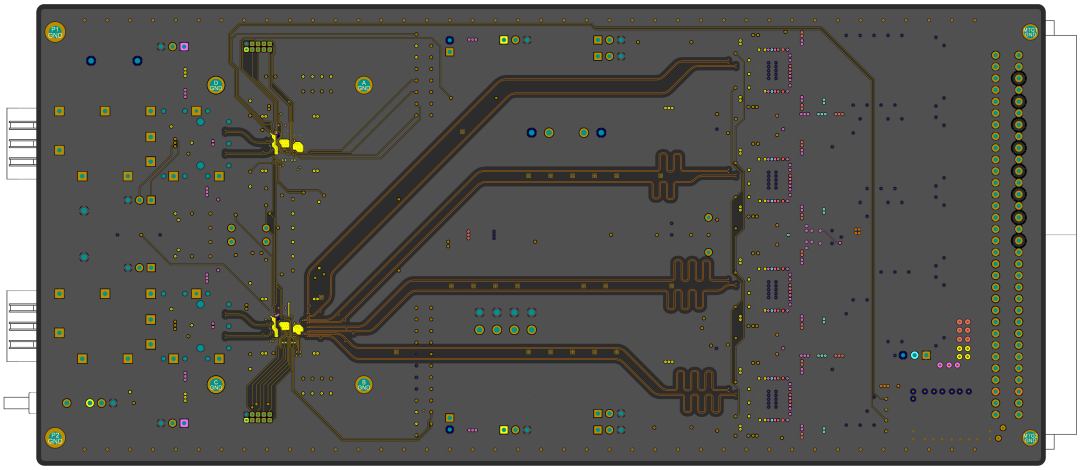

Figure 6.27: E-Links **B** and High-Speed striplines routed on layer 4.

Another challenge faced during the development concerns the fanout of the lpGBT, which is embedded in a square BGA package of 289 pins, spaced 0.50 *mm* each other. This characteristic turned out to be crucial in deciding the adoption of blind and buried vias, without which the routing would be unfeasible. The device supply rails were distributed in the planes below and connected by a combination of vias. Furthermore, it was decided not to coaxially stack the 4 to 7 through-hole vias with the outermost ones, to avoid PCB lamination problems and thus improve the product reliability. The via stitching technique has also been widely employed throughout the card, to allow a ruggedised grounding and improve the power dissipation.

Finally, despite the limited card power consumption, the chassis **GND** has only been dedicated to the dissipation of the heat. In fact, the copper plane surrounding the board edges is exposed and in contact with the containing metal box, facilitating the board cooling. An additional copper plane (see Fig. 6.30) has been located on the bottom layer for the same purpose since an aluminium plate supports the card in the same area. Before generating all the manufacturing files (*GERBER*, *NC Drill*, *Pick and Place*, *BOM*), the design rules were checked and the violations corrected. The finalised card can be seen in Fig. 6.29 and Fig. 6.30.

6.6 Board Start-Up and Configuration

The board devices automatically begin their initialisation procedures as soon as the power is applied. However, the lpGBTs require an additional configuration, to enable communication with the back-end by the optical link. It can be achieved either by direct access to their slave I^2C interfaces, which are always active, or by programming the internal e-Fuses in advance. This latter option is preferred for the commissioned system and is described in detail by the official device documentation (cf. procedure reported in the manual [S13]).

It is mandatory to configure the internal lpGBT registers dedicated to the high-speed link, including the PLL parameters and the locking requirements. When the desired values are correctly programmed, it is possible to assert the bits 1 (*pllConfigDone*) and 2 (*dllConfigDone*) of the register POWERUP2 [0x0fb], which enables the communication procedures. It has been proven that the lpGBTv1 is able to correctly lock the downlink stream even if the whole configuration registers are left in their default status. However, this was not the case for the lpGBTv0, which required tuning the settings. In any case, it is recommended to investigate the best operational setup, before programming a permanent configuration into the e-Fuses.

Finally, the settings of E-Links must be programmed (even by the optical interface), to enable communication between the BLMASICs and the lpGBTs, thus starting the transmission of measurements. This procedure can also be pre-loaded within the e-Fuses, to allow a ready-to-go booting. Contrary to the mandatory high-speed link configuration, the E-Link settings can be modified anytime. The following Tab. 6.28 reports an example of register configurations that were successfully adopted during the board validation. Refer to the official lpGBT documentation [S13] for the details about the function of each specific bit.

Register	Value	Comment
LpGBTv1 A only		
0x8E	**0x1A**	BLMASIC-0 Clock (ECLK16) Enabled at 80 MHz and 2 mA
0x8C	**0x1A**	BLMASIC-1 Clock (ECLK15) Enabled at 80 MHz and 2 mA
0x7E	**0x1A**	BLMASIC-2 Clock (ECLK8) Enabled at 80 MHz and 2 mA
0x7C	**0x1A**	BLMASIC-3 Clock (ECLK7) Enabled at 80 MHz and 2 mA
LpGBTv1 B only		
0x94	**0x1A**	BLMASIC-0 Clock (ECLK19) Enabled at 80 MHz and 2 mA
0x8E	**0x1A**	BLMASIC-1 Clock (ECLK16) Enabled at 80 MHz and 2 mA
0x84	**0x1A**	BLMASIC-2 Clock (ECLK11) Enabled at 80 MHz and 2 mA
0x7E	**0x1A**	BLMASIC-3 Clock (ECLK8) Enabled at 80 MHz and 2 mA
LpGBTv1 A & B		
0xDE	**0xC2**	BLMASIC-0 E-Link (Group 3 Channel 2) Phase and 100 ohm Termination
0xDC	**0xC2**	BLMASIC-1 E-Link (Group 3 Channel 0) Phase and 100 ohm Termination
0xDA	**0xC2**	BLMASIC-2 E-Link (Group 2 Channel 2) Phase and 100 ohm Termination
0xD8	**0xC2**	BLMASIC-3 E-Link (Group 2 Channel 0) Phase and 100 ohm Termination
0xCB	**0x58**	Enable Group 3 Channels 2 and 0 at 320 Mb/s and Fixed Phase
0xCA	**0x58**	Enable Group 2 Channels 2 and 0 at 320 Mb/s and Fixed Phase
0xFB	**0x06**	DLL / PLL Config Done bits

Figure 6.28: Configuration template for lpGBTv1 **A** and **B**, to enable the transmission of the measurements.

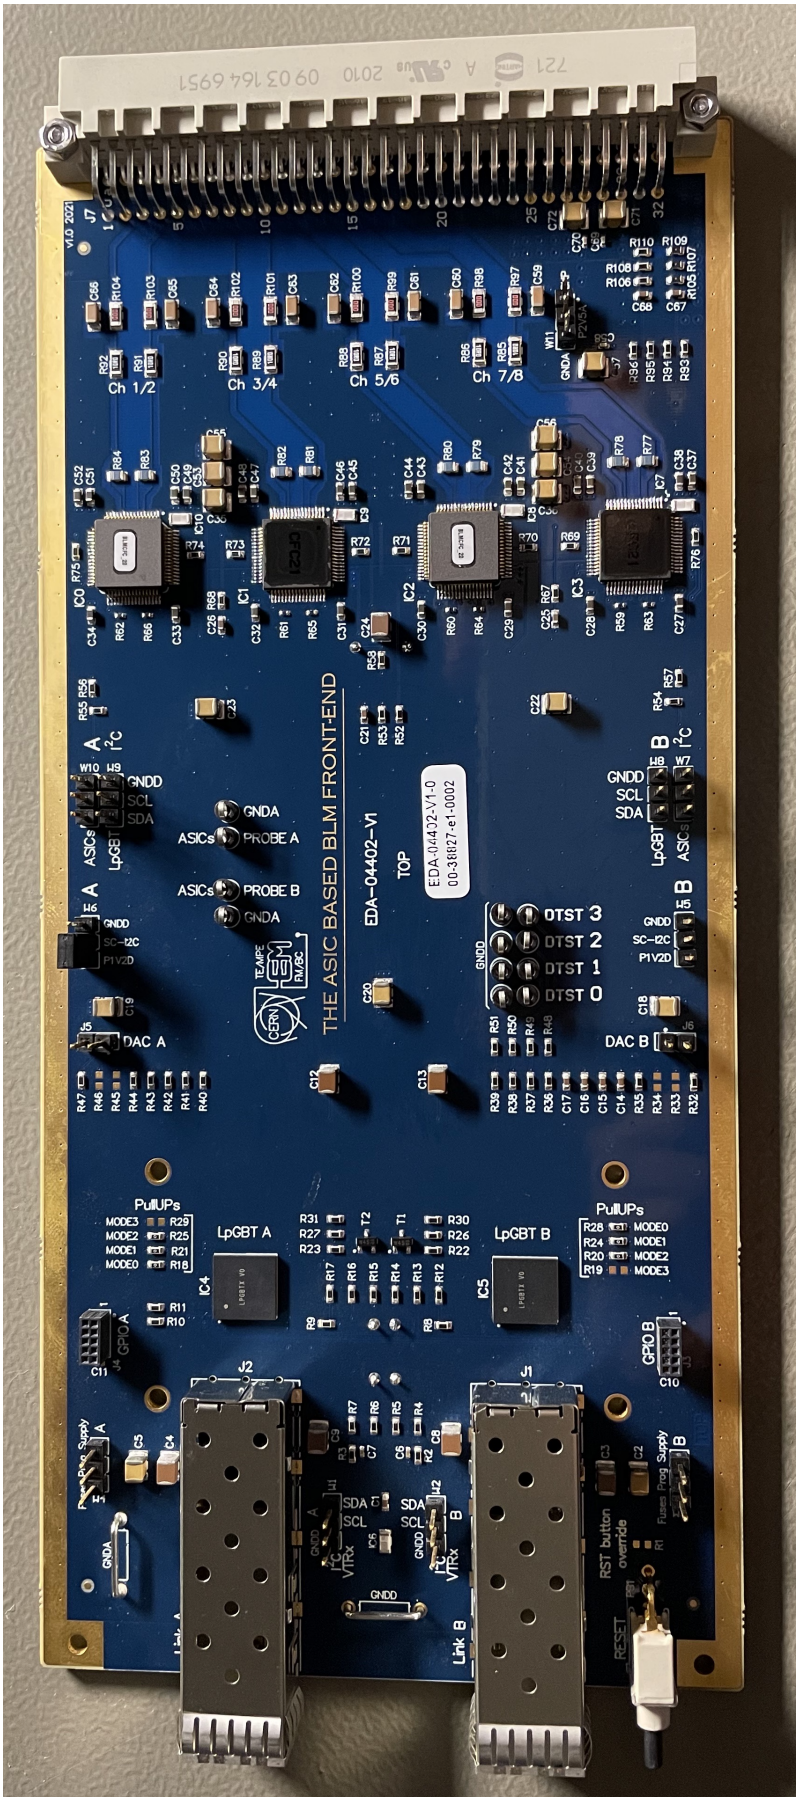

Figure 6.29: Top view of the assembled front-end. Including the BLMASICs v2 and v3, visible on the right.

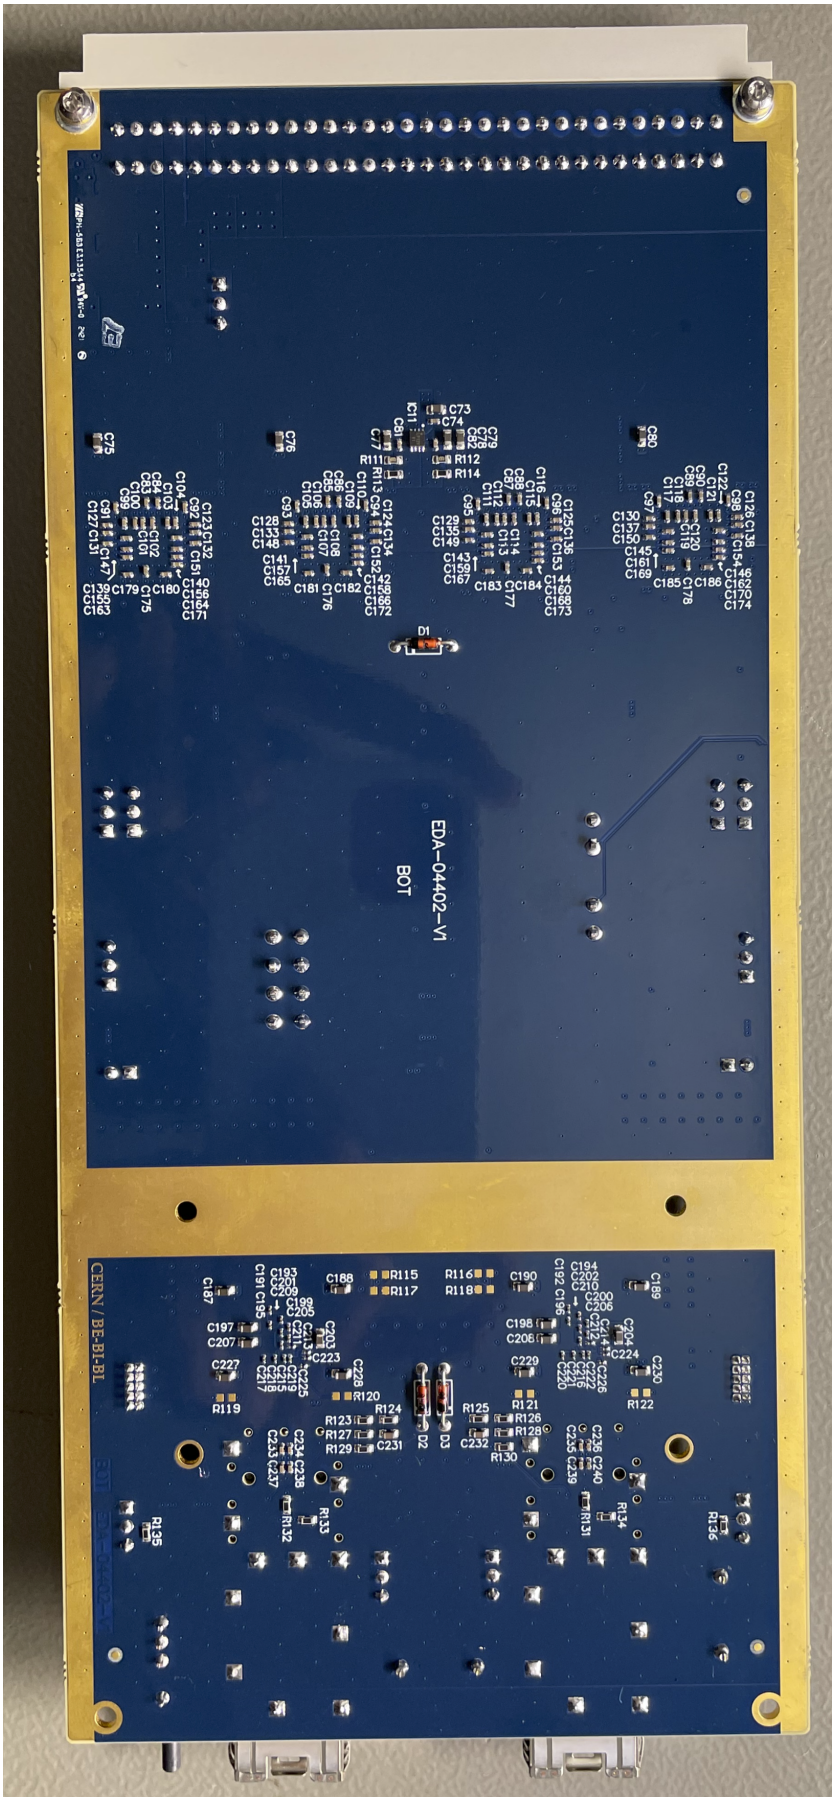

Figure 6.30: Bottom view of the assembled front-end.

Chapter 7

Back-End Development

The BLM readout electronics is covered, focusing on the development of the firmware to be installed within the back-end FPGA. The data acquisition architecture is illustrated and the function of the leading digital modules will be analysed in detail. These include the high-speed transceivers, the lpGBT core and the BLMASIC data retriever. Solutions to the major implementation challenges are proposed, such as the BLMASIC bitstream phase alignment, the redundancy management and the diagnostics. Furthermore, a feasibility study of novel floating-point filters is carried out. This reports promising outcomes that can be extended to applications beyond those discussed here. Finally, the modules developed for the overall system test-bench are illustrated, together with the performance considerations about the FPGA implementation.

7.1 Key Purposes

In agreement with the BLM readout electronic description (cf. Par. 2.3.2) and the specifications established for the new system (Par.s 2.5 and 2.6), the back-end involves the ensemble of hardware and firmware, required to:

1. establish a reliable communication with the front-end boards;
2. acquire the raw data produced by the BLMASIC converters;
3. process the data into beam loss measurements;
4. trigger the beam dump request in case of excessive losses;
5. verify the correct monitoring function and the system sanity.

The board which will carry out such tasks is the VFC-HD [45], which proved to include suitable resources. Thus, the work focused on developing and testing the VHDL codes, which are then programmed within the embedded FPGA device. The whole design has been developed exploiting Quartus Prime Standard [S44], version 20.1.1.

The tasks 1 and 2 were of utmost interest at this prototype stage, since their implementation is entirely novel. The data processing was described in detail in Chapter 5

and the analysis of its implementation feasibility will be displayed in the following Par. 7.8. The measurement thresholding modules will be derived from an adaptation of the previous firmware, the features of which were presented in Par. 2.3.2. Hence, this manuscript will not address such an argument, being part of the works of integration into the existing plant. Finally, an overview of the diagnostics is reported in Par. 7.9.

7.2 Data Readout Architecture Overview

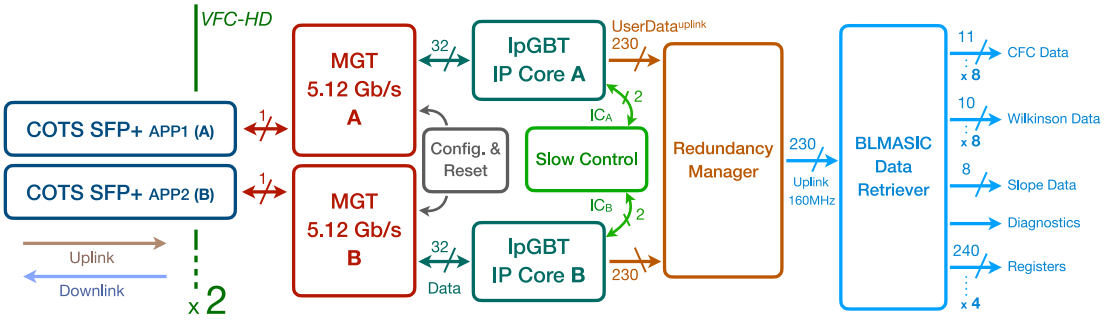

Figure 7.1: Architecture of the data readout chain.

In Fig. 7.1, there are the digital modules and the related connections necessary to manage the communication with the front-end, acquire the transmission streams and extract the data produced by the BLMASICs. The optical bitstream, generated by the front-end VTRx devices, is here digitised by commercial multi-gigabit transceivers. Due to the non-standard and asymmetric bit-rates (Uplink: 5.12 *Gbps* and Downlink: 2.56 *Gbps*), an investigation was carried out together with the VTRx developers, to select the devices suitable for the application. Specifically, 10 *Gbps* SFP+ plus modules have been used, implementing a single-mode optical link (LC connector) that operates at the same wavelengths of the VTRx. The modules are installed within the board general purpose slots (*Application* SFPs) and they are interfaced with the FPGA Current Mode Logic (CML) transceivers through differential lanes. The VFC-HD can host up to 4 SFP modules, thus it can be connected to 2 front-end boards employing the redundancy (cf. the Unit presented in Par. 2.6). In fact, the acquisition chain shown in Fig. 7.1 is actually doubled for the parallel monitoring of 16 BLM detectors.

The shown **MGT** blocks concern the Intellectual Property (IP) cores providing access to the FPGA transceiver primitives. These perform the data serialisation (downlink) or parallelisation (uplink), to convert the fast bitstreams into buses composed of multiple bits and operating at a slower clock frequency. Such a solution is mandatory for handling high throughput data inside the FPGA logic, since the latter is optimised to operate with a high level of parallelisation rather than very fast clocks.

Subsequently, the raw data are exchanged with the lpGBT IP core, which is the firmware counterpart of the hardware device. It includes all the processing logic required for encoding or decoding the stream, performing the FEC and handling the

communication session. The downlink is employed to provide the clock to the front-end and to write into the configuration memory of all the embedded devices (cf. Par. 6.4.5). On the other hand, the configuration readout and the measurement data are received through the uplink. In particular, the latter are propagated to the rest of the logic modules by the *userData* (uplink) bus.

A redundancy manager module is also included in the design to deal with the parallel data streams, thus improving the availability of the system. Several methodologies can be exploited to identify which of the links, between **A** or **B**, should be selected. In this way, a single userData bus will be used to feed the successive BLMASIC data retriever block, allowing a transparent operation in case of link failure. This last module performs the parallelisation of the BLMASIC bitstream, decodes the 8b/10b [70] frame (cf. Par. 4.7) and retrieves the content of its registers. Finally, the same module includes the logic necessary to collect and organise the measurements to be later processed.

7.3 FPGA Transceivers

The interface between the fabric logic and the high-speed SFP modules is represented by the transceivers embedded into the Arria V FPGA. Specific IP cores, supplied by the manufacturer, are used to integrate such interfaces within the design. They also provide effective tools for achieving the desired configuration, like the GUI wizard shown in Fig. 7.2. The block reported displays the nets involved in the transceiver primitive, with the inputs on its left side and the outputs on the right. On the other hand, the menu reflects the configuration employed in this project. It is possible to notice that RX or TX data paths have been enabled, both operating at 5120 $Mbps$. In fact, the actual rate asymmetry is managed by the lpGBT IP core, simplifying the setup even for devices which do not natively support such a solution in their primitives. Precisely, the downlink bits are duplicated just before the transmission, emulating a 2.56 *Gbps* link.

The logic of transceivers is enabled with just the basic features, such as the Physical Medium Attachment (PMA) module (comparable to that used in computer networks [S45]). This is used for the clock recovery (in the uplink) and the data serialisation, or parallelisation, by a factor of 16. Furthermore, the Physical Coding Sublayer (PCS) logic has also been enabled. This is employed to remap the 16 *bit* PMA bus over a 32 *bit* one, thus allowing a data clock of 160 MHz. Such a frequency represents the best compromise for this FPGA generation. The frame alignment and decoding features, usually performed by the PCS, are not enabled here since they are all implemented inside the lpGBT IP. This latter is connected to the transceiver primitive by the 32 *bit* interfaces, the data validity flags, the clocks and the uplink bit-slip control.

Finally, the PLL embedded into the FPGA high-speed interfaces has been driven by a 160 MHz reference clock, which is sourced by an external programmable Si5338N oscillator, produced by Skyworks [S46].

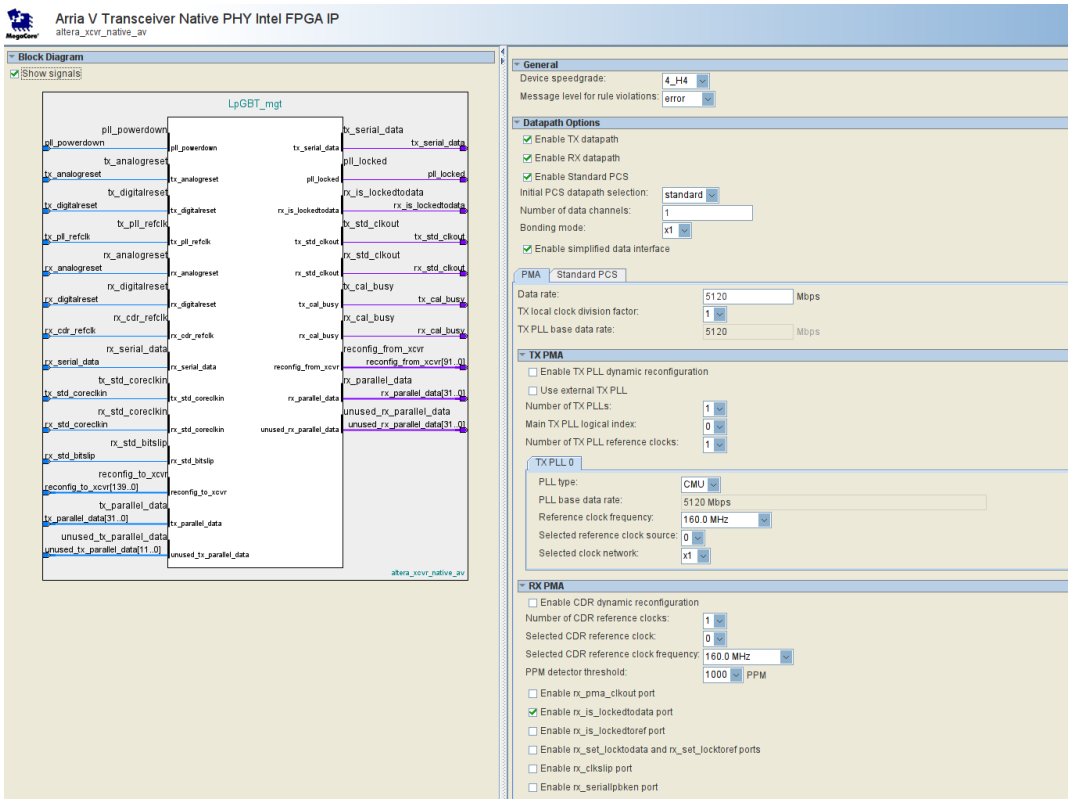

Figure 7.2: Multi-Gigabit Transceiver (**MGT**) configuration for operating with the lpGBT.

7.4 The lpGBT IP Core

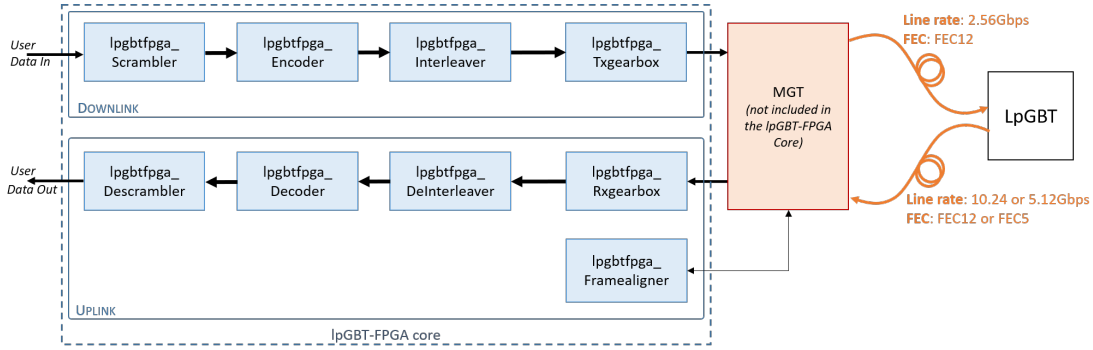

Figure 7.3: Internal architecture of the lpGBT IP core.

This IP core is divided into two modules, handling the uplink and the downlink sessions, respectively. They are provided by the lpGBT development team (EP-ESE Group [S47]), which is in charge of supporting the code repository and documentation [S48]. From the user point of view, the modules require a basic configuration of parameters, which set the bit-rate, the employed FEC, the transceiver bus and clock morphologies. Then, it performs its operations transparently (see Fig. 7.3) and the exposed interfaces can be connected to the rest of the logic. During normal operations, the digitiser measurements will be presented at the userData output bus and the validity signal

uplinkClkOutEn will be used to strobe the updates. Indeed, such a large bus (230 *bit*) includes the data coming from all the E-Links, like the transmissions of BLMASICs (cf. Par 6.4.3). It is clocked at $f_{bus} = 160\ MHz$ and the strobe is asserted every 4 clock cycles, corresponding to a duty cycle $\delta = 25\ \%$.

According to the module implementation, 16 *bit* of the userData bus are assigned to an entire E-Link group. Since the connections of BLMASICs are spread over 2 groups (cf. Tab. 6.3), a total of 32 *bit* is employed for their readout (8 *bit* for each device). Furthermore, due to the acquisition oversampling $OS_{frame} = 4$ (cf. Par. 6.4.3), 2 new bits are actually received at each bus strobe. In fact, the BLMASIC data bit-rate BR_{ADC} is obtained by

$$BR_{ADC} = \frac{\delta \times f_{bus} \times 8\ bit}{OS_{frame}} = 80\ Mbps\ . \tag{7.1}$$

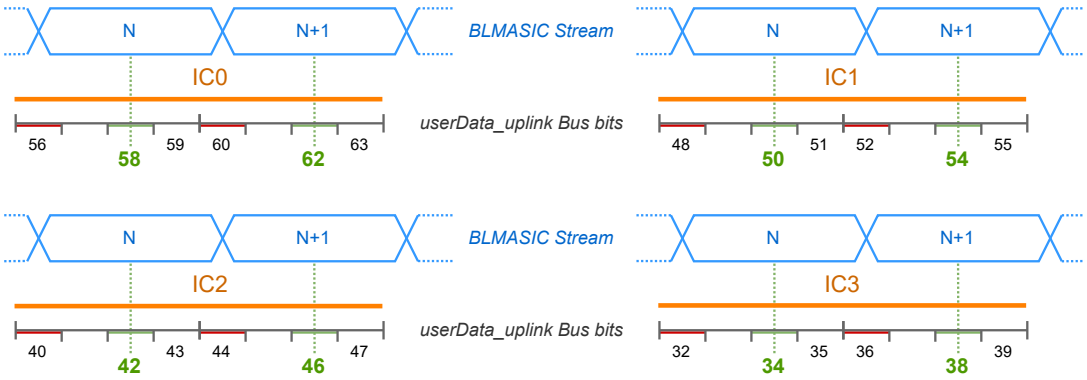

Figure 7.4: Data stream allocations on the *userData_uplink* bus for the 4 BLMASICs.

In Fig. 7.4 are reported the specific userData bit indexes b_{index} that are allocated to each BLMASIC converter, highlighting their designators (IC0 - IC3). Note that such indexes remain fixed in time. For instance, the picture representing the successive bus strobe will result the same as Fig. 7.4, except for $(N, N+1) \Longrightarrow (N+2, N+3)$. Furthermore, it is shown how the bit phase should be aligned to the sampling in optimal conditions. In the example, the bus bits represented in green (like b_{34}, b_{38}) will hold the higher confidence values, since they are likely sampled at the maximum eye opening (cf. Par. 4.7). Instead, those in red (like b_{32}, b_{36}) will be affected by the utmost uncertainty, being located in the proximity of the transitions. A real data example is shown in Fig. 7.5 and its zoomed version Fig. 7.6, which were obtained by using the Signal Tap tool, provided by Quartus. The clock driving the logic analyser is shared with userData.



Figure 7.5: Real userData_uplink bus values, showing the transmissions of BLMASICs.

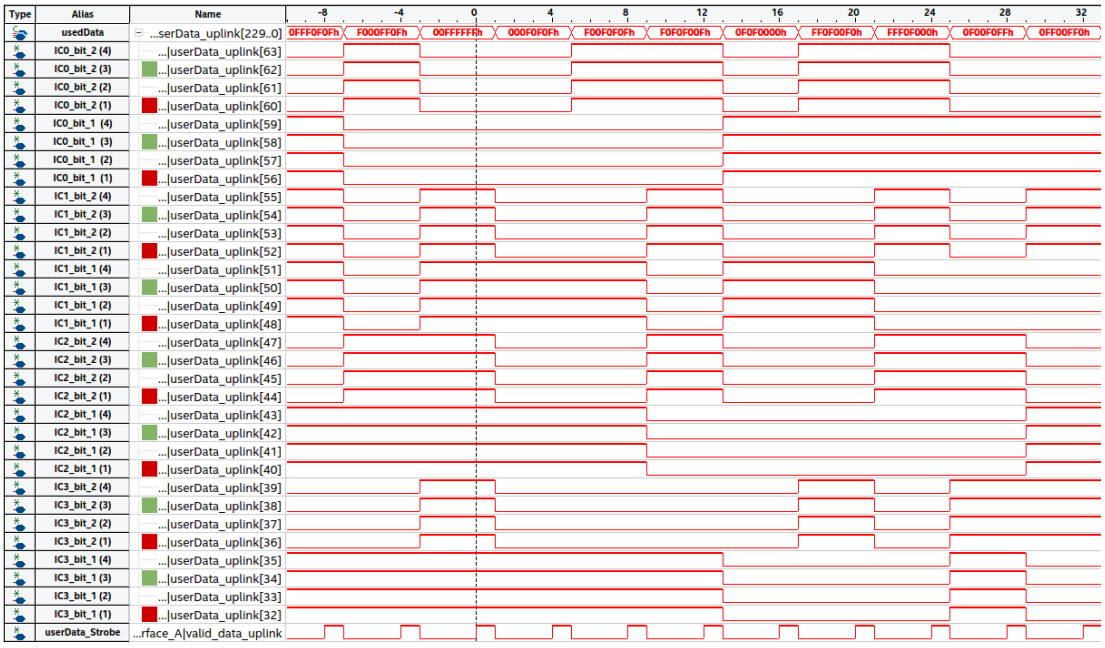

Figure 7.6: Zoomed version of Fig. 7.5.

In both figures, the optimal sampling bits and the uncertainty ones are highlighted with the same colours as before. In Fig. 7.5 it is possible to roughly distinguish the BLMASIC payloads from the periodic synchronisation commas, having patterns reflecting those presented in Par. 4.7. As expected, some sampling discrepancies are visible in the uncertainty bits. In the zoomed picture, it is also possible to notice the bus strobe signal (last row), which is asserted for 1 clock cycle over 4.

7.5 Phase Alignment

By configuring the lpGBT registers, the data sampling phase was fixed while monitoring the userData bus waveforms. Such a procedure aims to realise the conditions shown in Fig. 7.4. However, it is not guaranteed that the programmed values would provide the optimal result on different boards or in long-term operations (see considerations of Par. 6.4.3). Thus, a phase tracking logic should be implemented at this stage.

One possible method can force the data transition to be always located on a fixed index like, for instance, the bit b_{60} of the IC0 data. This can be achieved by monitoring the adjacent bits (b_{59} and b_{61}) values and comparing them with the one sampled in b_{60}, obtaining the flags

$$\begin{cases} b_{early} & = \overline{b_{59} \oplus b_{60}} \\ b_{late} & = \overline{b_{60} \oplus b_{61}} \end{cases} . \tag{7.2}$$

Since the data are DC balanced, it is expected to observe very frequent transitions between b_{59} and b_{61}. If the edges are synchronous to the b_{60} sampling, the average duty cycle of b_{early} should be approximately equal to the b_{late} one (50 %). If b_{late} is asserted more often than b_{early}, the delay of the data must be incremented, otherwise reduced. Such a mechanism is widely employed for the phase detectors included in PLLs, (cf. Alexander Phase Detector [93]). The data delay circuit is integrated within the lpGBT E-Link logic and can be configured by remote, using the optical interface. This solution can realise an efficient phase locking system, requiring a relatively simple logic and providing fixed indexes for the optimal sampling bits. However, it can be affected by a relevant latency since the phase alignment control involves both the front-end and the back-end sides. Furthermore, the method's effectiveness can be severely compromised by the limited data delay range of the lpGBT. In fact, a test on the real system revealed that the maximum delay may not cover the entire bit period (12.5 ns).

Because of these considerations, a second option has also been studied to implement a stand-alone phase tracking module, not involving the front-end. It is guaranteed that the BLMASIC data are driven by a clock synchronous to the one used for their sampling. Thus, no relevant phase drifts are expected during the operations.



Figure 7.7: Data sampling without phase correlation in the best and the worst conditions.

Under this assumption, it is possible to develop a logic to detect where the transitions occur, regardless of their position within the observation windows (cf. again Fig. 7.4).

The implementation could exploit the oversampling to locate the data belonging to the same bit eye. Since the phase is generally unknown at system start-up, different scenarios can occur. As reported in Fig. 7.7, the best condition involves 2 groups of 3 samples each sharing the same value. In this case, the optimal sampling index is certain and it is possible to read correctly the values N (b_{58}) and $N+1$ (b_{62}). On the contrary, the worst condition makes identifying the optimal sample more difficult. But, the logic can be programmed to recognise such a situation and it is still guaranteed that two adjacent bits are equal, at least. They must reside in the middle of the N word. Thus, it is possible to sample N (i.e. in b_{59}) and exploit b_{56} or the b_{63} to acquire $N-1$ or $N+1$, respectively. It is not important which of those bits is actually chosen, because the whole data stream will be read anyway. Nevertheless, it is necessary to prioritise one of the choices, to avoid the selection logic jumping from one to the other.

Finally, the actual eye diagram of the streams must be taken into account to validate the assumptions made in this study. The results of the BLMASIC characterisation (cf. Fig. 4.11) reported $T_{rise} \approx T_{fall} \approx 850\ ps$, which are about 3.7 times shorter than the oversampling period. Thus, a properly scaled waveform would look as in Fig. 7.8, showing narrower transition regions than what is represented in previous figures.

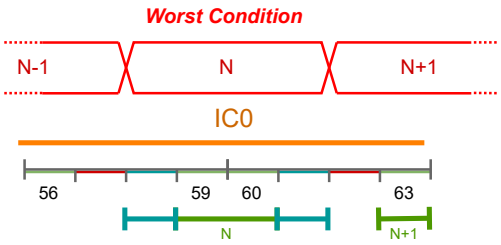

Figure 7.8: Realistic eye sampling in the worst condition.

Here, it is much more common to identify a cluster of three identical samples amongst b_{58} to b_{60} or b_{59} to b_{61}. Thus, if the system can identify it, the optimal sampling locations can be easily chosen. Clearly, it cannot be assumed that data would change at every eye, like in the figures. Hence, the logic must detect the different conditions, otherwise some may lead to an inability to resolve the sampling phase. For instance, if the system detects no transitions, it should hold the previous phase setting.

This last solution requires a more complex logic and it needs to be extensively validated. However, it has the significant advantage of being front-end independent and can potentially eliminate the locking latencies. Obviously, the lpGBT data delay feature must be disabled in this case.

7.6 Redundancy

The redundancy management logic is necessary to handle the doubled data stream received by the back-end (cf. again the Fig. 7.1). This enables the system to perform its protection role, even in case of a single link disconnection. Such an event, for instance, may be due to a faulty component (lpGBT, VTRx, SFP transceivers), a problem with the optical connections or the sporadic case of an unlock of PLLs. Anyhow, the module must undertake the correct measures to ensure an uninterrupted flow of data towards the processing chain. Its development is out of the scope of this work, since many solutions have already been validated and employed in the BLETC card (cf. Par. 2.3.2) for the currently operational system. However, it has been ensured that the new architecture could potentially support them. In addition, the new BLMASIC characteristics and the FEC features provided by the lpGBT pave the way for further possibilities.

The legacy module implements a logic that stores a certain amount of data in the memory, then it also performs a comparison between the two channels, before proceeding with the propagation. In fact, the optical transmission and the successive readout modules introduce a slight delay between the streams. Therefore, it is necessary to integrate some sort of FIFO logic before selecting one of these by a multiplexer. Otherwise, the output data samples may not be contiguous. In addition, the new system provides detailed information about how many bits, and which ones, have been corrected by the FEC (*dataCorrected* bus) in each link. Hence, it can be possible to adopt predictive techniques, to withdraw the channel that is likely to fail and select the backup one.

Finally, a different approach could involve relocating the redundancy manager after the BLMASIC data retriever block, thus implementing a higher-level redundancy. This requires duplicating several modules, but it may dramatically simplify the logic. In fact, the delay between the fast data streams (userData buses) is usually much smaller than the current measurement period, leading to a single sample skid on the results. Furthermore, it is not necessary to manage the contiguity between the fast buses, as required in the previous solution, because they would feed different measurement retriever modules.

7.7 BLMASIC Data Retriever

This module includes the logic necessary to retrieve the content of the frames from the raw userData bus (presented in Par. 7.4), which were generated by the BLMASICs. Furthermore, it establishes a new level of abstraction in the FPGA firmware, enabling simplified access to the status of converters and measurements. The following description is referred to one out of four readout logics (one for each BLMASIC) since they are all alike. It involves two main sub-modules, which are represented in Fig. 7.9, including their ports.

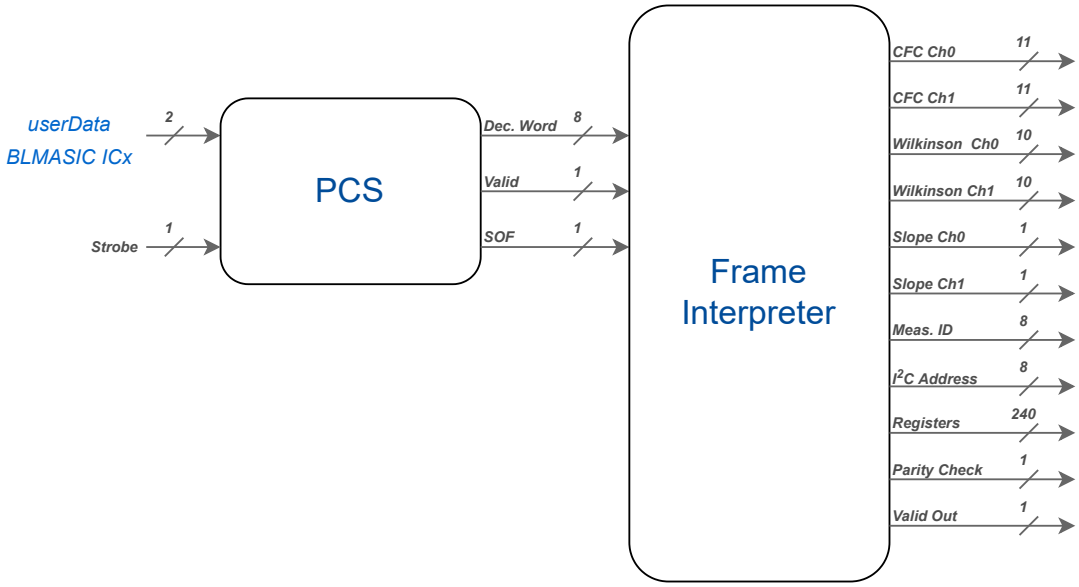

Figure 7.9: Simplified architecture of the data retriever module.

