UNIVERSITÀ
DEGLI STUDI
DI PADOVA

Sede Amministrativa: Università Degli Studi di Padova
Dipartimento di Ingegneria dell'Informazione

CORSO DI DOTTORATO DI RICERCA IN: Ingegneria dell'Informazione
CURRICOLO: Scienza e Tecnologia dell'Informazione
CICLO XXXI

Design and Development of a Multi-Purpose Input Output Controller Board for the SPES Control System

Tesi redatta con il contributo finanziario dell'Istutito Nazionale di Fisica Nucleare - Laboratori Nazionali di Legnaro

Coordinatore: Ch.mo Prof. Andrea Neviani
Supervisore: Ch.mo Prof. Matteo Bertocco
Co-Supervisore: Ch.mo Dr. Marco Bellato

Dottorando: Davide Pedretti

Design and Development of a Multi-Purpose Input Output Controller Board for the SPES Control System.

In collaboration with:

National Institute for Nuclear Physics (INFN)

Legnaro National Laboratories (LNL)

SPES Project

Davide Pedretti

Padova, September 2018

Contents

List of Figures

List of Tables

List of Abbreviations

ADC	**A**nalog to **D**igital **C**onverter
ATX	**A**dvanced **T**echnology e**X**tended
ALPI	**A**cceleratore **L**ineare **P**er **I**oni (Ion Linear Accelerator)
ASCII	**A**merican **S**tandard **C**ode for **I**nformation **I**nterchange
BGA	**B**all **G**rid **A**rray
BPM	**B**eam **P**rofile **M**onitor
CAN	**C**ontroller **A**rea **N**etwork
CERN	**C**onseil **E**uropéen pour la **R**echerche **N**ucléaire
CEX	**C**OM **EX**press
COM	**C**omputer-**O**n-**M**odule
CPU	**C**entral **P**rocessing **U**nit
DAC	**D**igital to **A**nalog **C**onverter
ECR	**E**lectron **C**yclotron **R**esonance
EEPROM	**E**lectrically **E**rasable **P**rogrammable **R**ead **O**nly **M**emory
EMC	**E**lectro**M**agnetic **C**ompability
EMI	**E**lectro**M**agnetic **I**nterference
eMMC	**e**mbedded **M**ulti **M**edia **C**ard
ENOB	**E**ffective **N**umber **O**f **B**it
EPICS	**E**xperimental **P**hysics and **I**ndustrial **C**ontrol **S**ystem
FEC	**F**ront**E**nd **C**omputer
FFT	**F**ast **F**ourier **T**ransform
FIFO	**F**irst **I**nput **F**irst **O**utput
FMC	**F**PGA **M**ezzanine **C**ard
FPGA	**F**ield **P**rogrammable **G**ate **A**rrays
GTP	**G**igabit **T**ransceiver **P**ort
I^2C	**I**nter-**I**ntegrated **C**ircuit bus
IBERT	**I**ntegrated **B**it **E**rror **R**ate **T**ester
IC	**I**ntegrated **C**ircuit

IEEE	**I**nstitute of **E**lectrical and **E**lectronics **E**ngineers
INFN	**I**stituto **N**azionale di **F**isica **N**ucleare (National Institute for Nuclear Physics)
IOC	**I**nput **O**utput **C**ontroller
IP	**I**nternet **P**rotocol
ISOL	**I**sotope **S**eparation **O**n-**L**ine
LAN	**L**ocal **A**rea **N**etwork
LNL	**L**egnaro **N**ational **L**aboratories
LSB	**L**east **S**ignificant **B**it
LVDS	**L**ow **V**oltage **D**ifferential **S**ignaling
MAC	**M**edia **A**ccess **C**ontrol
MAV	**M**oving **AV**erage filter
MCP	**M**icro-**C**hannel **P**late
MCU	**M**icro**C**ontroller **U**nit
PCB	**P**rinted **C**ircuit **B**oard
PCIe	**P**eripheral **C**omponent **I**nterconnect **e**xpress
PDN	**P**ower **D**elivery **N**etwork
PLC	**P**rogrammable **L**ogic **C**ontroller
PoE+	**P**ower **o**ver **E**thernet **P**lus
PRBS	**P**seudo**R**andom **B**inary **S**equence
PSD	**P**ower **S**pectral **D**ensity
PWM	**P**ulse **W**idth **M**odulation
RAM	**R**andom **A**ccess **M**emory
RTD	**R**esistance **T**o **D**igital
RF	**R**adio **F**requency
RIB	**R**adioactive **I**on **B**eam
RTL	**R**egister **T**ransfer **L**evel
SATA	**S**erial **A**dvanced **T**echnology **A**ttachment
SFP	**S**mall **F**orm-factor **P**luggable
SINAD	**SI**gnal-to**N**oise **A**nd **D**istortion ratio
SNR	**S**ignal to **N**oise **R**atio
SPES	**S**elective **P**roduction of **E**xotics **S**pecies
SPI	**S**tandard **P**eripheral **I**nterface
TCP/IP	**T**ransmission **C**ontrol **P**rotocol / **I**nternet **P**rotocol
TDR	**T**ime **D**omain **R**eflectometer
TIA	**T**rans**I**mpedance **A**mplifier

UART	**U**niversal **A**synchronous **R**eceiver **T**ransmitter
UDP	**U**ser **D**atagram **P**rotocol
USB	**U**niversal **S**erial **B**us
VHDL	**V**HSIC **H**ardware **D**escription **L**anguage
VHSIC	**V**ery **H**igh **S**peed **I**ntegrated **C**ircuit
VME	**V**ersabus **M**odule **E**urocard
VNA	**V**ector **N**etwork **A**nalyzer
VGA	**V**ideo **G**raphics **A**rray
WB	**W**ish**B**one system on chip architecture
JTAG	**J**oint **T**est **A**ction **G**roup

Sommario

Questo lavoro di dottorato è stato svolto presso ai Laboratori Nazionali di Legnaro (LNL), uno dei quattro laboratori nazionali dell'Istituto Nazionale di Fisica Nucleare. La missione principale dei LNL è la ricerca di base nella fisica e astrofisica nucleare, sostenuta da un'importante ricerca relativa alle tecnologie emergenti. La tecnologia e l'innovazione sono i punti di forza che sostengono la ricerca scientifica, lo sviluppo industriale, e guidano il progresso della società in cui viviamo. Le attività di ricerca nell'ambito dell'elettronica e dell'informatica sono essenziali per lo sviluppo del sistema di controllo del progetto SPES (produzione selettiva di specie esotiche). SPES è il progetto piú importante e rappresenta il futuro dei laboratori di Legnaro. Si tratta di una infrastruttura di tipo ISOL (separazione di isotopi in linea), di seconda generazione, il cui obiettivo è quello di generare e accelerare un fascio di ioni radioattivi dedicato alla ricerca nel campo della fisica nucleare (studio di nuclei in condizioni estreme), e ad applicazioni sperimentali in diversi campi della scienza come la produzione di particolari radionuclidi per la medicina nucleare che saranno utili per la diagnosi e la cura di patologie oncologiche. Il progetto del sistema di controllo di SPES sarà basato su tecnologie innovative che consentiranno di monitorare e controllare dispositivi tra loro molto diversi e che eseguono funzioni differenti che vanno dall'acquisizione e visualizzazione dei dati, condivisione dei dati in rete, memorizzazione delle informazioni, operazioni di sorveglianza, diagnostiche e trasporto del fascio. Le applicazioni in tempo reale e la gestione della strumentazione comunemente utilizzata nei sistemi di controllo distribuiti di acceleratori di particelle sono spesso accompagnate dallo sviluppo di sistemi embedded. In questo contesto, il dottorato di ricerca proposto descrive il progetto e la realizzazione di una scheda elettronica di controllo (IOC) multi funzione capace di controllare quasi tutte le apparecchiature coinvolte nel trasporto del fascio di ioni radioattivi. L'idea di base di questo lavoro è quella di estendere il controllo a livello di singola apparecchiatura o piccoli gruppi di dispositivi senza rinunciare alla modularità e alla standardizzazione dell'elettronica. Il risultato del

lavoro di dottorato è un computer embedded multi-funzione progettato con tecnologie all'avanguardia che diventerà lo standard, a livello hardware, su cui si baserà la nuova generazione di computer di frontend del sistema di controllo distribuito di SPES. La scheda IOC, con a bordo uno dei computer-on-module della famiglia COM Express di tipo 6, integra tutte le funzionalità di un computer commerciale e in aggiunta è equipaggiata con un dispositivo programmabile sul campo (FPGA) e alcune periferiche non standard dedicate ad applicazioni di controllo specifiche. L'utente finale potrà sfruttare questa scheda elettronica come un qualunque pc commerciale, oppure, potrà sfruttare le potenzialità della FPGA per le elaborazioni digitali dei dati in tempo reale, per il trasferimento dei dati ad alta velocità su fibra ottica, per chiudere anelli di controllo a larga banda e per avere tempi di risposta agli stimoli in ingresso dal campo deterministici e molto brevi.
Il documento apre con una introduzione sul progetto SPES e sull'architettura e le tecnologie utilizzate nel sistema di controllo di un acceleratore di particelle, prima di descrivere la progettazione, prototipizzazione e validazione della scheda IOC dando particolare risalto alle attività in cui il mio contributo è stato fondamentale. La tesi si chiude descrivendo l'integrazione della scheda IOC nel sistema di diagnostiche di fascio di SPES. Le misure del profilo trasversale e della corrente di fascio eseguite sul campo dimostrano che la nuova elettronica estende la sensibilità di corrente a pochi pA.