7.7.1 Physical Coding Sublayer

For each converter, the userData bus (cf. Fig. 7.4) yields a pair of new bits at each strobe. The **PCS** logic (in Fig. 7.10) provides the parallelisation of the resulting stream, identifies the frame morphology and decodes the 8b/10b (cf. Par.s 3.3.8 and 4.7).

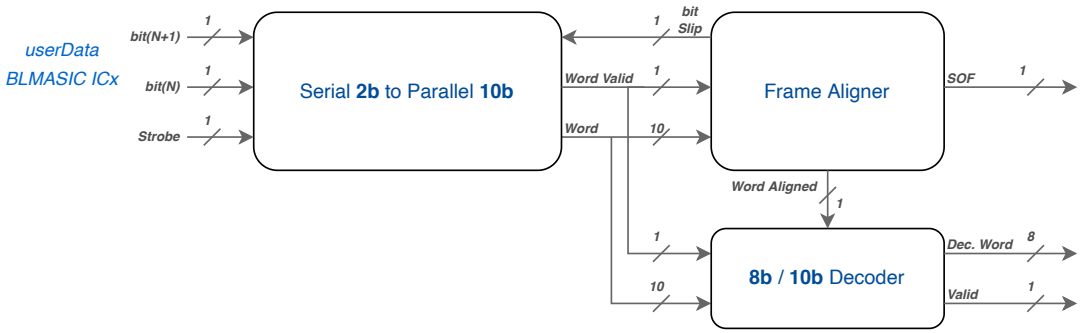

Figure 7.10: Physical Coding Sublayer included in the BLMASIC data retriever.

Data Paralleliser

The first task involves converting the 2 *bit* input bus into a 10 *bit* one, which can contain a whole 8b/10b encoded symbol. Furthermore, its output bits must be correctly aligned to the latter, so the decoder can interpret it. The paralleliser is based on a 2 *bit* shift register, with a depth of 6, shown in Fig. 7.11. The shifter inputs are directly fed by the userData bits and the registers share the clock and strobe control with the bus. Hence, the data are moved from one register to the next one, whenever a new update is presented.

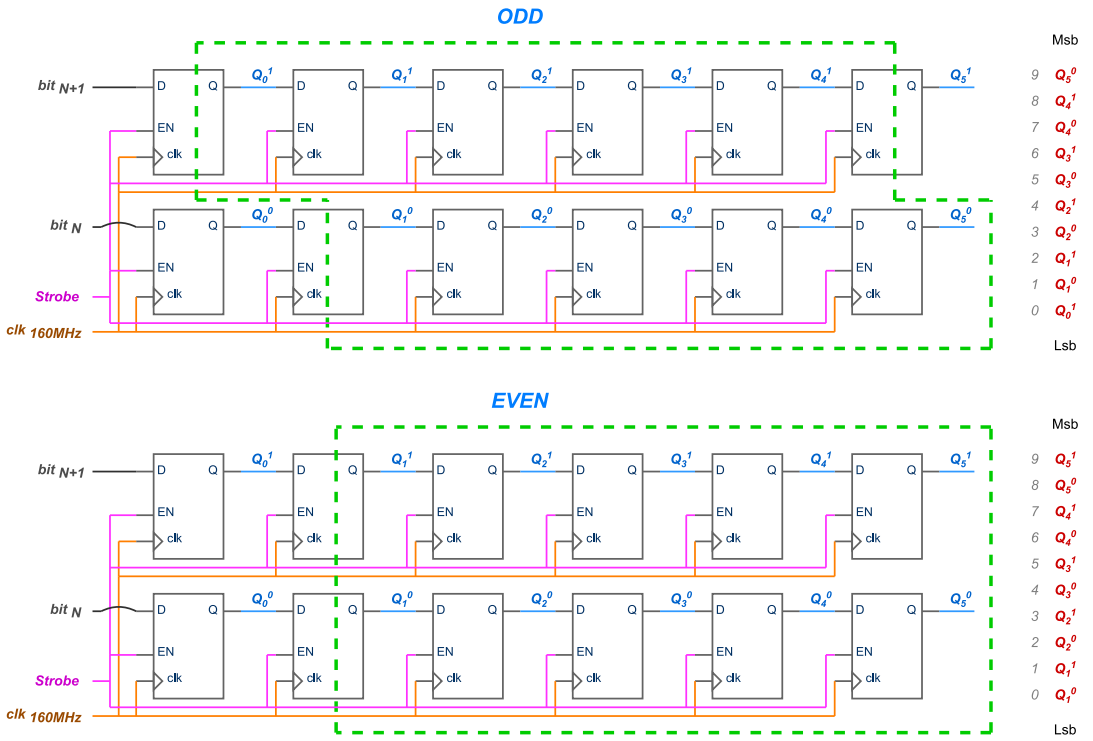

Figure 7.11: Shift register employed in the PCS data paralleliser. The two possible selections of outputs (**ODD**, **EVEN**) are reported.

The word alignment is achieved by bit-slip control circuitry, not shown in the figure. Its behaviour can be easily understood by imagining the system taking periodic 10 *bit* snapshots of the data flowing. In particular, the snapshot is represented by the outputs of registers and is propagated to the successive modules with a period 5 times larger than the bus strobe one. From the schematics of Fig. 7.11, it can be noticed that the outputs are actually updated at the strobe rate, simplifying the logic. In fact, the propagation of results is commanded by an additional logic, which asserts the *word valid* signal once every 5 strobes. Furthermore, a phase control mechanism for the snapshot bits is required to enable the correct alignment of symbols. By only modifying the generation of *word valid* pulses, achieving a 2 *bit* resolution step in phase is possible. However, to further reduce it to a single bit, the multiplexing of the outputs between two possible configurations is required (see Fig. 7.11).

The output phase value is stored into a 4 *bit* register, which can be incremented by acting on the bit-slip control signal. The 3 most significant bits are compared with a counter running synchronously to the data, to generate *word valid* pulses with the desired phase. Meanwhile, the least significant bit drives the multiplexing of the outputs. By exploiting such a solution, it is possible to shift the 10 *bit* output codes with a resolution of a single bit, thus correctly aligning the 8b/10b symbols.

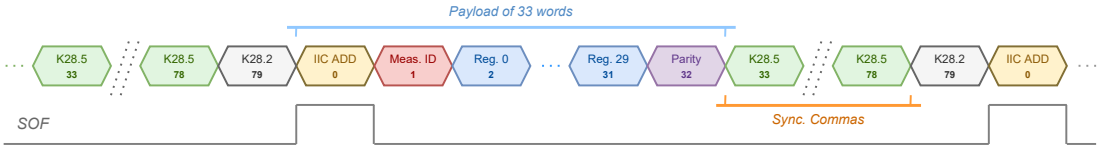

Figure 7.12: BLMASIC frame structure and the generated SOF signal.

Frame Aligner and Decoder

The 10 *bit* words, sourced by the paralleliser, feed both the frame aligner and the decoder modules as shown in Fig. 7.10. The former is employed to detect whether 8b/10b symbols are recognised and it shifts the paralleliser phase otherwise. The methodology of identification exploits the property stating that 5 consecutive 1s or 0s in the bitstream can only occur within the *K28.5* synchronisation commas, which are **0b1100000101** or the bitwise negated **0b0011111010**. Thus, a finite state machine (FSM) has been developed to identify such a pattern, during its *SEARCH* state. As soon as the comma is recognised, the FSM enters in the *LOCK* state and asserts the *word aligned* flag. Otherwise, if the pattern is not identified for the entire duration of the frame (80 words), the FSM asserts the bit-slip for a clock cycle (*BITSLIP* state) and it returns back to *SEARCH*. Additional control is performed during *LOCK*, to keep checking that a minimum number of *K28.5* commas are still detected within the frame period. If this is not the case, for instance due to the front-end disconnection, the FSM is reset to the *SEARCH* state, *word aligned* is de-asserted and the whole procedure repeated. Finally, the same module is designed to locate the sequence of control symbols {*K28.5*, *K28.2*}, which is followed by the beginning of the data payload (cf. Par. 3.3.8). This allows asserting the Start of Frame flag (**SOF** in Fig. 7.12), which is used subsequently to synchronise the content acquisition.

These operations are carried out in parallel with the 8b/10b decoding, performed by the block shown at the bottom of Fig. 7.10. New decoded bytes are propagated by asserting the PCS output validity, whenever the paralleliser produces a new word, and the described logics guarantee the correct alignment.

Such processes have been validated using the Signal Tap analyser, the waveforms of which are reported in Fig. 7.13. Before the trigger event (sample 0), a reset of the BLMASIC IC0 has been forced to visualise the word and the frame locking mechanism. It can be observed how the frame aligner module drives the bit-slip control, to modify the 10 *bit* words phases. In fact, both the *K28.5* and *K28.2* symbols have been recognised before the lock, but not synchronously with the *word validity* signal. Thus, the waveform shows that the phase is shifted by 3 *bit*, before the locking condition is achieved. After that, the aligner FSM goes into the *LOCK* state, the word aligned flag is asserted and the SOFs are generated. Such conditions enable the 8b/10b decoder, thus it is possible to distinguish the decoded *K28.5* comma value (*0xBC*) in the resulting data (*8 bit Word* row in Fig. 7.13).

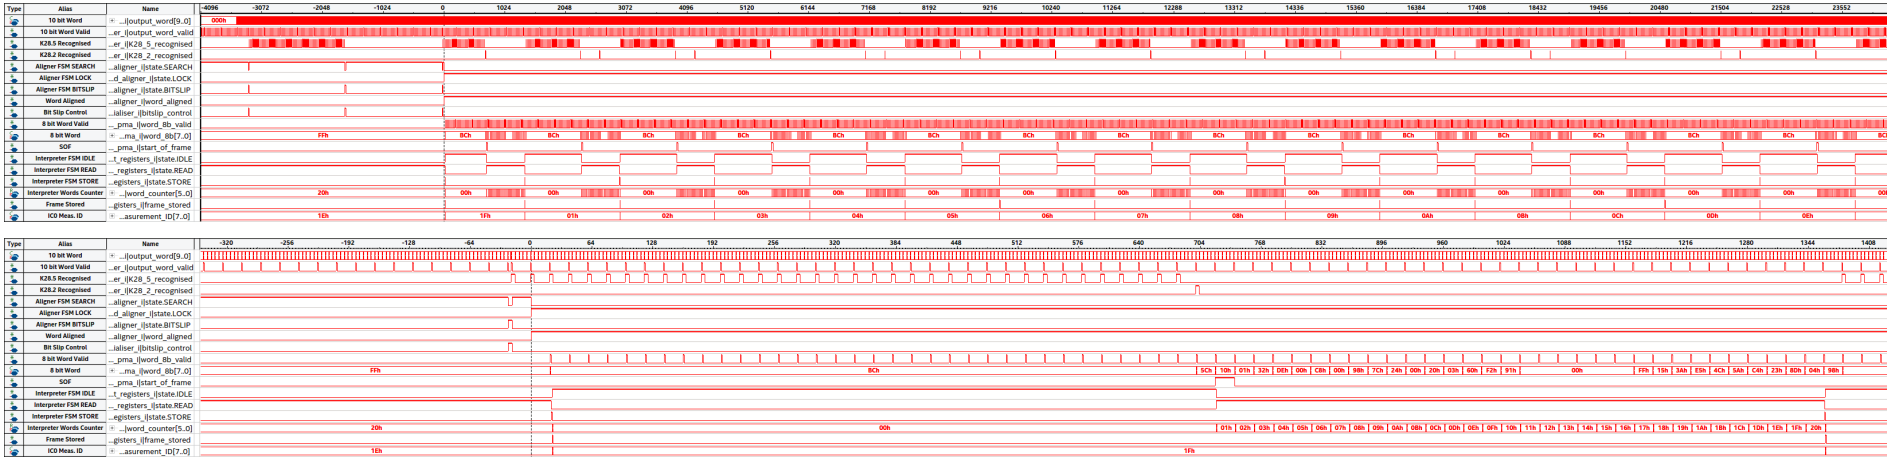

Figure 7.13: BLMASIC data retriever module operations, including the frame alignment and the measurement ID capture. The bottom figure is the zoomed version of the top one, close to the trigger instant.

Finally, it can be noticed that the *K28.5* symbols are only recognised outside the payloads, while 10 *bit* words that are coincidentally equivalent to the symbol *K28.2* are sometimes found in the data. This agrees with the encoding standard [70], since the uniqueness of bit representation is only guaranteed for the comma symbols (like *K28.5*). Such a statement emphasises the importance of detecting the sequence {*K28.5*, *K28.2*} before asserting the SOF, instead of just *K28.2*.

7.7.2 Frame Interpreter

Once the PCS provides the decoded stream of BLMASIC bytes, the payload information is easily extracted. Such a goal is performed by a dedicated FSM, paired with an array of registers, which is used to store incoming data. In particular, the FSM waits for the SOF signal to be asserted during the *IDLE* state, then it begins incrementing the array pointer as the new bytes are received, in order to fill the storage (*READ* state). When all the payload bytes are written, the FSM reaches its last state *STORE* and the entire frame is copied into a shadow array of the same size. This solution has been introduced to simplify the subsequent data processing, which retrieves the data from the shadow copy, remaining untouched during the *READ* state. After the *STORE*, a specific flag is also set for a clock cycle, to trigger the data processing and to provide the strobe signal for the module outputs. The frame acquisition mechanism can be observed at the bottom of Fig. 7.13. The interpreter FSM begins operations after the PCS properly aligns the frame, so it starts to alternate *IDLE* and *READ* states. The word counter value is also reported, incrementing up to 32, before the frame stored flag is asserted. The 33th word is the parity check byte, which is not stored. Instead, it is directly verified during the *READ* state. The incrementing measurement ID byte is shown in the last row of the figure. As with all the propagated results, this latter is delayed by one frame, explaining why the unexpected value *0x1F* is displayed before the new data arrival. Indeed, it belonged to the acquisition before the reset.

Some received values, like the I^2C interface address, the measurement ID, the auxiliary Wilkinson ADC data and the configuration registers, are directly passed to the module outputs, for onward use. The parity check is computed during the *READ* state, in parallel with the data storing, adopting the methodology described in Par. 3.3.8. At the same time, the main CFC measurements are pre-treated before their propagation. In particular, the 10 *bit* registers containing the counters of the threshold crossing (positive and negative) are reassembled from the stored bytes. Thus, running differences are iteratively computed between the new and the previous counts, to obtain the number of threshold breaches that occurred during the last acquisition period. The results of the positive and the negative differences are finally added together, to get the total integrator waveform inversions (cf. Par. 3.2). It can be proved that the overflow of the BLMASIC counters is transparently managed by the binary differences in two's complement, if the output is also truncated to 10 *bit*. Obviously, the total addition

requires an extra bit, so the CFC measurement output is represented on 11 *bit*.

The integrator slope direction is also a valuable information for the combination of CFC and auxiliary ADC measurements (cf. Par. 5.2). It can be read by directly querying the dedicated converter register (cf. Appendix C) or derived from the data of CFC counters. This last solution involves the comparison of the differences between the positive and the negative crossings of thresholds. If the number of positive ones is higher than the negative, then the integrator output is currently falling, otherwise it is rising. Instead, if the values are equal, the slope calculated at the previous iteration is held. After many tests performed on the real system, it has been decided to delay the propagation of the CFC total counts and the slope information outputs by one frame (10 μs). This is necessary for a correct synchronisation to the auxiliary ADC data, thus enabling the combination of the measurement sources according to the technique described in Par. 5.2.

Finally, all the described modules have been quadruplicated to operate independently on the data produced by the four BLMASICs. Since each frame interpreter propagates its results by asserting a dedicated *Valid Out* signal (visible in Fig. 7.9), these are combined in an additional logic, to generate a unique validity flag. In particular, the latter is asserted for a single clock cycle after updating the results of all the BLMASICs.

7.8 Processing Implementation

The complete back-end firmware must involve the measurement processing and the successive thresholding required for the accelerator protection. An investigation was carried out during the development, to identify which of these tasks can be improved by an upgrade or retained from the previous system. In particular, it is convenient that the thresholding logics (cf. Par. 2.3.2) would be migrated to the new FPGA, exploiting the features already well integrated into the BLM plant. On the other hand, data processing can benefit from the advantages of novel methodologies, which provide higher precision measurements and faster responses. Their theoretical bases have been extensively discussed in Chapter 5, and the computer implementation allowed a quicker analysis of their performance in view of the future FPGA implementation.

The research activity was focused on the feasibility study of the IIR filter implementation (cf. Par. 5.5.2). In particular, it has been highlighted that filters characterised by very low cutting frequencies require a high accuracy tuning of parameters. Furthermore, their stability can be influenced by the numeric quantisation and it must be validated case by case. A floating point (FP) processing enables us to overcome these difficulties, but its implementation could be challenging.

The Arria V device does not include dedicated DSP blocks for the FP computation, but the manufacturer provides all the IP cores able to fulfil this need, in compliance with the IEEE-754 standard [94]. In particular, the *ALTFP_ADD_SUB* and *ALTFP_MULT*

modules were examined for the computation of the additions and multiplications. The main issue concerning the employment of such FP cores concerns the influence of their inherent latency on the behaviour of the filter. In fact, the delays set into its structure (cf. Direct Form II in Fig. 7.14) are strictly constrained by the target frequency response characteristics. The additional computation latencies were not considered in the theoretical study (cf. Par. 5.5.2) since they were all assumed null. However, their effect cannot be disregarded at this stage. For instance, a first-order IIR is described by the difference equation

$$y_k = b_0 \; x_k - a_1 \; z^{-1} \; y_k \; . \tag{7.3}$$

However, assuming that the product operations add latency of D clock cycles, it becomes

$$y_k = b_0 \; z^{-D} \; x_k - a_1 \; z^{-(D+1)} \; y_k \; , \tag{7.4}$$

which describes a filter of order $(D+1)$, having different characteristics respect to (7.3), because of the supplementary poles. The configuration of IP cores allows setting the desired latencies, which may even be removed in some cases, making the logic completely combinatorial. In this way, the value of D vanishes, thus the filter implementation will exactly correspond to the structure of Fig. 7.14. However, such a solution could dramatically limit the maximum operation frequency since the FPGA timing constraints become more difficult to be accomplished. In particular, if a high data throughput is required, advanced pipelining techniques [95], [96] must be adopted to compensate for the effect caused by the additional poles. They involve the introduction of properly tuned zeros into the transfer function or the interleaving of channels, if required.

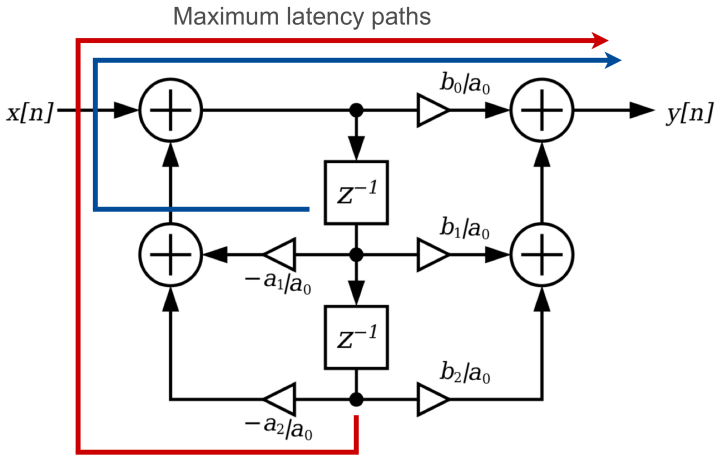

Figure 7.14: Maximum data latency paths for a second order IIR filter (Direct Form II).

Nevertheless, because of the BLM specifications, such issues are significantly facilitated, since the sampling frequency of measurements is much lower than the computing clock one. In particular, each input channel is sampled at 100 kHz, while the raw data processing will be carried out at 100 MHz. This means that up to 1000 clock cycles can be exploited to produce the filtered output sample.

To analyse the maximum latency affecting the filter implementation, it is possible to refer to Fig. 7.14, the structure already presented in Par. 5.5.2. Furthermore, the factor $1/a_0$ has been included in the other coefficients, to reduce the number of required multipliers by one. As stated above, the addition and multiplication operations are implemented using the FP cores. On the other hand, the central column delays (z^{-1}) are realised by shift registers, fed at the same rate as the input. Their depth must be set accordingly to the desired FP representation, for example 64 *bit*, if the double precision is employed.

As soon as the new input x_k is introduced, the computation is started and carried on until all the internal signals are propagated throughout the required operations and the output y_k is produced. At this stage, the structure nodes are stable and the logics are disabled, pending the new input sample. The longest paths are highlighted with colours in Fig. 7.14 and both of them involve 3 additions and 2 products. The typical operation latencies for the employed FPGA are reported in the manufacturer FP cores manual [97]. In particular, the double precision adder requires 12 clock cycles and the multiplier 7 cycles. Furthermore, their maximum operating frequencies are 251 MHz and 186 MHz, respectively, which are well inside the targets. Thus, following the above considerations, it is possible to obtain the maximum structure latency, corresponding to 50 clock cycles. Since all the operations are performed concurrently, such a value also corresponds to the total number of cycles required for the overall computation. The obtained results have demonstrated that the filter logic is employed for only the 5 % of the available time (50 over 1000 cycles), even if the computation is carried out in double precision. This represents a remarkable achievement, considering that the floating point provides a more effective processing and better precisions with the high dynamic range BLM measurements.

If the overall data readout is considered, interleaving the total 16 channels into the same logic significantly reduces the employed FPGA resources. In fact, this would increase the computation duty cycle up to 80 %. In doing this, the number of FP cores will remain the same, while the registers must be replicated for each channel. Furthermore, multiplexers can be adopted to switch the different register banks, the input and output ports. The number of allocated resources has been estimated using the official documentation data [97] and the synthesis process results. In particular, the required FPGA primitives are reported in Tab. 7.1, including the percentage use of the target device.

Finally, it must be pointed out that the complete system will require the instantiation of 12 filters (cf. Par. 2.3.2), tuned with the parameters presented in Tab. 5.3. Thus, the number of resources is multiplied accordingly, obtaining: ALMs 69.6 %, DSP Blocks 26.4 % and Registers 25.2 %. These results are still promising, but further optimisation may be necessary to preserve more device area for the other tasks.

Function	ALMs	DSP Blocks	Registers
Adders	3544	0	4500
Multipliers	1695	20	2810
Multiplexers	1280	0	0
Data Buffers	0	0	2048
Total	6519	20	9358
Utilisation	5.8 %	2.2 %	2.1 %

Table 7.1: Arria V (model **5AGXMB1G4F40C4N** [S49]) deployment of resources for the 16 channels IIR filter implementation. Double precision floating-point is used.

7.9 Diagnostics

The back-end must provide different levels of diagnostics to validate the monitoring operations, alerting the staff about anomalies and providing the tools to investigate their causes. Initially, this is achieved by querying the extensive features included in the front-end boards to which it is connected (cf. Par. 6.4.8). The commands and the values can be exchanged by using an auxiliary slow control IP core (named *GBT-SC*), which is included in the lpGBT related repository [S48]. This is directly connected to the session management modules, described in Par. 7.4, by using a 2 *bit* bus (*ICData*). The architecture provides for the diagnosis to be either launched automatically, to periodically monitor certain parameters, or requested by the operator through the VME bus (cf. Par.s 2.3.2 and 2.5).

By exploiting the described logics, the back-end can access the internal registers of lpGBTs and, consequently, all the other front-end components. These include the sensors, the BLMASICs and the VTRx transceivers (cf. Par. 6.4.8). Some critical values, such as the detector biasing voltage, will be regularly monitored by iterated requests, because their role is mandatory for machine protection and the beam dump must be triggered in case of fault. Although the front-end architecture allows direct reading of the critical quantities by both the lpGBTs, a cooperation with the redundancy manager (cf. Par.s 7.2 and 7.6), can speed up queries for auxiliary signals as well. Indeed, the triggering of cross I^2C requests between the two lpGBTs may be necessary to allow complete board monitoring in case of a single optical link failure (cf. Par. 6.4.5).

Additional diagnostic features are also provided by exploiting the mainstream transmissions. In fact, the continuous reception of the configuration registers of BLMASICs allowed for validating the devices in the radiation environment. In particular, custom modules developed to carry out SEU statistics were widely employed during the irradiation campaign performed at the Paul Scherrer Institut (PSI) (cf. the results in Par. 8.10). They periodically check the content of converter registers, reporting any corruptions. If even a single bit is modified, the counter is incremented and the diagnostic module will adopt the new content as the model for the successive comparison. In view

of the irradiation sessions, additional running counters were also implemented to log: parity and 8b/10b decoding errors, continuity of samples and number of corrections applied by the FEC on the optical link. Specific software has also been developed to visualise in real time and store the diagnostic reports. The example shown in Fig. 7.15 highlights three recognised issues, involving errors in the BLMASIC transmission and the corruption of registers. The first one (Meas. ID 2) was triggered by a forced reset of all the converter devices during the acquisition. Instead, the others (Meas. ID 5 and 7) were caused by external I^2C register configurations of IC0 and IC3, respectively.

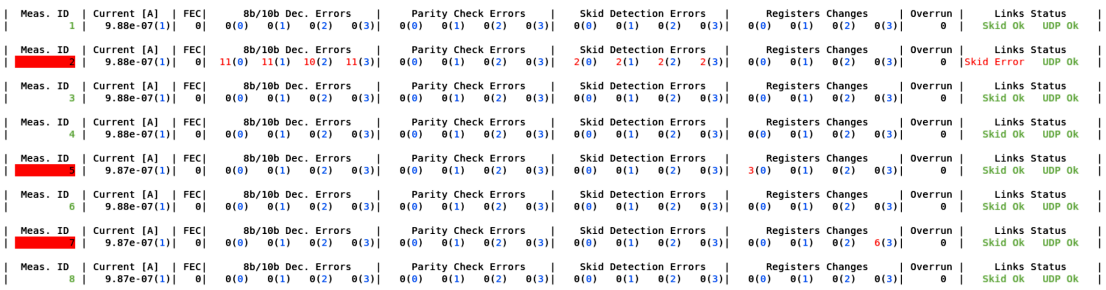

Figure 7.15: Screenshot of the diagnostic software terminal.

7.10 Test-Bench Data Transmission

In order to evaluate different processing chains by the computer-based prototype (cf. Par. 5.8), a module for real-time data transmission has been developed. A study was conducted to determine the interface that would provide sufficient bandwidth, an effective way of uploading measurements to the computer and be economical in terms of development time. Thus, the choice focused on the 1000Base-T Gigabit Ethernet interface (IEEE 802.3ab [S45]), which was exploited to realise an ad-hoc network involving the back-end board and the test-bench management computer. Due to the specific use case, it was not necessary to implement the numerous protocols involved in standard local-area networks, but only the bare essentials. In particular, a custom FPGA module has been designed to generate User Datagram Protocol (UDP) [98] packets. The UDP module is based on an FSM that manages the data acquisition from a cross-domain clocking FIFO and propagates the produced packets toward the Media Access Controller (MAC) [S45], provided by the manufacturer. The FIFO is required to transfer the BLMASIC measurements and the diagnostic information from the 160 MHz to the 125 MHz clock domain, which is the standard frequency for the Ethernet logic. The PCS/PMA modules, included in the Quartus IP catalogue, were configured to meet the physical link requirements.

The produced UDP packets need to be encapsulated in the lower level Internet and Ethernet protocols (EtherNet/IP) [S45], which include the details about the packet morphology, the sender and the recipient (both IPv4 and MAC addresses). These parameters will remain fixed in the assembled test-bench, so the same values are transmitted at every iteration. The same happens for the IPv4 header Cyclic Redundancy Check (CRC) value, which is calculated using the Wireshark tools [S50] and hard-coded into the firmware. On the other hand, the UDP payload includes 16 acquisition samples of 64 *byte* each, composed of measurement and diagnostic data. Each of them is organised according to the following C++ structure, adopted in the processing software presented in Par. 5.8.

```
struct BLMASIC_data_struct{

    // -------------------------------
    // Measurements Data
    // -------------------------------

    // 8 x channels CFC data
    uint16_t CFC[8];

    // 8 x channels Wilkinson data
    uint16_t Wilkinson[8];

    // -------------------------------
    // Diagnostic running counters
    // for each BLMASIC device (IC0 to IC3)
    // -------------------------------

    // 8b/10b decoding errors
    uint8_t decoding_errors_counter[4];

    // abnormal configuration register changes
    uint8_t registers_changes_counter[4];

    // wrong BLMASIC frame parity
    uint8_t parity_errors_counter[4];

    // non-contiguous packet reception
    uint8_t skid_detection_counter[4];

    // UDP transmitted samples
    uint32_t tx_sample_counter;

    // FEC12 corrected frames
    uint16_t FEC_counter;

    // UDP transmission overruns
    uint8_t overrun_counter;

    // -------------------------------
    // Payload Padding
    // (all bits set to 1)
    // -------------------------------
    uint8_t set_bit_padding[9];
};
```

Finally, an FSM to receive UDP data has also been integrated within the same module, to enable an essential control of the board operations (acquisition enabling, diagnostic reset). It is programmed to accept a payload made by a pair of unsigned 32 *bit* integers, which indicate the desired register address and the value to be configured, respectively. Thus, a simple interpreter was programmed to validate the requests and perform the commands. In Fig. 7.16, an example is shown in which the computer asks for a reset of the diagnostic counters (first two packets) and the start of data transmission. Then, it is possible to notice the packets sent by the UDP module on the *MAC Tx Data* bus.

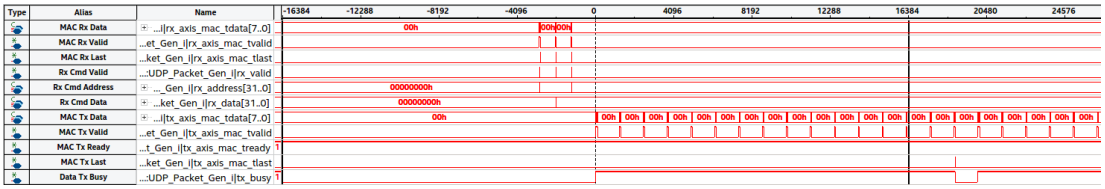

Figure 7.16: UDP transactions, including the commands received by the computer (top rows) and the transmitted data (bottom).

7.11 Performance Considerations

The VFC-HD firmware developed during this research has laid the foundations for the renewed BLM readout system. However, further work will be required to implement the described signal processing and to adapt the new logic interfaces to the legacy modules. The whole measurement signal paths have been clocked at 160 MHz, except the modules dedicated to the UDP data transmission (cf. Par. 7.10), to accelerate the prototype development and proceed with the overall system validation. Although this frequency is relatively high for the employed FPGA, the positive timing slacks, reported in Tab. 7.2, show the result of a successful implementation.

Timing	**Slow** 85° C	**Slow** 0° C	**Fast** 85° C	**Fast** 0° C
Setup Slack	0.092	0.222	2.076	2.606
Hold Slack	0.176	0.052	0.092	0.080
Recovery Slack	1.101	1.222	3.206	3.325
Removal Slack	0.463	0.373	3.206	3.325
min. Pulse Width	0.195	0.195	0.195	0.195

Table 7.2: Timing margins of the FPGA firmware used in the test-bench.

Nevertheless, it will be a good practice to reduce this frequency domain to only those operations that strictly require it, since the number of FPGA resources deployed in the final firmware could likely make it difficult to achieve the same results.

In particular, cross clocking domain FIFOs can be introduced before all the buses characterised by a slow duty cycle, such as after the lpGBT cores ($\delta_{userData} = 25$ %) or after the BLMASIC data retriever ($\delta_{measurements} < 0.1$ %). This procedure will reduce the bus frequencies while maintaining the same data throughput. A 100 MHz clock is expected to represent a good compromise between the processing capabilities and the FPGA timing requirements. But, it may eventually be reduced to 80 MHz, maximising the duty cycle of the filters (cf. Par. 7.8).

Chapter 8

Results

The assembly of the overall system test-bench is shown, describing the involved equipment. The power consumption of the novel front-end is verified to comply with the design specifications and a strategy is established to validate its features. Downstream from a procedure of measurement calibration, the analogue conversion performance is illustrated in detail. The extent of the collected information is summarised through the introduced merit factors, which prove that the instrument precision is within the specifications. In particular, currents down to $1.32\ pA$ *are detectable, with an error of* $\epsilon_r \leq 10\ \%$*, which decreases to* $\epsilon_r \leq 1\ \%$ *for currents above* $3.05\ pA$*. Furthermore, the same study demonstrates the benefits brought by the novel processing filters, confirming the theory. A custom chart will be provided for identifying the optimal processing tuning, given a neighbourhood of currents and a target precision. The digitiser accuracy is investigated and a corrective function is suitably defined to limit the systematic deviation at currents* $> 10\ \mu A$*. Furthermore, at least up to* $10\ mA$*, no ambiguities can arise from injected currents out of the design operating range. Finally, the proton irradiation test performed at PSI (* $200\ MeV$*) is detailed. It shows no major performance drifts or SEE consequences up to at least* $\approx 3.3\ kGy$ *of TID.*

8.1 Overall System Assembly

During the research activity, the new hardware was developed and produced. This allowed for the assembly of a pair of prototyping test-benches, involving the same components that will be installed during the commissioning phase. The setup is shown in Fig. 8.1 and it presents an architecture similar to those used during the BLMASIC validation (cf. Chapter 4). However, the analysis is here focused on the system level, emulating a scenario closer to the real one. The front-end has been assembled within a prototype of the new crate (EDA-03916), which will be employed both in SPS [S51] and HL-LHC, with some geometric differences. It is also equipped with the power transformer (IPU, EDA-03918) and the voltage regulator (PSU, EDA-03919), the rails of which have been adapted to cope with the new requirements (cf. Par. 6.4.1).

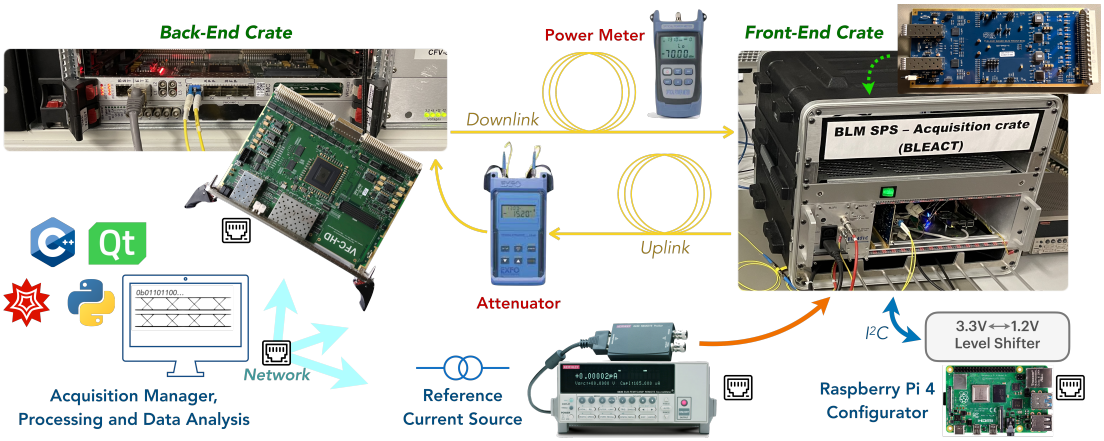

Figure 8.1: System level test-bench architecture.

On the back-end side, a commercial ELMA [S52] crate hosts the VFC-HD card, which has been programmed with the test-bench firmware described in Par. 7.10. The connection of the two domains is established by two pairs of optical fibres, enabling the up/downlink data paths between the front-end board and the VFC-HD. Actually, in the test-bench only a single pair has been assembled, being still sufficient for accessing all the front-end features, whether it is connected to the link A or B (cf. Par.s 6.4.4 and 7.2). However, both the fibre pairs have been installed during the redundancy functionality test and the evaluation of the board power consumption (cf. Par. 8.2). Moreover, the employment of some attenuators has been required, since such a setup involves much shorter connections than those needed for the final installation. This is especially true for the uplink, due to the high output power (2 dBm) produced by the front-end VTRx transmitter (cf. Par. 6.4.4).

The front-end configuration is provided by a Raspberry Pi 4 [S53] embedded computer, which runs a custom software to launch I^2C transactions. These can be controlled by the network via TCP packets, which enable remote operations. A voltage level shifter ($3.3\ V \leftrightarrow 1.2\ V$) has been installed between the Raspberry and the front-end to adapt their digital interfaces. Furthermore, as seen in the previous validation architectures (cf. Chapter 4), a bench-top reference current source has been installed and connected to the front-end, to carry out the analysis of its conversion performance. The available 8 input channels have been tested individually.

Finally, a dedicated computer is integrated into the setup to allow the overall system management. In particular, it executes the processing software described in Par. 5.8, acquires the data from the VFC-HD board and stores the results of measurements. Furthermore, the same computer can control the I^2C configuration and the current source. This is obtained thanks to additional custom-developed software, which also provides the possibility of programming automatised measurement sessions.

8.2 Front-End Power Consumption

The front-end board power consumption was estimated by monitoring the electrical characteristics of the supplying rails for about 6 h. During the whole time, both the redundant links were kept activated and several test input currents were injected into the converters, to emulate typical operating conditions. The rail voltages were set to the nominal design values and constantly regulated by the sense feedback of a bench-top power supply. Meanwhile, the values of currents have been recorded with a period of 1 s.

Power Rail	Current [mA]	Power [mW]
1.2V	543	652
2.5V	420	1050
5.0V	24	120

Table 8.1: Power consumptions of the front-end board at full load.

No major fluctuations have been detected during the test and the maximum standard deviation of 1.5 mA has been reported for the 1.2 V rail. The averaged consumption measurements are presented in Tab. 8.1. The current drawn from the first two rails in the table above is well below the maximum allowed by the PSU unit, corresponding to 4 A. Also, the 5.0 V line is inside the limit, given by 3 A. Furthermore, a total power consumption of 1.822 W has been reported, including the major contribution brought by the optical transceivers (2.5 V rail). Such a value is well inside the specifications and ensures low operating temperatures, which also has considerable advantages in terms of noise. In particular, at the ambient temperature of 27 ^oC, the thermal camera reported the maximum readings shown in Tab. 8.2.

Component	Natural Convection	Crate Fan Ventilation
lpGBTs	36.2 ^oC	32.6 ^oC
BLMASICs	34.3 ^oC	30.6 ^oC

Table 8.2: Temperature of components (27 ^oC ambient).

8.3 Analysis Strategy

During the development of the hardware electronics and the related firmware, analyses at different levels were performed to validate the features of several specific components. This made it possible to gradually proceed with assembling such a complex project, reporting any issues that might have compromised its optimal success. In particular, all the basic functionalities have been validated in the previous stages, such as: the proper system powering, the BLMASIC operations, the digital data transmission, the front-end

configuration and the measurement processing. Hence, this latest test-bench has finally offered a comprehensive view of the upgraded system, at least as a high dynamic range current monitor. This is because the protection and the threshold logic will be inherited from the former BLM.

The present investigation has therefore focused on an exhaustive validation of the conversion performance in terms of precision and response speed. In particular, the absolute current errors σ_r and the relative ones ϵ_r have been measured downstream of the full data chain, beginning with the raw BLMASIC data and terminating with the different processing methods (cf. Chapter 5). To enable a targeted study of the different current ranges and filter parameters, test-bench automation was employed to manage the entire measurement session, which took about a week of continuous data acquisition. According to the options presented in Par.s 5.6 and 5.7, the automation plan has been set by defining:

- the type of employed filter (AVG, IIR Eq.Loss or Eq.Rise),
- the different RS_X tuning,
- the sweep minimum and maximum currents,
- the number of sweep currents,
- the number of samples to be acquired for each current,
- the settling time after each current configuration,
- the time interval between each sample,

the values of which are reported in Tab. 8.2.

Signal	Sweep Features			Acquisition Settings			Diagnostics
	Currents Tested	Min Current [A]	Max Current [A]	Samples Acquired per Current	Settling Time [s]	Sample Period [s]	Total Measurement Time [minutes]
Direct	1000	100n	1m	10000	0.10	10u	3.4
RS0	1000	10n	1m	10000	0.10	10u	3.4
RS1	1000	10n	1m	10000	0.10	10u	3.4
RS2	1000	1n	100u	10000	0.10	20u	5.1
RS3	1000	1n	100u	10000	0.10	20u	5.1
RS4	1000	100p	10u	10000	0.11	190u	33.6
RS5	1000	100p	10u	10000	0.13	190u	33.9
RS6	500	10p	1u	10000	0.30	1.83m	155.1
RS7	500	10p	1u	10000	0.92	1.83m	160.3
RS8	250	1p	100n	10000	3.38	18.3m	776.7
RS9	250	1p	10n	10000	13.21	18.3m	817.6
RS10	150	1p	10n	10000	52.53	18.3m	588.9
RS11	100	1p	1n	10000	209.82	18.3m	654.8

Figure 8.2: Measurement plan for each filter tuning.

It can be noticed that each analysed sweep range does not cover the entire device dynamics. In fact, a limitation of the current minimum has been imposed on the fastest tunings (like RS0-RS3), since they cannot carry out sufficiently precise measurements below such values. Instead, although the slow tunings (RS9-RS11) are able to provide high precision on the full range, their characterisation has been focused on the lower one only, to limit the long acquisition times while maintaining fine sweep steps. In fact, the theory proves that, up to full scale, these cases provide lower σ_r than their faster versions (cf. Par 5.7). In any case, precise measurements of weak currents are still challenging and must be carefully treated.

For each current, the number of acquired samples is fixed to 10^4 for all tunings, representing an optimal compromise between good statistics and a reasonable duration of the data analysis computations. The *settling time* concerns the period the system has to wait between the configuration of the injected current and the beginning of data saving. This quantity has been calculated by selecting the longest amongst the source stabilisation, the electronic delay and the processing step transient (cf. Par. 5.5), which is generally dominant. In this way, the data saving begins when the transient has been completed at the 99.9 % at least. Thus, σ_r will not be significantly influenced.

The novel processing has been developed to operate at the sampling frequency of 100 kHz and it does not apply any decimation of the outputs, contrary to the currently operational BLM (cf. Par. 2.3.2). This minimises the over-threshold detection latency, optimising the accelerator protection. However, during the batch data analysis, it is not necessary to employ such a high rate for data saving while slower filter tunings are selected. In fact, filter bandwidth can be so narrow (< 10 mHz, cf. Par. 5.7) that the higher frequency components of the noise will no longer affect σ_r. Hence, the decimation of the output samples will not entail any disadvantages for this characterisation.

Following such remarks, specific sample periods (cf. Tab. 8.2) have been set according to the desired measurement duration. At the early acquisitions, it was observed that σ_r could be noticeably higher, whether peaks (or valleys) occurred in the integrator waveform, that is for $s_k^{CFC} > 0$ (cf. Par.s 3.2.1 and 5.2). The causes of this phenomenon can be related to an inaccurate calibration of the auxiliary ADC thresholds or other non-ideal behaviours of the converter (cf. Par. 4.8.7). Thus, for each current, the acquisition duration has been imposed long enough to involve at least 3 waveform reversing ($s_k^{CFC} \geq 3$). Such a solution allows for avoiding underestimating the measurement error, regardless of the specific cause of the degradation.

In conclusion, the time required to run the entire sweep acquisition is reported in the latest column of Tab. 8.2. It is possible to observe that the cases involving slow tunings take several hours to be completed, due to the long settling times and the conditions imposed. The total session length was 54 h, which must be tripled (6.75 $days$) to allow the complete characterisation of the 3 different processing methods. The details about each tuning and the corresponding filter parameters were reported in Tab. 5.3.

8.4 Data Calibration

The data produced by the BLMASICs and propagated up to the processing software are not calibrated, but the system provides both hardware and software features to achieve this goal. Thus, different steps must be performed before carrying out the analysis of the precision and the commissioning of the instrument. Furthermore, the real-time processing employed in this research (cf. Par. 5.8) exploits the combination of CFC and auxiliary ADC data (cf. Par. 5.2), which must be preliminarily tuned. For the sake of simplicity, the front-end operates in its default configuration during measurements.

The first calibration step involves identifying the operating range for the auxiliary ADC (cf. Par. 3.3.7). In fact, its raw data are represented over 10 *bit*, but the span which actually contains the triangular waveform is narrower, corresponding to the total V_{swing} of about 2 V (cf. Par.s 3.3.5 and 3.2.1). So, the ADC data have been monitored for a long time, tracking the maximum and minimum values. Furthermore, the procedure is iterated for the 8 front-end inputs to obtain better statistics about different BLMASICs. Since no major differences have been detected amongst the channels, the characterisation was carried on by exploiting the averages of the maxima and minima. They are:

- $\overline{ADC_{max}} = 583$, corresponding to the integrator output of 0.96 V,
- $\overline{ADC_{min}} = 11$, corresponding to the integrator output of -0.96 V.

Momentarily, assuming that the system response is linear with the input current, the subsequent calibration steps involve finding a gain coefficient G_0 and a baseline offset b, which optimise the accuracy of the results m_k (hence also r_k, cf. Chapter 5), that is

$$m_k = G_0 \; m_k^{RAW} + b \; , \tag{8.1}$$

where m_k^{RAW} denotes the dimensionless stream of values produced by the front-end. The device linearity was investigated by employing the CFC counter data (cf. Par. 5.2). In particular, the post-processed measurements were compared with the intensities of the injected reference currents, to identify a linear relation between them. It has been observed that such behaviour is well approximated between 100 pA and 1 μA, while the outer ranges are affected by a deflection. Unfortunately, this degrades the accuracy of the instrument, although it does not influence its precision (cf. Par. 5.1). Hence, a possible mitigation mechanism can be sought and it will be described in Par. 8.8.

To carry out a precise calibration of G_0 and b values, a linear fitting [99] was computed amongst the measurements and the corresponding reference currents, exploiting the data acquired for the characterisation (cf. Par. 8.3). This technique minimises the absolute error power over the whole acquired sweep by definition, thus maximising the accuracy of r_k. It was decided to use the results achieved by RS6 / RS7 tunings (AVG method), which are characterised by high precision and were acquired in sweeps centred in the linearity zone.

Thus, the computed coefficients yield:

- $G_0 = 10.787\ nA$,
- $b = 91\ pA$.

The first one represents the ratio between the raw data produced by the system and the real injected currents, that is $\propto f_s\ V_{swing}\ C$ (cf. Par. 3.2.1). Instead, b is the measurement bias and corresponds to the current added to the input by the calibration DAC (cf. Par. 4.8.6) and additional leakages. Both the results agree with the estimations carried out during the design.

In order to accommodate any considerable mismatches amongst the various input channels, or the different front-end boards, three hardware calibration features can be exploited. The first one concerns adjusting the CFC quantiser thresholds by the external pins **VREFA** and **VREFB**, which directly affect G_0. In the present board prototype, these have been tied to the $1.2\ V$ analogue supply rail, leading to $\pm\ 0.96\ V$ thresholds. However, improved solutions can be adopted to enable better repeatability of the results across multiple boards (cf. Par. 6.4.9).

The second feature involves the tuning of the auxiliary ADC range, which can be regulated by the dedicated registers (cf. the 5 *bit* value *WADCRBe* in Appendix C). This action will actually modify the $\overline{ADC_{max}}$ and $\overline{ADC_{min}}$, affecting both the data combination process and the gain G_0. Nevertheless, since an analogous action can also be achieved by operating on the firmware or the software sides, like in this case, it is recommended to keep the registers in their default status. This is up to the preferred calibration strategy and either the solutions can still be implemented after the hardware installation.

Finally, the calibration DAC can set the desired measurement bias, to ensure that a minimum positive current is always injected into the converter (cf. Par. 3.3.6). In fact, its output can be fine-tuned by the *PLEAKBe* and *NLEAKBe* registers and it is also possible to modify the operating range, by replacing the external precision resistor (cf. Par. 4.8.6). Furthermore, a comparison procedure can verify the DAC bias accuracy. Concretely, it is possible to run the converter while an external reference current is injected and the DAC is disabled. Hence, the measurements are acquired and an averaged result is computed. After that, the external reference is disconnected and the DAC re-enabled, configuring its registers to obtain an output similar to the previous one. Thus, this bias calibration method employs the converter to adjust its own embedded DAC and can be adopted even if the converter itself is inaccurate. It should be noted that such a procedure assumes that, throughout the test, the intrinsic device leakage is constant. If the estimation of the actual leakage magnitude is required, the DAC must be disabled and a decreasing sweep of reference currents must be injected externally until the integrator waveform stops (cf. analogue probing described in Par.s 3.3.8 and 4.8). Hence, the latest configured current will represent an approximation of it. This characterisation reported a value of $\approx 2\ pA$.

8.5 Conversion Performance

At this stage, the most significant analysis carried out concerns the validation of the system conversion capabilities. In particular, the final test-bench setup (see Fig. 8.1) has been employed according to the strategy presented in Par. 8.3, to measure a variety of reference currents and calculate the corresponding error. The overall characterisation data, obtained by the three described filtering methodologies, are graphically summarised in Fig.s 8.3, 8.4 and 8.5, which are reported in the following pages. The plots on the left represent the conversion characteristic, namely the current measured by the system as a function of the injected one. Specifically, the absolute minimum / maximum of the acquired samples are plotted for each injected current, highlighting the respective enclosed interval. Instead, the measurement relative errors ϵ_r are reported on the right. These provide an effective insight into the system precision, which is strongly dependent on the intensity of the input. However, the absolute errors σ_r are also reported for completeness in Fig. 8.6.

These results deliver a detailed description of the converter behaviour over a large variety of cases. In fact, the plots can be used to foresee ϵ_r and estimate the measurement outcomes under certain operative conditions, precisely: the type of employed filter, the selected tuning and the range of the input currents. Higher errors have been reported while measuring the lowest currents of each sweep since the signal has approximately the same magnitude as the noise in those cases. Then, a decreasing trend is noticeable for increasing currents, the shape of which differs significantly between the IIRs and the AVG (see examples in Fig.s 8.7 and 8.8). In particular, it has been observed that the IIR slope is usually steeper than its counterpart, allowing it to achieve a higher precision over larger current spans. Furthermore, when high currents are measured, plateaux of minimum error can be easily detected for most of the tunings. Their values do not appreciably change amongst the filter types, meaning that the converter operates at its maximum achievable resolution. This performance limit cannot be enhanced by modifying the spectral characteristics of the filter, since the conversion noise is superimposed on the useful information within the passband (cf. Par.s 4.5 and 5.3). However, such an issue does not entail relevant drawbacks for the goals, because the measurement precision is still sufficiently high ($\epsilon_r < 0.1$ %).

The detection of the maximum and minimum samples, acquired at each current, provided a direct overview of the range of measured values. Although such limits are helpful for roughly estimating the distribution of outcomes, a deeper analysis shows that it often deviates from the Normal one. For example, while measuring the lowest currents of each sweep, a quasi-Gaussian distribution was identified with a variance much smaller than the estimated error σ_r. This occurs because a few percentages of samples are spread almost uniformly over an interval large as far as 10 times the mentioned variance. Furthermore, the outliers sporadically reach 0, when the measured currents are of the same order as the noise ($\epsilon_r \sim 100$ %). The null outliers can be seen

in the plots, where only the maximum samples are drawn. This effect is due to the data combination method employed before the filtering (cf. Par. 5.2). Nevertheless, it does not lead to an underestimation of the losses in the long term, since each null sample will be balanced by the successive ones to sustain the average. The plotted ϵ_r considers all these contributions, providing a worst-case estimation of the device precision. A detailed study of the probability distribution of samples is out of the goals of the present work.

Finally, the exploitation of the slowest tuning RS11 has allowed for investigating the maximum resolution capabilities of the developed system. It proved that it is possible to carry out a measurement of currents lower than 10 pA with $\epsilon_r < 1$ %, which is largely inside the design requirements (cf. Par. 2.6). Specifically, down to 1.32 pA it is reported an error of $\epsilon_r \leq 10$ %, which decreases to $\epsilon_r \leq 1$ % for currents above 3.05 pA. On the other hand, intensities above 1 μA can still be monitored with an error $\epsilon_r < 10$ %, even employing the fastest tuning RS0. This ensures that even prompt beam loss spikes can be detected with the desired precision. A complete overview of the achievable performance will be described in the next Par. 8.6.

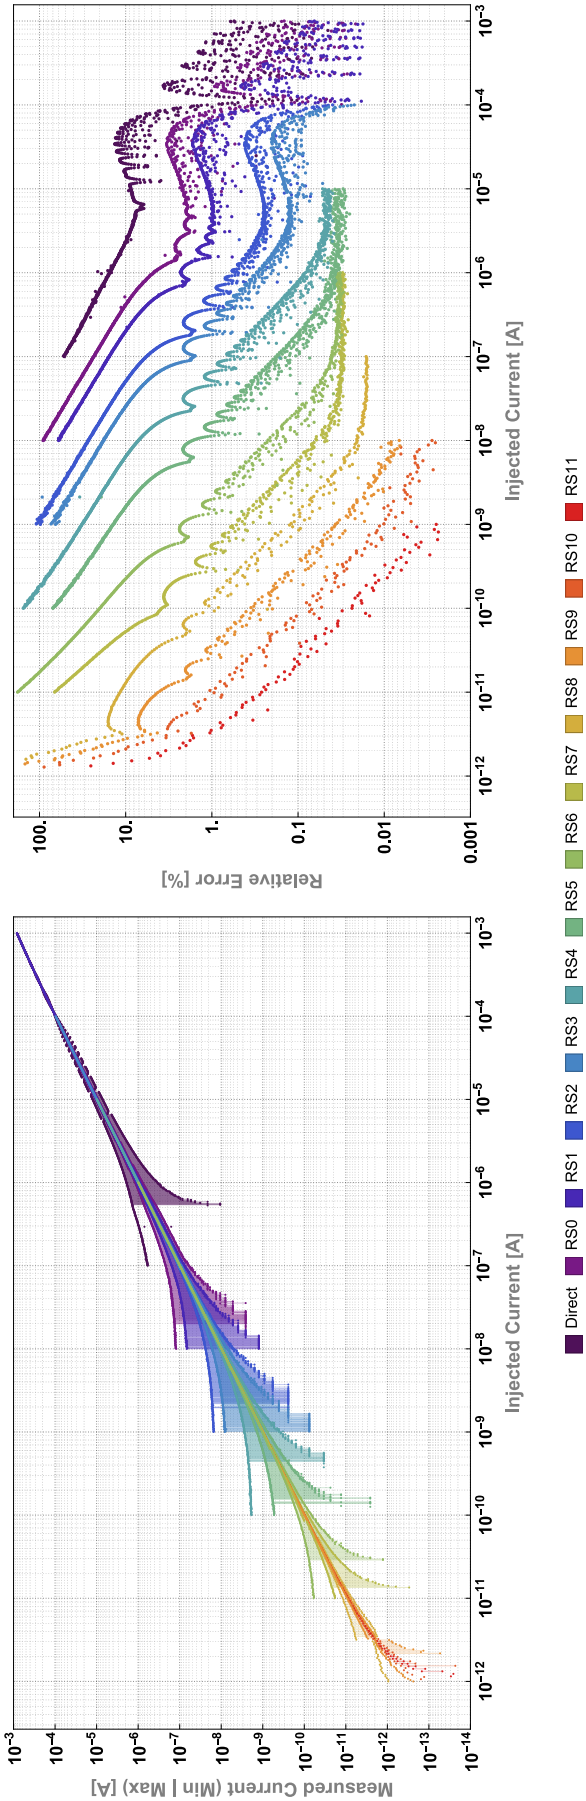

Figure 8.3: Conversion characteristic and relative error ϵ_r achieved by different **AVG** filter tunings.

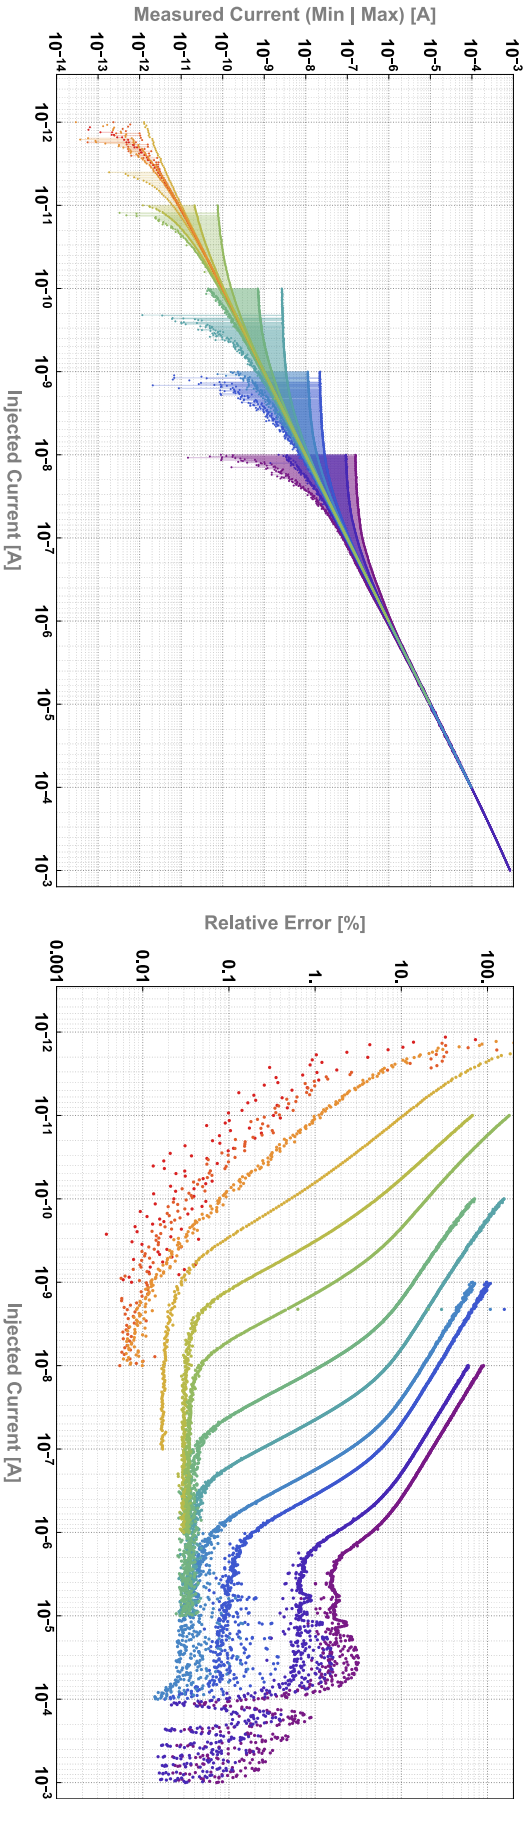

Figure 8.4: Conversion characteristic and relative error ϵ_r achieved by different **IIR Eq.Loss** filter tunings.

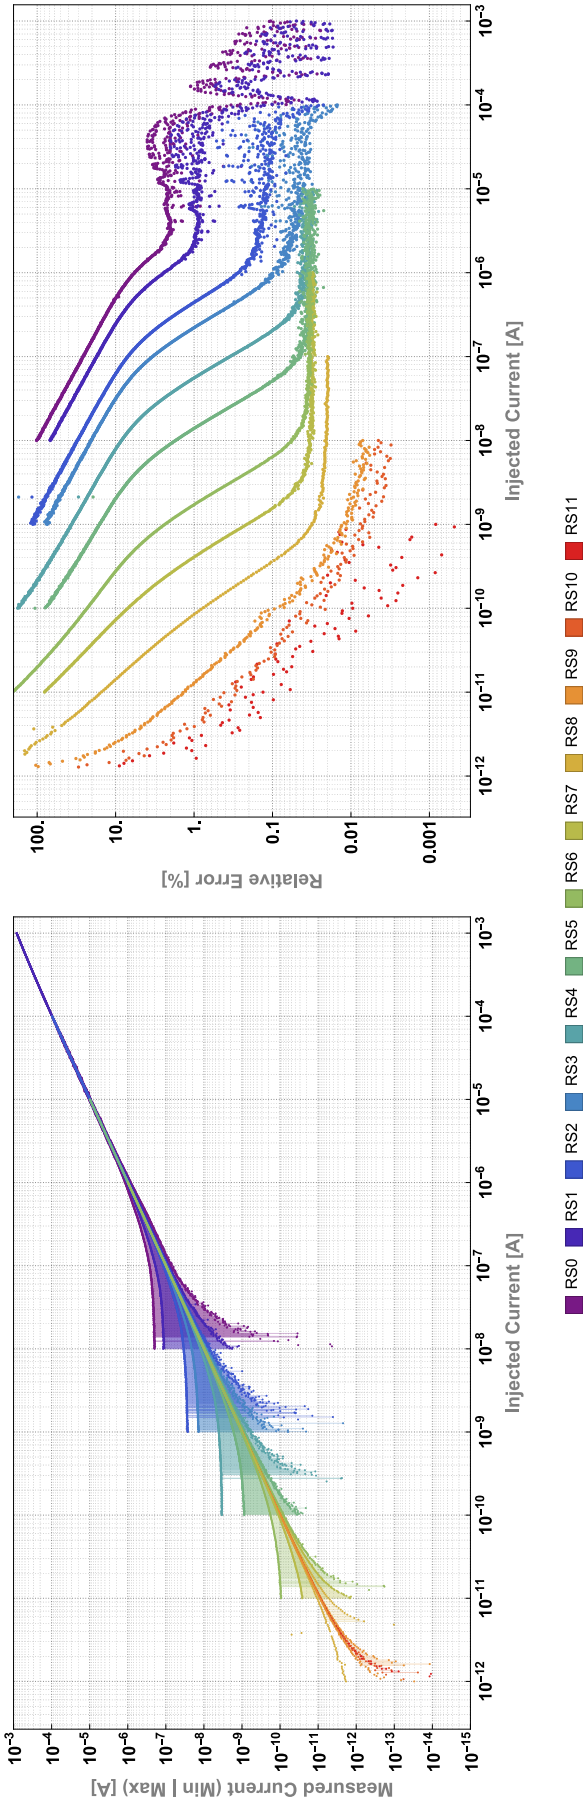

Figure 8.5: Conversion characteristic and relative error ϵ_r achieved by different **IIR Eq.Rise** filter tunings.

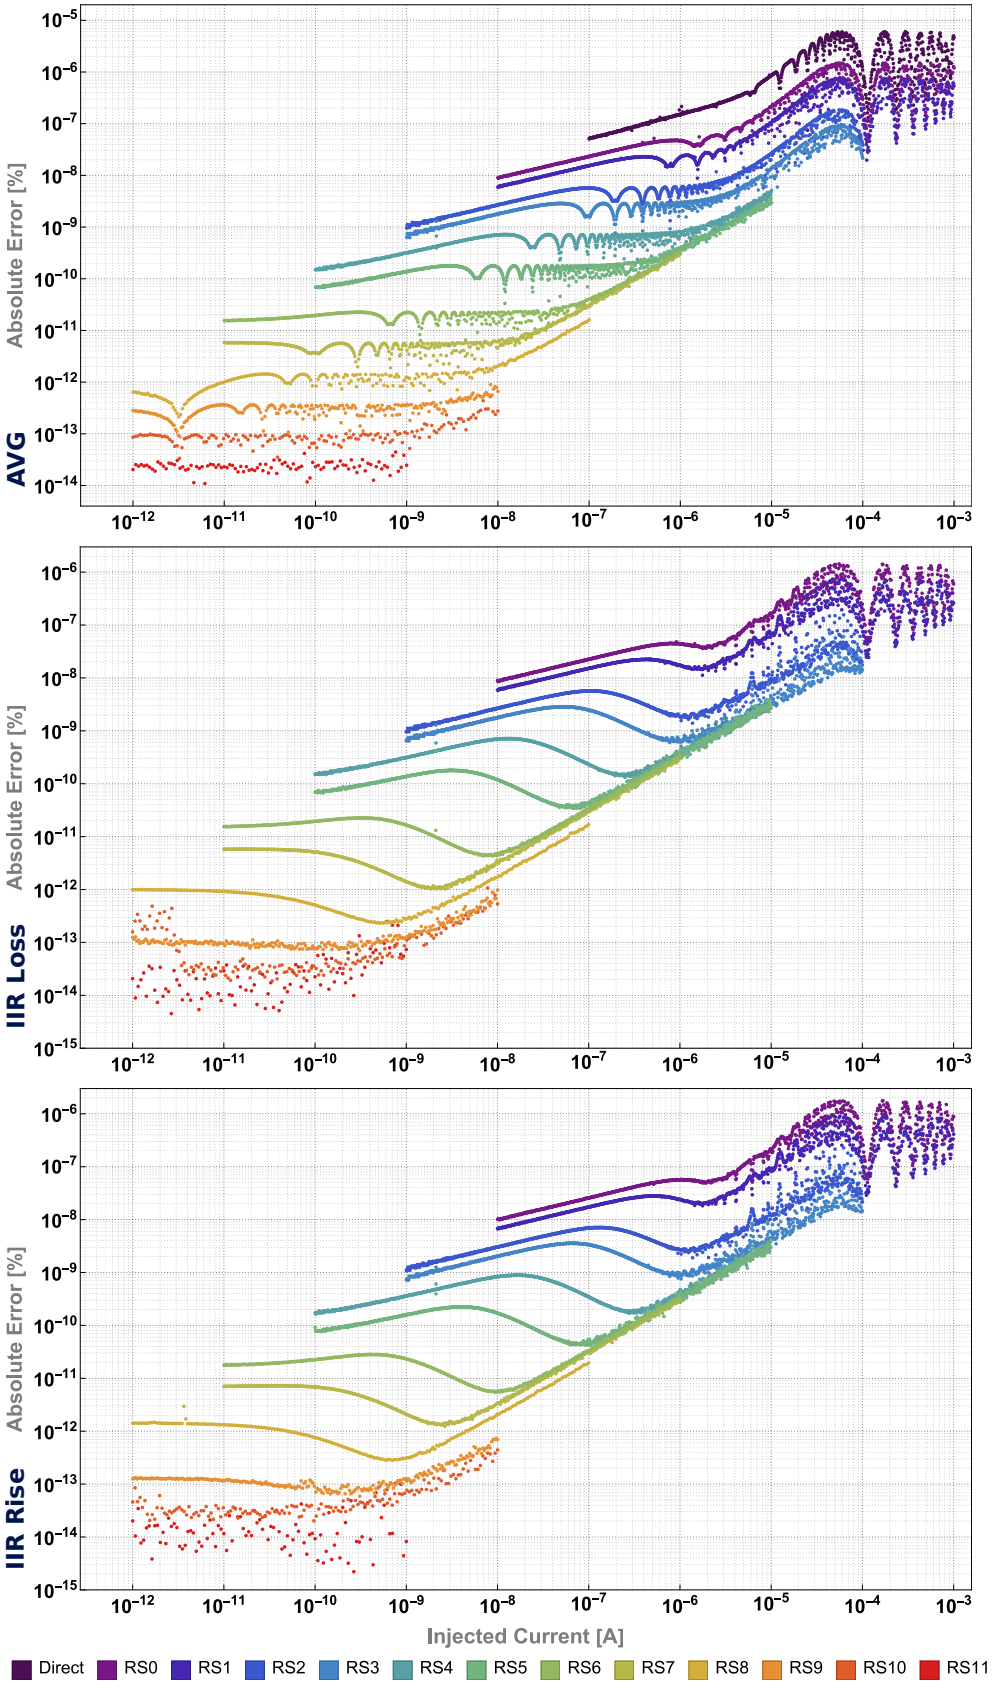

Figure 8.6: Absolute error σ_r reported for different filters and tunings.

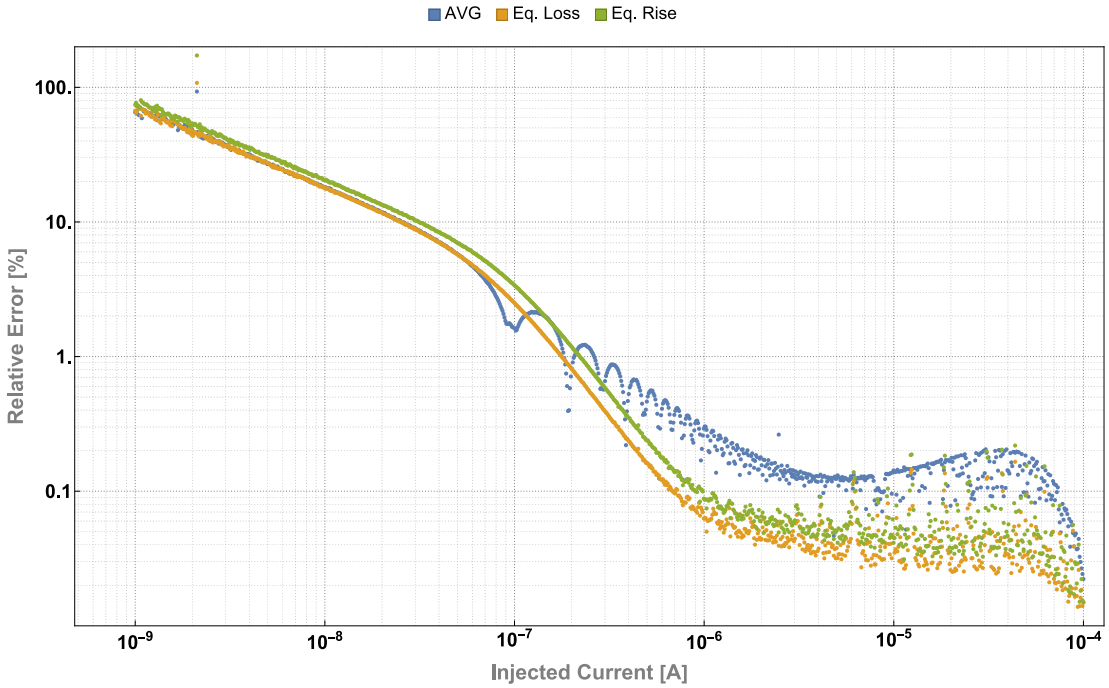

Figure 8.7: Error ϵ_r comparison with the different filter types at **RS3** tuning.

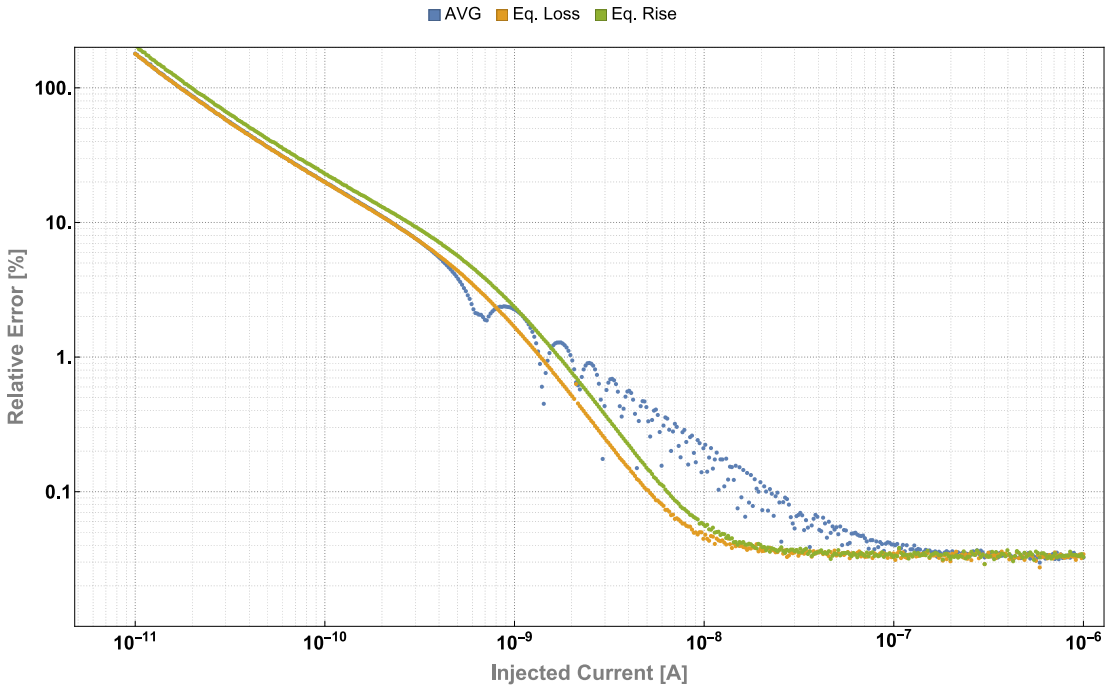

Figure 8.8: Error ϵ_r comparison with the different filter types at **RS6** tuning.

8.6 Merit Factors

The results presented in Par. 8.5 show a detailed analysis of the converter performance under specific circumstances. Nevertheless, the system capabilities can be synthesised in fewer key specifications (presented in Tab. 8.9) by introducing appropriate merit factors. This also facilitates comparing performance amongst the different processing methods, to identify which is the best suitable for the goals.

The first class of factors reports the minimum currents that the system can measure with the target error ϵ_r^T (equal to 50 %, 10 % and 1 %) for the various filter tunings RS_X. The investigation was carried out by comparing the calculated ϵ_r data with the desired target ϵ_r^T. Due to the non-monotonicity of ϵ_r as a function of the current, the maximum value of this variable that best approximates the target ϵ_r^T (from below) is reported. This means that currents weaker than those shown in the table could possibly be estimated with higher precision, but the measurements of all the major ones are guaranteed to deliver errors $\epsilon_r \leq \epsilon_r^T$.

Across the different tunings, it can be observed that the performance of the IIR Eq.Loss is approximately equivalent to the AVG filter, when an error of 50 % is imposed as the target. However, both the IIR filters drastically extend the effective range of measurements, while higher precisions are required (10 % and 1 %). The results confirm that the system operates within the established specifications (cf. Par. 2.6), allowing for measuring currents below 10 pA with an error of 10 % for all the processing techniques. In particular, the best achievable performance for weak signals is obtained by the IIR Eq.Rise, which allowed to monitor currents down to 3.05 pA with $\epsilon_r = 1$ % (that is $\sigma_r = 31\ fA$). The presented theory does not explain why this filter performs unexpectedly better than the Eq.Loss version. However, a possible explanation can be sought in the quantisation effects carried by the IIR implementation, which begins to dominate over the conversion errors while employing very low cutting frequencies. In fact, the coefficients used in the direct form structure (cf. Par. 5.5.2 and Tab. 5.1) tend to $a_0 \to 1,\ a_1 \to -2,\ a_2 \to 1,\ b_0 \to 0,\ b_1 \to 0,\ b_2 \to 0$, which represent a degenerate case of filter having a double pole in $z = 1$. Thus, even minimal changes in the coefficients may produce relevant discrepancies in the numerical calculation result.

Whereas the presented figures provide a local evaluation of the device performance, an additional merit factor can be introduced to score the processing methods over the entire sweep. In particular, for the application concerned, it is interesting to evaluate the average error $\overline{\epsilon_r}$, obtainable while measuring a generic current in the input range. Assuming on it a uniform probability density, such a value can be estimated by

$$\overline{\epsilon_r} = \frac{1}{I_N - I_1} \sum_{k=2}^{N} \epsilon_r^{(k)} \left(I_k - I_{k-1}\right) \quad \left(I_k > I_{k-1}\right), \tag{8.2}$$

where I_k represents each current in the acquired sweep, N the total number of currents and $\epsilon_r^{(k)}$ is the relative error computed by the I_k current measurements.

Signal	50% Error Achievable Current [A]			10% Error Achievable Current [A]			1% Error Achievable Current [A]			Mean Expected Error [‰]			Required Representation Bits		
	AVG	IIR Eq. Loss	IIR Eq. Rise	AVG	IIR Eq. Loss	IIR Eq. Rise	AVG	IIR Eq. Loss	IIR Eq. Rise	AVG	IIR Eq. Loss	IIR Eq. Rise	AVG	IIR Eq. Loss	IIR Eq. Rise
Direct	109n			60.1u			570u			14.80			14		
RS0	27.6n	26.0n	32.8n	521n	390n	503n	84.9u	82.0u	170u	3.81	3.25	4.17	16	16	15
RS1	14.0n	14.0n	17.2n	202n	196n	246n	65.1u	63.7u	69.0u	1.94	1.42	1.88	17	17	16
RS2	3.59n	3.43n	4.53n	50.9n	49.7n	63.4n	534n	357n	444n	2.49	1.14	1.54	19	19	18
RS3	2.14n	2.14n	2.27n	25.5n	25.5n	31.4n	264n	179n	223n	1.28	0.47	0.62	19	19	19
RS4	485p	502p	611p	6.48n	6.34n	7.98n	65.7n	44.9n	55.9n	0.88	0.65	0.73	21	21	21
RS5	160p	159p	184p	2.14n	2.14n	2.14n	16.5n	11.3n	14.0n	0.45	0.40	0.42	23	23	23
RS6	35.6p	35.6p	40.9p	225p	225p	277p	2.02n	1.39n	1.72n	0.49	0.44	0.47	25	25	25
RS7	12.9p	13.2p	15.9p	53.9p	57.7p	72.7p	580p	326p	410p	0.34	0.32	0.34	27	27	26
RS8	2.19p	3.18p	4.00p	13.3p	10.6p	14.6p	134p	64.2p	84.7p	0.25	0.22	0.39	29	29	28
RS9	1.81p	1.50p	1.50p	2.82p	2.43p	2.62p	36.2p	11.1p	13.3p	0.28	0.12	0.05	30	30	30
RS10	1.54p	1.74p	1.28p	2.10p	2.86p	1.54p	11.1p	3.90p	4.41p	0.37	0.42	0.12	30	30	30
RS11	1.32p	1.32p	1.32p	1.52p	1.63p	1.32p	3.76p	4.04p	3.05p	0.16	0.31	0.03	30	30	30

Figure 8.9: Measurement merit factors for the different employed processing. Optimal results are highlighted in green.

Note that, under the made assumptions, the equation (8.2) can also be successfully adopted when the measurements are acquired over a logarithmic sweep, like in this case. The values reported in the table confirm the theory anticipated during the design of filters (cf. Par. 5.7) and they agree with the consideration stated for Fig.s 8.7 and 8.8. In particular, a noticeable reduction of the estimated error $\overline{\epsilon_r}$ is obtained by employing the novel proposed techniques, which translates to higher conversion precision over the ranges of interest.

Finally, the overall analysis has been performed by employing a double precision floating point representation (IEEE-754 [94] at 64 *bit*). However, the values of current measurements can be exhibited in more condensed fixed point numbers. The presented characterisation allows calculating the minimum number of required bits for such a representation, in order to prevent the loss of information in all possible cases. These may be interpreted as merit factors for the different tunings and filters, since having more significant bits implies enhanced dynamics. However, it is important to recall that this information should not be used as a benchmark of precision, contrary to the other factors. For instance, the measurement of a low current could be reported as a number having many leading zeros and few rightmost significant bits. On the contrary, a high current would involve a large number of non-zero bits, the leftmost of which are the only ones actually significant for the result, due to the increased absolute error. In fact, in such a complex system, both σ_r and ϵ_r are dependent on the input signal, thus the number of significant bits fluctuates accordingly. The required number of bits Q_{req} is computed by the ratio between the system full scale FS (1 mA) and the minimum detectable currents I_{min} (target $\epsilon_r^T = 50$ %). In particular, it holds

$$Q_{req} = \left\lceil Log_2 \frac{FS}{I_{min}} \right\rceil \ , \tag{8.3}$$

where the ceiling of the logarithm result is computed.

8.7 Optimal Parameters

An alternative presentation of the results can be achieved by exploiting the data acquired during the characterisation sessions (cf. Par. 8.3). In particular, it concerns the realisation of charts which allow for the identification of the optimal tuning for a given input current neighbourhood and a target error ϵ_r^T. Thus, such tools deliver more practical guidelines about employing the different filters and tunings, even for applications beyond BLM purposes.

As it was introduced in Par. 5.5 and shown in the analogue characterisation results (cf. Par. 8.5), filters with lower cutting frequencies provide better precision at the expense of a slower step response. The BLM system must be able to detect as quickly as possible the events that might compromise the safety of the accelerator (cf. Par. 2.1).

Thus, the optimal criterion concerns the identification of the fastest response tuning that delivers the desired target precision. Thus, a data analysis software was developed to correlate ϵ_r^T at the injected currents, similarly to the procedure adopted for the computation of the merit factors (cf. Par. 8.6). This led to a two-dimensional graph, which correlates each tuning (in colours) to regions in the space of measuring currents and target errors. Clearly, different charts have been obtained for the three filter types and they are shown in Fig.s 8.10, 8.11 and 8.12 (in the following pages).

Such representations should be exploited by identifying a neighbourhood of the current of interest and the desired measurement relative error. Hence, the colour of the area denoted by these contours indicates the most effective tuning to be employed. Of course, the real case would involve rather extended spans for those variables, particularly the injected currents, so the parallel monitoring of several tunings is eventually required. This explains why all the possible tunings are generally monitored by the BLM (cf. Par. 2.3.2). However, exceptions can be considered to minimise the required computational resources whenever the solicitations of specific loss ranges are foreseen.

By observing the charts, the presence of white-coloured regions can be noticed at the left and bottom right of the plane. The former refer to target performance unattainable by the device, since a too high precision is demanded while measuring weak signals. Instead, the latest ones are due to the lack of data, because some of the sweeps were limited to currents lower than the full scale, according to the reasons described in Par. 8.3. It is very likely that the tunings RS2 to RS4 actually cover these areas.

A comparison amongst the results of the AVG and the IIR filters shows that the coloured bands of these latter have a steeper descent and are slightly translated towards the left. For instance, the RS8 band of the AVG reports an error $\epsilon_r^T = 0.1$ % between $\approx 1\ nA$ and $\approx 6\ nA$, while the same target is achieved by the IIRs in the range $\approx 300\ pA$ to $\approx 1\ nA$. This agrees with the considerations made about the merit factors (cf. Par. 8.6), especially concerning high precision measurements, and it can be interpreted in two different ways. Assuming that the demanded precision remains the same, the IIR filters allow monitoring weaker currents than the equivalent tuning of AVG. Alternatively, the IIRs can measure the same current more rapidly than the AVG counterpart since faster step response tunings can be employed. For example, the IIR RS8 tuning can be used to detect currents in the neighbourhood of 600 pA with $\epsilon_r^T = 0.1$ %, but the slower RS9 has to be selected if the processing is carried out with the AVG filter.

Finally, the IIR filtering reduces the surface of the unattainable white area on the left-hand side, further pushing the performance of the system. In addition, their bands are less irregular than those observable for the AVG, as a consequence of a smoother frequency response curve. This simplifies the identification of the optimal tuning.