Le basi di questo lavoro di dottorato risalgono al mio arrivo ai LNL nel Marzo 2014, quando ho avuto l'opportunità di progettare una schedina di controllo a bassi consumi e basso costo, basata sul modulo Intel Edison. Questo semplice sistema embedded, progettato per controllare le lenti elettromagnetiche necessarie a focalizzare e guidare il fascio nella camera a vuoto, ha confermato i vantaggi che si hanno nel progettare sistemi embedded basati su computer-on-module. Questi moduli rendono accessibili le tecnologie e i processori di ultima generazione per la progettazione di computer embedded. Alla fine del 2014 ispirato dall'esito positivo del progetto della schedina di controllo basata sul modulo Intel Edison, e dalla realizzazione di prototipi basati su schede di supporto commerciali per i moduli COM Express, ho iniziato a lavorare al progetto e alla realizzazione della scheda elettronica IOC multi funzione, successivamente finalizzata durante i tre anni del dottorato di ricerca.
Durante la prima parte del dottorato ho studiato e definito l'architettura della scheda, riassumendo in uno schema a blocchi le funzionalità e le periferiche necessarie per soddisfare i requisiti dei principali sottosistemi di controllo ai LNL. In

seguito, utilizzando programmi di progettazione elettronica professionali ho convertito lo schema a blocchi in uno schema elettrico comprensivo di oltre 2000 componenti. Successivamente, ho esportato lo schema elettrico nell'ambiente di sviluppo del layout assieme ai vincoli elettrici, meccanici e costruttivi della scheda. In questa fase ho curato particolarmente l'analisi di integrità di segnale e di potenza e la reperibilità sul mercato dei componenti attivi e passivi, mentre ho potuto contare sull'esperienza di un professionista esterno per lo sbroglio del circuito stampato. Prima di generare i file di produzione ho eseguito simulazioni post-layout sulle linee di trasmissione differenziali a impedenza controllata e ho analizzato i modi risonanti dei piani di alimentazione della scheda per ridurre la probabilità di crosstalk, disadattamenti di impedenza e modi risonanti che possono generare problemi difficilmente riproducibili durante la fase di prototipizzazione.
Durante la seconda parte del dottorato ho concentrato i miei sforzi sulla caratterizzazione del prototipo. I test funzionali sono stati eseguiti contemporaneamente allo sviluppo del codice VHDL per la programmazione della FPGA. L'obietivo primario di questo firmware è quello di gestire la comunicazione a basso livello con le periferiche collegate alla FPGA e mappare in modo trasparente, nello spazio di indirizzamento del processore, le risorse interne dei dispositivi sotto controllo. L'attività di validazione hardware è stata supportata da misure accurate di jitter, del numero di bit effettivi dei convertitori analogico digitali, di rumore sulle principali tensioni di alimentazione, e dei diagrammi ad occhio dei link seriali ad alta velocità. I problemi emersi durante ai test dei prototipi sono stati corretti nelle due successive revisioni dei circuiti stampati. Ho dedicato l'ultima fase del dottorato di ricerca alla qualifica del primo lotto di produzione di 20 schede IOC e alla loro integrazione nel sistema di diagnostiche di fascio di SPES. Le misure del profilo di fascio eseguite sul campo e l'estensione della sensibilità di corrente a pochi pA confermano che la scheda elettronica progettata è una soluzione affidabile per standardizzare, a livello hardware, il controllo di diverse apparecchiature nel complesso degli acceleratori del progetto SPES. Questa scheda sostituirà la tecnologia VME in diverse applicazioni e sarà la base su cui implementare un sistema di trasporto di fascio automatico e di qualità, fondamentale per il successo delle attività di ricerca ai LNL. L'installazione in campo della scheda elettronica rappresenta una soddisfazione personale enorme e corona questi anni di duro lavoro durante ai quali ho trasformato quella che nel 2014 era solo un'idea, in un computer embedded pienamente funzionante.

Abstract

This PhD work has been carried out at the Legnaro National Laboratories (LNL), one of the four national labs of the National Institute for Nuclear Physics (INFN). The mission of LNL is to perform research in the field of nuclear physics and nuclear astrophysics together with emerging technologies. Technological research and innovation are the key to promote excellence in science, to excite competitive industries and to establish a better society. The research activities concerning electronics and computer science are an essential base to develop the control system of the Selective Production of Exotic Species (SPES) project. Nowadays, SPES is the most important project commissioned at LNL and represents the future of the Lab. It is a second generation Isotope Separation On-Line (ISOL) radioactive ion beam facility intended for fundamental nuclear physics research as well as experimental applications in different fields of science, such as nuclear medicine; radio-pharmaceutical production for therapy and diagnostic. The design of the SPES control system demands innovative technologies to embed the control of several appliances with different requirements and performing different tasks spanning from data sharing and visualization, data acquisition and storage, networking, security and surveillance operations, beam transport and diagnostic. The real time applications and fast peripherals control commonly found in the distributed control network of particle accelerators are accompanied by the challenge of developing custom embedded systems. In this context, the proposed PhD work describes the design and development of a multi-purpose Input Output Controller (IOC) board capable of embedding the control of typical accelerator instrumentation involved in the automatic beam transport system foreseen for the SPES project. The idea behind this work is to extend the control reach to the single device level without losing in modularity and standardization. The outcome of the research work is a general purpose embedded computer that will be the base for standardizing the hardware layer of the frontend computers in the SPES distributed control system. The IOC board is a Computer-on-Module (COM) carrier board designed to host

any COM Express type 6 module and is equipped with a Field Programmable Gate Array (FPGA) and user application specific I/O connection solutions not found in a desktop pc. All the generic pc functionalities are readily available in off-the-shelf modules and the result is a custom motherboard that bridges the gap between custom developments and commercial personal computers. The end user can deal with a general-purpose pc with a high level of hardware abstraction besides being able to exploit the on-board FPGA potentialities in terms of fast peripherals control and real time digital data processing.
This document opens with an introductory chapter about the SPES project and its control system architecture and technology before to describe the IOC board design, prototyping, and characterization. The thesis ends describing the installation in the field of the IOC board which is the core of the new SPES diagnostics data readout system. The beam profile and current measured under real beam conditions prove that this embedded controller extends the current sensitivity to the pA range outperforming the legacy VME technology. The document gives emphasis to those activities where my contribution has been extensive and considerable.

The foundations of this PhD work date back to my arrival at the LNL, in March 2014, when I had the opportunity to design a custom Intel Edison carrier board whose primary aim is to embed the control of small groups of power supply boxes necessary to operate the electromagnetic lenses in the SPES accelerators complex. The project of this low cost and low power embedded board validated the COM approach that makes cutting-edge technology easily accessible for creating low power and cost-effective smart embedded solutions, ideal to replace the legacy technology in many control subsystems at LNL. In late 2014, inspired by the positive outcome of the Intel Edison carrier board project and by prototypes built around commercial COM Express carrier boards, I started working on the design and development of the multi-purpose IOC board, afterward finalized during the three years of PhD studies.
During the first part of the work, I successfully studied and defined the IOC board architecture and I conceptually fixed the on-board features and peripherals intended to satisfy the requirements of the main control subsystems at LNL.
Then, using professional tools, I turned the conceptual design into a detailed electrical schematic that integrates over 2000 components, and, I transferred the electrical schematic to the layout environment together with the collection of all the electrical and physical constraints, the board stack-up and the board outline information. I took care of the signal and power integrity analysis, the power and

thermal management, bill of material, manufacturability and testability aspects whilst I could count on the expertise of a third party company for the place and route of the printed circuit board. Prior to releasing the production files, I performed signal and power integrity post-layout simulations in order to assess the probability of crosstalk, impedance mismatches, plane resonances as well as not optimal thermal distribution problems difficult to be reproduced and debug.
Upon arrival of the first IOC board prototype, I focused on the hardware debug and characterization. The functional tests have been carried out alongside the FPGA firmware development necessary to handle the low-level device access protocol and to provide the end user with a transparent mapping of the process variables of the peripherals under control. The hardware validation has been supported by accurate performance measurements, like the data digitization accuracy, clock jitter, voltage ripple, and high speed data eye diagrams. Several minor issues highlighted during the tests have been fixed with the two subsequent hardware releases which fell mainly under my responsibility.
I dedicated the last months of the PhD to the qualification of the first production lot of twenty IOC boards and their integration in the SPES beam diagnostic data readout and signal processing system. The positive outcome of the beam profile measurements performed under real beam conditions and the extension of the beam current measurement resolution to the pA range prove that the IOC board is a reliable solution to standardize the control of several appliances in the SPES accelerators complex where it will be embedded into physical equipment, or in their proximity, and will control and monitor their operation replacing the legacy VME technology. The installation in the field of the IOC board represents a great personal reward and crowns these years of busy time during which I turned what was just an idea in 2014, into a working embedded computer today.

Acknowledgments

First of all, a special thanks to Dr. Marco Bellato who has made this work possible. He taught me important lessons and gave me all the support that I needed to successfully finalize this PhD project.

I am obliged to Prof. Matteo Bertocco, my referee at the University of Padua. He gave me precious advice and suggestions to improve this document.

Many thanks to Dr. Damiano Bortolato, Mr. Roberto Isocrate, and Dr. Claudio Fanin for their help, their valuable tips and insightful comments necessary to pursue the achievements.

A special thanks to Diego, Jesus, Stefano, Enrico, Davide, Fabio, Giovanni, Marco and all my friends and colleagues at the Legnaro National Laboratory. All of them welcomed me into the control group and gave me precious suggestions. With them all, I shared good moments that made the busy time running away smoothly.

Grazie di cuore alla mia famiglia e alla mia ragazza Paola per avermi sostenuto durante questi anni e per la loro vicinanza. A loro voglio dedicare questo lavoro di tesi.

Chapter 1

The SPES Control System

1.1 Introduction

Particle accelerators play a leading role in the understanding of the universe. We just know about 5% of the universe matter and there are good reasons to think that the study of high energy particle collisions can take us to answer many questions still unanswered. In 1895, Roentgen, while experimenting with a cathode radiation discovered the x-ray. This discovery had a huge impact on the world and laid the foundations of particle accelerator physics.
The X-ray tube shown in Figure 1.1, is a very simple particle accelerator where electrons generated at the cathode are accelerated against the anode by an electrical field. The anode converts the energy of incident electrons into x-ray. Anode and cathode are kept under high vacuum. Starting from this simple model, more and more complex accelerators have been built to propel charged particles to nearly light speed and at very high energies in order to study atomic nuclei far away from stability conditions.