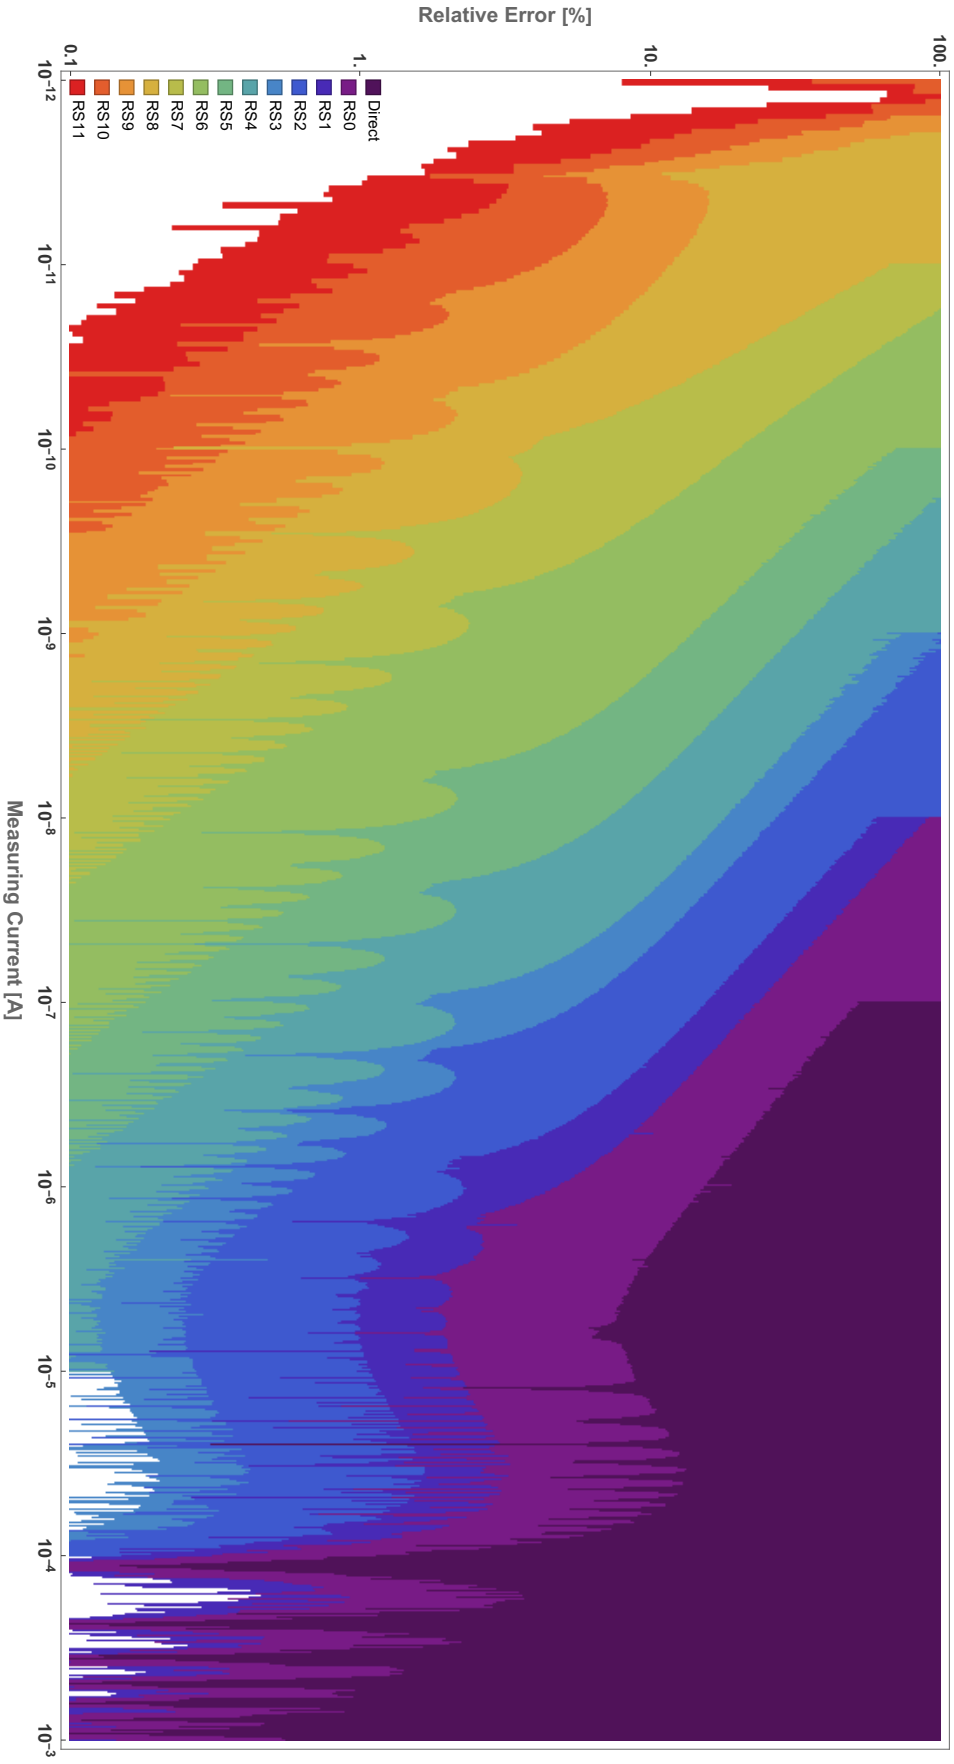

Figure 8.10: Measurement precision chart for the **AVG** filter.

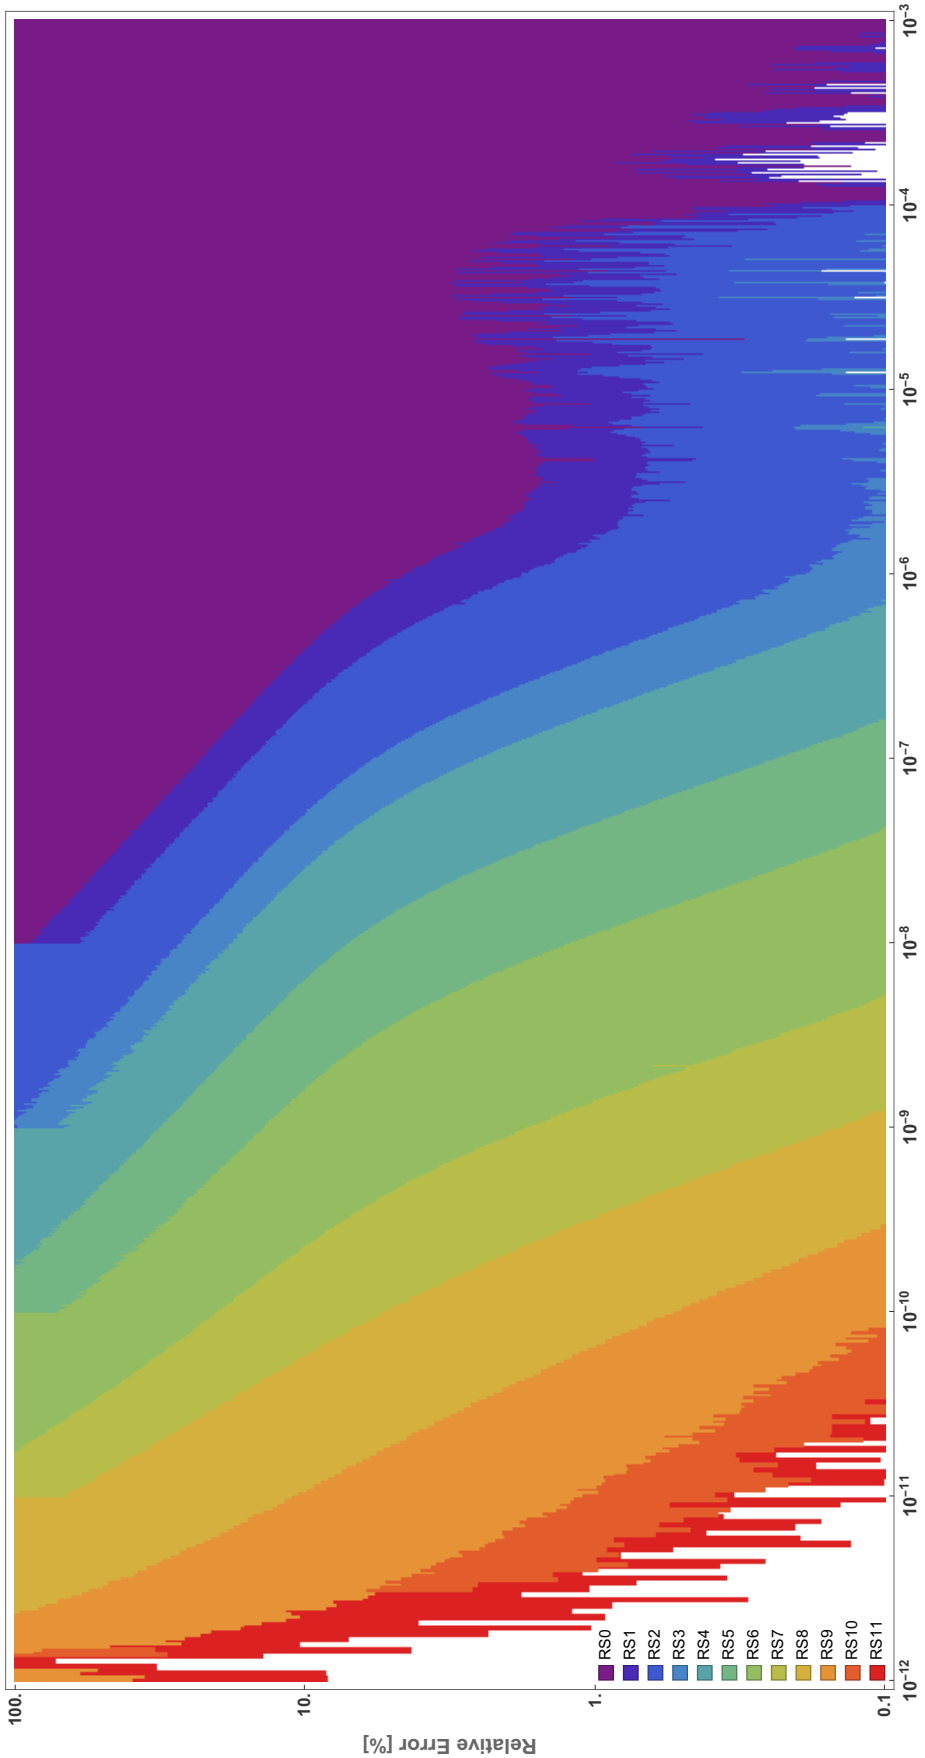

Figure 8.11: Measurement precision chart for the **IIR Eq.Loss** filter.

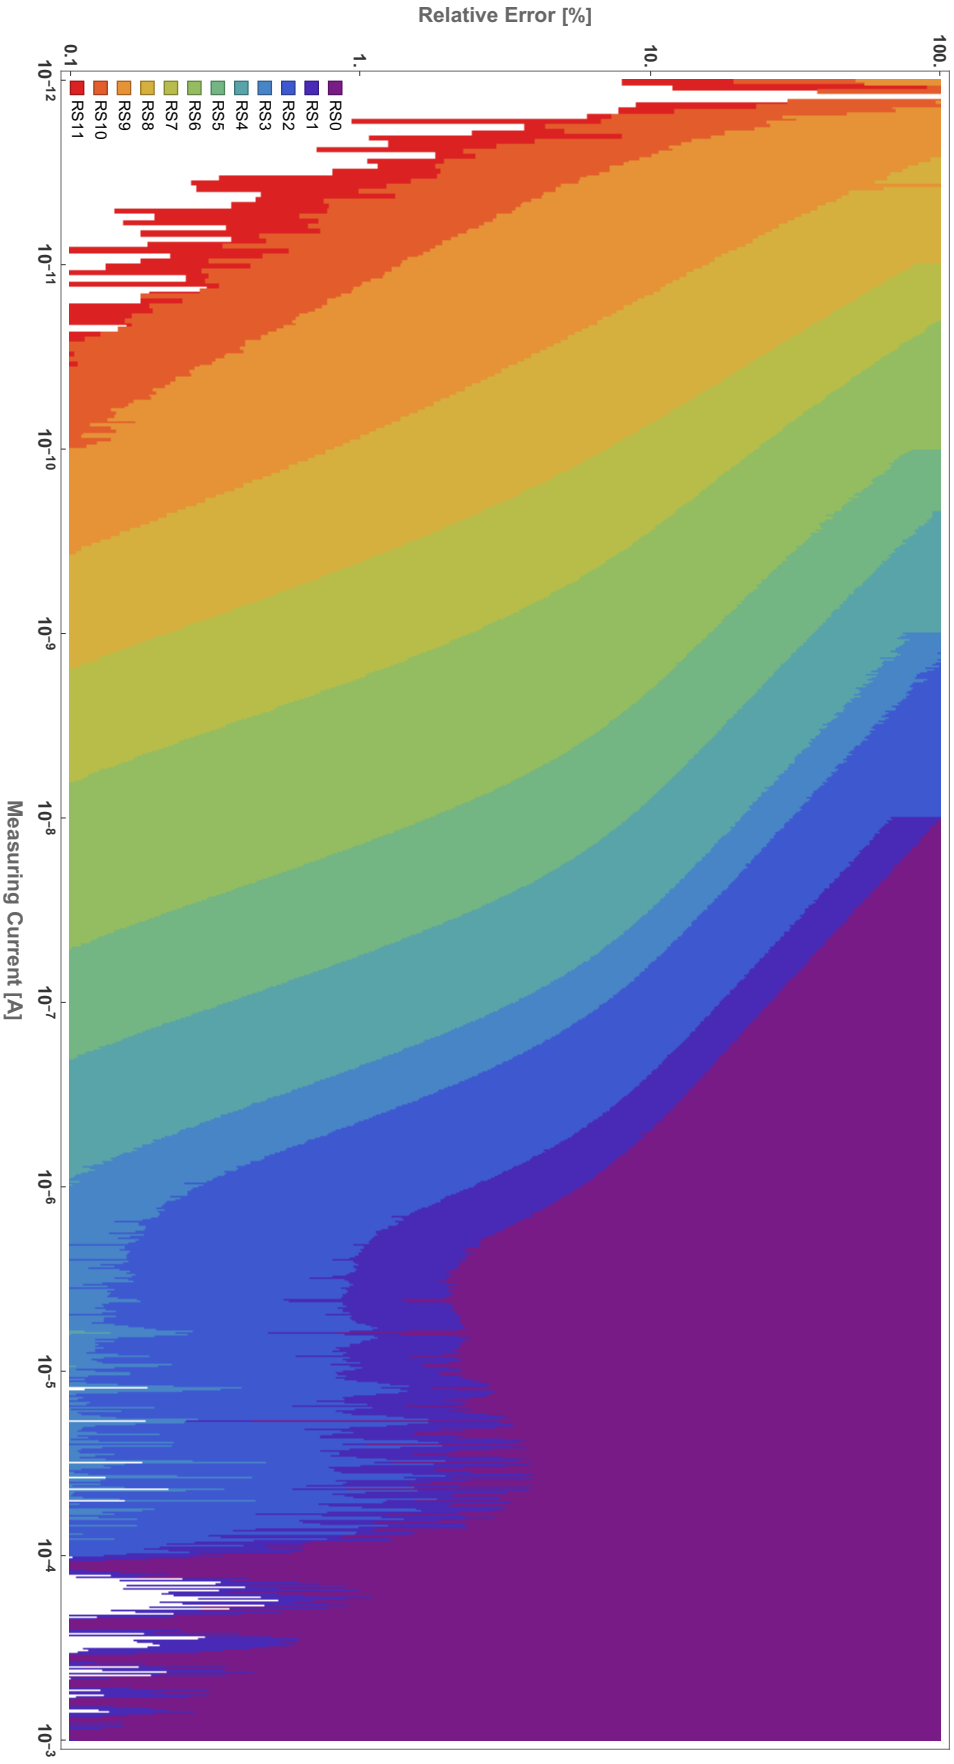

Figure 8.12: Measurement precision chart for the **IIR Eq.Rise** filter.

8.8 Accuracy Considerations

A considerable amount of time invested in this research was devoted to optimising and characterising the conversion precision of the system. Indeed, this is one of the main requirements for the success of the project, according to the considerations made in Par. 5.1. However, the validation of accuracy is also a critical point in this context, since the entire measurement chain must produce realistic estimates of the beam loss over the broadest possible range of magnitudes.

The analysis was based on comparing the performed measurements and the corresponding list of reference currents, injected by the bench-top sources (Keithley [S23], models 6430 and 6221). The data acquired during the precision characterisation (cf. Par. 8.3) have been preliminary calibrated, following the methodologies presented in Par. 8.4. Thus, to evaluate the accuracy, one needs the ratio between the results r_k and the input currents I_k, namely the relative gain coefficient

$$G_k = \frac{r_k}{I_k} \, , \tag{8.4}$$

which is calculated and plotted over the whole range of interest (see Fig. 8.13). Note that r_k already contains the calibration factor G_0. Wherever the ratio is equal to 1, the device response is accurate, meaning that the gain coefficient G_0 and the baseline offset b are optimally calibrated. Otherwise, they should be corrected to fulfil the goal. Usually, the offset b is dominated by the effect of the calibration DAC. Furthermore, it has a value of the order of 100 pA, so its impact on the measurements is limited to the lower range of the dynamics. Consequently, the main deviations to be corrected are described by G_k.

While G_k has been initially assumed constant for the entire measurement range, the results in Fig. 8.13 highlight that it is actually a function $G(I)$ of the input current I, supposed to varies continuously. In fact, above 10 μA, a significant drift of the response is observed, degrading the accuracy of the outcomes. In particular, the full-scale current (1 mA) is actually measured by the instrument as ≈ 829 μA, which is the 17.1 % less and does not comply with the specifications. Nevertheless, the characterisation assessed that the precision achieved in the neighbourhood of the full-scale is high ($\epsilon_r < 0.6$ %), even if direct raw data are considered. Thus, it is possible to introduce a corrective function $F(I)$ to enhance the accuracy, namely

$$F(I) = \frac{1}{G(I)} \, . \tag{8.5}$$

As observed in Fig. 8.13, the weak currents are also affected by the degradation of features, but corrections are not recommended in this case. In fact, the measurement precision is here generally lower, hence the application of an emphasis function may not be as effective as in the previous situation. Specifically, the relative errors may become

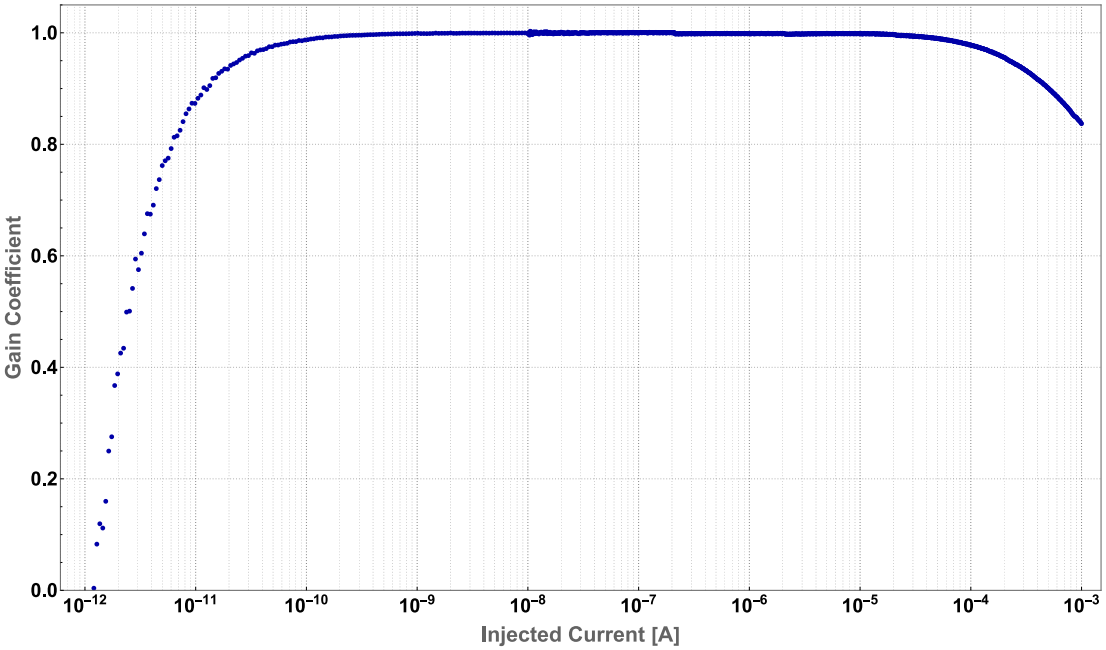

Figure 8.13: The relative gain coefficient G_k is reported for a sweep of input currents. The employed data correspond to the AVG sweeps RS1 and RS10 combined.

so high that the benefits expected in terms of accuracy are actually outweighed by the limited precision. On the other hand, the correction can be applied for sufficiently high currents, estimating a modelled fitting function $G^{>}(I)$ from the characteristic response. It was observed that a good approximation of the trend is achieved by an expression of the form

$$G^{>}(I) = \frac{1 + a\,I + b\,I^2}{1 + c\,I + d\,I^2 + e\,I^3}\ , \tag{8.6}$$

where all the parameters were calculated to minimise the least squares between the $G^{>}(I)$ and the real AVG RS1 data, as an example. The values obtained are

- $a = 5962.93$,
- $b = 9.08392 \times 10^6$,
- $c = 6146.28$,
- $d = 1.09181 \times 10^7$,
- $e = 1.1232 \times 10^9$.

Thus, the fitting curve shown in Fig. 8.14 is obtained. Consequently, the function $F^{>}(I) = G^{>}(I)^{-1}$ can be employed to correct the measurements in the considered range, improving their accuracy. The deviation reached at the full scale is now reduced to 1.8 % and all the computed ratios $F^{>}(r_k)\ r_k\ /\ I_k$ are reported in Fig. 8.14.

The real-time implementation of this correction may be inefficient in the back-end FPGA device, due to the extensive employment of arithmetic operations on the samples.

Thus, it would be convenient to include in the device memory a pre-calculated table of the factors $F^{>}(r_k)$, or even the corrected measurements $F^{>}(r_k)\ r_k$. The latter table would suppress the multiplications, but the monotonicity of r_k as a function of I_k must be guaranteed. This procedure can be effectively applied when the data are represented in fixed-point, for instance, after the filtering processing (cf. Par. 8.6). Finally, it is remarkable that such a methodology introduces a deterministic noise, which could be possibly mitigated by resorting to the dithering technique [100].

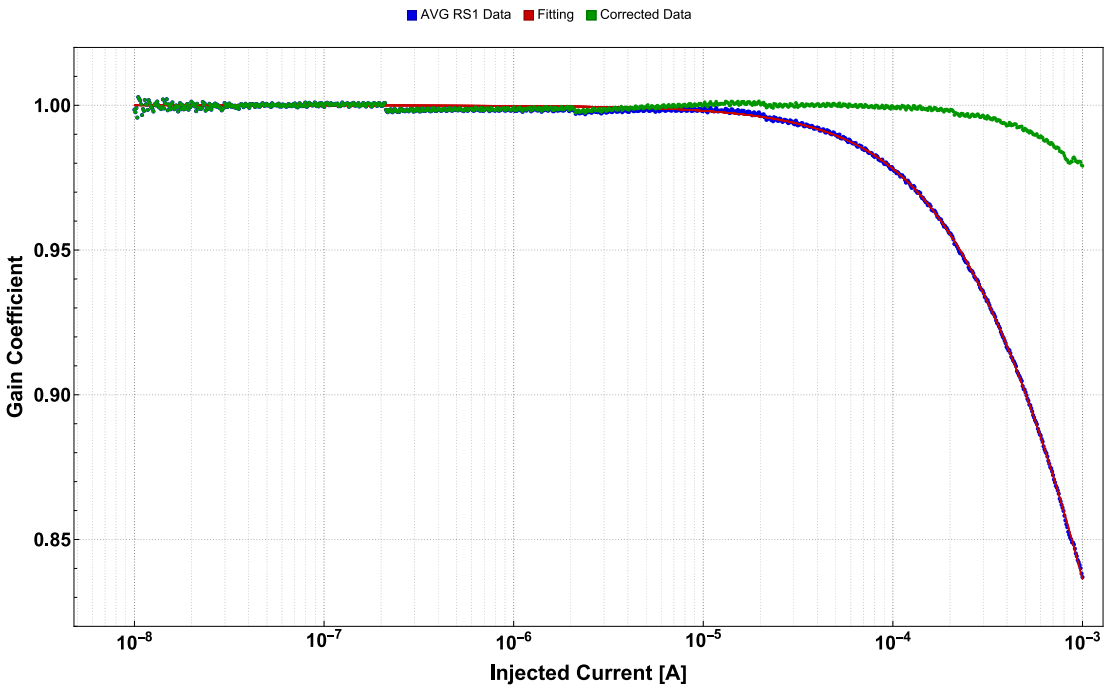

Figure 8.14: Comparison of the measured relative gain coefficient (blue) with its fitting function (8.6) (red), in the 10 nA to 1 mA range. The corrected gain values are represented in green. All the data were obtained by the AVG RS1 sweep.

8.9 Full-Scale Reliability

Although the device performance was examined in detail up to the full-scale specification of 1 mA, it has been investigated for currents above such a limit. In fact, whether huge beam loss events occur in the accelerator, the device response shall not be less than the full-scale value. In particular, the raw converter data were acquired up to 10 mA and corrected by the function described in Par. 8.8. The results are shown in Fig. 8.15, highlighting the maximum and minimum values as it was done for the study in Par. 8.5. It can be qualitatively observed that the system linearity is kept up to 1 mA, as required by the specification. Later, the measurements increase to a maximum, at the input current of 1.38 mA, before decreasing to a plateau value of about 1.1 mA. This behaviour is mainly due to the limitations introduced by the CFC integrator bandwidth on its highest waveform oscillation frequencies.

The investigation ensured that the plateau of measurements is maintained at least for a decade above the full scale. Thus, it can be concluded that, up to this limit, no ambiguities can arise from injected currents out of the design operating range.

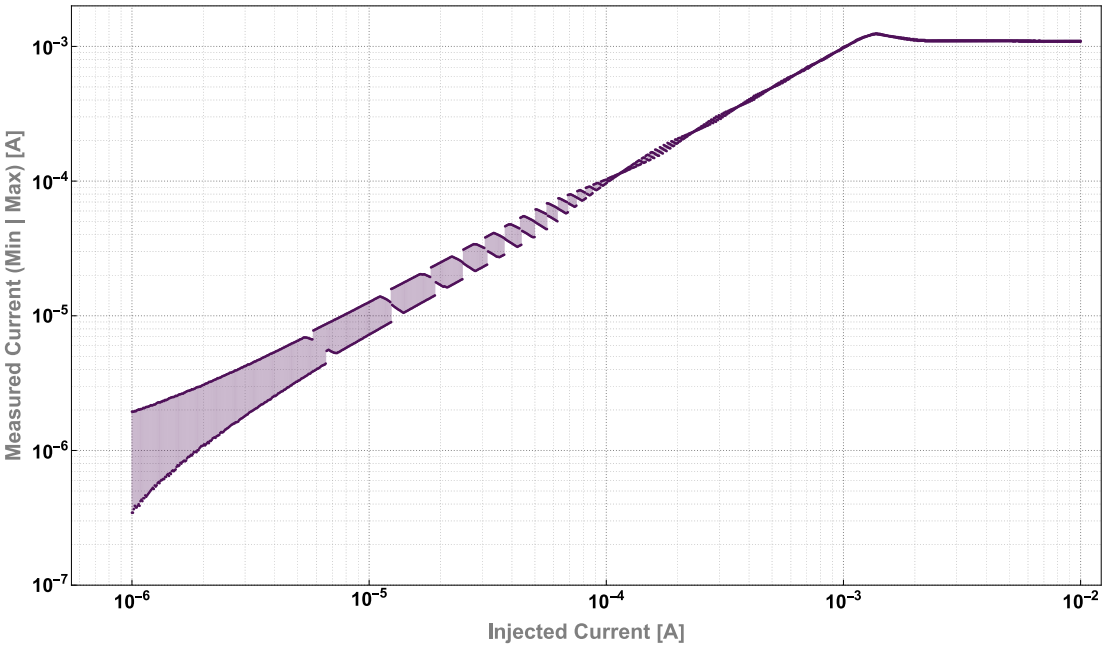

Figure 8.15: Characteristic above the full scale (up to 10 mA), using unprocessed data (Direct).

8.10 Radiation Test

An essential milestone required prior to the release of the BLM upgrade concerns the validation of the front-end ability to work in a radiation environment. In fact, the fulfilment of such a specification involved a substantial fraction of the development time, representing one of the major distinguishing features with respect to commercial solutions. A variety of parameters characterise ionising radiation fields, which can affect electronics in many ways (cf. Appendix A). Thus, several scenarios were identified to better emulate the final working conditions. Furthermore, the team discussed the strategy for the irradiation campaign in view of an effective validation within a reasonable time. The test targets focused on the inspection of the performance drifts due to the TID, to which the front-end will be subjected, and its resilience to SEEs as a consequence of its exposure to high-energy beam fluxes.



Figure 8.16: Aerial view of Paul Scherrer Institute (PSI).

8.10.1 Irradiation Facilities

The considered irradiation facilities have included the CERN High Energy Accelerator Mixed-field (CHARM) [S54] and the Proton Irradiation Facility (PIF) [S55], which is located at the Paul Scherrer Institute (PSI) (Villigen, Switzerland) [S6]. The features provided by the former site have been considered well suited for the present research, due to the possibility of evaluating the overall system behaviour under conditions very closely resembling the operational ones. However, very long acquisition times would have been required to achieve the desired TID (order of some months). In addition, since no collimation solutions are available in this laboratory, a targeted analysis of the individual BLMASIC would have been more complex. Hence, it has been decided to perform such a validation campaign at the PSI (in Fig. 8.16), irradiating the front-end board directly with a proton beam.

Originally, the PIF (in Fig. 8.17) was constructed by the cooperation between the European Space Agency (ESA) [S56] and the PSI for testing spacecraft components under a controlled radiation field. Thus, it has become a landmark for the scientific community that employs systems which operate in such harsh conditions. The facility is part of the PROSCAN project for cancer proton therapy. Its beam is sourced by the compact superconducting cyclotron COMET (COmpact MEdical Therapy cyclotron) [101], [102], which can be fully configured to achieve the desired features. The PIF transfer line is operative only on weekends since the working days are reserved for biomedical applications.

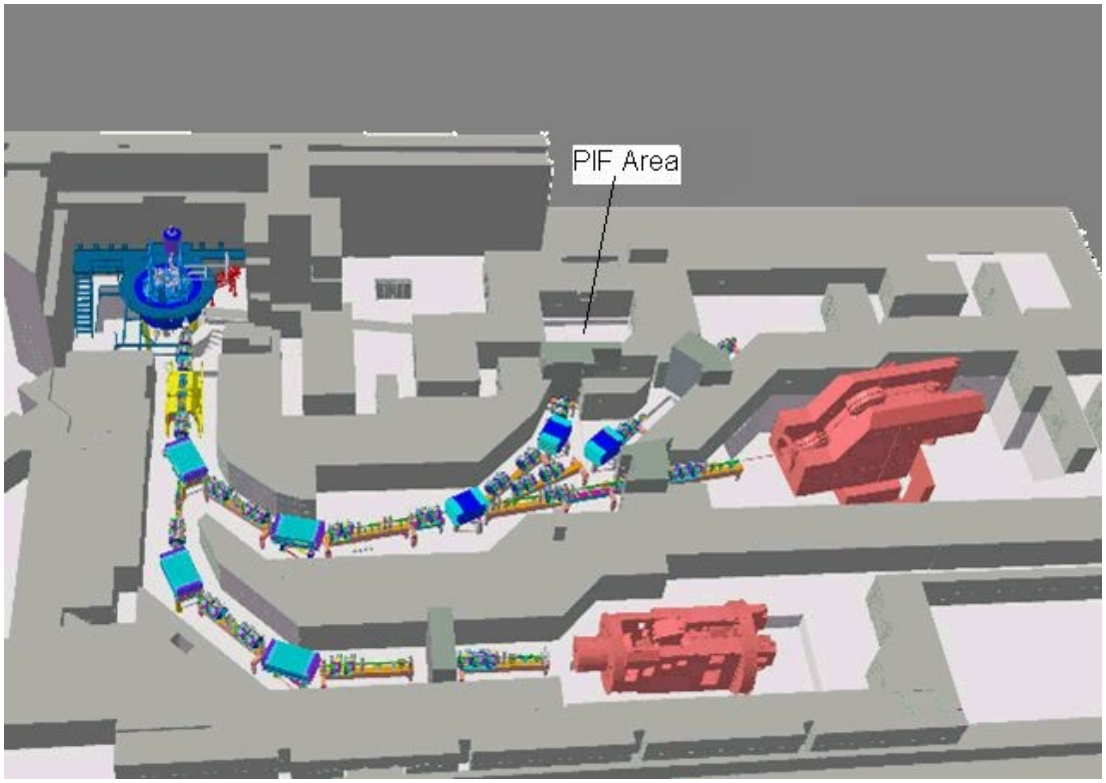

Figure 8.17: Rendering of the PROSCAN project layout, including the proton beam transfer lines and the containing building structure. The COMET accelerator is shown in the top left and the PIF area is designated. The related accessing maze is also visible between this latter and the entering room (on the centre top).

8.10.2 Test-Bench Installation

During the evening of 20th May 2022, a version of the test-bench described in Par. 8.1 was installed at the PIF. The bare front-end board was fixed on the target holder frame (see. Fig. 8.18) and its connections were propagated outside the experimental area through a 20 m umbilical cable. The hosting crate, employed in the laboratory setup, was not used during this campaign since the analysis was only focused on the board. Instead, the remaining setup components, namely the power supply and the back-end crate, were placed at the beginning of the accessing maze, protected from the radiation by concrete walls. While, the computer receiving the data was installed in the facility control room. Finally, the configuration Raspberry Pi and the level shifter box were positioned about 1.50 m below the target and shielded by 50 mm thick lead blocks. Such a solution has been adopted to avoid excessive signal integrity degradation due to the transmission of the I^2C bus over longer distances.

The power supply was configured to drive the electronics in the experimental area. Its recording capabilities allowed logging the values of the voltages and the currents throughout the test duration. Furthermore, since undesired voltage drops were reported on the cabling, sense feedbacks were exploited to control the supplying rails at the delivery location. Both the back-end FPGA and the acquisition software were equipped

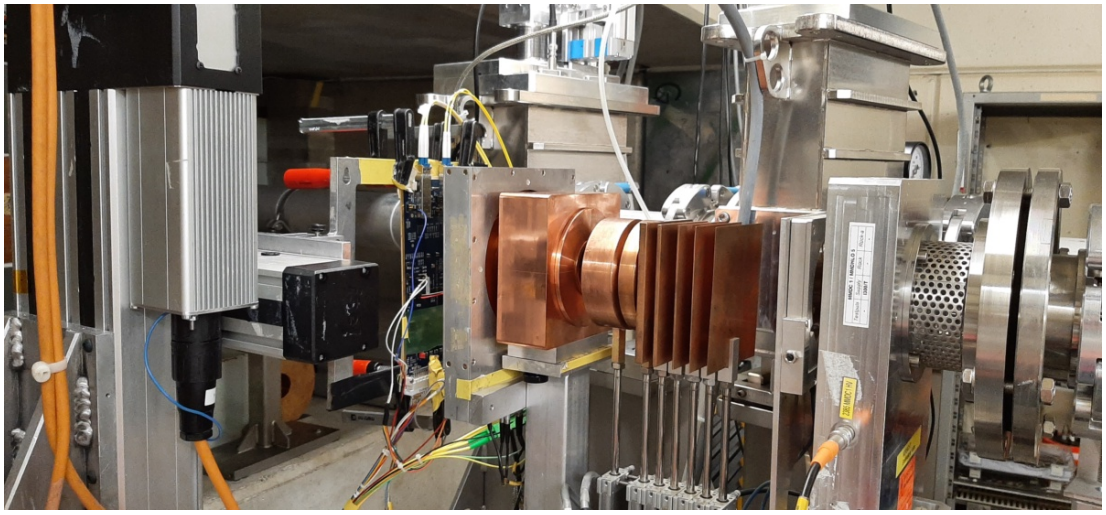

Figure 8.18: Front-End board assembled as a target at the Proton Irradiation Facility of PSI.

with the diagnostic features presented in Par. 7.9 and a set of currents were injected into the inputs for the planned test. In particular, their intensities for the different channels and the type of sources are reported in Tab. 8.3.

Input Channel	**Approximate Injected Current [A]**	**Current Source**
Ch.1	10 μA	External reference
Ch.2	100 pA	Internal calibration DAC
Ch.3	700 μA	Locally mounted resistor
Ch.4	100 μA	””
Ch.5	10 μA	””
Ch.6	1 μA	””
Ch.7	100 nA	””
Ch.8	10 nA	””

Table 8.3: Test currents injected into the front-end during the irradiation test at PSI.

8.10.3 Irradiation Session

The analyses focused on detecting SEE occurrences (cf. Appendix A), monitoring current measurement drifts, in relation to the TID, and evaluating the overall system behaviour under radiation exposure. Using the developed diagnostic framework (cf. Par. 7.9), it has been possible to visualise and store the front-end data and carry out the setup of beam parameters accordingly. In particular, a preliminary cycle was started with low beam intensity, to detect the possible presence of major issues. Then, it was gradually increased up to the maximum achievable particle flux.

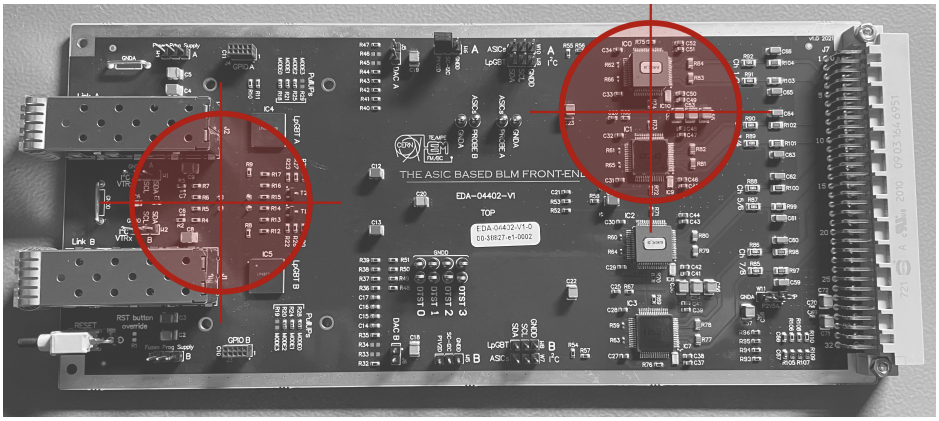

Figure 8.19: Beam pointing locations and collimation area.

To allow a focused evaluation of performance, the beam was collimated (50 mm) on the BLMASICs IC0 and IC1, as shown in Fig. 8.19, while the rest of the card has only been subjected to the scattered radiation. Instead, on the last run, the beam was pointed on the area of transceivers, moving the card by using the XY positioning system. Due to the limited time allocated for the test, only a minor dose was delivered to this zone. Nevertheless, it was deemed necessary to identify possible issues affecting the data serialisation and transmission.

The irradiations started at about *23:30* of May 20th and finished in the early morning of the 21st. In Tab. 8.4 are reported the details about each specific run. The beam energy was set to 200 MeV and the flux, calculated at its maximum current of 10 nA, was 3.7×10^8 $p\ cm^{-2}\ s^{-1}$, with p denoting the proton.

Run	Start Time	Stop Time	Beam Current [nA]	TID [Gy]	Target Area
1	*23:34*	*00:25*	1	50	IC0 / IC1
2	*00:30*	*00:46*	3	50	””
3	*00:51*	*01:14*	7	200	””
4	*01:17*	*02:39*	10	1k	””
5	*02:41*	*04:00*	10	1k	””
6	*04:02*	*05:22*	10	1k	””
7	*05:31*	*05:52*	10	200	Transceivers

Table 8.4: Specifications of the irradiation runs and collimation targets.

8.10.4 Radiation Test Results

Power Supply Monitoring

In Fig.s 8.20 and 8.21, the trends of voltages and currents during the runs are reported for the powering rails. They correspond to those provided by the back panel connector under standard operations within the crate (cf. Par. 6.4.1), with the only exception of the 7.0 V one (shown in purple), which was actually used to supply the P5V line (cf. Par. 6.4.9). Such a higher voltage was set to correctly operate the linPOL12V regulator even without the need for sense feedback. This does not lead to any difference in the front-end performance. Instead, the 5.0 V bus (in green) was only used to supply the Raspberry Pi.

Due to the successful employment of the voltage regulation at the delivery point, no significant voltage changes were reported during the whole session, hence the devices were always correctly powered. The current consumption also remained well inside the margins, however some features can be remarked. It is possible to notice that the 1.2 V rail fluctuates into an interval of $\approx 3\ mA$. This is an intrinsic behaviour of the board neither related to radiation exposure nor to an abnormal condition. The root cause can be sought in the internal circuits of the BLMASICs or the lpGBTs.