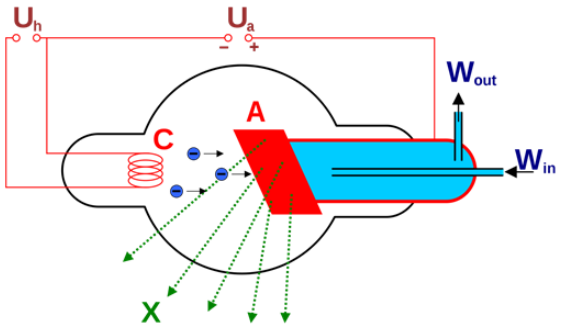

Figure 1.1: X-Ray tube.

1.1.1 Binding Energy and the Mass Defect

The nucleus is composed of a certain number of protons (Z) and neutrons (N). The atomic number Z represents the nuclear charge. Nuclei with the same Z may exist with a different number of neutrons (the isotopes). At least 3000 nuclides have been experimentally characterized and they are all represented in the Segrè chart, Figure 1.2. Stable nuclides require approximately equal numbers of protons and neutrons even if, as you can see in the Sagrè chart, the number of neutrons increases faster than the number of protons. Neutrons play an essential role in stabilizing nuclides, binding together protons inside the nucleus, that otherwise would be repelled by Coulomb force. Starting from stable nuclides one can obtain unstable isotopes by adding or removing neutrons. In an unstable atom, the nucleus changes by giving off radiations; those atoms are said to be radioactive. In other words, radioactive isotopes tend to decay to a more stable configuration emitting energy under radiation form. During radioactive decay the principle of conservation of nucleon number apply; the total number of nucleons (neutrons + protons) must be the same before and after the decay. Radioactive isotopes are often called radioisotopes. All elements with an atomic number greater than 83

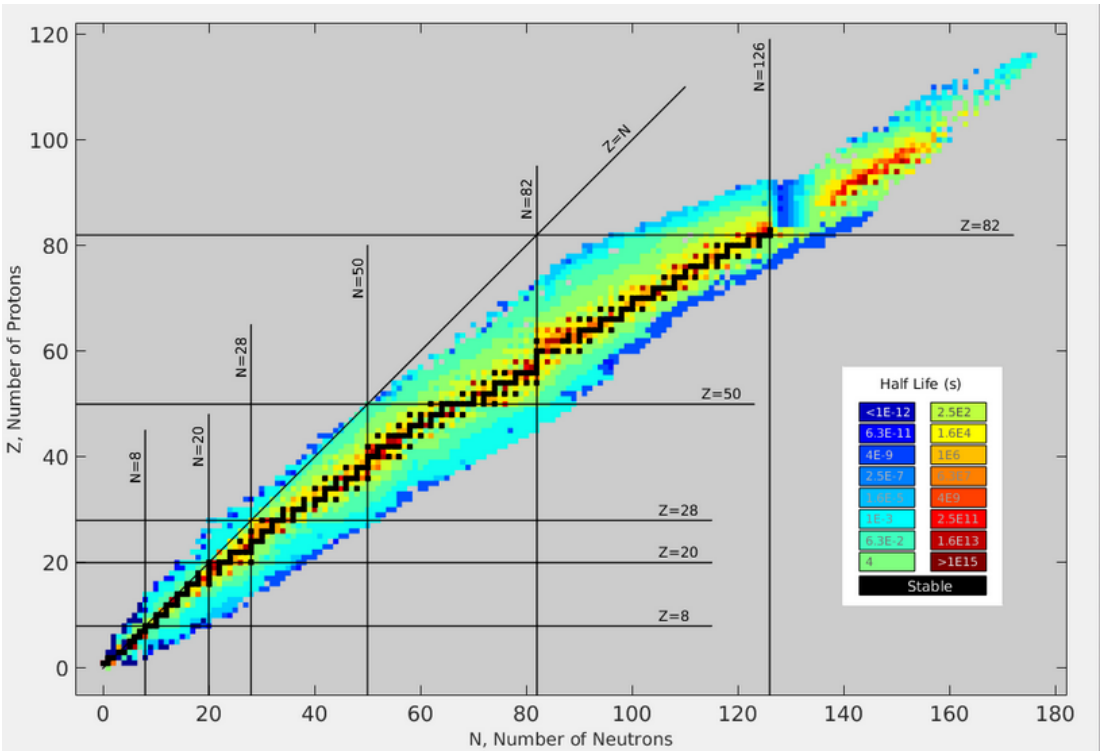

Figure 1.2: Stability valley.

are radioisotopes meaning that these elements have unstable nuclei and are radioactive. Radioisotopes are widely used in medicine, industry and are of particular interest for scientific research. In order to generate a radioisotope from a stable atom, we must somehow counteract the binding energy of the nucleus. The energy of a particle accelerator ranges from tens of MeV, for small machines, up to several TeV in facilities like the Conseil Européen pour la Recherche Nucléaire (CERN). In both cases, the magnitude of these energies is relevant if compared with the binding energy, defined below. The mass of protons and neutrons are often given in terms of an atomic mass unit, where one atomic mass unit (u) is defined as 1/12th of the mass of a carbon-12 atom. The proton mass is 1.0073 u. The neutron mass is 1.0087 u. The carbon-12 atom has a mass of 12.000 u, and yet it contains 12 objects (6 protons and 6 neutrons), each with a mass greater than 1.000 u. Protons and neutrons have larger mass when they are separated than when they are bound together into a nucleus and this is true for all nuclei. The missing mass is known as the *mass defect* and the binding energy can be deduced from here using the Einstein's equation (1.1).

$$E = mc^2 \tag{1.1}$$

In any nucleus, there is an amount of binding energy equal to the mass defect and is essentially the equivalent energy needed to split the nucleus into individual protons and neutrons. The binding energy is measured in MeV and is considerably larger than the energy associated with the binding energy of electrons in the atom. The binding energy per nucleon in the nucleus is on the order of 8 MeV. Therefore, to activate a nuclear reaction an amount of energy equal to or greater than 8 MeV must be transferred to the nucleus. This is the goal of a particle accelerator which increases the kinetic energy of ions before letting them impact each other (collider) or against a target. In the last case, the target nuclei are bombarded by charged particles, for example protons, that activate nuclear reactions of different types depending on a number of parameters including the type of bombarding particle and the energy of the projectiles. Analyzing data generated by these events, scientists study the structure of the subatomic world and the laws of nature governing it. If the incoming projectile has more than 8 MeV of energy, the collision causes other particles to be ejected from the target nucleus. Controlling the target, the beam type and energy it is possible to produce a specific radionuclide.

1.2 The Production of Radioactive Ion Beams at LNL

Legnaro National Laboratories (LNL) [1] is one of the four national labs of the National Institute for Nuclear Physics (INFN). Nuclear physics and nuclear-astrophysics research is the driving force behind the development of particle accelerators at LNL, together with their applications in health therapy and diagnostic techniques. The Selective Production of Exotic Species (SPES) [2] project represents the future of the laboratory and is being commissioned with the primary purpose of producing intense exotic beams; with exotic species, we refer to those nuclides that have no stability. So far, all the experiments performed at LNL were based on stable beams whilst SPES is a second generation Isotope Separation On-Line (ISOL) facility [3]. An ISOL facility is essentially a proton driver which generates a primary proton beam that collides against a target and the nuclear reaction products of interest are afterwords extracted, ionized and selected with an electromagnetic mass analyzer coupled in series. The intense Radioactive Ion Beam (RIB) of short-lived exotic nuclei so generated will be re-accelerated using a superconducting linear accelerator and made available for fundamental nuclear physics research. Indeed, the products of the collisions of those RIBs with stable targets, are rare neutron-rich nuclei similar to those generated in advanced stellar stages and not present on Earth due to their short lifetime. The investigation on such systems is a new frontier of physics for extending our knowledge of nuclei at extreme conditions and to give basic information for the study of nuclear structure, the formation of matter and stellar evolutions. Moreover, these RIBs will be exploited in interdisciplinary fields such as astrophysics, nuclear medicine (radio-pharmaceutical production for therapy and diagnostic), alongside the development of an intense neutron source to test radiation effects on electronics in space and experimental cancer treatments.

1.3 SPES Integration in the LNL Accelerators Complex

Figure 1.3 gives an overview of the SPES facility. The layout plays an important role while transporting short-lived particles from the source to the experimental halls. In blue is the part of the facility completely new, dedicated to the neutron-rich ion beam generation and includes all the equipment foreseen by the ISOL

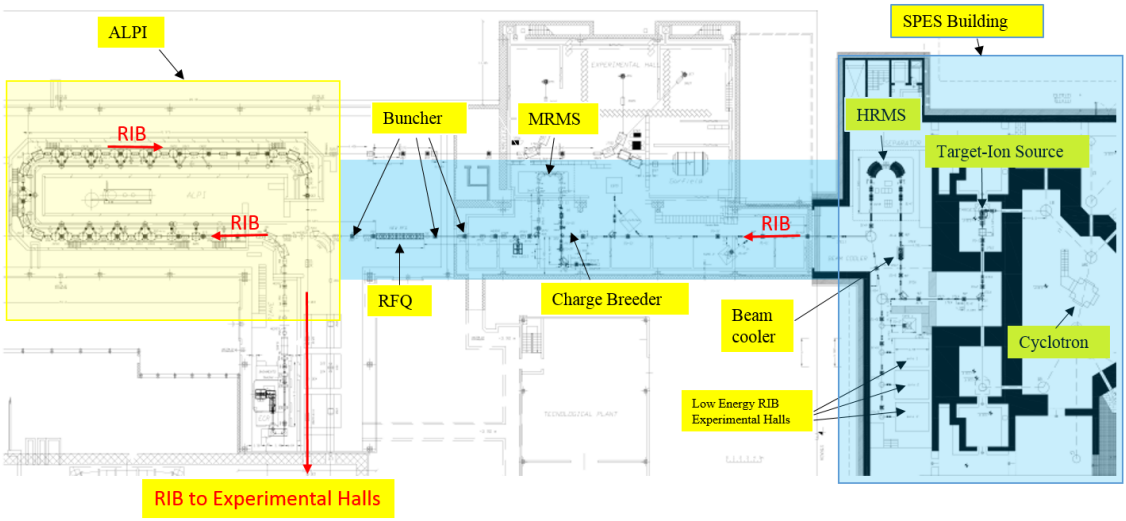

Figure 1.3: The SPES facility layout.

technique. The primary proton beam is generated with the cyclotron provided by the BEST Cyclotron Systems Inc., a member of TeamBest^{TM}. The 70 MeV cyclotron is a particle accelerator in which negative hydrogen ions (H^-) accelerate outwards from the center along a spiral path until they reach the stripper foils, enabling the extraction of currents up to 750 µA of protons in the energy range 35 to 70 MeV. Those energies exceed the binding energy of target nuclei and can be exploited to activate the target.