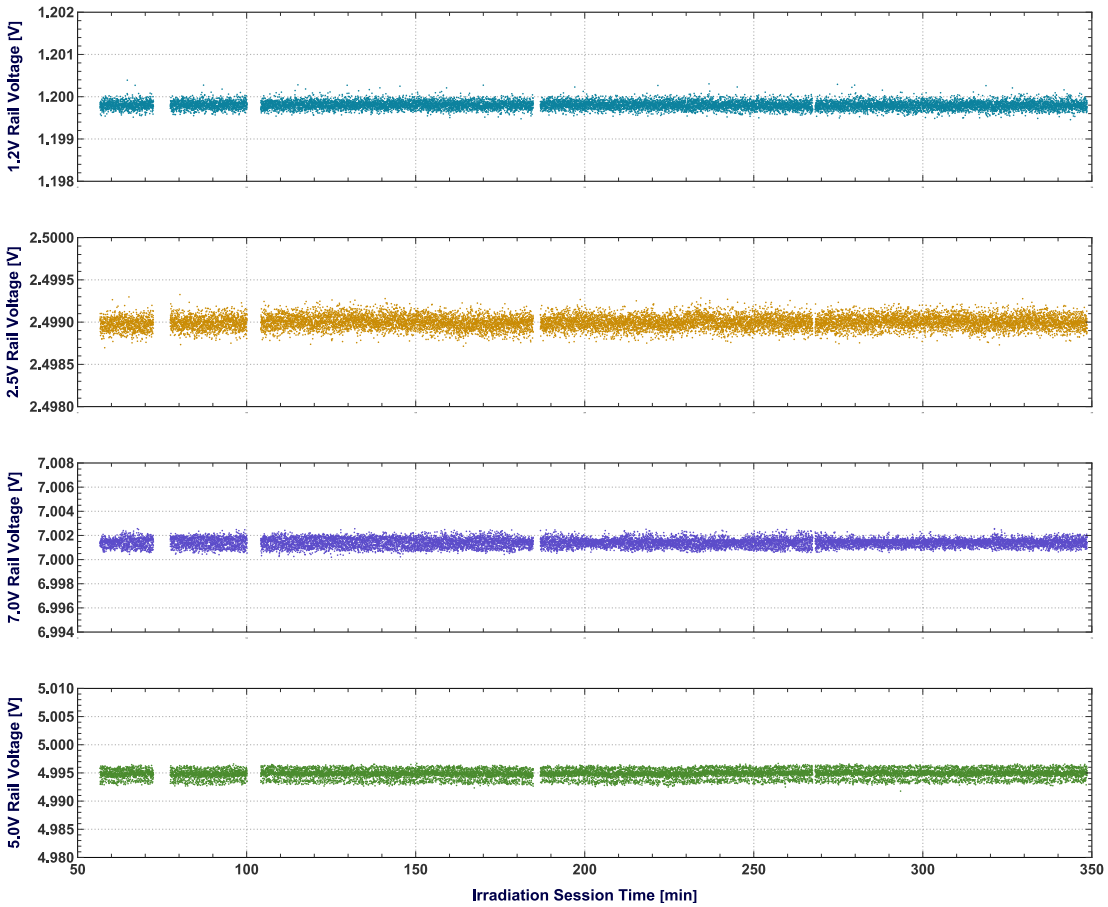

Figure 8.20: Voltage measurements of the supply rails during the irradiation session.

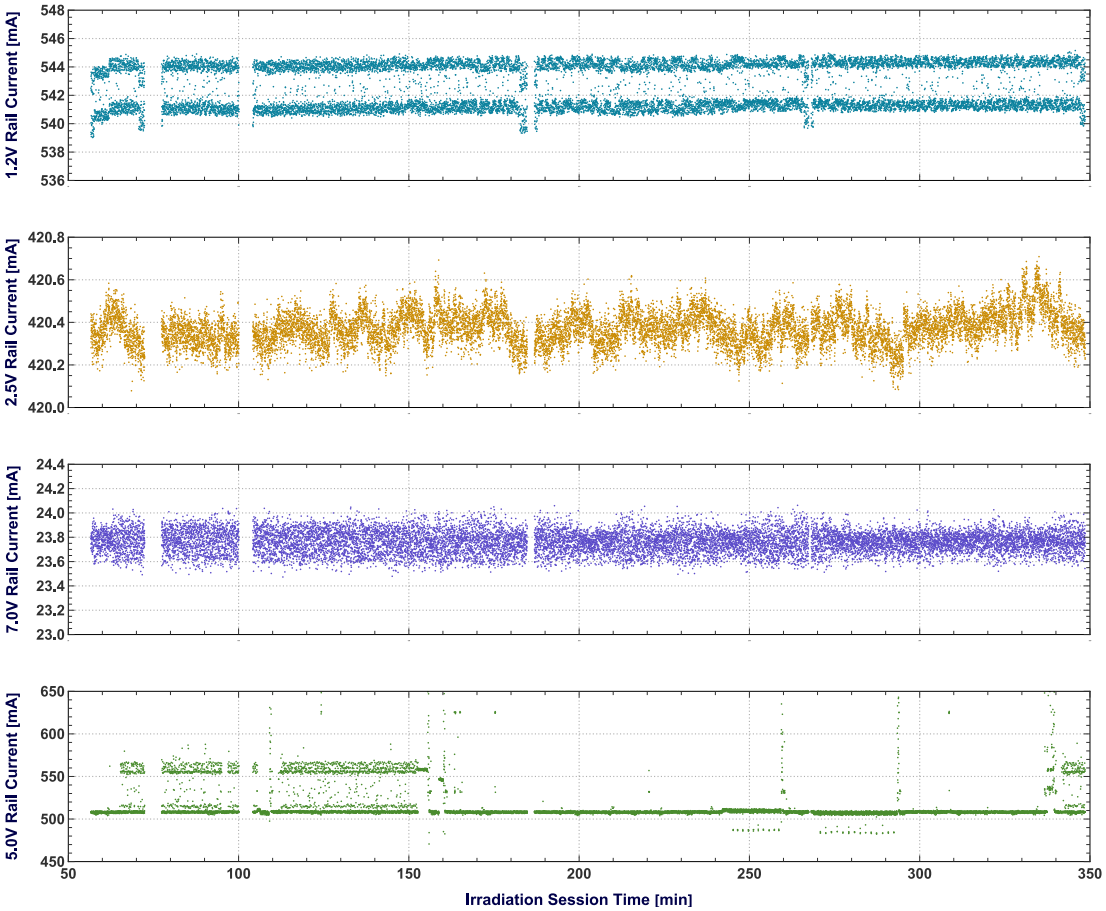

Figure 8.21: Current measurements of the supply rails during the irradiation session.

Instead, for the same rail, it is possible to notice that small jumps of $\approx 1\ mA$ are displayed between the end of a run and the beginning of the next one (white spaces). These are probably due to the photocurrent effect induced by the incident beam. In fact, the data acquisitions were configured to always include a short time during which the beam shutter was closed, facilitating the recognition of the irradiation periods. In any case, this phenomenon is not related to the electronics of the board and does not pose any threats to the application.

Of particular interest is the only rapid current rise that certainly occurred during the irradiation (after 60 min) for the 1.2 V rail. This might represent a case of a single event latch-up which, however, caused no further detectable effects.

Finally, relevant consumption changes were reported for the rail supplying the Raspberry Pi. In particular, higher currents were measured while it was committed to perform I^2C transactions towards the front-end devices, which is reasonable. Furthermore, at least 6 cases in which the current reached 650 mA were observed during the irradiation periods. It turned out that, in these instants, the device got stuck into deadlock states due to the radiation exposure, thus the reboot was forced. Although the characterisation of this single board computer is beyond the scope of the research, such results carry a relevant scientific implication. In particular, it has been demonstrated

that devices not built with radiation tolerant features quickly fail in such harsh environments, even if they are heavily shielded. This is further supported by observing that the Raspberry Pi was located rather sideways to the beam trajectory, hence it was subjected to a much lower flux than the hardware developed in this work.

Digital Validation

Exploiting the diagnostic tools presented in Par.s 7.9 and 7.10, the reliability of the configuration memory of BLMASICs and the integrity of the data propagation, up to the acquisition computer, have been validated with excellent outcomes. In particular, no errors were detected throughout the test in the 8b/10b coding and the frame parity check, involving the communication between the BLMASICs and the lpGBTs. The absence of skids guaranteed the contiguity of samples and no unexpected changes in the registers were observed in any BLMASICs. This means that, even if SEUs could have potentially occurred, the triple redundancy of the digital logic mitigated all the consequences.

The uplink optical data transmission has also been monitored by reading the diagnostic report provided by the lpGBT IP core (cf. Par. 7.4). Indeed, its embedded FEC system declares any correction performed on the received bits, but no occurrences were detected. Instead, the acquisition software recognised a single case of a UDP packet loss. This event is not radiation-induced, since both the back-end and the computer were placed in protected areas (cf. Par. 8.10.2). Also, it is guaranteed that this did not influence the diagnostic reports, since all the faulty cases are counted within the back-end FPGA, the results of which are regularly transmitted to the computer. Finally, such a missing packet certainly did not affect the current measurements, because the raw data are filtered over a larger time window.

Current Measurement Drifts

The test currents were monitored during the session and processed by an IIR filtering (RS8 tuning). Their percentage drifts, with respect to the mean values, are shown in Fig. 8.22. The channels CH1 to CH4 belong to the devices directly irradiated (IC0 and IC1). An overall mild increase is observed in them because of the absorbed $\approx 3.3\ kGy$ TID. In particular, this can be easily recognised in CH3 and CH4. On the other hand, in CH1 and CH2, the photocurrent is neatly detected, in analogy to that seen for the consumption measurements presented in Par. 8.10.4.

Remarkably, the maximal measurement deviations never exceed the 0.5 %, which is an outstanding achievement accordingly to the design specifications.

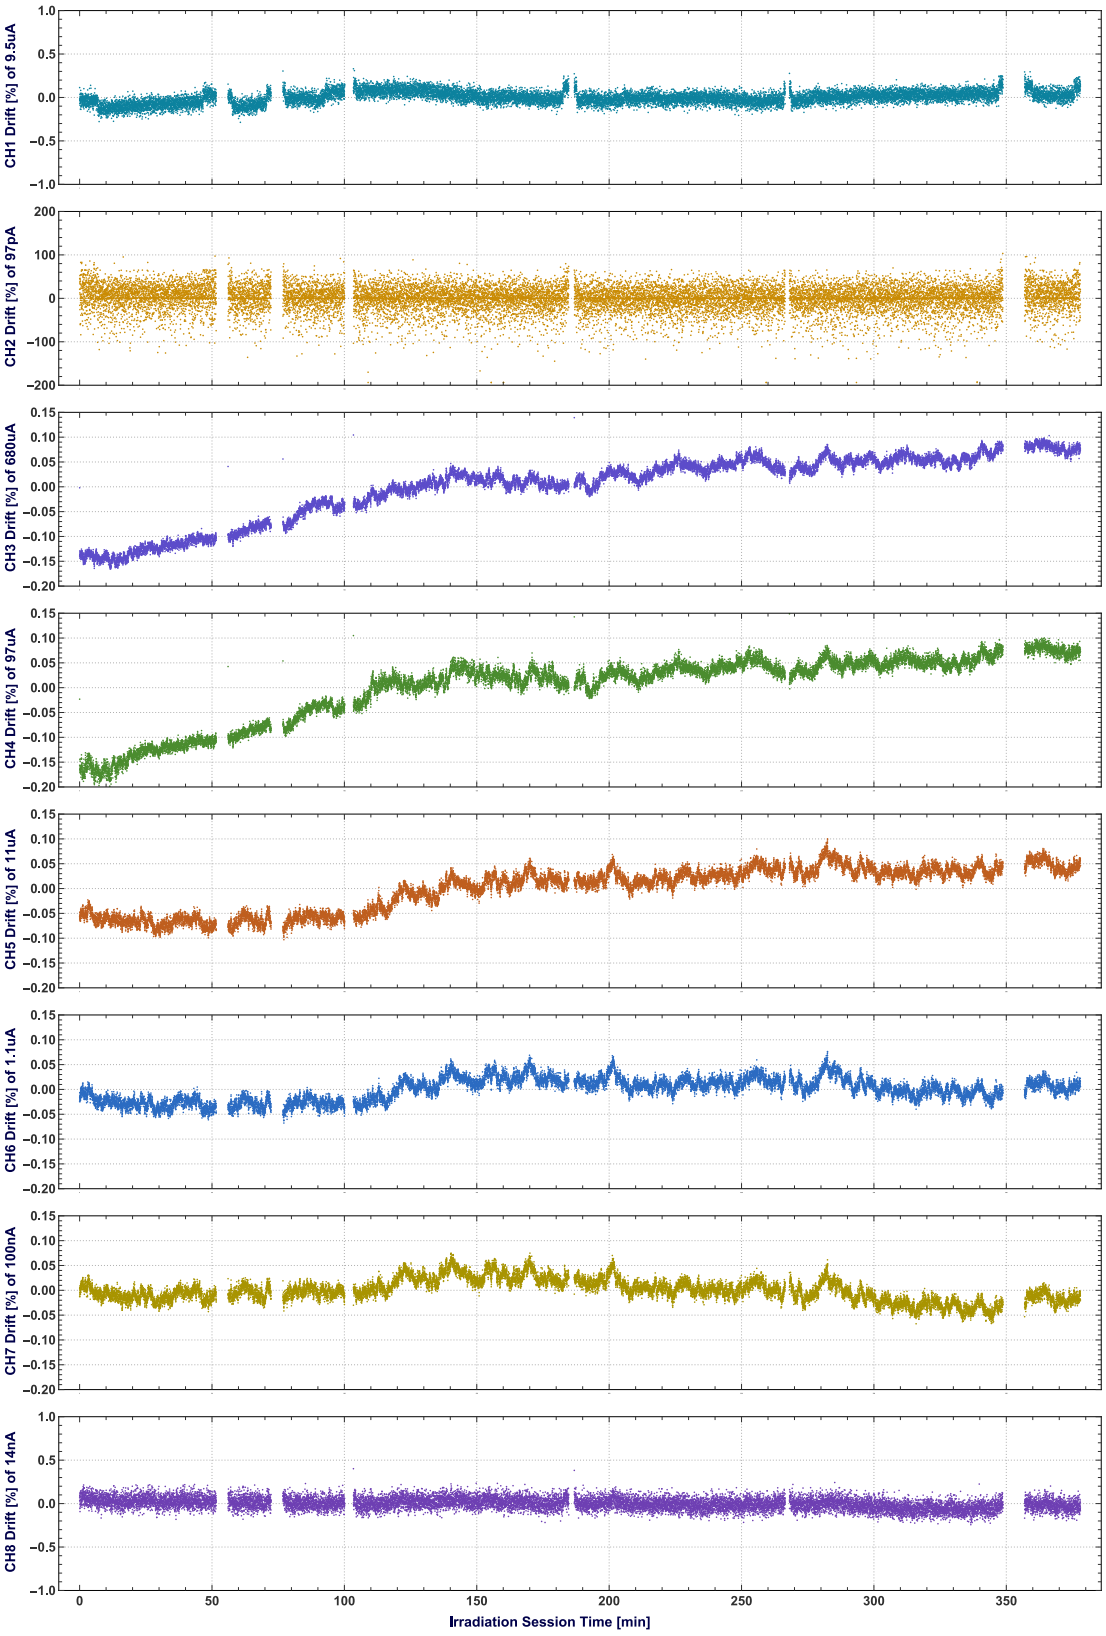

Figure 8.22: Percentage drifts of the current measurements during the irradiation session.

Chapter 9

Conclusions

The HL-LHC accelerator will significantly increase the beam loss radiation throughout the tunnel, exposing the superconductive magnets, as the new Triplets of ATLAS and CMS experiments, to a higher risk of quenching and the whole critical equipment to a harsh working environment. For these reasons, the BLM system performance needs to be upgraded in terms of resolution and response speed, enhancing its ability to detect potential hazardous scenarios promptly. Such a goal has demanded the development of radiation-tolerant digitisers, which must be located as close as possible to the detection spots. In particular, an improvement factor of more than 6 was assessed as the TID target (3 kGy) with respect to the state of the art, considering a 20 years lifetime.

Since such requirements are not achievable by employing commercial products, a completely new front-end board was engineered and validated in this work. The activity also involved supporting the design and performing the characterisation of an application-specific integrated circuit (the *BLMASIC*), which represents the crucial component devoted to the digital conversion of the detector signals. This agrees with the strategy of exploiting in-house designed ASICs. In such a manner, the number of components populating the card scaled by a factor of ≈ 5, translating into benefits concerning reliability, power consumption and reduction of costs to about 1/3, in comparison with the former solution. Moreover, the ASIC-based design enables more precise planning of spare procurement and obsolescence. The new electronics are compatible with the latest solutions provided for the beam instrumentations at CERN, the optical data transmission and the readout back-ends. Their firmware was also largely developed during the reported research. Finally, all the hardware products were optimised to simplify future commissioning and embed a broader range of diagnostic features, to effectively supervise the BLM system as part of its "sanity checks".

The delivered instrument has been optimised to fulfil the BLM specifications for the HL-LHC project. However, the flexibility of the employed architecture allows it to be effectively adapted to applications beyond the commissioning target. In fact, designers of modern particle accelerators for fundamental research, medical and technological purposes strive to boost either the beam brightness, or the luminosity. Hence, they must

cope with increasing losses and the associated damaging potential. Therefore, having electronics that can effectively monitor and operate in harsh radiation environments will increase the reliability and availability of their systems. Furthermore, the developed digitisers pave the way to miniaturised conversion solutions, which can eventually be integrated within the detector. This prevents the need for bulky font-end crates. In addition, the sensitive analogue transmission is drastically reduced and the optical data stream is generated on site. Such a concept could greatly simplify the maintenance of the BLM system in very large accelerators, like in the foreseen FCC [103]–[105], and it will allow for the plant robotisation [106].

A detailed frequency domain analysis of the data stream and the signal processing was undertaken to optimise the system conversion performance. In particular, by exploiting the theoretical foundations of Signal Theory and using specific numerical procedures, novel data treatment solutions were proposed. These are designed to enhance the precision of the loss current measurements and the response speed to abrupt transitions. Precisely, the technique of running sums, employed in the previous BLM filtering, was compared with a more advanced IIR one. Then, the assumptions and the deriving benefits have been successfully validated on real data.

As filtering of base-band signals is popular in scientific applications, the presented methods can be effectively generalised and reproduced to gain the same kind of improvements. Actually, as shown in Par. 5.9, it is also possible to accurately estimate the outcome error for a wide range of filters and inputs, without performing extensive measurement campaigns. This directly derives from the simple rules of composition among the spectral amplitudes of the data and the frequency response of the applied filter. In conclusion, the novel proposed methodology shows that the targeted acquisition of the only measurement noise spectrum can be exploited for a flexible and automatised error prediction. This significantly reduces the test-bench occupancy time and, consequently, increases its availability.

Furthermore, the convenience of adopting IIR filters for the treatment of measurements carries a general value, since their improved spectral characteristics allow for sharper conversion noise suppression, with respect to the usual moving average approach. This feature is particularly demanded in locations where high background noise is combined with detection cross-talk or beam collision debris, which makes it arduous to establish appropriate protection thresholds. A processing prototype was implemented in C++ to validate the proposed methods, applying a multi-threaded software architecture. It readily carried out a detailed characterisation of the system conversion performance, without the need to involve an FPGA logic for such an aim at this stage. In addition, different embedded systems can execute the same code on stand-alone processors, enabling the filter realisation to be used for generic applications. For the final implementation, an analysis of the minimal required logic resources was undertaken, proving the feasibility of such methods, even if a double precision floating-point representation is exploited.

Different analysis levels were performed to validate the basic features of the critical components. They include the power supply, the converter calibration, the data transmission and the configuration of the boards. For all of these, successful responses have been obtained. This led to the assembly of the overall system test-bench, which is very similar to the definitive configuration that will be installed at HL-LHC. Leveraging this, a standardised strategy of measurements provided the complete characterisation of the analogue conversion performance. In concrete, for 54 h, suitable current sweeps were injected into the system and the measurements were acquired and processed, to validate the converter resolution for different filtering techniques and tunings. The front-end power consumption was also measured at full load, obtaining a value of 1.822 W, which is well inside the upper bound of the design. Similarly, the temperatures of lpGBT and BLMASIC devices were monitored by a thermal camera, reporting readings always below 40 ^oC at an environmental temperature of 27 ^oC. The analogue characterisations show that currents down to 1.32 pA were detectable by the system, with an error of $\epsilon_r \leq 10$ %, fulfilling the design specifications. Furthermore, the minimum intensity of 3.05 pA was measured with high precision ($\epsilon_r \leq 1$ %), using the equivalent rise time IIR filter. Concerning the high-intensity loss currents, values above 1 μA can still be monitored with an error $\epsilon_r < 10$ %, employing the fastest tuning (RS0) with the rise time of 40 μs. This ensures that even prompt beam loss spikes can be precisely detected. Instead, by direct use of the raw data (period of 10 μs), the minimum detectable current is 109 nA (at $\epsilon_r \leq 50$ %), increased to 60.1 μA if demanding an error $\epsilon_r \leq 10$ %. Finally, precision map charts were displayed, to facilitate the identification of the best tuning given a specific neighbourhood of currents and the target error.

For measurements with a target error of 50 %, the performance achieved using the equivalent loss IIR filters proved similar to the AVG one. However, the former dramatically improves the capabilities when higher precisions are required ($\epsilon_r < 10$ % or < 1 %). The mean measurement error was also estimated over the whole sweep range, providing a score assessing the best performing filter for the various tunings. In all the examined cases, it turned out to be the IIR. In conclusion, although the AVG solution still delivers excellent results, the IIRs provide a lower average error for equivalent step responses, according to the introduced metrics.

The analysis of the conversion accuracy reported that the instrument suffers from degradation of characteristics for currents above 10 μA. In fact, the full-scale of 1 mA is actually measured as ≈ 829 μA, the 17.1 % less, not complying with the specifications. Nevertheless, in this range, the response is monotonic and the precision achieved close to the full scale is still high enough ($\epsilon_r < 0.6$ %), even if direct raw data are considered. Hence, a function has been introduced to correct the data and improve the accuracy. It may be adopted to emphasise the instrument response and reduce the systematic deviation down to 1.8 %. Although the linearity is lost above 1.38 mA, the measurements reach a plateau for a decade above the full scale. So, at least up to 10 mA, no ambiguities can arise from injected currents out of the nominal operating range.

In a preliminary development stage, the BLMASIC converters were subjected to X-Ray radiation at CERN for a TID of 1 MGy, not reporting any issues affecting its digital and analogue circuit behaviours. Later, the complete front-end board was exposed to a proton beam (200 MeV) irradiation at the PIF experimental area of PSI, reaching the TID of ≈ 3.3 kGy. During the test, its power consumption, the drift of the current measurements, the quality of the data transmission and the occurrence of SEEs were constantly monitored. The only remarkable anomaly was represented by a case of a single event latch-up, which did not lead to further detectable consequences. On the other hand, the maximal current measurement drift never exceeded the 0.5 %, which constitutes an outstanding achievement according to the targets. Such features will finally allow the BLM system to comply with the desired accuracies in areas with exceptionally high TID, thanks to a localised installation of the front-ends. This will be crucial not only for LHC and its upgraded version HL-LHC, but also for the SPS accelerator, where the performance of the existing system is limited.

Based on the above results, it is possible to claim that the developed system prototype is sufficiently mature to be considered the best candidate for the commissioning of the BLM upgrade in view of HL-LHC. Larger scale production of the BLMASIC and the front-end is currently under discussion, while the readout electronics and the optical links have already been established. In this perspective, the installation of the discussed devices will soon be realised on the current accelerator plant. Thus, the digitisers will be calibrated by exploiting the real beam loss measurements. Furthermore, such an activity will facilitate the know-how transfer to all the project maintainers.

The upcoming period activities will also involve: 1) the migration of the signal processing into the FPGA logic, 2) the integration of the legacy features into the back-end firmware, such as the thresholding module, 3) the completion of the diagnostic blocks and 4) the development of the system user interface. Concerning the hardware, a few remarks are in order to improve the accuracy of the digitisers. First, as explained in Par. 6.4.9, the voltage which sets the CFC quantiser thresholds should be provided by a stable reference. For instance, a secondary linPOL12V regulator can be employed, instead of directly tying the device reference pins to the power supply rail. Second, the anomaly observed in the CFC waveform (cf. Par. 4.8.7) should be investigated in more depth, before proceeding with its possible employment in the accelerator, or tackling the problem by designing a new BLMASIC version. Currently, the most accepted opinion about this phenomenon concerns the existence of parasitic effects within the integrated circuit. Hence, it is suggested to repeat the simulations by considering a larger class of causes, eventually correcting the device layout to minimise their influence. Since an upgrade of the front-end hosting crate is planned alongside the BLM electronic enhancement, a final observation involves the possibility of providing direct access to the I^2C buses from the back-panel connector. This will allow the complete board configuration and diagnostics by remote, even if both the optical links are down.

When all these issues are overcome, both the BLEIC front-end board and the hosted BLMASIC converters will be manufactured in large volumes, to optimise the costs and to conclude favourable contracts with the microelectronic foundries. Thus, the provision of components will probably be vast enough to cover not only the needs of the CERN upgrades and the maintenance of the future system, but also those of external projects requiring the same kind of technology.

Appendices

Appendix A

Radiation Effects in Electronics

The electronic circuits can be disturbed by radiation-induced effects. These must be considered when designing mission critical systems requiring high reliability, long operation times and exposure to radiations, like in the case of high-energy physics instrumentations.

The protons circulating in LHC beams can escape from the closed orbit, leaving the boundaries of the accelerator and being lost into local materials. There are several beam loss scenarios:

- Intentional beam losses, due to the collision of beams in experiments;
- Direct beam losses, when beams strike collimation systems or absorbers;
- Beam / gas interactions, when circulating beams hit residual gas.

Considering a basic MOSFET device as test target, it is possible to classify the effects of beam loss radiation into three principal types:

- *Total Ionising Dose* (TID),
- *Single Event Effects* (SEE),
- *Displacement Damage* (DD).

A brief description of the phenomena involved is reported in the following paragraphs.

A.1 Total Ionising Dose (TID)

An incident particle can cause direct or indirect ionisation of the semiconductor. Indirect ionisation occurs when a primary particle causes nuclear reactions that then lead to ionisation. These sequence of events may occur:

1. a charged particle passes through a semiconductor material and interacts with the semiconductor;
2. electron–hole pairs are generated along the path of the particle because of the energy that it released;

3. generated holes remain trapped in the integrated circuit oxides (such as passivation, gate oxides, etc.).

In both cases, the holes created within the oxides cause changes to the device characteristics. The absorbed, or total ionising, dose is a cumulative radiation effect measured in Grays [Gy]. This is integrated over time, causing gradual degradation of the semiconductor performance.

All active components are potentially susceptible to TID and several adverse effects are due to it. For example, if gate oxides collect positive charges, the N-MOS semiconductors suffer from a switch-on voltage reduction as the gate is progressively activated. On the contrary, the P-MOS devices exhibit an increase in the switch-on voltage, as the positive charges progressively inhibit the channel establishment. Furthermore, changes in timing margins and increases in the leakage currents and the total power consumption are induced. Eventually, N-MOS devices will fail so that they are permanently activated, and P-MOS devices are permanently inactivated.

A.2 Single Event Effects (SEE)

A single event effect is instantaneous and it has a certain probability (cross-section) of occurring at every particle interaction. In almost all cases, the root consequence of a SEE is a current pulse. For instance:

1. a charged particle passes through a semiconductor and interacts;
2. electrons are generated by the particle as it traverses the material;
3. electrons are highly mobile, they flow through the MOSFET and are collected at the reverse biased junction;
4. these electrons create a current pulse shortly after the interaction.

The impacts and outcomes on the circuit lead to two typical soft failures:

- Single Event Transient (SET): transient created on the signal path. One potentially significant impact concerns the consequent violation of set-up and hold timings.
- Single Event Upset (SEU): the initial particle interaction causes a stored bit to change in value. This erroneous value generally persists until it is re-written or the system is reset. The effect of an SEU on the system level depends upon the location and the function of the flipped bit.

Furthermore, in complementary metal oxide semiconductor (CMOS) circuitry, if a particle interaction deposits sufficient charge, then the parasitic structure can cause hard failures, such as:

- Single Event Latch-Up (SEL): a low-impedance short-circuit is created between the supply rails and the ground. This requires the power cycle of the device for the latch-up to be cleared. An SEL can be catastrophic if it is not mitigated, due to localised heating of the component in the high current state.
- Other hard failures, such as single event burnout (SEB) and single event gate-rupture (SEGR).

A.3 Displacement Damage (DD)

Displacement damages occur when particles interact with the silicon lattice. The case example can be described as follow:

1. an ideal silicon lattice has regularly spaced atoms;
2. incident particles traverse the lattice;
3. the particle could probably strike and dislodge an atom;
4. this interaction can create Frenkel defects, which consist of a vacancy and an interstitial defect.

Damage to the lattice is proportional to the integrated flux of particles, therefore, DD gradually increases over time. The effects of DD generally reduce the device gain, which eventually leads to the complete failure of the device. DD mainly affects bipolar technologies, while CMOS ones are generally less susceptible. The DD is a cumulative effect and is measured as the total number of 1 MeV equivalent neutrons that have passed through a cm^2 surface area. The term non-ionising energy loss (NEIL) is also used in the subject to estimate the DD.

Appendix B

Ionisation Chamber Assembly

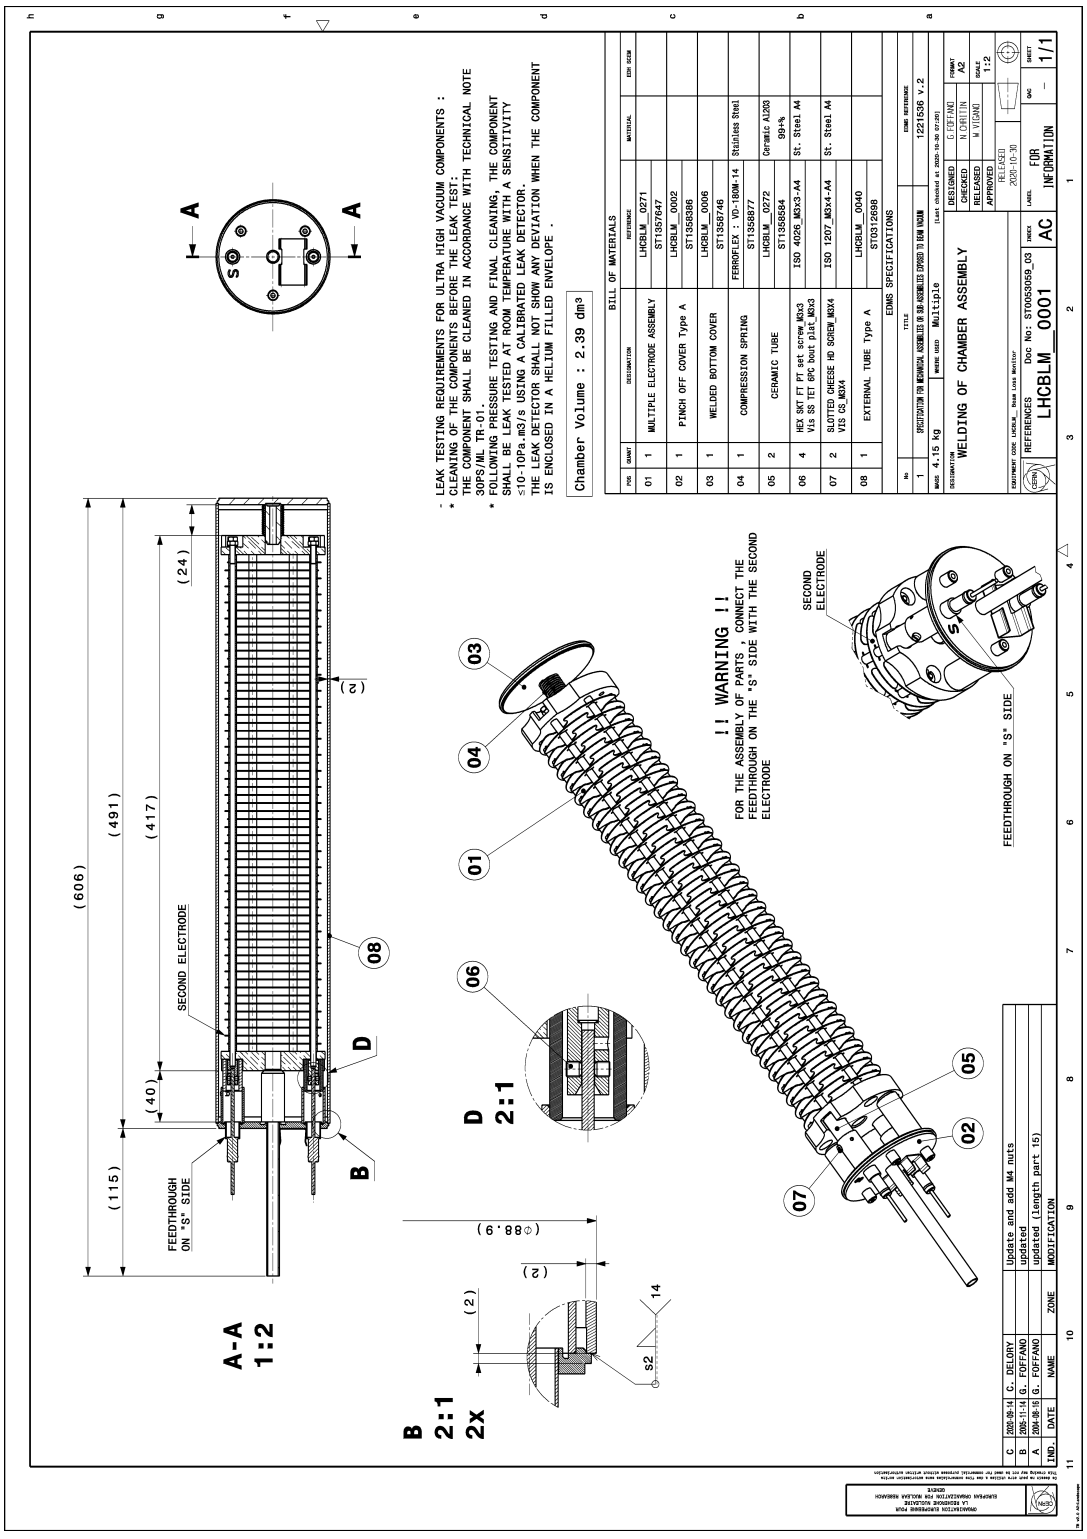
- LEAK TESTING REQUIREMENTS FOR ULTRA HIGH VACUUM COMPONENTS :
* CLEANING OF THE COMPONENTS BEFORE THE LEAK TEST:
THE COMPONENT SHALL BE CLEANED IN ACCORDANCE WITH TECHNICAL NOTE
30PS/ML TR-01.
* FOLLOWING PRESSURE TESTING AND FINAL CLEANING, THE COMPONENT
SHALL BE LEAK TESTED AT ROOM TEMPERATURE WITH A SENSITIVITY
≤10-10Pa.m3/s USING A CALIBRATED LEAK DETECTOR.
THE LEAK DETECTOR SHALL NOT SHOW ANY DEVIATION WHEN THE COMPONENT
IS ENCLOSED IN A HELIUM FILLED ENVELOPE .
Chamber Volume : 2.39 dm3
BILL OF MATERIALS
WELDING OF CHAMBER ASSEMBLY
LHCBLM_0001
!! WARNING !!
FOR THE ASSEMBLY OF PARTS , CONNECT THE
FEEDTHROUGH ON THE "S" SIDE WITH THE SECOND
ELECTRODE
SECOND ELECTRODE
FEEDTHROUGH ON "S" SIDE
FEEDTHROUGH ON "S" SIDE
SECOND ELECTRODE
A-A
1:2
B
2:1
2x
D
2:1
(606)
(491)
(417)
(115)
(40)
(24)
(2)
(ϕ88.9)
S2
14

Appendix C

BLMASIC Registers

The BLMASIC memory map is reported in Tab.s C.1 to C.3 and the configuration of the digital and the analogue probe registers is after detailed in Tab.s C.4 to C.6.

It should be noticed that, concerning the analogue probes, the specific monitoring signal (centre column) can be connected to its dedicated output pin by setting the corresponding bit. Moreover, it is neither recommended nor useful to enable more than one signal pointing to the same pin, even if they belong to different channels. For instance, the differential probing of the CH0 integrator output is obtained by configuring: $CFCFE\_CH0\_TESTBITe = 0x03$ and $CFCFE\_CH1\_TESTBITe = 0x00$. While, for the integrator of CH1, it implies: $CFCFE\_CH0\_TESTBITe = 0x00$ and $CFCFE\_CH1\_TESTBITe = 0x03$. Finally, it can be remarked that the polarity of differential probes is eventually inverted between the two channels.

Register ADDR	Memory Bit	Register Name	Read/ Write	Comment	Default Binary Value	Default Decimal Value
0	0	BLM_SLVSTXBIAS	R/W	Transceivers current. 0: 1.4mA; 1: 1.7mA; 2: 2.2mA; 3: 3mA	0	2
	1				1	
	2	blm_reset	R/W	Full Reset bit 1/3 (Active High)	0	0
	3	blm_FEreset	R/W	Analogue Circuit Reset bit 1/3 (Active High)	0	0
	4	testClockOutSel[1:0]	R/W	Selects the digital test output probe	1	3
	5				1	
	6	blmA_bypass8b10b	R/W	bypasses the 8b10b encoding for Ch.A (Active High) bit 1/3	0	0
	7	blmB_bypass8b10b	R/W	bypasses the 8b10b encoding for Ch.B (Active High) bit 1/3	0	0
1	8	testClockOutSel[3:2]	R/W	Selects the digital test output probe	0	2
	9				1	
	10	blm_reset_bar	R/W	Full Reset bit 2/3 (Active Low)	1	1
	11	blm_Fereset_bar	R/W	Analogue Circuit Reset bit 2/3 (Active Low)	1	1
	12	testClockOutSel[5:4]	R/W	Selects the digital test output probe	1	1
	13				0	
	14	blmA_bypass8b10b_bar	R/W	bypasses the 8b10b encoding for Ch.A (Active Low) bit 2/3	1	1
	15	blmB_bypass8b10b_bar	R/W	bypasses the 8b10b encoding for Ch.B (Active Low) bit 2/3	1	1
2	16	testClockOutSel[6]	R/W	Selects the digital test clock output	0	0
	17	----------	R	Empty register. Not used. Contains zero.	0	0
	18	blm_reset	R/W	Full Reset bit 3/3 (Active High)	0	0
	19	blm_FEreset	R/W	Analogue Circuit Reset bit 3/3 (Active High)	0	0
	20	----------	R	Empty register. Not used. Contains zero.	0	0
	21	----------	R	Empty register. Not used. Contains zero.	0	0
	22	blmA_bypass8b10b	R/W	bypasses the 8b10b encoding for Ch.A (Active High) bit 3/3	0	0
	23	blmB_bypass8b10b	R/W	bypasses the 8b10b encoding for Ch.B (Active High) bit 3/3	0	0
3	24	CFCFE_CH0_PLEAKBe[9:0]	R/W	Leakage compensation DAC (positive current injection). Channel 0.	0	200
	25				0	
	26				0	
	27				1	
	28				0	
	29				0	
	30				1	
	31				1	
4	32				0	
	33				0	
	34	CFCFE_CH0_NLEAKBe[7:0]	R/W	Leakage compensation DAC negative current injection). Channel 0.	0	0
	35				0	
	36				0	
	37				0	
	38				0	
	39				0	
5	40				0	
	41				0	
	42	CFCFE_CH0_TRIM20Ue[5:0]	R/W	Adjustment of main bias generator. Channel 0.	0	38
	43				1	
	44				1	
	45				0	
	46				0	
	47				1	
6	48	CFCFE_CH0_TRIM60UBGAe[2:0]	R/W	Adjustment of the bias for Band Gap amplifier (self reference current source). Channel 0.	0	4
	49				0	
	50				1	
	51	CFCFE_CH0_WADCRBe[4:0]	R/W	Adjustment of the bias for the WILKINSON ADC range. Channel 0.	1	15
	52				1	
	53				1	
	54				1	
	55				0	
7	56	CFCFE_CH0_CFCCOMPBe[2:0]	R/W	Adjustment of the bias for the CFC comparators. Channel 0.	0	4
	57				0	
	58				1	
	59	CFCFE_CH0_WADCCOMPBe[2:0]	R/W	Adjustment of the bias for the WILKINSON ADC comparator. Channel 0.	0	4
	60				0	
	61				1	
	62	CFCFE_CH0_TESTBITe[11:0]	R/W	Control register for the analogue probe output multiplexer. Channel 0.	0	0
	63				0	
8	64				0	
	65				0	
	66				0	
	67				0	
	68				0	
	69				0	
	70				0	
	71				0	
9	72				0	
	73				0	
	74	CFCFE_CH1_PLEAKBe[5:0]	R/W	Leakage compensation DAC (positive current injection). Channel 1.	0	200
	75				0	
	76				0	
	77				1	
	78				0	
	79				0	

Figure C.1: BLMASIC register map, addresses from 0 (0x00) to 9 (0x09).

10	80	CFCFE_CH1_PLEAKBe[9:6]	R/W	Leakage compensation DAC (positive current injection). Channel 1.	1	200
	81				1	
	82				0	
	83				0	
	84	CFCFE_CH1_NLEAKBe[7:0]	R/W	Leakage compensation DAC negative current injection). Channel 1.	0	0
	85				0	
	86				0	
	87				0	
11	88				0	
	89				0	
	90				0	
	91				0	
	92	CFCFE_CH1_TRIM20Ue[5:0]	R/W	Adjustment of the main bias generator. Channel 1.	0	38
	93				1	
	94				1	
	95				0	
12	96				0	
	97				1	
	98	CFCFE_CH1_TRIM60UBGAe[2:0]	R/W	Adjustment of the bias for Band Gap amplifier (self reference current source). Channel 1.	0	4
	99				0	
	100				1	
	101	CFCFE_CH1_WADCRBe[4:0]	R/W	Adjustment of the bias for the WILKINSON ADC range. Channel 1.	1	15
	102				1	
	103				1	
13	104				1	
	105				0	
	106	CFCFE_CH1_CFCCOMPBe[2:0]	R/W	Adjustment of the bias for the CFC comparators. Channel 1.	0	4
	107				0	
	108				1	
	109	CFCFE_CH1_WADCCOMPBe[2:0]	R/W	Adjustment of the bias for the WILKINSON ADC comparator. Channel 1.	0	4
	110				0	
	111				1	
14	112	CFCFE_CH1_TESTBITe[11:0]	R/W	Control register for the analogue probe output multiplexer. Channel 1.	0	0
	113				0	
	114				0	
	115				0	
	116				0	
	117				0	
	118				0	
	119				0	
15	120				0	
	121				0	
	122				0	
	123				0	
	124	----------	R	Empty register. Not used. Contains zero.	0	0
	125	----------	R	Empty register. Not used. Contains zero.	0	0
	126	----------	R	Empty register. Not used. Contains zero.	0	0
	127	----------	R	Empty register. Not used. Contains zero.	0	0
16	128	----------	R	Empty register. Not used. Contains zero.	0	0
	129	----------	R	Empty register. Not used. Contains zero.	0	0
	130	----------	R	Empty register. Not used. Contains zero.	0	0
	131	----------	R	Empty register. Not used. Contains zero.	0	0
	132	----------	R	Empty register. Not used. Contains zero.	0	0
	133	----------	R	Empty register. Not used. Contains zero.	0	0
	134	----------	R	Empty register. Not used. Contains zero.	0	0
	135	----------	R	Empty register. Not used. Contains zero.	0	0
17	136	----------	R	Empty register. Not used. Contains zero.	0	0
	137	----------	R	Empty register. Not used. Contains zero.	0	0
	138	----------	R	Empty register. Not used. Contains zero.	0	0
	139	----------	R	Empty register. Not used. Contains zero.	0	0
	140	----------	R	Empty register. Not used. Contains zero.	0	0
	141	----------	R	Empty register. Not used. Contains zero.	0	0
	142	----------	R	Empty register. Not used. Contains zero.	0	0
	143	----------	R	Empty register. Not used. Contains zero.	0	0
18	144	----------	R	Empty register. Not used. Contains zero.	0	0
	145	----------	R	Empty register. Not used. Contains zero.	0	0
	146	----------	R	Empty register. Not used. Contains zero.	0	0
	147	----------	R	Empty register. Not used. Contains zero.	0	0
	148	----------	R	Empty register. Not used. Contains zero.	0	0
	149	----------	R	Empty register. Not used. Contains zero.	0	0
	150	----------	R	Empty register. Not used. Contains zero.	0	0
	151	----------	R	Empty register. Not used. Contains zero.	0	0
19	152	----------	R	Empty register. Not used. Contains zero.	0	0
	153	----------	R	Empty register. Not used. Contains zero.	0	0
	154	----------	R	Empty register. Not used. Contains zero.	0	0
	155	----------	R	Empty register. Not used. Contains zero.	0	0
	156	----------	R	Empty register. Not used. Contains zero.	0	0
	157	----------	R	Empty register. Not used. Contains zero.	0	0
	158	----------	R	Empty register. Not used. Contains zero.	0	0
	159	----------	R	Empty register. Not used. Contains zero.	0	0

Figure C.2: BLMASIC register map, addresses from 10 (0x0A) to 19 (0x13).

20	160	blmLockConfidence[6:0]	R	Confidence that the selected lpGBT is locked; max is good, mid range is locking/unlocking, zero is no lock	0	0
	161				0	
	162				0	
	163				0	
	164				0	
	165				0	
	166				0	
	167	BLM_blmSel	R/W	Select lpGBT_A: 0; lpGBT_B 1;	0	0
21	168	BLM_B_lpgbt_ready	R	lpGBT_B is ready	0	0
	169	BLM_A_lpgbt_ready	R	lpGBT_A is ready	0	0
	170	BLM_VERSONID[3:0]	R	Version identifier	X	FE Dependent
	171				0	
	172				1	
	173				0	
	174	----------	R	Empty register. Not used. Contains zero.	0	0
	175	----------	R	Empty register. Not used. Contains zero.	0	0
22	176	CFCFE_CH0_CFC_COUNTER_NEGEDGE[9:0]	R	CFC counter; number of negative edges in 800x80MHz cycles (resetless and stopless counter). Channel 0.	Measurements	
	177					
	178					
	179					
	180					
	181					
	182					
	183					
23	184	CFCFE_CH0_CFC_COUNTER_POSEDGE[9:0]	R	CFC counter; number of positive edges in 800x80MHz cycles (resetless and stopless counter). Channel 0.	Measurements	
	185					
	186					
	187					
	188					
	189					
	190					
	191					
24	192	CFCFE_CH1_CFC_COUNTER_NEGEDGE[9:0]	R	CFC counter; number of negative edges in 800x80MHz cycles (resetless and stopless counter). Channel 1.	Measurements	
	193					
	194					
	195					
	196					
	197					
	198					
	199					
25	200	CFCFE_CH1_CFC_COUNTER_POSEDGE[9:0]	R	CFC counter; number of positive edges in 800x80MHz cycles (resetless and stopless counter). Channel 1.	Measurements	
	201					
	202					
	203					
	204					
	205					
	206					
	207					
26	208	CFCFE_CH0_WILKINSON_10bCOUNTER[9:0]	R	Wilkinson counter; number of 80MHz counts between PHI3 and STOP (timeout is 700 counts). Channel 0.	Measurements	
	209					
	210					
	211					
	212					
	213					
	214					
	215					
27	216	CFCFE_CH1_WILKINSON_10bCOUNTER[9:0]	R	Wilkinson counter; number of 80MHz counts between PHI3 and STOP (timeout is 700 counts). Channel 1.	Measurements	
	217					
	218					
	219					
	220					
	221					
	222					
	223					
28	224	CFCFE_CH0_SIGN	R	sign from the first conversion	Measurements	
	225					
	226					
	227					
	228					
	229					
	230					
	231					
29	232	CFCFE_CH1_SIGN	R	sign from the first conversion		
	233					
	234					
	235					
	236	----------	R	Empty register. Not used. Contains zero.	0	0
	237	----------	R	Empty register. Not used. Contains zero.	0	0
	238					
	239					

Figure C.3: BLMASIC register map, addresses from 20 (0x14) to 29 (0x1D).

testClockOutSel[6:0]	Selected Digital Probe Output
0	1'b0
1	1'b1
2	blmSmClock
3	blmSel
4	BLM_B_lpgbt_ready
5	BLM_A_lpgbt_ready
6	BLM_resetB
7	CFCFE_CH0_COUNTER
8	CFCFE_CH0_SIGN
9	CFCFE_CH0_STOP
10	CFCFE_CH0_PHI3e
11	CFCFE_CH0_PHI2e
12	CFCFE_CH0_PHI1e
13	CFCFE_CH0_INITe_digital
14	CFCFE_CH0_INITe_analog
15	CFCFE_CH1_COUNTER
16	CFCFE_CH1_SIGN
17	CFCFE_CH1_STOP
18	CFCFE_CH1_PHI3e
19	CFCFE_CH1_PHI2e
20	CFCFE_CH1_PHI1e
21	CFCFE_CH1_INITe_digital
22	CFCFE_CH1_INITe_analog
23	SDFE_CH0_clk_to_test
24	SDFE_CH0_bit_to_test
25	SDFE_CH1_clk_to_test
26	SDFE_CH1_bit_to_test
27	&blmLockConfidence
28	blmLockConfidence[0]
29	blmLockConfidence[6] & blmLockConfidence[5]
30	CFCFE_CH0_SMSTATE[2:0]
31	
32	
33	CFCFE_CH1_SMSTATE[2:0]
34	
35	
36	----------
37	elinkclk_80MHz
38	blmA_elinkclk_80MHz
39	blmB_elinkclk_80MHz
40	blmA_elinkdataout_80Mbps
41	blmB_elinkdataout_80Mbps
42	clk8b10b
43	blmA_clk8b10b
44	blmB_clk8b10b
45	FEcontrol_sendMeas
46	blmA_FEcontrol_sendMeas
47	blmB_FEcontrol_sendMeas
48	blmA_sendMeas_inst_upsy
49	blmB_sendMeas_inst_upsy
50	FEcontrol_measCounter[0]
51	FEcontrol_measCounter[1]
52	FEcontrol_measCounter[2]
53	FEcontrol_measCounter[3]
54	FEcontrol_measCounter[4]
55	FEcontrol_measCounter[5]
56	FEcontrol_measCounter[6]
57	FEcontrol_measCounter[7]
58	blmA_rctxBusy
59	blmA_FEcontrol_sendMeas & blmA_rctxBusy
60	blmA_txsmState[0]
61	blmA_txsmState[1]
62	blmA_txsmState[2]
63	blmA_txsmState[3]

testClockOutSel[6:0]	Selected Digital Probe Output
64	blmB_rctxBusy
65	blmB_FEcontrol_sendMeas & blmB_rctxBusy
66	blmB_txsmState[0]
67	blmB_txsmState[1]
68	blmB_txsmState[2]
69	blmB_txsmState[3]
70	blmA_data8b10bEncoded_fortx[0]
71	blmA_data8b10bEncoded_fortx[1]
72	blmA_data8b10bEncoded_fortx[2]
73	blmA_data8b10bEncoded_fortx[3]
74	blmA_data8b10bEncoded_fortx[4]
75	blmA_data8b10bEncoded_fortx[5]
76	blmA_data8b10bEncoded_fortx[6]
77	blmA_data8b10bEncoded_fortx[7]
78	blmA_data8b10bEncoded_fortx[8]
79	blmA_data8b10bEncoded_fortx[9]
80	blmA_encoder_disparity
81	blmB_data8b10bEncoded_fortx[0]
82	blmB_data8b10bEncoded_fortx[1]
83	blmB_data8b10bEncoded_fortx[2]
84	blmB_data8b10bEncoded_fortx[3]
85	blmB_data8b10bEncoded_fortx[4]
86	blmB_data8b10bEncoded_fortx[5]
87	blmB_data8b10bEncoded_fortx[6]
88	blmB_data8b10bEncoded_fortx[7]
89	blmB_data8b10bEncoded_fortx[8]
90	blmB_data8b10bEncoded_fortx[9]
91	blmB_encoder_disparity
92	i2cActive
93	blmA_bypass8b10b
94	blmB_bypass8b10b
95	blm_reset
96	blm_FEreset
97	1'b0
98	1'b0
99	1'b0
100	1'b0
101	1'b0
102	1'b0
103	1'b0
104	1'b0
105	1'b0
106	1'b0
107	1'b0
108	1'b0
109	1'b0
110	1'b0
111	1'b0
112	1'b0
113	1'b0
114	1'b0
115	1'b0
116	1'b0
117	1'b0
118	1'b0
119	1'b0
120	1'b0
121	1'b0
122	1'b0
123	1'b0
124	1'b0
125	1'b0
126	1'b0
127	1'b0

Figure C.4: Table of correspondences amongst the *testClockOutSel* register values (left column) and the selected digital output probe.

CFCFE_CH0_TESTBITe[11:0]	Enabled Analogue Probe Output CH0	Probe Pin
0	CFC Integrator Output NEGATIVE	B_{TEST}
1	CFC Integrator Output POSITIVE	A_{TEST}
2	DC test 240mV (CFC Quantiser Bias)	A_{TEST}
3	DC test 240mV (WADC Quantiser Bias)	A_{TEST}
4	Wilkinson ADC (WADC) Output POSITIVE	A_{TEST}
5	Wilkinson ADC (WADC) Output NEGATIVE	B_{TEST}
6	DC Test 200mV (20uA General Bias)	A_{TEST}
7	DC Test 200mV (60uA Bias)	A_{TEST}
8	Resistor 10k to GND Rppolywo	A_{TEST}
9	Resistor 20k to GND Rppolywo	A_{TEST}
10	Resistor 10k to GND Hres	A_{TEST}
11	Resistor 20k to GND Hres	A_{TEST}

Figure C.5: Table of correspondences amongst the *CFCFE_CH0_TESTBITe* register values (left column) and the selected analogue probe for channel 0. The output pins are reported.

CFCFE_CH1_TESTBITe[11:0]	Enabled Analogue Probe Output CH1	Probe Pin
0	CFC Integrator Output NEGATIVE	A_{TEST}
1	CFC Integrator Output POSITIVE	B_{TEST}
2	DC test 240mV (CFC Quantiser Bias)	B_{TEST}
3	DC test 240mV (WADC Quantiser Bias)	B_{TEST}
4	Wilkinson ADC (WADC) Output POSITIVE	B_{TEST}
5	Wilkinson ADC (WADC) Output NEGATIVE	A_{TEST}
6	DC Test 200mV (20uA General Bias)	B_{TEST}
7	DC Test 200mV (60uA Bias)	B_{TEST}
8	Resistor 10k to GND Rppolywo	B_{TEST}
9	Resistor 20k to GND Rppolywo	B_{TEST}
10	Resistor 10k to GND Hres	B_{TEST}
11	Resistor 20k to GND Hres	B_{TEST}

Figure C.6: Table of correspondences amongst the *CFCFE_CH1_TESTBITe* register values (left column) and the selected analogue probe for channel 1. The output pins are reported.

Appendix D

BLMASIC Pinout

The following pages report the BLMASIC pinout list, with the numbering in the first column corresponding to the actual pins in the package.

Pin	Name	Type	Description
1	VSSPST	Power	*ESD Ground*
2	VCMA	Power / Analog Reference	*Common Mode Ch.A (0.6V)*
3	VSSA	Power	*Analog Ground*
4	VCMA	Power / Analog Reference	*Common Mode Ch.A (0.6V)*
5	INPUTA	Analog Input	*Input Channel A*
6	VCMA	Power / Analog Reference	*Common Mode Ch.A (0.6V)*
7	LDACA	Analog Output	*Calibration DAC Output for Ch.A (connect to INPUTA through an ext. 1GΩ Resistor)*
8	VDDA	Power	*Analog Supply (1.2V)*
9	VCMB	Power / Analog Reference	*Common Mode Ch.B (0.6V)*
10	VDDA	Power	*Analog Supply (1.2V)*
11	LDACB	Analog Output	*Calibration DAC Output for Ch.B (connect to INPUTB through an ext. 1GΩ Resistor)*
12	INPUTB	Analog Input	*Input Channel B*
13	VCMB	Power / Analog Reference	*Common Mode Ch.B (0.6V)*
14	VSSA	Power	*Analog Ground*
15	VCMB	Power / Analog Reference	*Common Mode Ch.B (0.6V)*
16	VSSPST	Power	*ESD Ground*
17	VCMB	Power / Analog Reference	*Common Mode Ch.B (0.6V)*
18	VDDPSTA	Power	*ESD Analog Supply (1.2V)*
19	VDDA	Power	*Analog Supply (1.2V)*
20	VSSA	Power	*Analog Ground*
21	VREFB	Analog Input Reference	*Quantiser Thresholds B Reference (1.2V)*
22	VCMWB	Power / Analog Reference	*Common Mode of Wilkinson ADC Ch.B (0.6V)*
23	V25B	Power	*Analog Supply (2.5V)*
24	BTEST	Analog Output	*Analog Probe Output B*
25	VCMWB	Power / Analog Reference	*Common Mode of Wilkinson ADC Ch.B (0.6V)*
26	VDDPST	Power	*ESD Digital Supply (1.2V)*
27	VSSPST	Power	*ESD Ground*
28	ADDR1	Digital Input, CMOS	*I²C Device Address Bit₁*
29	VDD	Power	*Digital Supply (1.2V)*
30	BSDA	Digital I/O, CMOS	*I²C B Data Signal*
31	BSCL	Digital Input, CMOS	*I²C B Clock Signal*
32	VSS	Power	*Digital Ground*

Figure D.1: BLMASIC Pinout (1 to 32).

33	BRDY	Digital Input, CMOS	*Ready Signal from lpGBT B*
34	VSS	Power	*Digital Ground*
35	VDD	Power	*Digital Supply (1.2V)*
36	BECLKP	Digital Input, CLPS 400mV CM	*80 Mhz Clock B from lpGBT Positive*
37	BECLKN	Digital Input, CLPS 400mV CM	*80 Mhz Clock B from lpGBT Negative*
38	BTXP	Digital Output, CLPS 400mV CM	*Data Bit-Stream B 80Mbps Positive*
39	BTXN	Digital Output, CLPS 400mV CM	*Data Bit-Stream B 80Mbps Negative*
40	VSS	Power	*Digital Ground*
41	VDD	Power	*Digital Supply (1.2V)*
42	ATXN	Digital Output, CLPS 400mV CM	*Data Bit-Stream A 80Mbps Negative*
43	ATXP	Digital Output, CLPS 400mV CM	*Data Bit-Stream A 80Mbps Positive*
44	AECLKN	Digital Input, CLPS 400mV CM	*80 Mhz Clock A from lpGBT Negative*
45	AECLKP	Digital Input, CLPS 400mV CM	*80 Mhz Clock A from lpGBT Positive*
46	VSS	Power	*Digital Ground*
47	VDD	Power	*Digital Supply (1.2V)*
48	ARDY	Digital Input, CMOS	*Ready Signal from lpGBT A*
49	RSTB	Digital Input, CMOS	*General Reset*
50	ASCL	Digital Input, CMOS	*I²C A Clock Signal*
51	ASDA	Digital I/O, CMOS	*I²C A Data Signal*
52	DTSTOUT	Digital Output, CMOS	*Digital Output Probe*
53	ADDR0	Digital Input, CMOS	*I²C Device Address Bit$_0$*
54	VSSPST	Power	*ESD Ground*
55	VDDPST	Power	*ESD Digital Supply (1.2V)*
56	VCMWA	Power / Analog Reference	*Common Mode of Wilkinson ADC Ch.A (0.6V)*
57	ATEST	Analog Output	*Analog Probe Output A*
58	V25A	Power	*Analog Supply (2.5V)*
59	VCMWA	Power / Analog Reference	*Common Mode of Wilkinson ADC Ch.A (0.6V)*
60	VREFA	Analog Input Reference	*Quantiser Thresholds A Reference (1.2V)*
61	VSSA	Power	*Analog Ground*
62	VDDA	Power	*Analog Supply (1.2V)*
63	VDDPSTA	Power	*ESD Analog Supply (1.2V)*
64	VCMA	Power / Analog Reference	*Common Mode Ch.A (0.6V)*

Figure D.2: BLMASIC Pinout (33 to 64).

Appendix E

Front-End Schematics

The front-end board schematics are reported in the following pages, according to the description presented in Chapter 6.

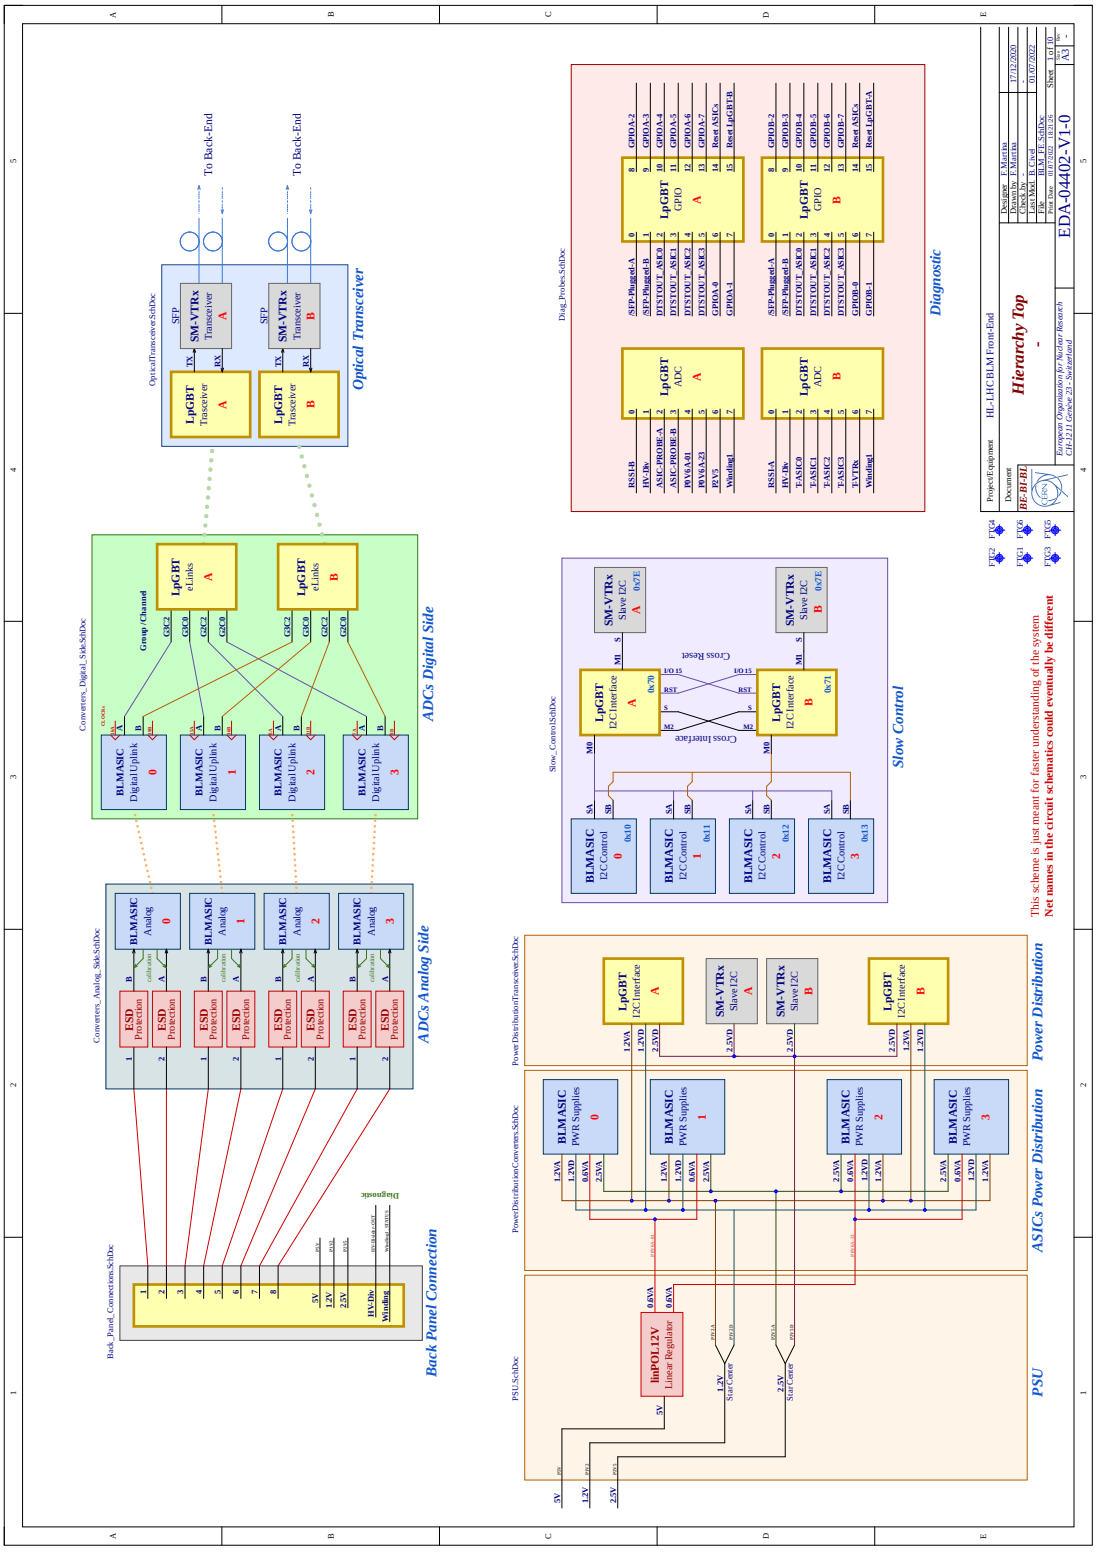
Optical Transceiver
ADCs Digital Side
ADCs Analog Side
Back Panel Connection
Diagnostic
Slow Control
ASICs Power Distribution
Power Distribution
PSU
This scheme is just meant for faster understanding of the system
Net names in the circuit schematics could eventually be different
Hierarchy Top
EDA-04402-V1-0
HL-LHC BLM Front-End

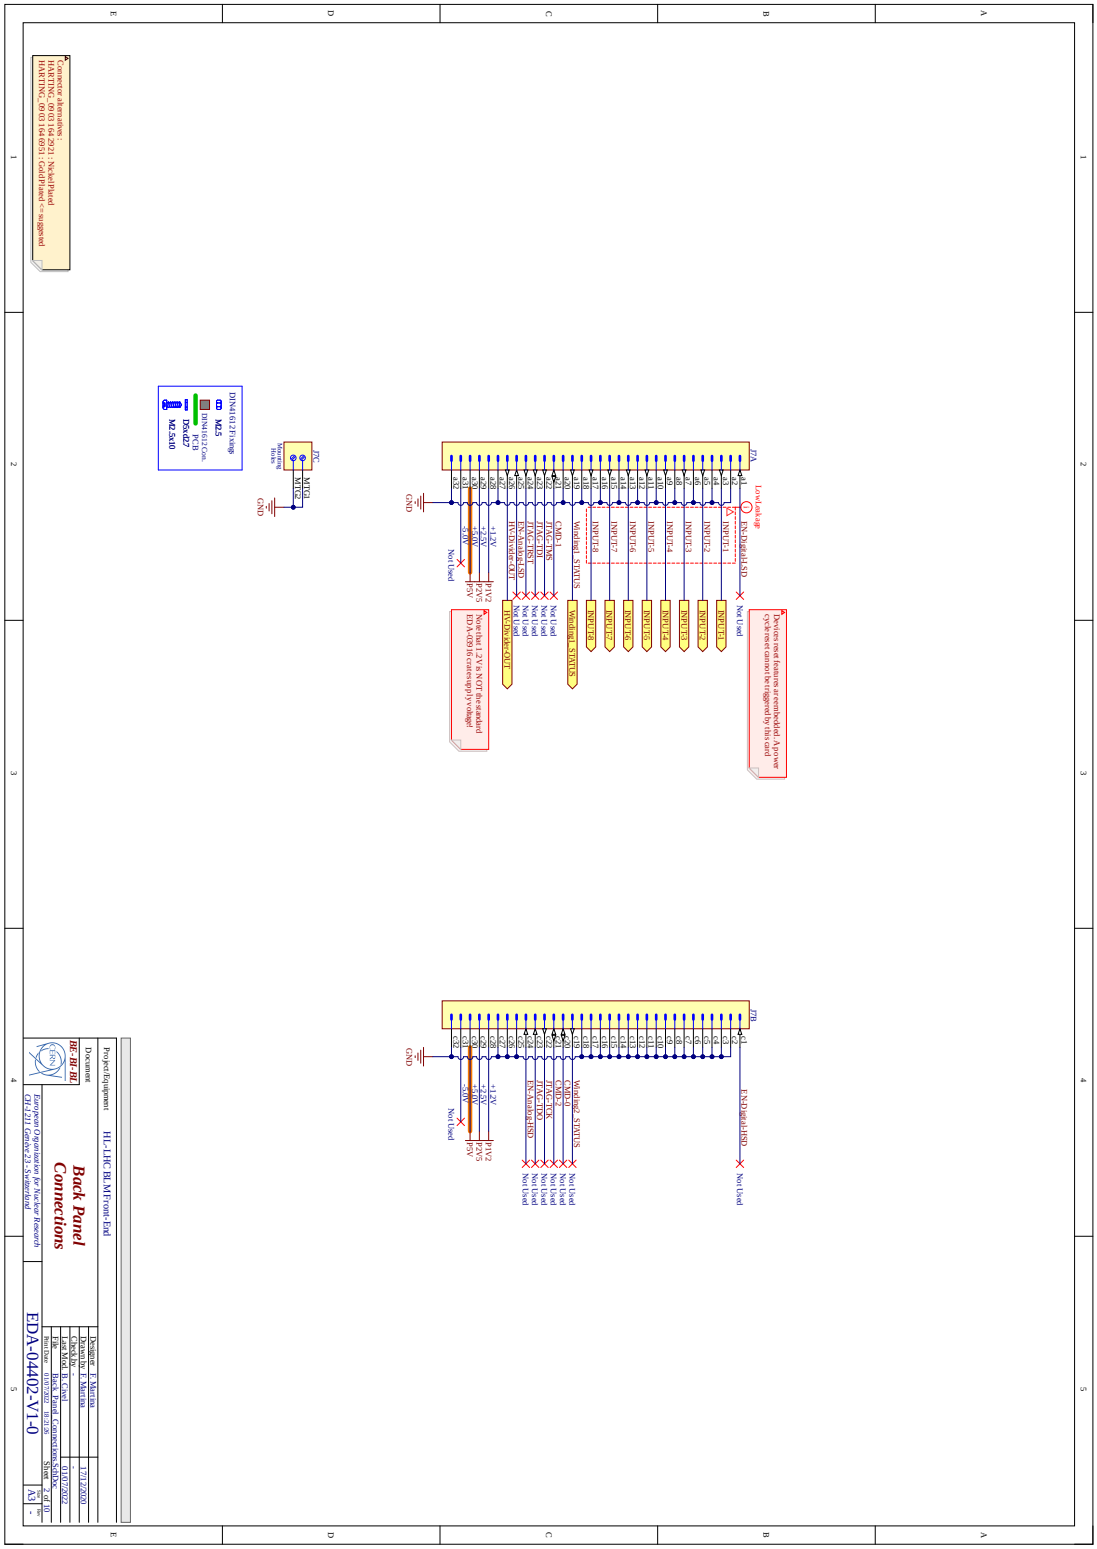
Back Panel
Connections
EDA-04402-V1-0
Project/Equipment
HL-LHC BL M Front-End
Designer E. Martins
Drawn by E. Martins
Check by B. Cisel
Date Mod. 01/07/2022
File Back_Panel_Connections.SchDoc
Sheet 2 of 10
A3
Document
BE-BI-BL
European Organisation for Nuclear Research
CH-1211 Genève 23 - Switzerland
Connector alternatives :
HARTING: 09 03 164 2921 : Nickel Plated
HARTING: 09 03 164 6951 : Gold Plated <= suggested
DIN41612 Fixings
M2.5
DIN41612 Cons.
PCB
Dsx.dZ7
M2.5x10
Mounting Holes
MTG1
MTG2
GND
J7A
J7B
Low Leakage
EN-Digital-LSD
EN-Digital-HSD
INPUT-1
INPUT-2
INPUT-3
INPUT-4
INPUT-5
INPUT-6
INPUT-7
INPUT-8
Wireline1_STATUS
Wireline2_STATUS
CMD-1
CMD-2
JTAG-TMS
JTAG-TCK
JTAG-TDI
JTAG-TDO
JTAG-TRST
EN-Analog-LSD
EN-Analog-HSD
HV-Divider-OUT
+1.2V
P1V2
+2.5V
P2V5
-5.0V
Not Used
Not that 1.2 V is NOT the standard EDA-4430 BI card supply voltage!
Not used connections are reserved for Low Leakage cycle reset cannot be triggered by this card

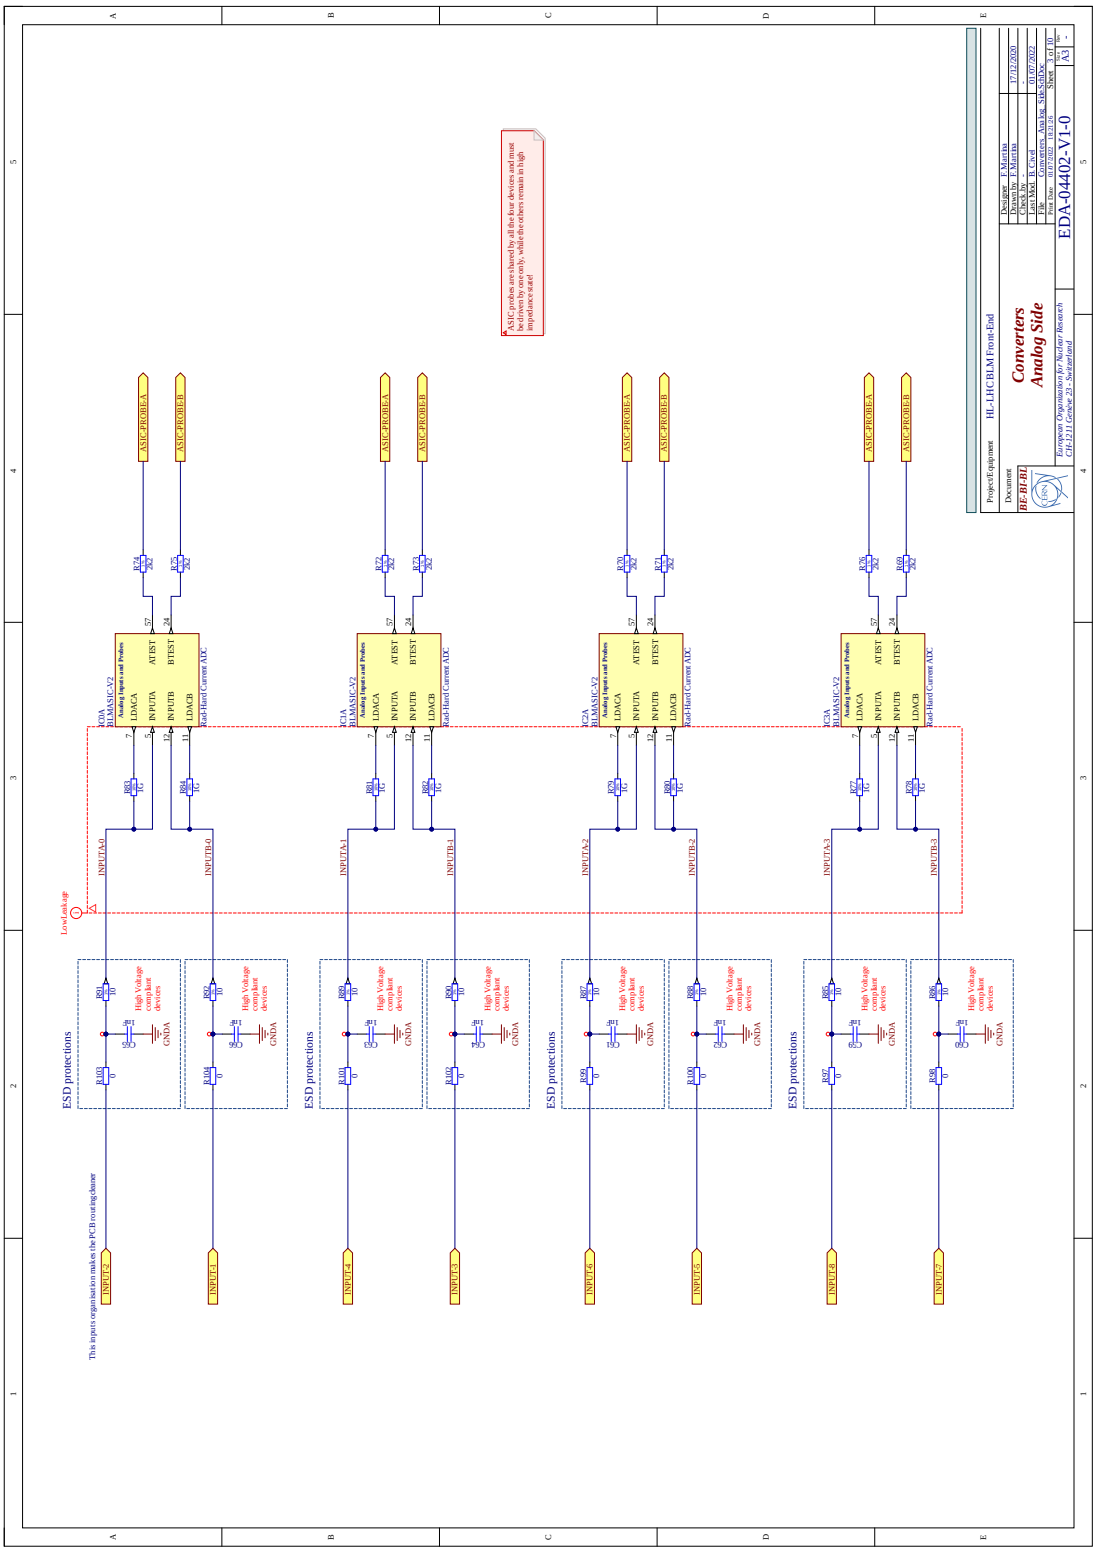
ASIC probes are shared by all the four devices and must be driven by one only, while the others remain in high impedance state!
Project/Equipment HL-LHC BLM Front-End
Document BE-BI-BL
Converters
Analog Side
European Organisation for Nuclear Research
CH-1211 Genève 23 - Switzerland
Designer F. Martina
Drawn by F. Martina
Check by -
Last Mod. B. Civel
File Converters_Analog_Side.SchDoc
Print Date 01/07/2022 18:21:26
Sheet 3 of 10
A3
EDA-04402-V1-0
ASIC-PROBE-A
ASIC-PROBE-B
ESD protections
High Voltage compliant devices
GNDA
Low Leakage
This inputs organisation makes the PCB routing denser

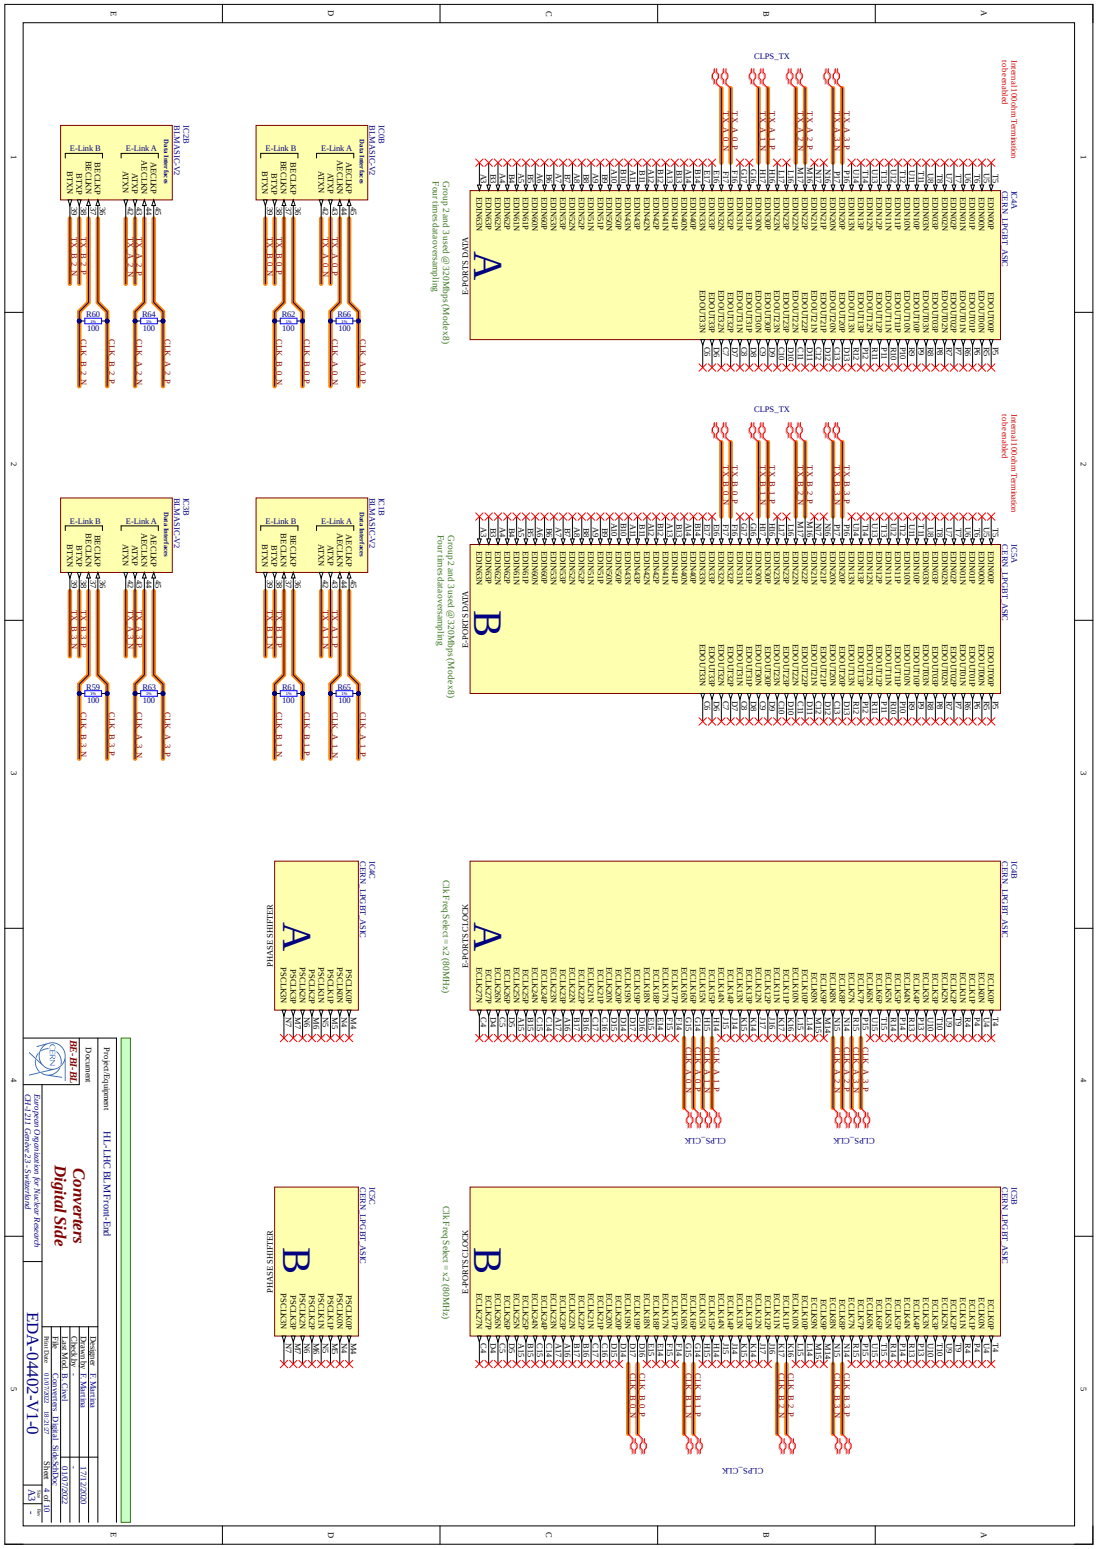
Converters
Digital Side
EDA-04402-V1-0
Project/Equipment HL-LHC BLM Front-End
Document BE-BI-BL
European Organization for Nuclear Research
CH-1211 Genève 23 - Switzerland
Designer E. Martina
Drawn by E. Martina
Check by
Last Mod. B. Civel
File Converters_Digital_Side.SchDoc
Print Date 01/07/2022 10:21:27
Sheet 4 of 10
A3
Internal 100ohm Termination to be enabled
CLPS_TX
CERN_LPGBT_ASIC
Group 2 and 3 used @ 320Mbps (Mode x8)
Four times data oversampling
Clk Freq Select = x2 (80MHz)
PHASE SHIFTER
CLPS_CLK
E-Link A
E-Link B
BLMASIC-V2