Figure 1.4 shows the cyclotron installed at LNL. The proton beam extracted by the cyclotron is driven to the target-ion source chamber where collides with 7 uranium carbide UC_x thin disks with a diameter of 40 mm and a thickness of 0.8 mm each. The radioactive isotopes produced in the SPES target by uranium fission, diffuse inside the target material and then effuse through the transfer line in the direction of the ion source, where they acquire the 1^+ charge state needed for their extraction. The fission rate in the target is 10^{13} fission per seconds of about 42 different elements plus neutrons. The neutron flux can be so intense that the bunker walls thickness range from 1 to 4 meters and are designed to accommodate the radioprotection requirements. The RIB is extracted from the target-ion source chamber by the ion extractor electrode and, then, it passes through a Wien Filter [4] for a preliminary separation that decreases as much as possible the contaminants (unwanted nuclides). Continuing in the vacuum chamber, the RIB temperature is reduced by a Beam Cooler in order to decrease the energy dispersion and the emittance. The last instrument in the SPES building is a High-Resolution Mass Separator (HRMS), foreseen to purify the desired beam removing isobar ions com-

Figure 1.4: The 70 MeV cyclotron installed at LNL.

ing from the target.

The so achieved low energy single species RIB can be delivered to the user for low energy experiments, or, it can be post-accelerated by exploiting the existing ion linear accelerator (ALPI) and transported to the experimental halls. To increase the transport efficiency and for a more efficient post-acceleration, an Electron Cyclotron Resonance (ECR) [5] based charge breeding produces a multiple charge state RIB. To avoid beam contamination induced by the impurities presents inside the charge breeder, a Medium-Resolution Mass Separator (MRMS) has been accurately chosen and installed downstream of the charge breeder. The transmission line used to inject the RIB to the ALPI superconducting linear accelerator is equipped with a sequence of bunchers (beam "packetizer"), a Radio Frequency Quadrupole (RFQ) that pre-accelerates the beam to maximize the transport efficiency in ALPI, magnetic dipoles and quadrupole lenses. ALPI represents the SPES post-accelerator stage and increases the RIB energy up to 20 MeV in order to counteract the binding energy of any stable target installed in the experimental halls.

1.4 The SPES Control System Architecture

The control system is a key component for the operation of the machine and integrates into a homogeneous architecture many subsystems that have different levels of complexity and safety requirements, providing a comprehensive means to control and monitor the operation of all the devices involved in the beam transport system. In addition, the control system provides the tools for the remote test and diagnosis of the machine fault conditions. At the beginning of the nineties, advances in computer technology and the increasingly challenging accelerator data acquisition and control requirements led to the development of a standard control system architecture, called *"standard model"*, afterward adopted by several laboratories with few variations [6]. The standard architecture together with standardized hardware and software components provides much greater flexibility easing the expansion and automation of the system. Moreover, the standard model facilitates the installation of new high-performance hardware platforms, the software portability, and resources sharing between laboratories and/or industries.
The SPES control system architecture, summarized in Figure 1.5, complies with the "standard model". The control room hosts the operator's interfaces and database management, provided by Linux workstations and servers. The Local Area Network (LAN) provides the connection between the distributed FrontEnd Computers (FEC) and the operator. The SPES control network will exploit the 802.3 standard Ethernet using the Transmission Control Protocol / Internet Protocol (TCP/IP) or eventually the User Datagram Protocol (UDP). The frontend computers distribute the intelligence and the computational power improving the system reliability, flexibility, and performances, opening the possibility to implement high bandwidth closed loop controls. By moving the data acquisition, data buffering, data processing, interlock and alarms handling and closed-loop controls in the proximity of the I/O devices the performances increase and the data throughput in the LAN is reduced. Signal conditioning and I/O interfaces are the lowest level of the control system and support the physical communication between the FEC and the equipment under control. The FEC and I/O modules, object of this PhD thesis, provide real time support to realize the local feedback on specific components, fast response to outside stimuli and interface to a wide variety of devices and accelerator instrumentation mainly grouped into five subsystems.
The vacuum subsystem produces a high grade of vacuum, below 10^{-7} mbar, necessary to remove obstacles from the particles paths and to avoid RIB contamination.

Various types of magnets are needed to drive and keep the beam well focused in the vacuum chambers. Among them, we have dipole magnets used to bend the trajectory of the beam following the Lorentz principle, quadrupole magnets that act like lenses to focus the beam gathering the particles together, steerer magnets that change the direction of the main lobe of the particle beam. The magnetic fields are generated by permanent magnets or by electrical current flowing through the magnet coils, depending on the strength of the magnetic field to be achieved. The acceleration of charged particles is obtained with radio-frequency cavities that exploit electromagnetic alternating fields (at the frequencies of 80 MHz and 160 MHz) to boost the energy of the particles crossing the cavity. The phase, amplitude, and frequency of the electromagnetic field must be finely tuned and these operations fall under the Radio Frequency (RF) control subsystem [7].
The ALPI cavities operate in superconducting conditions below the critical temperature, very close to the absolute zero. For this purpose, the cryogenics subsystem

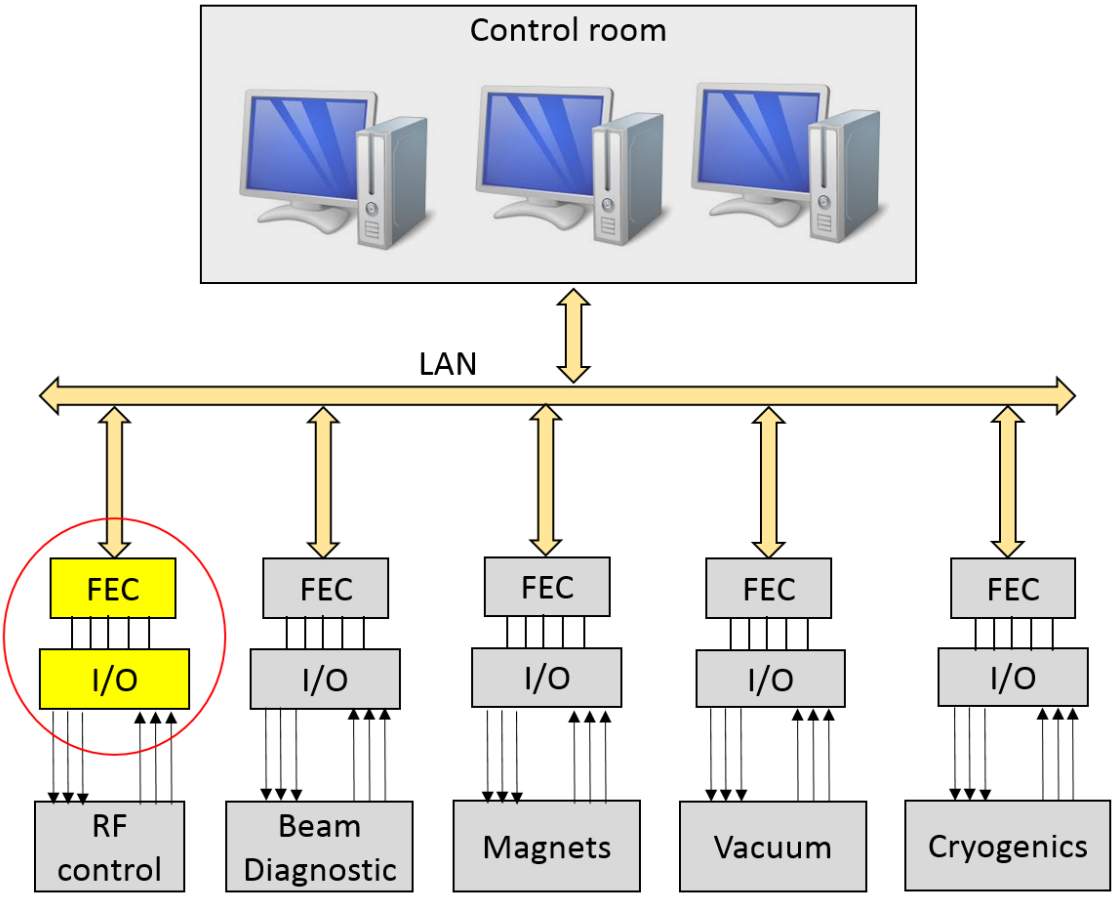

Figure 1.5: The SPES control system architecture.

cools the cavities exploiting liquid helium (He). In superconducting regime, the niobium (Nb) used to manufacture the cavities offers very low resistance to alternate currents. The consequent advantages are in terms of cost saving to run the machine and moreover the maximum electromagnetic field available in superconducting regime is significantly higher.
Another important subsystem integrated into the distributed control network of the SPES project is the beam diagnostics. The diagnostic system provides a feedback to the operator on the beam current and on the longitudinal and transverse beam position inside the vacuum pipe. These beam measurements are essential to reach nominal transport efficiency as well as to detect abnormal operations of the accelerator.

1.5 The SPES Accelerator Technology

Innovative technologies are exploited across all levels of the control system architecture. Figure 1.6 gives an overview of the control system technology required to address always more demanding tasks in terms of networking, data visualization, data storage, data management, computational power and to embed the control of a wide variety of instruments with very different requirements. Networking is the base to build distributed control systems, distributed data pools and distributed computational power. In SPES optical fibers are used to implement long distance backbones at 1 Gb/s, while 100/1000 Mb/s switches and standard copper cables can be used to link the components belonging to the same subsystem. The Experimental Physics and Industrial Control System (EPICS) has been chosen as main control system framework to standardize the operator interface, the data access and management among subsystems [8]. The basic concept underlying the EPICS architecture is a distributed database. All the relevant process variables are collected from the field by the frontend computers. Each FEC hosts a partition of the database and runs an EPICS software input output controller application that implements the channel access server-client protocol necessary to share the process variables between subsystems and the operator. EPICS supports soft real time feedback; to monitor a process, a minimum of 10 Hz update is necessary to offer to the operator an acceptable real time perception. Whilst, high-bandwidth control loops belong to the frontend computer technology. At the lowest layer in the control system architecture, depending on their application, the FECs are represented by a wide range of commercial modules and custom hardware platforms.

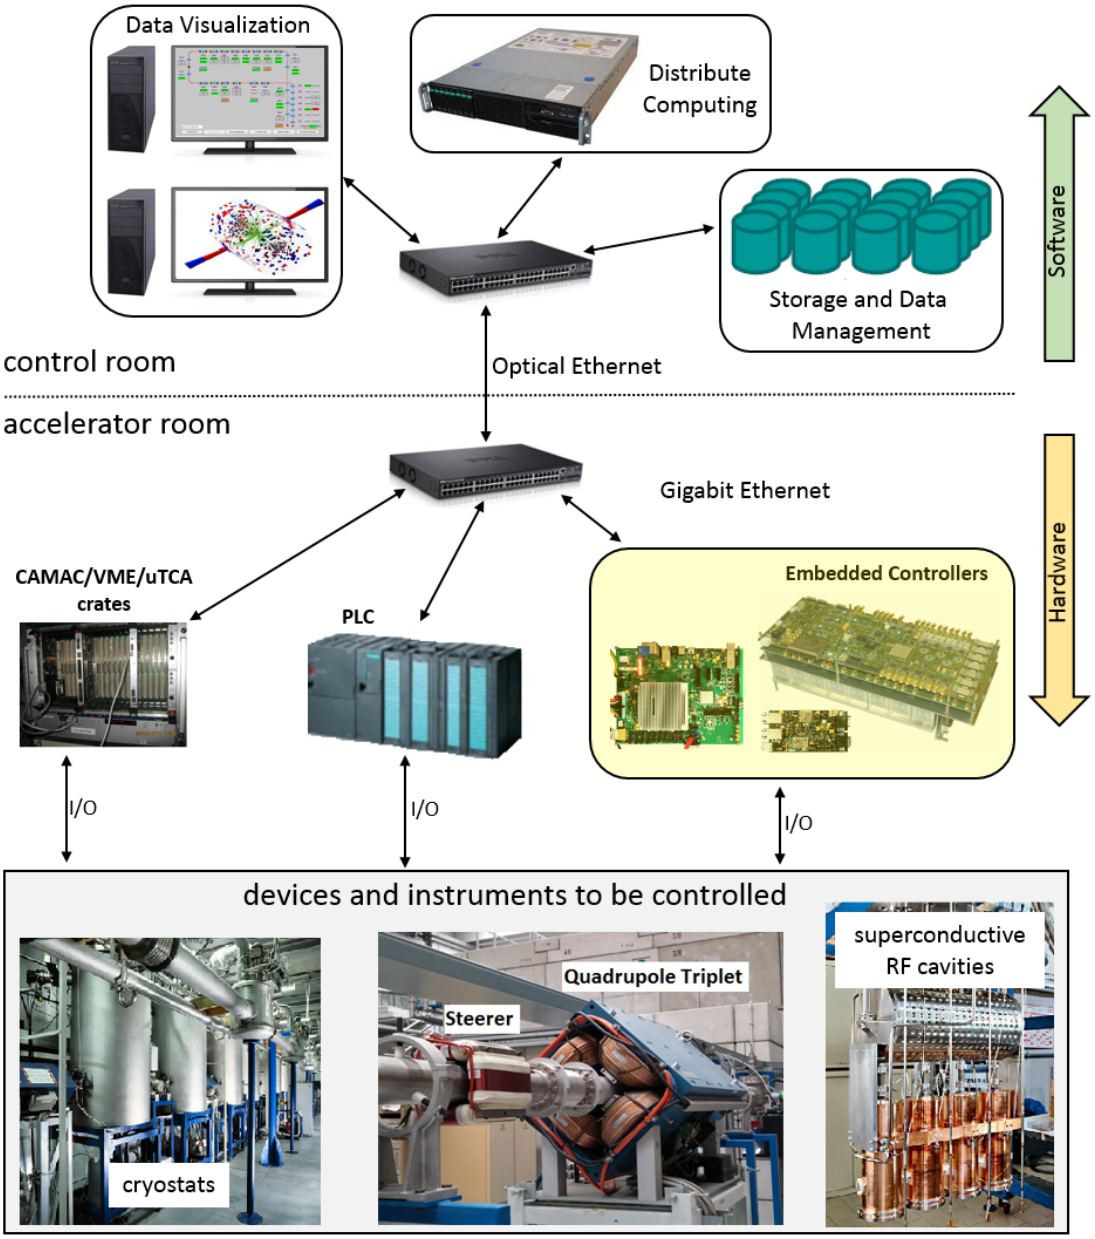

Figure 1.6: SPES control system technology.

Programmable Logic Controllers (PLCs) are an example of commercial hard real time controllers that integrate a Central Processing Unit (CPU) together with input and output devices in a compact and reliable solution capable of operating in harsh environments. In SPES, the PLCs are used for the vacuum and radio-protection subsystems, where the reliability is the primary concern.
The Computer-Aided Measurement And Control (CAMAC), the Versabus Module Eurocard (VME) and more recently the Advanced Telecommunications Computing

Architecture (ATCA) are all example of standard bus and modular crate electronics for data acquisition and control that are vastly used in particle physics. These standard bus allow the data exchange between plug-in modules and still represent an effective way to standardize the hardware layer of the control system. A large number of commercial I/O cards and single-board computers matching these form-factors are available from independent manufacturers.
Alongside these commercial solutions, the real time applications, fast peripherals control and the necessity to cope with tight synchronization requirements commonly found in the distributed control system of particle accelerators are accompanied by the challenge of developing custom embedded controllers endowed with leading semiconductor industry products like the Field Programmable Gate Arrays (FPGA). These custom developments are no longer limited to the most demanding equipment or non standard peripherals control not supported by commercial modules, but, the trend in our laboratory, is to design smart and cost-effective embedded solutions to replace the legacy technology in many applications. At the same time, the complexity of these control networks and the need for maintainability and modularity call for hardware standardization. In this context, the aim of this PhD research activity has been the design and development of a multi-purpose Input Output Controller (IOC) board capable of embedding the control of many equipment involved in the automatic beam transport system of the SPES project, becoming the base for hardware standardization at LNL where it will replace legacy VME technology in several applications.

Chapter 2

The Multi-Purpose Input Output Controller

The SPES distributed control system foresees the interconnection of application specific controllers enriched with computational and networking capabilities that exchange information to perform complex tasks. So far at LNL, the approach followed to embed the control of different devices has been the adoption of a VME based field bus containing a CPU board and booting a real time operating system where required, plus various application specific I/O modules for interfacing with the devices under control. The development of dedicated controller boards based on Digital Signal Processors (DSP) and FPGA technology has always been restricted to specific applications that realize complex control functions or with severe constraints in term of response time and synchronization. The Internet of Things (IoT) era and the Computer-On-Module (COM) approach represent a wind of change at this regard, making cutting-edge technology easily accessible and lowering the entry barrier for prototyping and creating custom hardware platforms endowed with off-the-shelf connectivity solutions, high-density memory capability, smart user interfaces as well as low power computational resources. On my arrival at LNL, in 2014, there was a strong intention to modernize the SPES control system extending the control reach to small groups of devices exploiting custom embedded computers to replace the legacy VME technology in several applications. The importance of the hardware standardization to facilitate the post-production maintenance, obsolescence, and support aspects led us to the design of a multi-purpose Input Output Controller (IOC) board that is general enough to satisfy the control and monitoring requirement of almost all equipment installed in the SPES facility, becoming the base for hardware standardization at LNL.