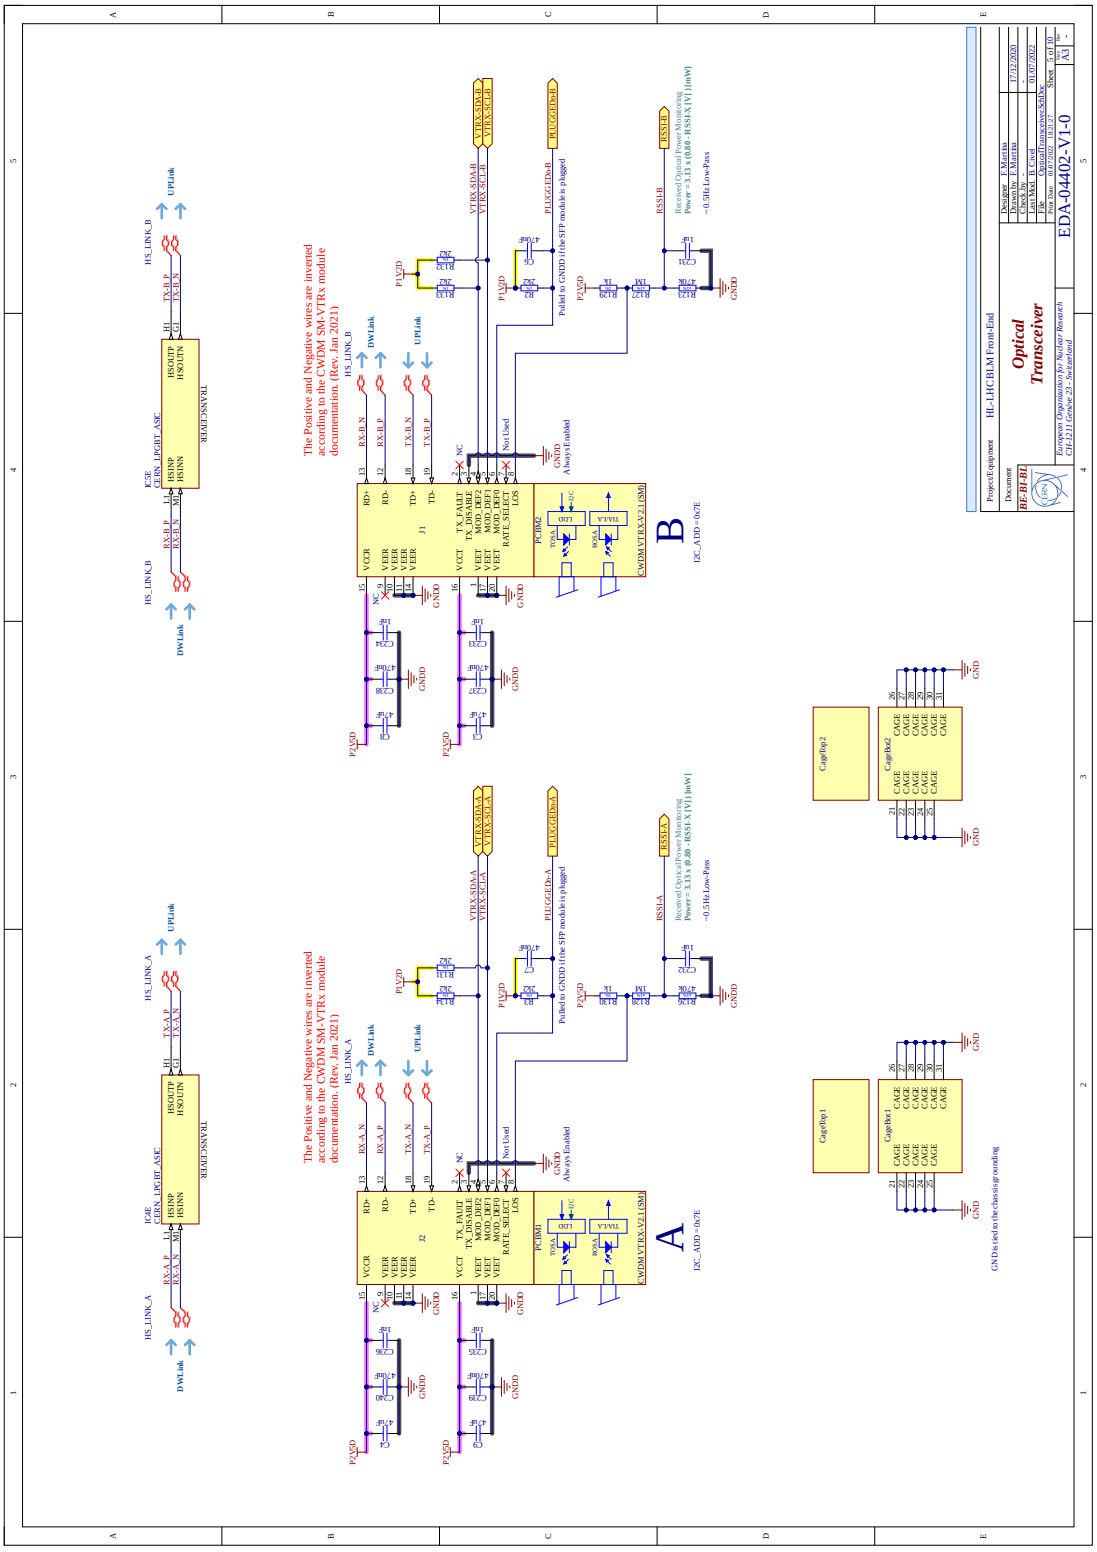
Optical Transceiver
EDA-04402-V1-0
HL-LHC BLM Front-End
The Positive and Negative wires are inverted according to the CWDM SM-VTRx module documentation. (Rev. Jan 2021)
GNDs tied to the chassis grounding

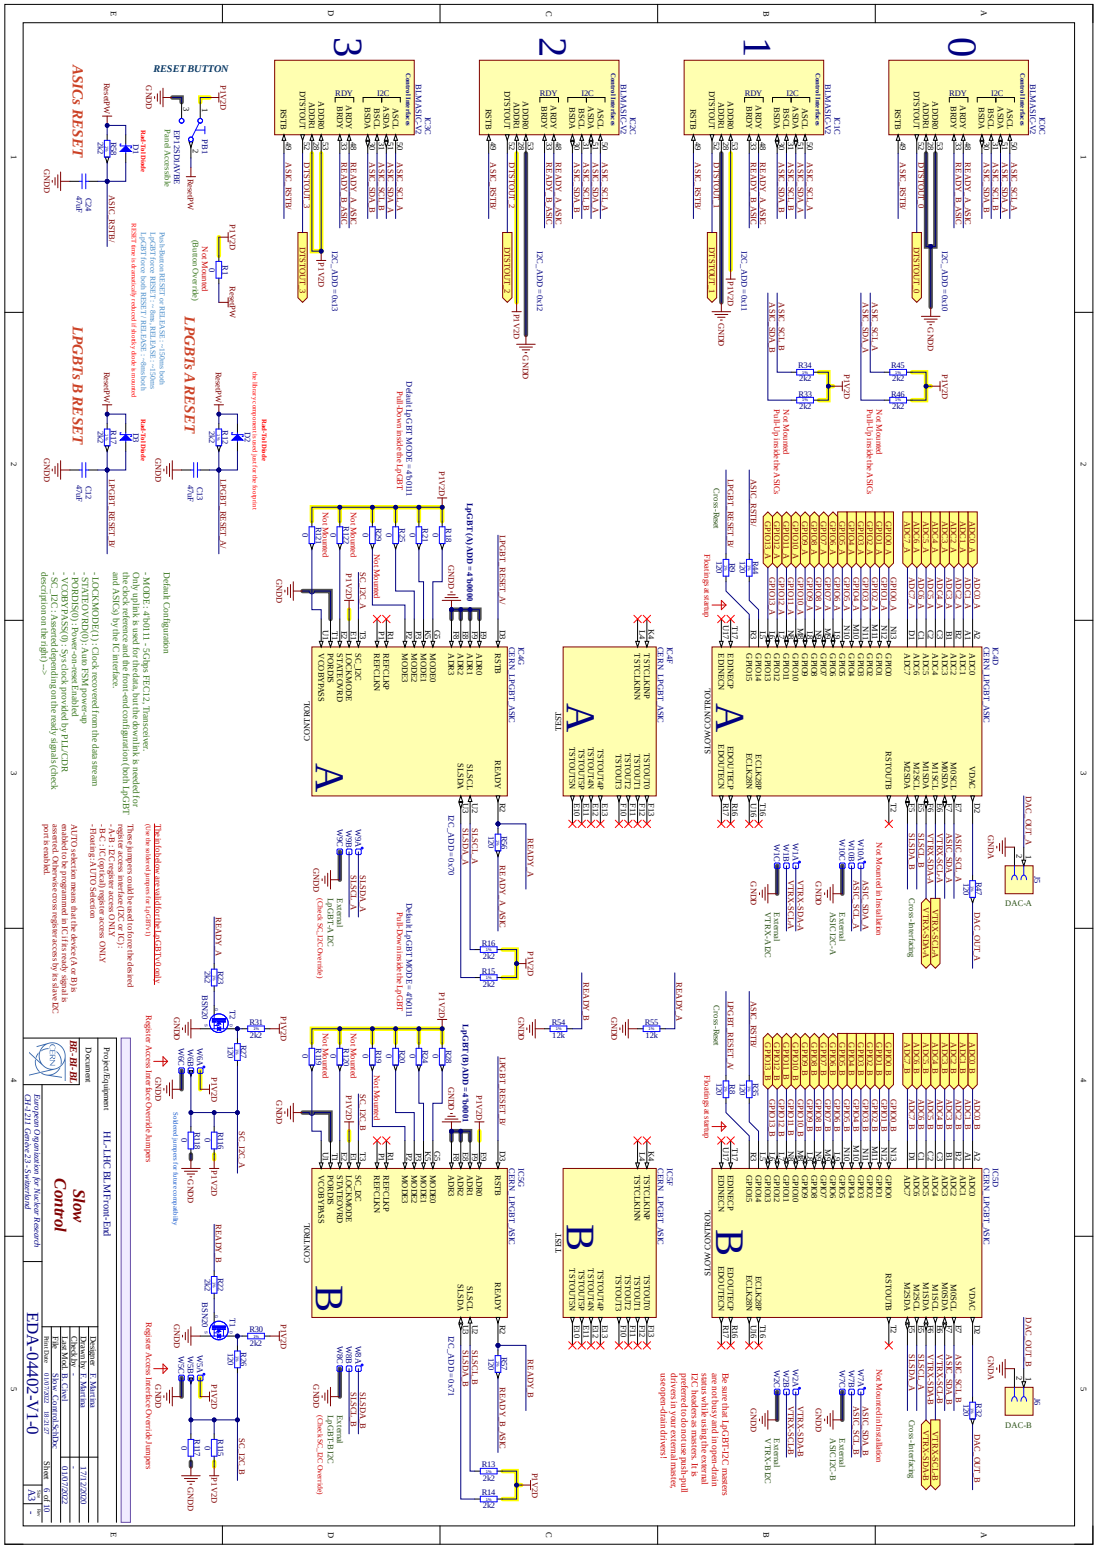
RESET BUTTON
ASICs RESET
LPGBTs A RESET
LPGBTs B RESET
Slow Control
EDA-04402-V1-0

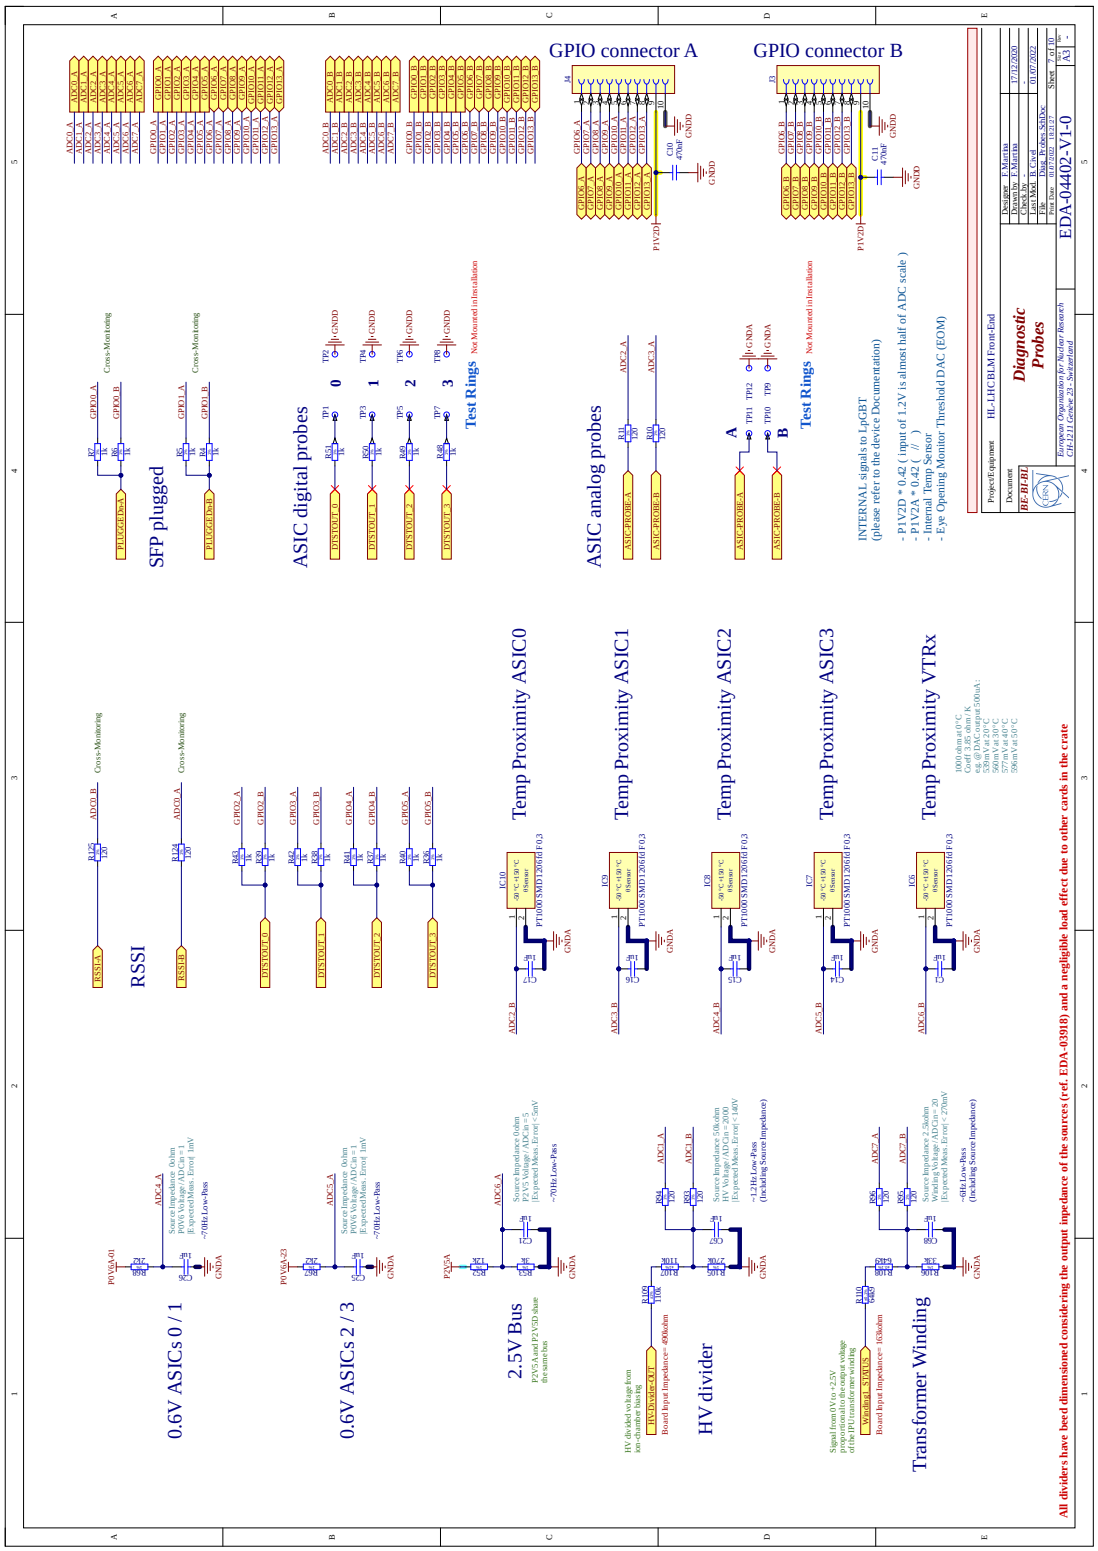
GPIO connector A
GPIO connector B
SFP plugged
ASIC digital probes
ASIC analog probes
Test Rings
RSSI
0.6V ASICs 0 / 1
0.6V ASICs 2 / 3
2.5V Bus
HV divider
Transformer Winding
Temp Proximity ASIC0
Temp Proximity ASIC1
Temp Proximity ASIC2
Temp Proximity ASIC3
Temp Proximity VTRx
INTERNAL signals to LpGBT
(please refer to the device Documentation)
- P1V2D * 0.42 (input of 1.2V is almost half of ADC scale)
- P1V2A * 0.42 (//)
- Internal Temp Sensor
- Eye Opening Monitor Threshold DAC (EOM)
Diagnostic Probes
EDA-04402-V1-0
All dividers have beed dimensioned considering the output impedance of the sources (ref. EDA-03918) and a negligible load effect due to other cards in the crate

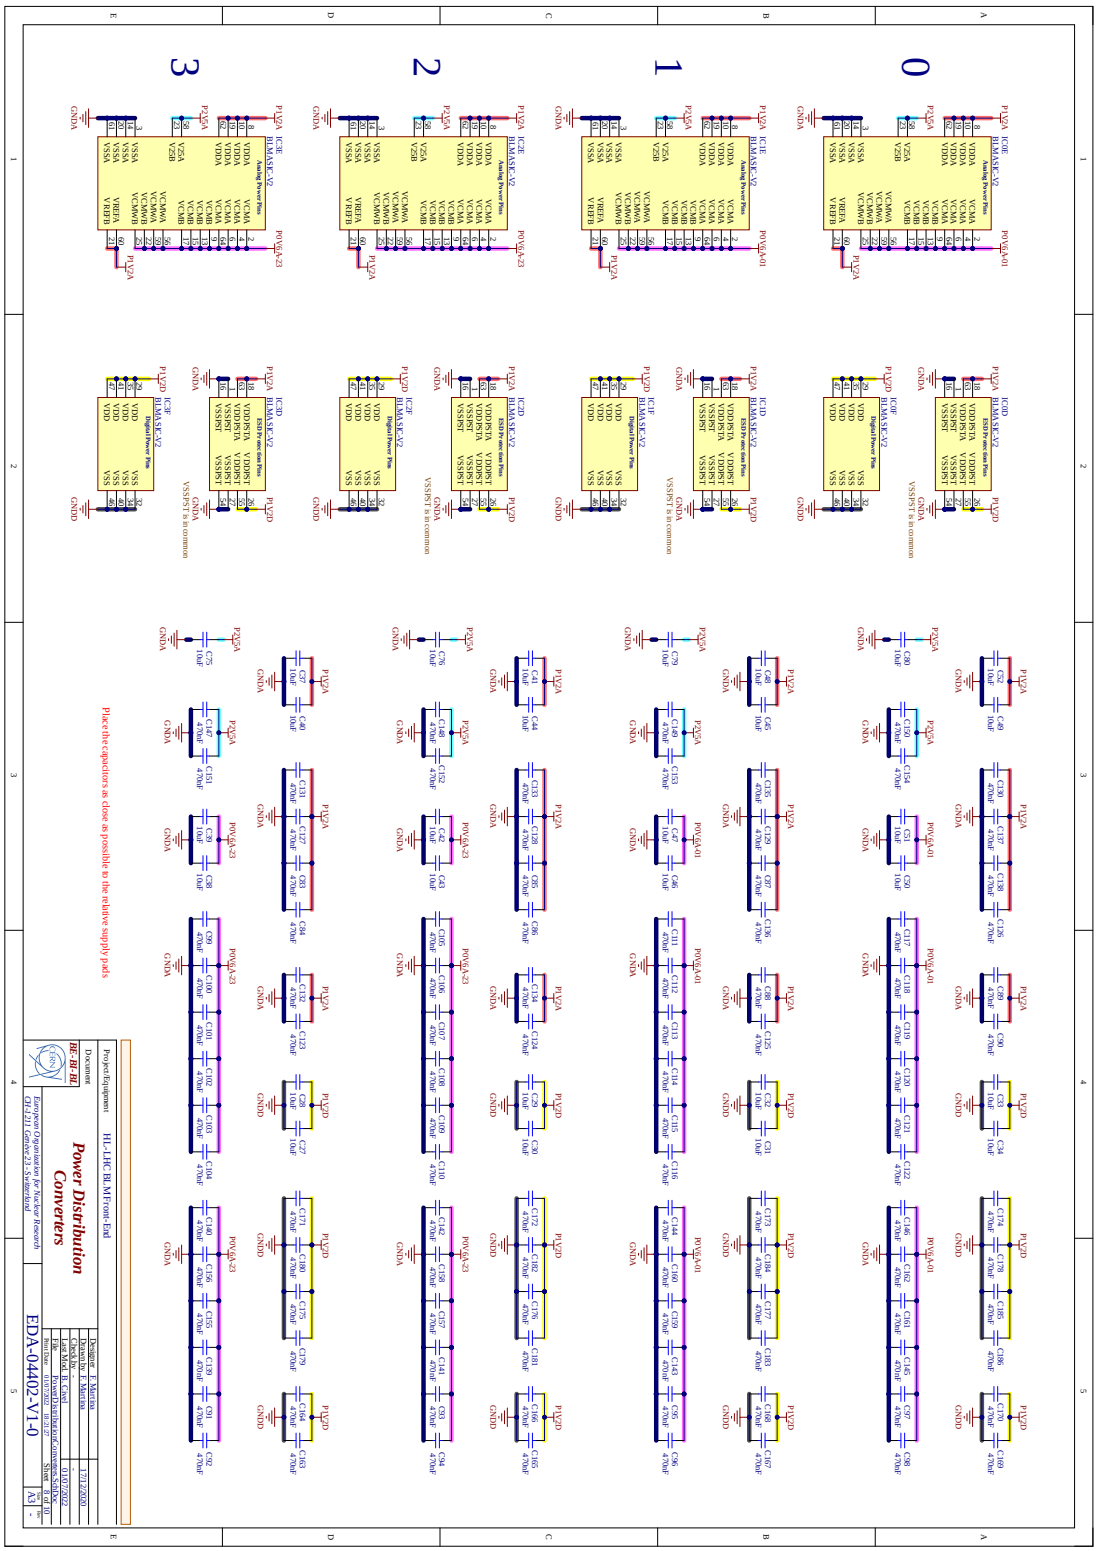
Power Distribution Converters
EDA-04402-V1-0
Project/Equipment HL-LHC BLM Front-End
Document BE-BI-BL
European Organisation for Nuclear Research
CH-1211 Genève 23 - Switzerland
Designer E. Martina
Drawn by E. Martina
Checked -
Last Mod. B. Civel
Date 01/07/2022
Project Date 01/07/2022
Sheet 8 of 10
A3
Place the capacitors as close as possible to the relative supply pads

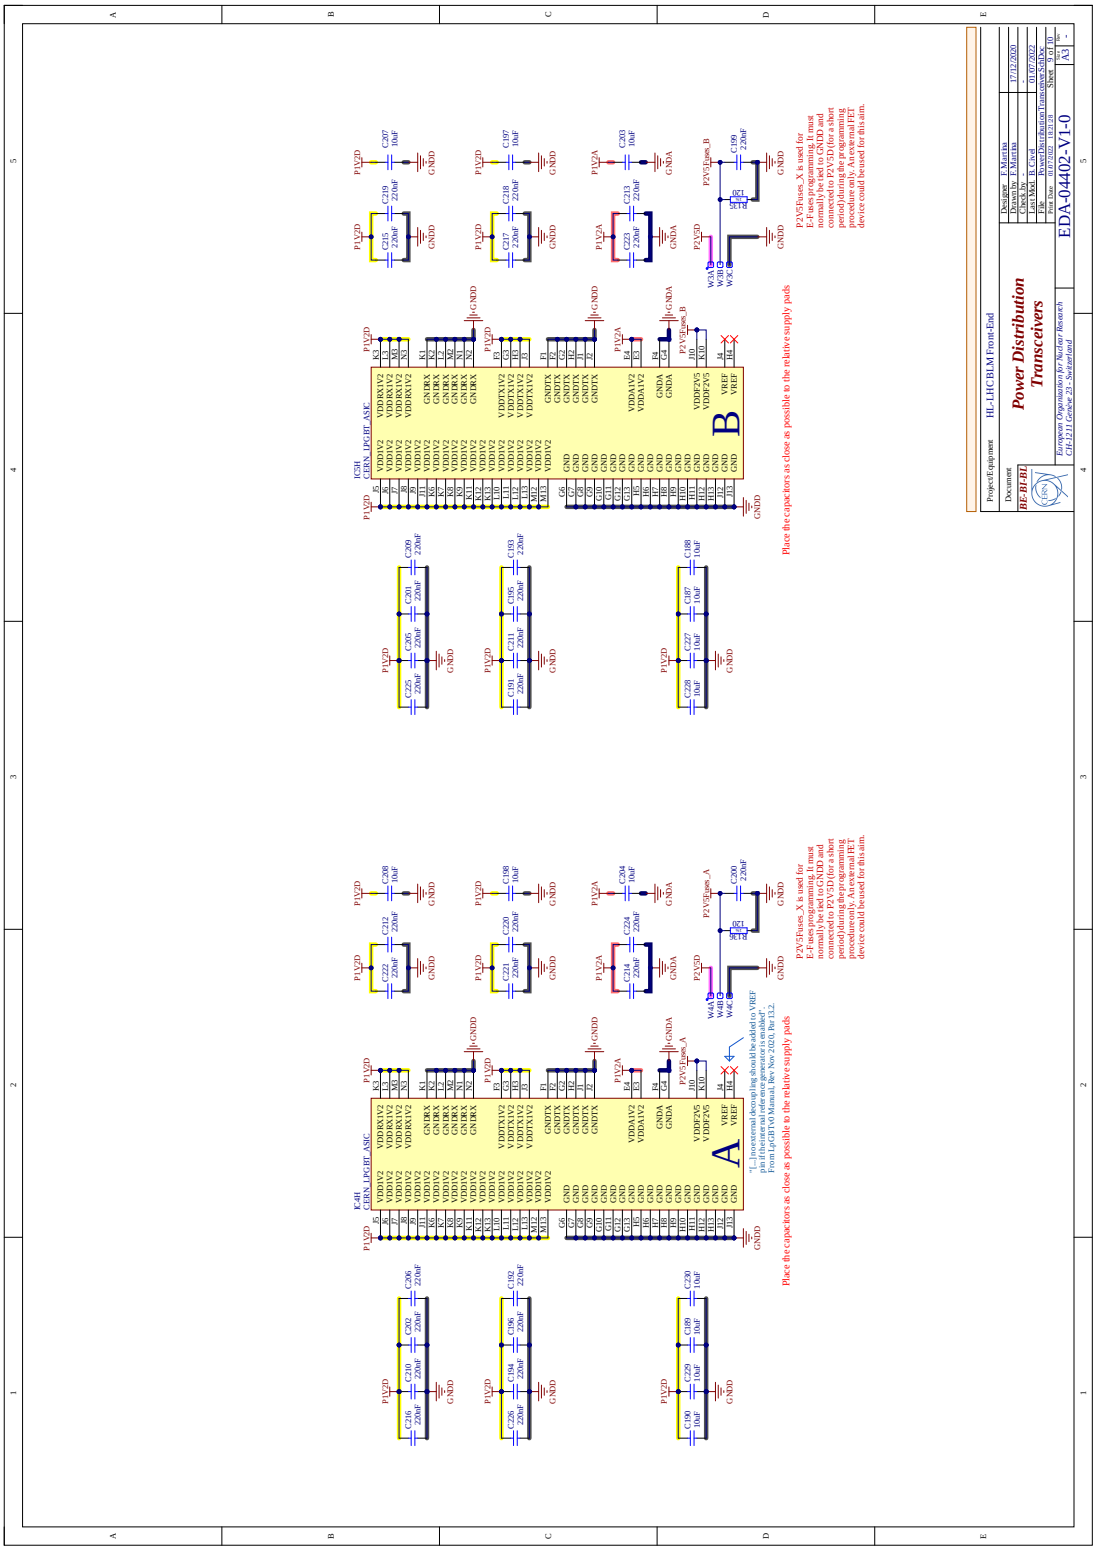
Power Distribution Transceivers
EDA-04402-V1-0
Place the capacitors as close as possible to the relative supply pads
Place the capacitors as close as possible to the relative supply pads

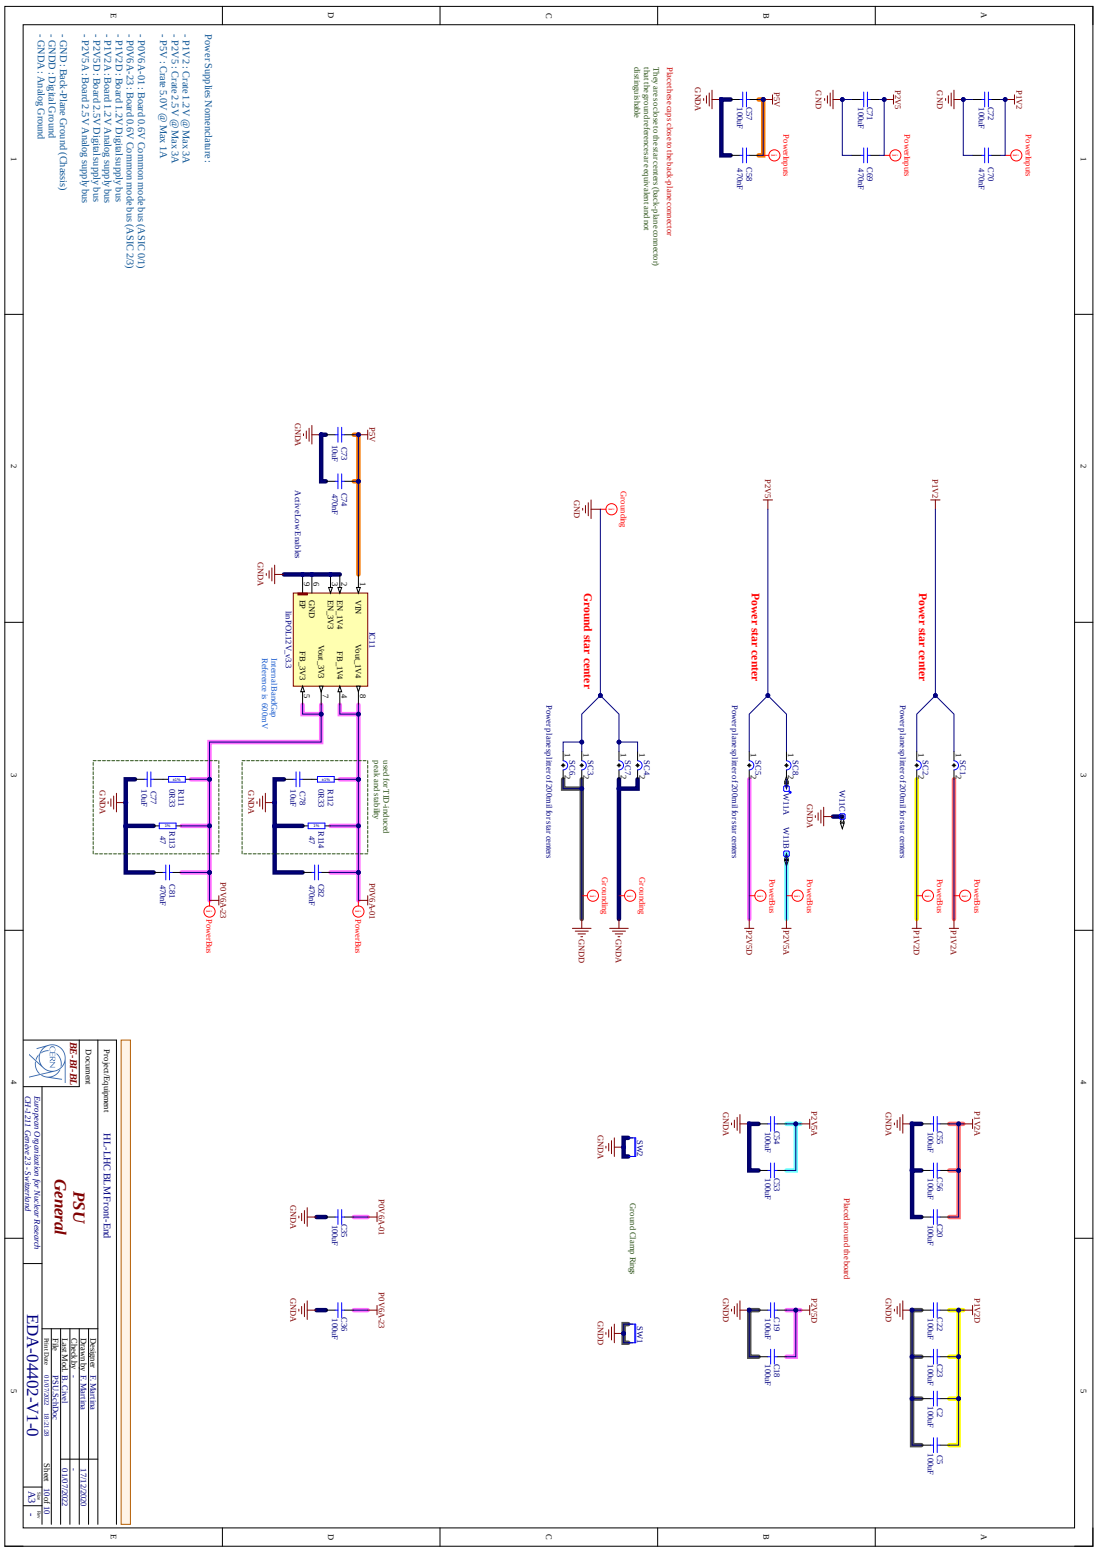
Power Supplies Nomenclature :
- P1V2 : Crate 1.2V @ Max 3A
- P2V5 : Crate 2.5V @ Max 3A
- P5V : Crate 5.0V @ Max 1A
- P0V6A-01 : Board 0.6V Common mode bus (ASIC 01)
- P0V6A-23 : Board 0.6V Common mode bus (ASIC 23)
- P1V2D : Board 1.2V Digital supply bus
- P1V2A : Board 1.2V Analog supply bus
- P2V5D : Board 2.5V Digital supply bus
- P2V5A : Board 2.5V Analog supply bus
- GND : Back-Plane Ground (Chassis)
- GNDD : Digital Ground
- GNDA : Analog Ground
Power star center
Power star center
Ground star center
Placed around the board
Ground Clamp Rings
Active Low Enables
Internal BandGap
Reference is 600mV
used for TID-induced peak and stability
Project/Equipment
HL-LHC BLM Front-End
Document
BE-BI-BL
PSU
General
European Organization for Nuclear Research
CH-1211 Genève 23 - Switzerland
Designer E. Martina
Drawn by E. Martina
Check by
Last Mod. B. Civel
File PSU.SchDoc
Sheet 10 of 10
EDA-04402-V1-0
A3

Appendix F

Front-End Layout

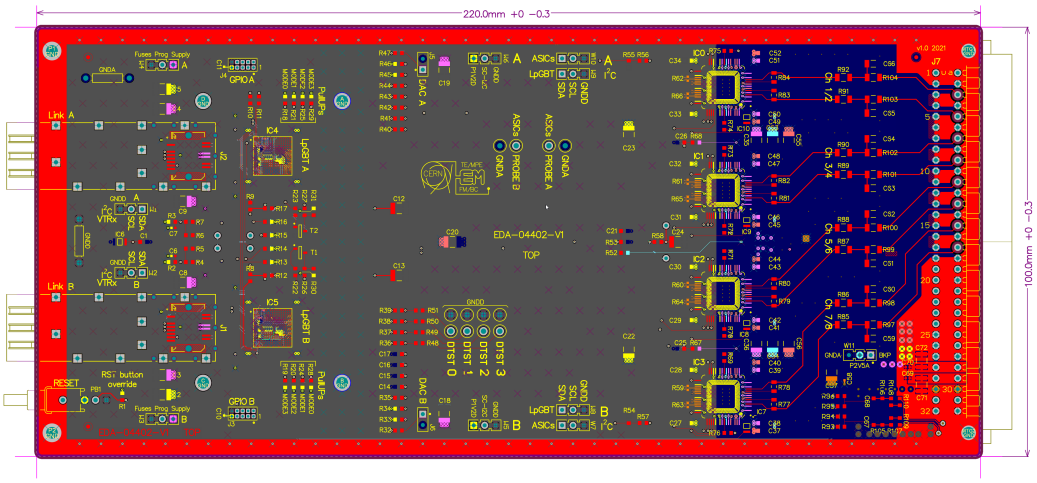

Figure F.1: Top layer. The 4 BLMASIC footprints are visible on the right, together with the routing of the input channels. On the left, the lpGBT devices and the related SFP cages can be distinguished. Analogue (blue), digital (grey) and chassis (red) ground polygons are visible.

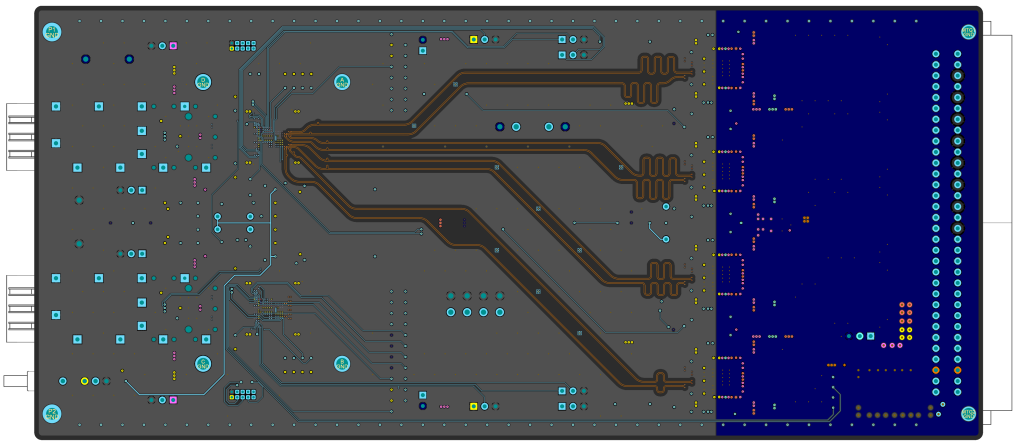

Figure F.2: Layer 2. Clock and data striplines amongst the BLMASICs and the lpGBT **A** device. Analogue (blue) and digital (grey) ground planes are visible.

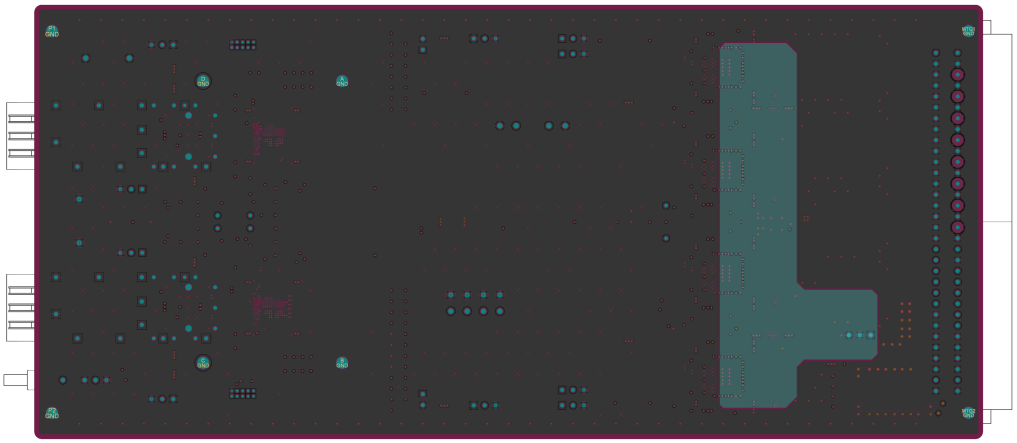

Figure F.3: Layer 3. The digital ground plane (grey) can be distinguished from the P2V5A analogue power plane.