2.1 The IOC Board Precursor

The Intel® Edison computing platform, shown in Figure 2.1, is an example of how the COM approach may change the way we look at embedded electronics. The Intel® Edison module foresees a 22 nm Intel® Silvermont Atom Processor, dual-core, 500MHz, x86 architecture. Moreover, this module integrates a 32-bit single-core Intel® Quark™ 100 MHz MicroController Unit (MCU), 1 GB of Random Access Memory (RAM) and 4 GB of embedded Multi Media Card (eMMC) flash storage. This tiny platform, slightly bigger than a coin, wraps a robust set of a typical desktop pc features and makes standard I/O solutions readily available in a 70 pins connector, alongside Bluetooth 4.0 and WiFi 802.11 a/b/g/n connections. The Intel® Edison module opens the possibility to create an embedded computer which outperforms the simple low cost microprocessor-based boards and bridges the gap between personal computers and custom embedded platforms.
At LNL, the Intel® Edison module is the core of a low-cost and low-power computer board, shown in Figure 2.2, exploited to control the power supplies of the electromagnetic beam lenses. On my arrival at LNL, most of my effort has been dedicated to the design of this carrier board which acts as an Ethernet to RS422 interface translator and is a good solution to control the magnets power supplies accessible via a serial interface [9]. This board implements two RS232/RS422/RS485 hardware configurable interfaces (one acting as a serial console by default), two Universal Serial Bus (USB) 2.0 type A ports/receptacles, a microSD card interface, it is Power over Ethernet Plus (PoE+) compatible, and features a fast Ethernet two ports network switch capability with the Ethernet over USB support enabled by default. In addition, 20 digital general purpose I/O, 3.3 V or 5 V selectable, are available to the end user on a flat cable connector. The possibility of bootstrapping the desired Linux distribution, or to build a custom Linux-based operating

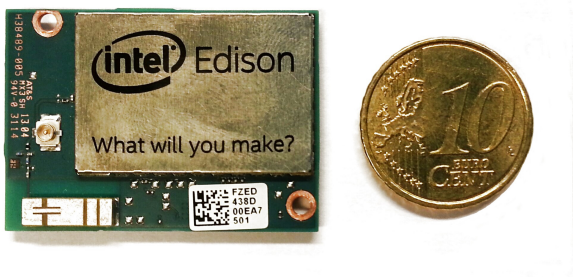

Figure 2.1: Intel® Edison computing platform.

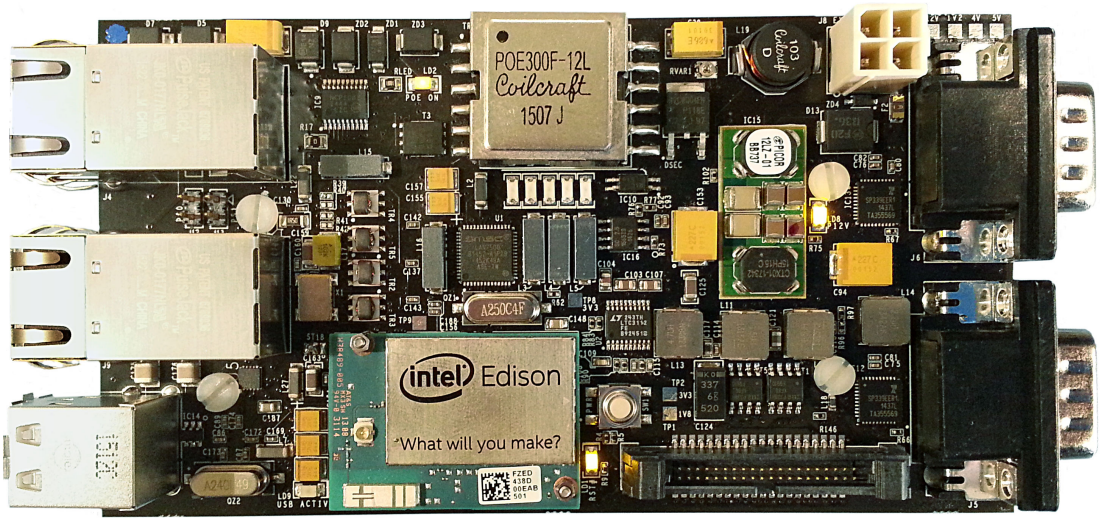

Figure 2.2: Intel® Edison carrier board.

system with the Yocto project tool-set [10], or to make a bootable micro SD card, represents a substantial advantage with respect to embedded operating systems. The porting of the EPICS application and corresponding database into the Edison platform has been straightforward. Real time peripheral control capabilities are provided by the integrated MCU. This low cost PoE+ computer board in din-rail mounting form factor, entirely designed at LNL, is a smart and cost-effective solution with respect to the legacy VME crate. At the same time, it is an example of ad hoc embedded controller targeted for the magnets power supply application. The on-board resources are limited and a different application with different control and monitoring requirements would probably demand a new design with different I/O solutions. This computer board validates the COM approach and it is the IOC board precursor, indeed, the considerations outcoming from the Intel Edison carrier board project constitute the guidelines for the design of the multi-purpose IOC board described in this thesis.

2.2 IOC Board Design Guidelines

The IOC board is thought to be general enough to satisfy the control requirements of different subsystems, becoming the base for hardware standardization at LNL. The main advantages that justify the development of the IOC board against commercial solutions are reflected in the following design guidelines.

- **Standardization**: the IOC board is based on the COM EXpress (CEX)

standard [11]. It is the highest performance COM standard and its production will not be discontinued in the coming years. Common connectors and mounting holes are maintained between different subsequent revisions allowing newer COM Express modules to be backward compatible and improving the carrier performance as they become available on the market. This increases the custom carrier board lifetime. Almost all standard pc peripherals are integrated in the COM Express module and made available to the user via two 220-pin connectors.

- **Modularity**: it is possible to extend the IOC board functionalities with commercial solutions. The two commercial form factors of interest are the Peripheral Component Interconnect express (PCIe) [12] and the Low Pin Count FPGA Mezzanine Card (LPC FMC).

- **Flexibility**. Alongside all generic pc features the IOC board is endowed with more application specific resources and peripherals. It targets different applications in our laboratory and must be general enough and on-field customizable in order to satisfy simple slow control monitoring operations as well as complex data processing requiring a high-level of computational power.

- **FPGA**. The carrier board hosts an FPGA to provide support for those application specific peripherals not commonly supported by a desktop pc. The end user may exploit the FPGA to perform fast peripherals control operations, data acquisition, data buffering and data processing, synchronization operations, and high bandwidth control loop, freeing the Operating System from real time tasks.

- **Hardware abstraction layer**. The possibility of booting any standard Linux distribution, or to run Windows out of the box, together with the Intel x86-64 architecture represent substantial advantages for the software development. The user can deal with a general purpose pc with a high level of hardware abstraction. The software development is independent of the underlying hardware and the software portability from legacy platforms currently in use at LNL is straightforward. The IOC board architecture minimizes the workload required to develop the gateway software necessary to interface the on-board resources to the control network.

2.3 IOC Board Main Features

A graphical representation of the IOC board main features is given by the high level block diagram shown in Figure 2.3. The main components and peripherals listed in the block diagram are intended to satisfy the requirements of the main control subsystems at LNL. At the heart of the multi-purpose controller, there are the COM Express module and the FPGA that transfer to the board computational power capabilities at different levels and with different approaches. It is a custom motherboard enriched with a Spartan-6 that assists the main processor in performing real time tasks, fast peripherals control, data acquisition, digital data pre-processing and data buffering. The peripherals subset linked to the FPGA is completely independent of the COM Express and the general purpose pc interfaces, hence, FPGA and COM Express can be used as independent units, also mutually exclusive, or tightly coupled as most of the applications at LNL require.

The IOC board populated with both the FPGA and the COM Express fits perfectly with a distributed control system architecture. In this topology, the COM

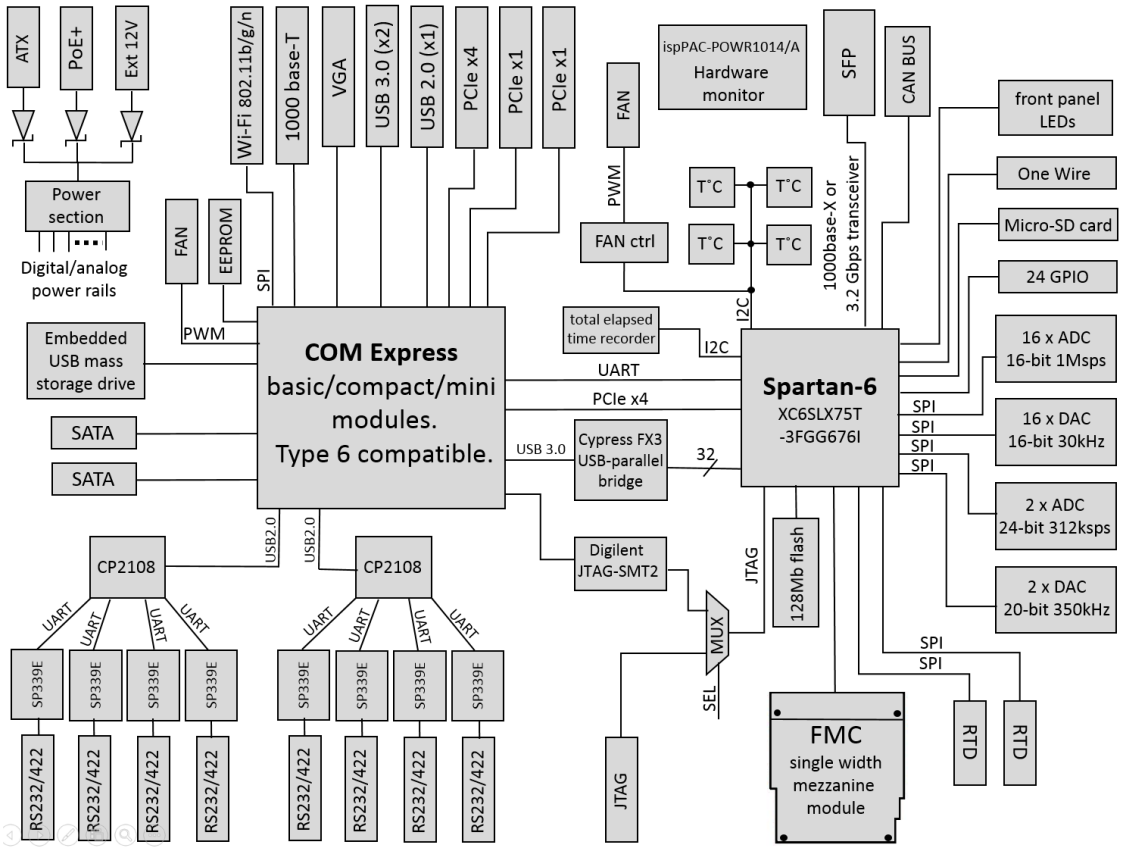

Figure 2.3: IOC board block diagram.

Express loads a database segment whilst the FPGA guarantees a fast response to input stimuli from the field. The communication between these two units is ensured by three independent buses whose choice of use depends on the data throughput and on the device driver support available:

- A simple serial Universal Asynchronous Receiver Transmitter (UART) link, that is a suitable choice for slow control operations and debug purposes.

- A PCIe 1.0 x 4 link that implements message-based point-to-point interconnections at 2.5 Gbps speed with Low Voltage Differential Signaling (LVDS). The PCIe endpoint core inferred into the FPGA can exploit lane widths from x1 to x4.

- USB 2.0/3.0 to 32-bit parallel bus bridge based on the Cypress CYUSB3014 Integrated Circuit (IC).

Device-to-device and device-to-cloud connectivity are of crucial importance for a custom platform acting as a local intelligent node in a distributed control system. The gigabit Ethernet controller is integrated into the COM Express module while the Wi-Fi connectivity is provided via a commercial module (GS2100MIP) that communicates with the host via a Standard Peripheral Interface (SPI). The Wi-Fi module supports the IEEE 802.11b/g/n standards and it is ideal for those applications requiring galvanic isolation.
The IOC board represents a practical solution to centralize the control of those applications not involved in safety or high-reliability processes where the control tasks are reduced to simple analog and digital I/O operations. In these centralized control topologies, the COM Express is not plugged into the carrier board. The FPGA handles the data acquisition and provides a transparent mapping of the relevant process variables to the remote central host. The digital and analog I/O resources are remotely accessible via an optical Ethernet implementation based on the UDP protocol, at the network layer, and supported by the Small Form-factor Pluggable (SFP) module directly coupled with the FPGA gigabit transceiver, at the physical layer. The central host is well represented by an open source server virtualization platform, like Proxmox VE [13], where each virtual machine has private virtualized hardware and network resources and run a EPICS segment database.

2.3.1 The COM Express Standard and General PC Peripherals

The COM Express® standard defines a family of small form factor computer on modules suitable for a wide range of applications. Figure 2.4 shows the top and bottom views of the compact CEX used in several applications at LNL. These modules can be used as standalone single board computers or, as we have done, as a processor mezzanine plugged into custom carrier boards. The CEX standard defines different module types and form factors. The type specifies the module pinout configuration and features available, whilst the form factor refers to the module size. The IOC board designed complies with the type 6 standard and can host the mini, compact or basic form factors, each available with different processors fulfilling the computational power requirement of varied applications [14]. In most applications at LNL, the data acquisition and real time processing are implemented by the FPGA, shrinking the CEX task to a database application. In this cases, the compact module form factor endowed with the low-power Intel Atom processor represents a suitable solution. The type 6 pinout closely matches the I/O features set of common motherboards. The final subset of these general purpose pc peripherals effectively available on the carrier board depends on the module plugged-in and on the carrier board design itself. In the IOC board, worth of notice is the presence of an analog Video Graphics Array (VGA) port, the 10/100/1000

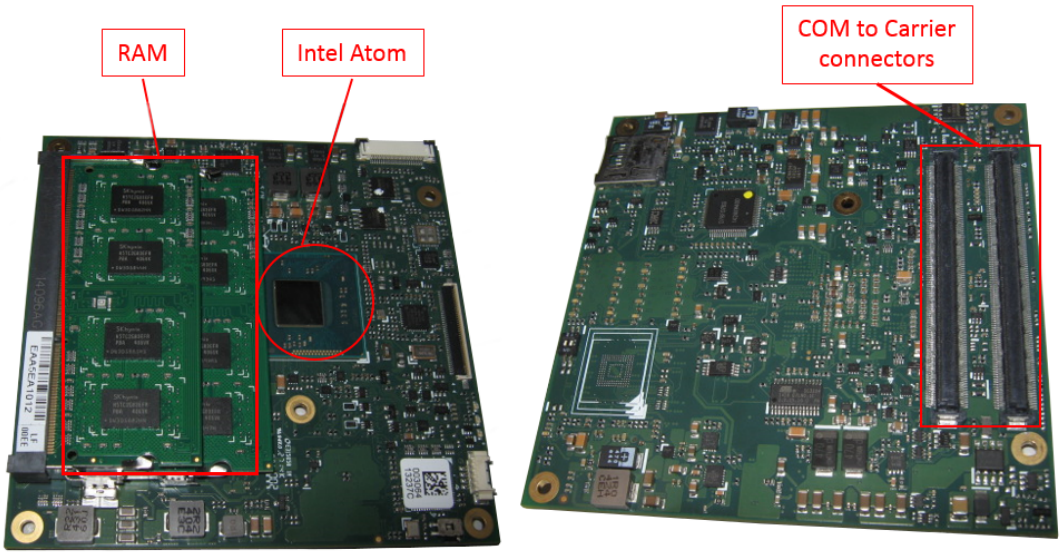

Figure 2.4: COM Express cExpress-BT type 6 module. Top and bottom views.

Gigabit Ethernet interface, three USB 2.0 type A receptacles among which one supports the super-speed data transfer (USB 3.0), two Serial Advanced Technology Attachment (SATA) 3 Gbps ports, one 4-wire FAN control connector, seven PCIe lanes: 4 lanes implement a PCIe x4 communication channel to the FPGA, and the remaining are linked to the three PCIe slots that open the possibility to extend the board functionalities with commercial modules. The IOC board does not exploit the audio, the Digital Display Interface (DDI) and the PCI Express Graphics (PEG) capabilities of the CEX module because the final installation on the field only requires the remote access to the console of the embedded computer. The IOC board is compatible with the Advanced Technology eXtended (ATX) power supply standard to ensure the wake-on-LAN capability when supported by the CEX module plugged in.

2.3.2 Serial Interfaces

Serial interfaces are commonly found in most embedded controllers since many electronic equipment can be fully controlled through the serial port. The IOC board is equipped with eight serial ports. Each of the eight serial links ends on a multiprotocol transceiver that supports RS-232, RS-485, RS-422 serial standards, hardware or software selectable by the user. These transceivers feature standardized logic levels and robust ElectroStatic Discharge (ESD) protection.
The RS-422 standard is vastly used in our laboratory because it allows data rates up to 10 Mbit/s through differential signaling over cables long hundreds of meters. The extension of the transmission range to hundreds of meters is key in preserving the electronics from radiation during the RIB transport and to ensure the accessibility of the IOC board during beam time. In the magnet power supply control application, the IOC board will replace its precursor: the Intel Edison carrier board. As shown in Figure 2.5, in those applications where the IOC board is devoted to the serial interfaces control, the carrier can be partially populated, e.g. non assembling the FPGA and the related peripherals, thus, significantly reducing the production costs. In the IOC board, the serial interfaces are implemented in eight 10-pin headers connected via flat cables to the standard D-subminiature (D-sub) 9-pin connectors installed in the IOC box rear panel. One IOC board serves up to 8 serial instruments in a point-to-point communication topology. Nevertheless, in RS-422 configuration, each serial port may drive a multi-drop bus extending the control range up to 10 slaves per serial channel. The final installation on the field

depends on the accelerator layout and on the cables layout.

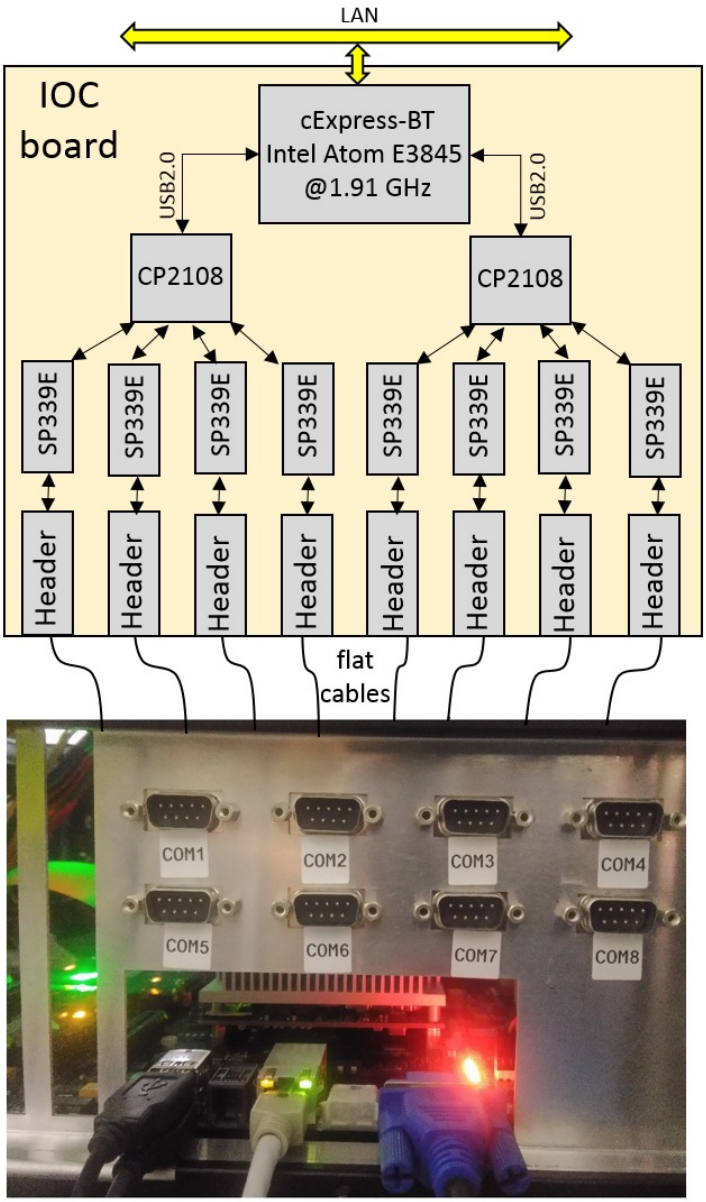

Figure 2.5: IOC box rear panel serial ports.