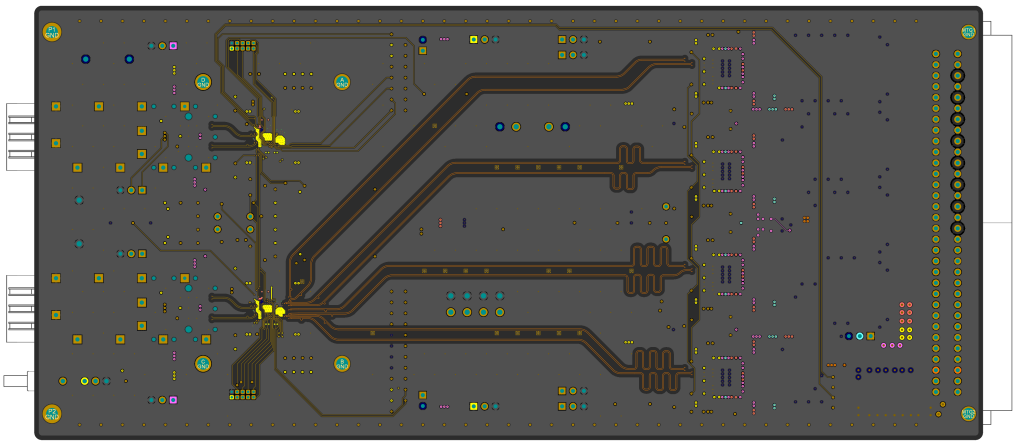

Figure F.4: Layer 4. Clock and data striplines amongst the BLMASICs and the lpGBT **B** device. The fast 5.12 *Gbps* striplines are visible on the left. The digital ground plane is in grey.

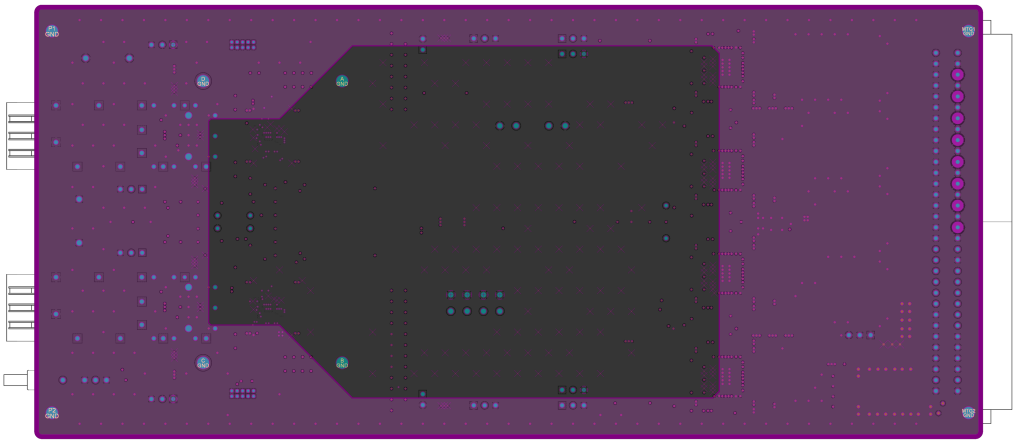

Figure F.5: Layer 5. The P2V5D digital power plane (violet) can be distinguished from the digital ground one (grey), which constitutes one of the stripline references.

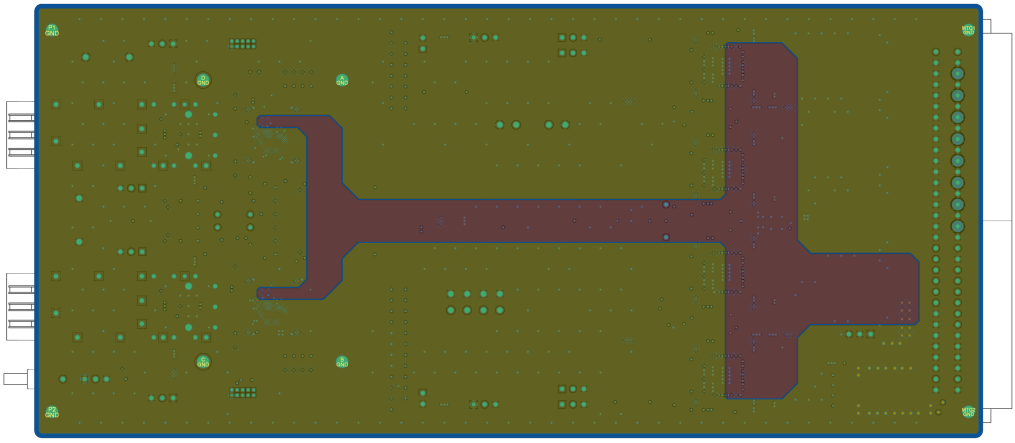

Figure F.6: Layer 6. The P1V2D digital power plane (yellow) and its analogue counterpart P1V2A (red) are visible.

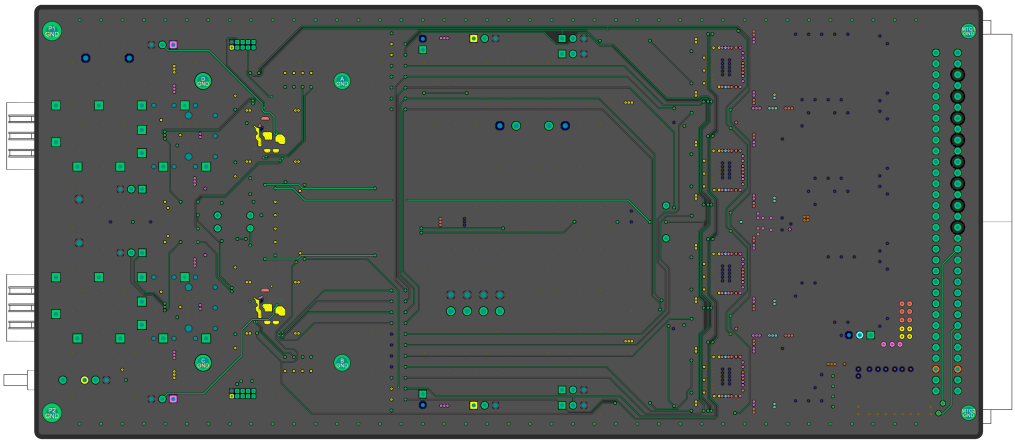

Figure F.7: Layer 7. The digital ground surrounds the slow control and diagnostic routing.

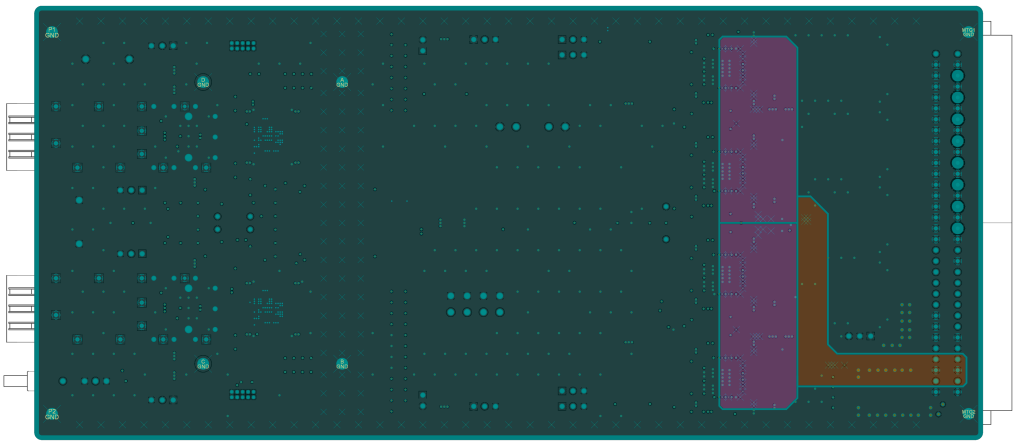

Figure F.8: Layer 8. The P5V polygon (orange) supplies the linPOL12V vias, from which the analogue common-mode references P0V6A-01 and P0V6A-23 (magenta) are derived and propagated towards the BLMASICs. The remaining area is filled by the chassis ground plane, used for heat dissipation.

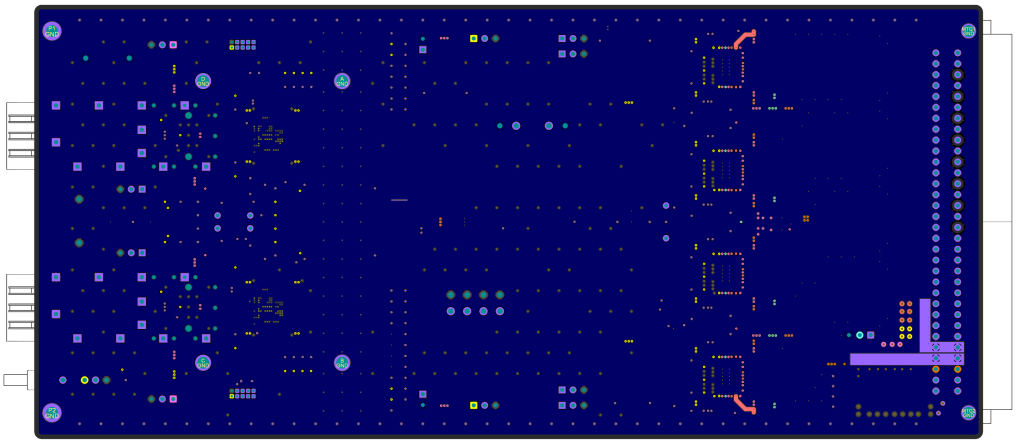

Figure F.9: Layer 9. Full analogue ground plane throughout the board and part of the power supply rails on the rear panel connector.

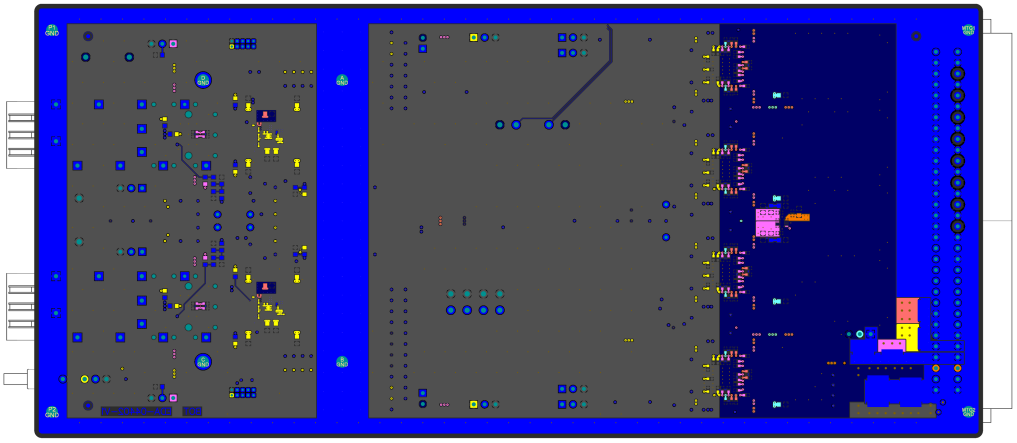

Figure F.10: Bottom layer. Analogue (dark blue), digital (grey) and chassis (light blue) ground polygons are visible. The right side includes the routing of the powering rails and the star centres. The back-panel connector pins, employed for the sensitive current inputs, can be distinguished due to the increased clearance with respect to the surrounding ground copper.

References

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," *Physics Letters B*, vol. 716, no. 1, pp. 1–29, 2012, ISSN: 0370-2693. DOI: https://doi.org/10.1016/j.physletb.2012.08.020.

[2] CMS Collaboration, "Combined results of searches for the standard model Higgs boson in pp collisions at $\sqrt{s} = 7$ TeV," *Physics Letters B*, vol. 710, no. 1, pp. 26–48, 2012, ISSN: 0370-2693. DOI: https://doi.org/10.1016/j.physletb.2012.02.064.

[3] B. Dehning, E. Effinger, J. Emery, *et al.*, "The LHC beam loss measurement system," in *2007 IEEE Particle Accelerator Conference (PAC)*, Jun. 2007, pp. 4192–4194. DOI: 10.1109/PAC.2007.4439980.

[4] European Strategy Group, "The European Strategy for Particle Physics Update 2013," *16th Session of European Strategy Council*, May 2013, CERN-Council-S-0106.

[5] O. Aberle, I. Béjar Alonso, O. Brüning, *et al.*, *High-Luminosity Large Hadron Collider (HL-LHC): Technical design report* (CERN Yellow Reports: Monographs). Geneva: CERN, 2020. DOI: 10.23731/CYRM-2020-0010.

[6] F. Martina and C. Zamantzas, "Testbed Development for the Characterisation of an ASIC for Beam Loss Measurement Systems," in *Proc. IBIC'20*, (Santos, Brazil), ser. International Beam Instrumentation Conference, JACoW Pub., Geneva, Switzerland, Nov. 2020, ISBN: 978-3-95-450222-6. DOI: 10.18429/JACoW-IBIC2020-THPP04.

[7] F. Martina *et al.*, "Methodology, Characterisation and Results from the Prototype Beam Loss Monitoring ASIC at CERN," in *Proc. IBIC'21*, (Pohang, Korea), ser. International Beam Instrumentation Conference, JACoW Pub., Geneva, Switzerland, Nov. 2021, pp. 294–298, ISBN: 978-3-95-450230-1. DOI: 10.18429/JACoW-IBIC2021-TUPP34.

[8] J. Wenninger, "LHC status and performance," in *Proceedings of Prospects for Charged Higgs Discovery at Colliders — PoS(CHARGED2018)*, vol. 339, 2019, p. 001. DOI: 10.22323/1.339.0001.

[9] C. Kurfuerst, *Quench Protection of the LHC Quadrupole Magnets*. 2010, vol. CERN-THESIS-2010, PhD Thesis Technische Univ. Wien.

[10] A. Priebe, K. Dahlerup-Petersen, B. Dehning, *et al.*, "Beam-induced quench test of LHC main quadrupole," *2nd International Particle Accelerator Conference 2011 San Sebastian, Spain, 4 - 9 Sep*, vol. CERN-ATS-2011-058, WEPC172, 2011.

[11] A. Priebe, B. Dehning, M. Sapinski, *et al.*, "Investigations of Quench Limits of the LHC Superconducting Magnets," *IEEE Transactions on Applied Superconductivity*, vol. 23, no. 3, pp. 4 701 205–4 701 205, 2013. DOI: 10.1109/TASC.2012.2236378.

[12] M. Sapinski *et al.*, "Beam-induced Quench Tests of LHC Magnets," in *Proc. 5th International Particle Accelerator Conference (IPAC'14), Dresden, Germany, June 15-20, 2014*, (Dresden, Germany), ser. International Particle Accelerator Conference, MOOCB01, Geneva, Switzerland: JACoW, Jul. 2014, pp. 52–55, ISBN: 978-3-95450-132-8. DOI: https://doi.org/10.18429/JACoW-IPAC2014-MOOCB01.

[13] M. Sapinski, T. Baer, M. Bednarek, *et al.*, "Beam induced quenches of LHC magnets," *Proceedings of IPAC'13, Shanghai, China*, pp. 3243–3245, 2013, THPEA045.

[14] M. Bajko, F. Bertinelli, N. Catalan-Lasheras, *et al.*, "Report of the Task Force on the Incident of 19th September 2008 at the LHC," CERN, Geneva, Tech. Rep., Mar. 2009.

[15] B. Dehning, E. Effinger, J. Emery, *et al.*, "The LHC beam loss measurement system," in *2007 IEEE Particle Accelerator Conference (PAC)*, 2007, pp. 4192–4194. DOI: 10.1109/PAC.2007.4439980.

[16] B. Dehning, E. Effinger, J. Emery, *et al.*, "LHC beam loss detector design: Simulation and measurements," in *2007 IEEE Particle Accelerator Conference (PAC)*, 2007, pp. 4198–4200. DOI: 10.1109/PAC.2007.4439981.

[17] B. Dehning, E. Effinger, J. Emery, *et al.*, "Quench Protection with LHC Beam Loss Monitors," *CERN-BE-2011-002*, no. https://cds.cern.ch/record/1348019, 4 p, 2010.

[18] C. Zamantzas, B. Dehning, E. Effinger, *et al.*, "The LHC beam loss monitoring system's data contribution to other systems," in *2007 IEEE Nuclear Science Symposium Conference Record*, vol. 3, 2007, pp. 2331–2335. DOI: 10.1109/NSSMIC.2007.4436612.

[19] B. Dehning, M. Dabrowski, E. Effinger, *et al.*, "Overview of LHC Beam Loss Measurements," *Proceedings of IPAC'11, Sep 4-9, San Sebastian, Spain, 2011*, pp. 2854–2856, 2011, THOAA03.

[20] S. Eitelbuss, "Screening and Reliability Testing of Beam Loss Monitor Electronics at CERN," Universität Stuttgart, Mar. 2019.

[21] V. Schramm, "Dependable System Development Methodology and Case Study for the LHC Beam Loss Monitoring System at CERN," Ph.D. dissertation, Universität Stuttgart, Feb. 2021.

[22] R. Filippini, B. Dehning, G. Guaglio, *et al.*, "Reliability Assessment of the LHC Machine Protection System," in *Proceedings of the 2005 Particle Accelerator Conference*, 2005, pp. 1257–1259. DOI: 10.1109/PAC.2005.1590726.

[23] G. Guaglio, "Reliability of the Beam Loss Monitors System for the Large Hadron Collider at CERN," Université Clermont-Auvergne, Dec. 2005.

[24] J. Emery, B. Dehning, E. Effinger, *et al.*, "First experiences with the LHC BLM sanity checks," *Journal of Instrumentation*, vol. 5, no. 12, pp. C12044–C12044, Dec. 2010. DOI: 10.1088/1748-0221/5/12/c12044.

[25] E. B. Holzer, B. Dehning, E. Effnger, *et al.*, "Beam loss monitoring for lhc machine protection," *Physics Procedia*, vol. 37, pp. 2055–2062, 2012, Proceedings of the 2nd International Conference on Technology and Instrumentation in Particle Physics (TIPP 2011), ISSN: 1875-3892. DOI: https://doi.org/10.1016/j.phpro.2012.04.110.

[26] C. F. Hajdu, T. Dabóczi, and C. Zamantzas, "On-line system supervision of beam loss monitoring systems with sinusoidal excitation," in *XXI IMEKO World Congress "Measurement in Research and Industry", Prague, Czech Republic*, 2015.

[27] C. F. Hajdu, C. Zamantzas, and T. Dabóczi, "Development of a method for continuous functional supervision of BLM systems," in *Proceedings of IBIC2016, Barcelona, Spain*, MOPG21, 2016, pp. 90–93. DOI: 10.18429/JACoW-IBIC2016-MOPG21.

[28] M. Kalliokoski, B. Auchmann, B. Dehning, *et al.*, "BLM threshold evolution and 2016 proposal," in *6th Evian Workshop on LHC beam operation, Evian Les Bains, France, 15 - 17 Dec 2015*, CERN, 2016, 191–196. 6 p.

[29] S. C. Wagner, "LHC machine protection system: method for balancing machine safety and beam availability," Ph.D. dissertation, ETH Zurich, 2010.

[30] M. Stockner, B. Dehning, C. Fabjan, *et al.*, "Measurements and Simulations of Ionization Chamber Signals in Mixed Radiation Fields for the LHC BLM System," in *2006 IEEE Nuclear Science Symposium Conference Record*, vol. 3, 2006, pp. 1342–1345. DOI: 10.1109/NSSMIC.2006.354151.

[31] D. Kramer, B. Dehning, E. B. Holzer, *et al.*, "Very high radiation detector for the LHC BLM system based on secondary electron emission," in *2007 IEEE Nuclear Science Symposium Conference Record*, vol. 3, 2007, pp. 2327–2330. DOI: 10.1109/NSSMIC.2007.4436611.

[32] V. Grishin, B. Dehning, E. Effinger, *et al.*, "A Family of Gas Ionization Chambers and SEM for Beam Loss Monitoring of LHC and Other Accelerators," in *Proc. 26th Russian Particle Accelerator Conference (RUPAC'18), Protvino, Russia, 01-05 October 2018*, (Protvino, Russia), ser. Russian Particle Accelerator Conference, TUZMH03, Geneva, Switzerland: JACoW Pub., Nov. 2018, pp. 44–48, ISBN: 978-3-95450-197-7. DOI: doi:10.18429/JACoW-RUPAC2018-TUZMH03.

[33] M. Kalliokoski *et al.*, "Beam Loss Monitoring for Run 2 of the LHC," in *Proc. 6th International Particle Accelerator Conference (IPAC'15), Richmond, VA, USA, May 3-8, 2015*, (Richmond, VA, USA), ser. International Particle Accelerator Conference, MOPTY055, Geneva, Switzerland: JACoW, Jun. 2015, pp. 1057–1060, ISBN: 978-3-95450-168-7. DOI: https://doi.org/10.18429/JACoW-IPAC2015-MOPTY055.

[34] E. Effinger, B. Dehning, J. Emery, *et al.*, "The LHC beam loss monitoring system's data acquisition card," in *Proceedings of the Twelfth Workshop on Electronics for LHC and Future Experiments*, 2007, pp. 108–112. DOI: 10.5170/CERN-2007-001.108.

[35] E. Effinger, C. Zamantzas, G. Ferioli, *et al.*, "Single gain radiation tolerant LHC beam loss acquisition card," in *Proceedings of DIPAC'07, Venice, Italy*, THOAA03, 2007, pp. 319–321.

[36] C. Zamantzas, B. Dehning, E. Effinger, *et al.*, "The LHC Beam Loss Monitoring System's Real-Time Data Analysis Card," in *Proc. DIPAC'05*, (Lyon, France, Jun. 2005), POW019, JACoW Pub., Geneva, Switzerland.

[37] C. Zamantzas, "The real-time data analysis and decision system for particle flux detection in the LHC accelerator at CERN," Ph.D. dissertation, Brunel University, 2006.

[38] C. Zamantzas, B. Dehning, E. Effinger, *et al.*, "An FPGA Based Implementation for Real-Time Processing of the LHC Beam Loss Monitoring System's Data," in *2006 IEEE Nuclear Science Symposium Conference Record*, IEEE, vol. 2, 2006, pp. 950–954. DOI: 10.1109/NSSMIC.2006.356003.

[39] E. B. Holzer, B. Dehning, C. W. Fabjan, *et al.*, "Accuracy of LHC proton loss rate determination by the BLM system," in *2007 IEEE Nuclear Science Symposium Conference Record*, vol. 2, 2007, pp. 1679–1683. DOI: 10.1109/NSSMIC.2007.4437323.

[40] C. Zamantzas, B. Dehning, J. Emery, *et al.*, "Configuration and validation of the LHC beam loss monitoring system," in *Proceedings of the 9th European Workshop on Beam Diagnostics and Instrumentation for Particle Accelerators (DIPAC09), Basel, Switzerland*, TUPB31, 2009, pp. 240–242.

[41] B. Todd, "A Beam Interlock System for CERN High Energy Accelerators," Ph.D. dissertation, Brunel University, Nov. 2006.

[42] B. Puccio, A. Castañeda Serra, M. Kwiatkowski, *et al.*, "The CERN Beam Interlock System: Principle and Operational Experience," in *Proceedings of IPAC'10, Kyoto, Japan*, WEPEB073, Jun. 2010, pp. 2866–2868.

[43] S. K. Mohanan, H. Boukabache, D. Perrin, *et al.*, "Towards the next generation of CERN radiation monitoring front end ASICs," *Journal of Instrumentation*, vol. 17, no. 04, p. C04029, Apr. 2022. DOI: 10.1088/1748-0221/17/04/c04029.

[44] P. Moreira, "The LpGBT project status and overview," in *ACES 2016- Fifth Common ATLAS CMS Electronics Workshop for LHC Upgrades, 7-10 Mar, Geneva*, 2016.

[45] A. Boccardi, M. Barros Marin, T. Levens, *et al.*, "A Modular Approach to Acquisition Systems for Future CERN Beam Instrumentation Developments," in *Proceedings of ICALEPCS2015, Melbourne, Australia*, 2015, THHB2O02, 1103–1106. DOI: 10.18429/JACoW-ICALEPCS2015-THHB2O02.

[46] G. G. Venturini, B. Dehning, E. Effinger, *et al.*, "Systematic Study of Acquisition Electronics with a High Dynamic Range for a Beam Loss Measurement System," in *Proc. DIPAC'09*, (Basel, Switzerland, May 2009), TUPB33, JACoW Pub., Geneva, Switzerland, pp. 245–247.

[47] A. Sedra, K. Smith, T. Carusone, *et al.*, *Microelectronic Circuits* (Oxford series in electrical and computer engineering). Oxford University Press, 2020, ISBN: 9780190853549.

[48] J. Millman and A. Grabel, *Microelectronics* (Electrical & electronic engineering series). McGraw-Hill, 1987, ISBN: 9780070423305.

[49] C. Zamantzas, M. Alsdorf, B. Dehning, *et al.*, "System Architecture for Measuring and Monitoring Beam Losses in the Injector Complex at CERN," in *Proc. IBIC'12*, (Tsukuba, Japan), TUPA09, JACoW Pub., Geneva, Switzerland, 2012, pp. 347–350.

[50] W. Viganò, M. Alsdorf, B. Dehning, *et al.*, "10 orders of magnitude current measurement digitisers for the CERN beam loss systems," *Journal of Instrumentation*, vol. 9, no. 02, pp. C02011–C02011, Feb. 2014. DOI: 10.1088/1748-0221/9/02/c02011.

[51] E. Voulgari, "A Nine Decade Femtoampere Current to Frequency Converter," EPFL, Lausanne, Oct. 2017.

[52] S. Pavan, R. Schreier, and G. C. Temes, *Understanding delta-sigma data converters*. John Wiley & Sons, 2017.

[53] D. Zrilic, *Functional Processing of Delta-Sigma Bit-Stream*. Springer International Publishing, 2020, ISBN: 9783030476489.

[54] M. Kwiatkowski, M. Alsdorf, B. Dehning, *et al.*, "A Real-Time FPGA Based Algorithm for the Combination of Beam Loss Acquisition Methods used for Measurement Dynamic Range Expansion," in *Proc. IBIC'12*, (Tsukuba, Japan, Oct. 2012), TUPA31, JACoW Pub., Geneva, Switzerland, pp. 419–423.

[55] K. Rao and M. Swamy, *Digital Signal Processing: Theory and Practice*. Springer Singapore, 2018, ISBN: 9789811080814.

[56] M. Luise and G. Vitetta, *Teoria dei segnali* (Collana di istruzione scientifica). McGraw-Hill Companies, 2009, ISBN: 9788838665837.

[57] G. G. Venturini, B. Dehning, E. Effinger, *et al.*, "Design Specifications for a Radiation Tolerant Beam Loss Measurement ASIC," in *Proc. DIPAC'09*, (Basel, Switzerland, May 2009), TUPB32, JACoW Pub., Geneva, Switzerland, pp. 243–244.

[58] G. G. Venturini, F. Anghinolfi, B. Dehning, *et al.*, "Characterization of a Wide Dynamic-range, Radiation-tolerant Charge-digitizer ASIC for Monitoring of Beam Losses," in *Proc. IBIC'12*, (Tsukuba, Japan, Oct. 2012), MOPA12, JACoW Pub., Geneva, Switzerland, pp. 74–78.

[59] G. Venturini, F. Anghinolfi, and B. Dehning, "A 120dB dynamic-range radiation-tolerant charge-to-digital converter for radiation monitoring," *Microelectronics Journal*, vol. 44, no. 12, pp. 1302–1308, 2013, ISSN: 0026-2692. DOI: https://doi.org/10.1016/j.mejo.2013.08.020.

[60] G. Venturini, "Study of Acquisition Electronics with a High Dynamic Range for a Beam Loss Measurement System," Politecnico di Milano, Apr. 2010.

[61] W. Vigano, B. Dehning, E. Effinger, *et al.*, "Comparison of Three Different Concepts of High Dynamic Range and Dependability Optimised Current Measurement Digitisers for Beam Loss Systems," in *Proc. IBIC'12*, (Tsukuba, Japan, Oct. 2012), MOPA09, JACoW Pub., Geneva, Switzerland, pp. 66–70.

[62] L. Ratti, L. Gaioni, M. Manghisoni, *et al.*, "Investigating degradation mechanisms in 130 nm and 90 nm commercial CMOS technologies exposed to up to 100 Mrad ionizing radiation dose," in *2007 9th European Conference on Radiation and Its Effects on Components and Systems*, 2007, pp. 1–9. DOI: 10.1109/RADECS.2007.5205526.

[63] H. J. Barnaby, M. Mclain, and I. S. Esqueda, "Total-ionizing-dose effects on isolation oxides in modern CMOS technologies," *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, vol. 261, no. 1, pp. 1142–1145, 2007, The Application of Accelerators in Research and Industry, ISSN: 0168-583X. DOI: https://doi.org/10.1016/j.nimb.2007.03.109.

[64] W. Snoeys, T. Gutierrez, and G. Anelli, "A new NMOS layout structure for radiation tolerance," in *2001 IEEE Nuclear Science Symposium Conference Record (Cat. No.01CH37310)*, vol. 2, 2001, 822–826 vol.2. DOI: 10.1109/NSSMIC.2001.1009683.

[65] E. Shauly, "CMOS Leakage and Power Reduction in Transistors and Circuits: Process and Layout Considerations," *Journal of Low Power Electronics and Applications*, vol. 2, pp. 1–29, Dec. 2012. DOI: 10.3390/jlpea2010001.

[66] T. Vergine, M. De Matteis, A. Baschirotto, *et al.*, "A 32-channel 12-bits 65nm Wilkinson ADC for CMS central tracker," in *2014 10th Conference on Ph.D. Research in Microelectronics and Electronics (PRIME)*, 2014, pp. 1–4. DOI: 10.1109/PRIME.2014.6872683.

[67] R. J. Baker, *CMOS: circuit design, layout, and simulation*. John Wiley & Sons, 2019.

[68] R. Hogervorst, J. Tero, R. Eschauzier, *et al.*, "A compact power-efficient 3 V CMOS rail-to-rail input/output operational amplifier for VLSI cell libraries," *IEEE Journal of Solid-State Circuits*, vol. 29, no. 12, pp. 1505–1513, 1994. DOI: 10.1109/4.340424.

[69] P. Figueiredo and J. Vital, "Kickback noise reduction techniques for CMOS latched comparators," *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 53, no. 7, pp. 541–545, 2006. DOI: 10.1109/TCSII.2006.875308.

[70] A. X. Widmer and P. A. Franaszek, "A DC-Balanced, Partitioned-Block, 8B/10B Transmission Code," *IBM Journal of Research and Development*, vol. 27, no. 5, pp. 440–451, 1983. DOI: 10.1147/rd.275.0440.

[71] A. Dunkels, "Design and Implementation of the lwIP TCP/IP Stack," *Swedish Institute of Computer Science*, vol. 2, no. 77, 2001.

[72] M. Buckingham, *Noise in Electronic Devices and Systems* (Electrical and Electronic Engineering Series). E. Horwood, 1983, ISBN: 9780853122180.

[73] A. Van der Ziel, *Noise in Solid State Devices and Circuits*. Wiley, 1986, ISBN: 9780471832348.

[74] D. T. Marr, F. Binns, D. L. Hill, *et al.*, "Hyper-Threading Technology Architecture and Microarchitecture.," *Intel Technology Journal*, vol. 6, no. 1, 2002.

[75] P. P. Chu, *RTL hardware design using VHDL: coding for efficiency, portability, and scalability*. John Wiley & Sons, 2006.

[76] Y. Tsividis, *Operation and Modeling of the MOS Transistor* (Electrical engineering series). WCB/McGraw-Hill, 1999, ISBN: 9780070655232.

[77] M. Koh, W. Mizubayashi, K. Iwamoto, *et al.*, "Limit of gate oxide thickness scaling in mosfets due to apparent threshold voltage fluctuation induced by tunnel leakage current," *IEEE Transactions on Electron Devices*, vol. 48, no. 2, pp. 259–264, 2001. DOI: 10.1109/16.902724.

[78] G. Bishop, G. Welch, *et al.*, "An introduction to the Kalman filter," *Course notes of SIGGRAPH, Los Angeles, CA*, vol. 8, 2001.

[79] G. F. Welch, "Kalman filter," in *Computer Vision: A Reference Guide*. Cham: Springer International Publishing, 2020, pp. 1–3, ISBN: 978-3-030-03243-2. DOI: 10.1007/978-3-030-03243-2_716-1.

[80] N. Agrawal, A. Kumar, V. Bajaj, *et al.*, "Design of digital IIR filter: A research survey," *Applied Acoustics*, vol. 172, p. 107 669, 2021, ISSN: 0003-682X. DOI: https://doi.org/10.1016/j.apacoust.2020.107669.

[81] D. O'Loughlin, A. Coffey, F. Callaly, *et al.*, "Xilinx Vivado High Level Synthesis: Case studies," in *25th IET Irish Signals & Systems Conference 2014 and 2014 China-Ireland International Conference on Information and Communications Technologies (ISSC 2014/CIICT 2014)*, 2014, pp. 352–356. DOI: 10.1049/cp.2014.0713.

[82] C. Scarcella, M. B. Marin, S. Détraz, *et al.*, "Next generation of radiation tolerant single-mode optical links for accelerator instrumentation," in *Topical Workshop on Electronics for Particle Physics (TWEPP2018)*, vol. 343, 2018. DOI: 10.22323/1.343.0151.

[83] Hanks, C.L. and Battelle Memorial Institute Radiation Effects Information Center and United States., *Radiation Effects Design Handbook* (NASA contractor report v. 1,No. 1785). National Aeronautics and Space Administration, 1971.

[84] M. Brugger, "Radiation damage to electronics at the lhc," in *Proceedings of IPAC2012, New Orleans, Louisiana, USA*, THPPP006, JACoW Pub., Geneva, Switzerland, 2012, pp. 3734–3736.

[85] B. Todd and S. Uznanski, "Radiation Risks and Mitigation in Electronic Systems," CERN, Tech. Rep., Jun. 2015, 19 pages, contribution to the 2014 CAS - CERN Accelerator School: Power Converters, Baden, Switzerland, 7-14 May 2014, 19 p. DOI: 10.5170/CERN-2015-003.245. arXiv: 1607.01573.

[86] K. Iniewski, *Radiation effects in semiconductors*. CRC Press, 2018.

[87] I. Dawson, *Radiation effects in the LHC experiments: Impact on detector performance and operation* (CERN Yellow Reports: Monographs). Geneva: CERN, 2021. DOI: 10.23731/CYRM-2021-001.

[88] H. W. Wicklein and R. H. Dickhaut, *Transient Radiation Effects in Capacitors and Dielectric Materials*. Boeing Co., Seattle, WA (USA), Aug. 1961.

[89] ISO/IEC JTC 1 Information technology, "ISO/IEC 7498-1," International Organization for Standardization, Tech. Rep., 1994.

[90] I. S. Reed and G. Solomon, "Polynomial codes over certain finite fields," *Journal of the Society for Industrial and Applied Mathematics*, vol. 8, no. 2, pp. 300–304, 1960. DOI: 10.1137/0108018. eprint: https://doi.org/10.1137/0108018.

[91] Rigid Printed Board Cmte., "IPC-6012D," IPC International, Inc., Tech. Rep., 2015.

[92] Panasonic Corporation, "MEGTRON 6, Laminate R-5775, Prepreg R-5670, No.21041931," Tech. Rep., 2021.

[93] B. Razavi, "Challenges in the design high-speed clock and data recovery circuits," *IEEE Communications Magazine*, vol. 40, no. 8, pp. 94–101, 2002. DOI: 10.1109/MCOM.2002.1024421.

[94] IEEE, "IEEE Standard for Floating-Point Arithmetic," *IEEE Std 754-2019 (Revision of IEEE 754-2008)*, pp. 1–84, 2019. DOI: 10.1109/IEEESTD.2019.8766229.

[95] K. Parhi and D. Messerschmitt, "Pipeline interleaving and parallelism in recursive digital filters. I. Pipelining using scattered look-ahead and decomposition," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 37, no. 7, pp. 1099–1117, 1989. DOI: 10.1109/29.32286.

[96] M. Soderstrand, A. de la Sema, and H. Loomis, "New approach to clustered look-ahead pipelined IIR digital filters," *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, vol. 42, no. 4, pp. 269–274, 1995. DOI: 10.1109/82.378040.

[97] ALTERA, "Floating-Point IP Cores User Guide," Tech. Rep., 2016, UG-01058.

[98] J. Postel, "User datagram protocol," Information Sciences Institute, Tech. Rep., 1980, RFC 768.

[99] J. R. Taylor, *Error analysis*. Univ. Science Books, Sausalito, California, 1997.

[100] R. Etchenique and J. Aliaga, "Resolution enhancement by dithering," *American Journal of Physics*, vol. 72, no. 2, pp. 159–163, 2004. DOI: 10.1119/1.1613274.

[101] E. Pedroni, R. Bacher, H. Blattmann, *et al.*, "The 200-MeV proton therapy project at the Paul Scherrer Institute: Conceptual design and practical realization," *Medical physics*, vol. 22, no. 1, pp. 37–53, 1995. DOI: 10.1118/1.597522.

[102] M. Schillo, A. Geisler, A. Hobl, *et al.*, "Compact superconducting 250 MeV proton cyclotron for the PSI PROSCAN proton therapy project," in *AIP Conference Proceedings*, American Institute of Physics, vol. 600, 2001, pp. 37–39. DOI: 10.1063/1.1435191.

[103] FCC Collaboration, "FCC Physics Opportunities," *The European Physical Journal C*, vol. 79, no. 474, Jun. 2019. DOI: 10.1140/epjc/s10052-019-6904-3.

[104] FCC Collaboration, "FCC-ee: The Lepton Collider," *The European Physical Journal Special Topics*, vol. 228, Jun. 2019. DOI: 10.1140/epjst/e2019-900045-4.

[105] FCC Collaboration, "FCC-hh: The Hadron Collider," *The European Physical Journal Special Topics*, vol. 228, no. 4, pp. 755–1107, Jul. 2019. DOI: 10.1140/epjst/e2019-900087-0.

[106] M. Di Castro, M. L. B. Tambutti, M. Ferre, *et al.*, "I-tim: A robotic system for safety, measurements, inspection and maintenance in harsh environments," in *2018 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*, 2018, pp. 1–6. DOI: 10.1109/SSRR.2018.8468661.

Sitography

[S1] CERN Website: https://home.web.cern.ch.

[S2] LHC: https://home.cern/science/accelerators/large-hadron-collider.

[S3] High-Luminosity LHC: https://hilumilhc.web.cern.ch.

[S4] SY-BI-BL: https://sy-dep-bi.web.cern.ch/bl.

[S5] HL-LHC Radiation Levels: https://edms.cern.ch/document/2302154/1.0.

[S6] Paul Scherrer Institut (PSI): https://www.psi.ch/en.

[S7] ATLAS (A Toroidal LHC ApparatuS): https://atlas.cern.

[S8] ALICE (A Large Ion Collider Experiment): https://alice.cern.

[S9] CMS (Compact Muon Solenoid): https://cms.cern.

[S10] LHCb (Large Hadron Collider beauty): http://lhcb.cern.ch.

[S11] GEANT4 Simulation Toolkit: https://geant4.web.cern.ch.

[S12] Intel: https://www.intel.it.

[S13] lpGBT Device Documentation: https://lpgbt.web.cern.ch/lpgbt.

[S14] VFC-HD OH Repository: https://ohwr.org/project/vfc-hd/wikis/home.

[S15] EP-ESE-ME: https://ep-ese.web.cern.ch/structure/section-ese-me.

[S16] Taiwan Semiconductor Manufacturing Company: https://www.tsmc.com.

[S17] Analog Devices: https://www.analog.com.

[S18] Murata website: https://www.murata.com.

[S19] Phoenix PCB Provider: https://www.phoenixpcb.it.

[S20] Agilent: https://www.agilent.com.

[S21] Xilinx: https://www.xilinx.com.

[S22] Rohde & Schwarz: https://www.rohde-schwarz.com.

[S23] Keithley: https://www.tek.com/en/products/keithley.

[S24] MYiR: http://www.myirtech.com.

[S25] ARM: https://www.arm.com.

[S26] lwIP Stack: https://savannah.nongnu.org/projects/lwip.

[S27] MathWorks: https://mathworks.com.

[S28] Texas Instruments: https://www.ti.com.

[S29] Digilent: https://digilent.com.

[S30] Python: https://www.python.org.

[S31] Qt Framework: https://www.qt.io.

[S32] Qt Signals & Slots: https://doc.qt.io/qt-6/signalsandslots.html.

[S33] Wolfram Mathematica: https://www.wolfram.com/mathematica.

[S34] Intel HLS: https://www.intel.co.uk/content/www/uk/en/software/programmable/quartus-prime/hls-compiler.html.

[S35] Xilinx HLS: https://www.xilinx.com/products/design-tools/vivado/high-level-design.html.

[S36] Versatile Link: https://espace.cern.ch/project-versatile-link/public/default.aspx.

[S37] Altium Designer: https://www.altium.com.

[S38] AEC Documents: http://www.aecouncil.com/AECDocuments.html.

[S39] TDK Products: https://product.tdk.com.

[S40] Taiyo Yuden: https://www.yuden.co.jp/eu.

[S41] JEDEC: https://www.jedec.org.

[S42] DLA Land and Maritim: https://landandmaritimeapps.dla.mil.

[S43] CERN Regulators: https://espace.cern.ch/project-DCDC-new.

[S44] Intel FPGA Tools: https://www.intel.co.uk/content/www/uk/en/products/details/fpga/development-tools.html.

[S45] IEEE-802 Standards: https://www.ieee802.org.

[S46] Skyworks: https://www.skyworksinc.com.

[S47] EP-ESE: https://ep-ese.web.cern.ch.

[S48] lpGBT IP Core: http://lpgbt-fpga.web.cern.ch/doc/html.

[S49] Intel FPGA: https://www.intel.co.uk/content/www/uk/en/products/details/fpga.html.

[S50] Wireshark: https://www.wireshark.org.

[S51] SPS (Super Proton Synchrotron): https://home.cern/science/accelerators/super-proton-synchrotron.

[S52] ELMA: https://www.elma.com/en.

[S53] Raspberry Pi: https://www.raspberrypi.com.

[S54] CHARM Mixed Fields: https://kt.cern/technologies/charm-mixed-fields.

[S55] PIF at PSI: https://www.psi.ch/en/pif.

[S56] The European Space Agency (ESA): https://www.esa.int.