2.3.3 IOC Board Boot Options

The system boot process, that may be complex when considering various software and hardware implementations options in embedded systems, is not a concern for the IOC board. Indeed, the BIOS flash resides on the COM Express module and the outcome is a self-booting solution. The possibility of booting up Windows, Linux

or any operating system is an important advantage for an embedded computer board. The final user can counts on reliable standard operating systems, and, on a vast selection of software and drivers support. The boot options supported by the IOC board are:

- the two SATA-150 connectors and external Solid State Drive (SSD) disks.
- Network booting, useful to centralize the management of disk storage reducing the maintenance costs.
- On module micro-SD card and eMMC flash storage (up to 64 GB) when supported by the chosen CEX module.
- Embedded USB 2.0 mass storage drive. Figure 2.6 shows a tiny NAND flash memory that suits for low power embedded applications requiring a custom minimal operating system or a standard Linux distribution. More in general, any USB bootable flash drive is supported. These USB flash drive solutions are the most suitable to create system backup and to recover the system from failures like a desktop pc.

Figure 2.6: Embedded USB mass storage drive.

2.3.4 FPGA and Application Specific Interfaces

The IOC board design exploits the Spartan-6 XC6SLX75T device in a FGG676 package. The choice of the FPGA has been tailored as a tradeoff between cost, performance, power consumption, number of I/O, and availability on the market. The Spartan-6 family is built on a mature 45 nm process technology and provides leading integration capabilities and performance at low cost. The LXT version

chosen features power optimized high speed serial transceivers that implement PCIe endpoints and the optical Gigabit Ethernet interconnection. The Spartan-6 FPGA family is readily available on the market and Xilinx foresees shipments to at least 2027. This is a guarantee of the longevity of the IOC board that will be the base for hardware standardization in our laboratory. The available user general purpose I/O provided by the FGG676 package are 348, plus 8 high speed transceivers, enough to satisfy the I/O requirements of the on-board peripherals. The package pinout is compatible with three different devices of the Spartan-6 family: XC6SLX75T, XC6SLX100T, XC6SLX150T. The majority of applications at LNL will be addressed using the XC6SLX75T device.

Digital I/O

The IOC board features 24 general purposes single ended digital I/O linked to the FPGA through a bidirectional level translator. The logic level is 3.3 V or 5 V selectable, and, each I/O is endowed with a pull-up and a pull-down resistor to cope with different logic families. These I/O are available in a 34-pin connector header and can be exploited to quickly develop many standard digital interfaces like the Inter-Integrated Circuit bus (I^2C), UART, and SPI, or custom protocols, or to implement Pulse Width Modulations (PWM). These I/O can be exploited also as programmable industrial I/O if coupled with external isolators that feature the required 24 V level translation.

Analog I/O

The analog I/O channels make the IOC board a complete data acquisition system for real time applications. The custom carrier board includes:

- 16 single ended analog input channels, 16-bit resolution at 1 Msps. Each successive approximation Analog to Digital Converter (ADC) on-board has an independent signal conditioning circuit that features the single ended to differential conversion and extends the analog input range to $\pm 10V$.

- 16 single ended analog output channels, 16-bit resolution, 30 kHz output bandwidth (calculated using the maximum settling time specified by the manufacturer). The unipolar or bipolar output range of the Digital to Analog Converters (DAC) is software selectable from + 5 V, + 10 V, $\pm 5V$, $\pm 10V$, $\pm 10.8V$.

- 2 single ended analog input channels, 24-bit resolution, 312 ksps. These high-performance sigma-delta ADCs implements an internal finite impulse response filter and a pin selectable decimation rates of 64x, 128x, and 256x to address the most demanding data acquisition systems.

- 2 single ended analog output channels, 20-bit resolution, 350 kHz output bandwidth. These 1 ppm accuracy DACs operate in a gain of two configuration and thanks to the on-board signal conditioning circuitry feature a $\pm 10V$ output range.

Optical Link

The application specific SFP port opens the possibility to plug optical module transceivers to accomplish telecommunication and data communication applications. So far, this key feature has been exploited to implement a gigabit optical Ethernet interface. The SFP connector is linked to the FPGA transceiver and the maximum sustainable data rate is 3.125 Gbps.

CAN Bus

The Controller Area Network (CAN) is a message-based broadcast system designed to allow microcontrollers to communicate with each other in robust and flexible control networks with limited transmission ranges (tens of meters) [15]. Many industrial embedded computers feature CAN bus capabilities to monitor and operate instruments, and, since many devices involved in the beam transport system (especially in power distribution systems) include a CAN port, we decided to integrate the CAN bus peripheral into the IOC board. Moreover, since the connection of the COM Express to the CAN bus would require a commercial PCIe communication card with proprietary drivers, we decided to integrate into the IOC board a standalone CAN controller with SPI interface that communicates with the Spartan-6. This CAN controller chip, coupled with a CAN transceiver chip which satisfies the standard physical layer requirements, bridges the SPI bus to a CAN V2.0B, 1 Mb/s bus. In this way, in order to access a CAN device, the final user can decide either to exploit commercial PCIe solution when readily available or to integrate an SPI master core within the FPGA, depending on his expertise and convenience.

One Wire

One-wire is a serial protocol that exploits a single data line plus a ground reference for half-duplex bidirectional communication. It is the most economical solution to implement simple identification procedure, slow control operations, temperature measurements, manufacturer ID, status and calibration parameters monitoring. One-wire is a practical solution that may target simple applications and therefore we endowed the IOC board with a single channel one-wire master chip that bridges the I^2C master core running into the FPGA with any one-wire slave device.

MicroSD

The IOC board has a micro secure digital (micro-SD) card port that allows a micro-SD card to be connected to the FPGA. This solution is useful to preserve non-volatile information in those applications where the COM Express is not plugged into the carrier and the FPGA is the only on-board intelligent unit. Relevant information like calibration or application parameters, the Media Access Control (MAC) or Internet Protocol (IP) addresses, the IOC board identification code, can be loaded from a dedicated partition in the configuration Electrically Erasable Programmable Read Only Memory (EEPROM), but, the micro-SD card represents a further level of flexibility, opening the possibility to extend the memory capacity to several GB and to easily transfer data between two IOC boards on the desk. The micro-SD card interface will be handled by a custom or third party logic core instantiated within the FPGA, and an on-board high speed logic level translator provides the necessary signal level shifting.

FMC

The IOC board is a low-pin count FMC carrier board. The FMC slot has been designed in compliance with the ANSI/VITA 57.1-2008 standard [16]. Whilst almost all the FMC carrier boards on the market are available in PCIe or VME form factors, the IOC represents a compact solution that integrates into one single board the host processor, the FPGA and the LPC FMC connector suitable to extend the board functionalities using commercial or custom modules. Section 5.6 introduces a typical use-case of a commercial mezzanine card in the SPES control system.

Temperature Monitor and Fan Speed Control

The IOC board is able to monitor the local system temperature thanks to four on-board digital temperature sensors accessible from the FPGA via an I^2C bus. In addition, a fan speed controller provides two thermistor inputs and uses the temperature data to control a PWM output to adjust the speed of a cooling fan. The over-temperature flags are connected to a programmable hardware monitor chip that monitors the voltage levels of the main power rails inside the board and is able to start the soft-shutdown or hard-shutdown procedures. During the soft-shutdown procedure only the COM Express is forced into sleep mode, and, if the over-temperature alarm does not cease within a programmable time lapse, the hard-shutdown starts and switches-off the whole carrier board. The IOC board is also able to monitor the external temperature of any equipment installed in its proximity thanks to the two Resistance To Digital (RTD) converters optimized for platinum resistance temperature sensors.

2.3.5 The JTAG Chain and FPGA Remote Configuration

The Spartan-6 FPGA can load itself from an external SPI flash memory, or, the application specific configuration data (bitstream) can be loaded directly into the FPGA's internal memory by an external pc, via the standard 4-wire Joint Test Action Group (JTAG) configuration link. During the prototyping phase, a dedicated on-board connector provides a means of directly access the JTAG chain, from a local pc running the iMPACT software, using the Xilinx's USB platform cable [17][18]. Nevertheless, dedicated hardware support is necessary to preserve the

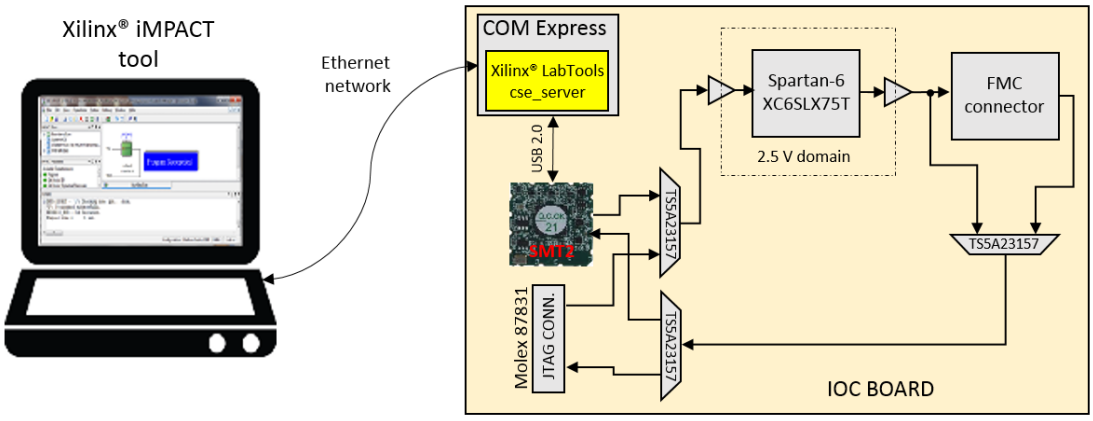

Figure 2.7: JTAG chain and FPGA remote reconfiguration and debug tools.

reconfiguration capability after the installation on the field when the electronics accessibility is a concern, or, sometimes, inhibited. Firmware may get corrupted during the reconfiguration phase itself and also stall situations due to bugs in the firmware may occur. Therefore it is advisable to be able to access the JTAG chain remotely via the Ethernet network. Moreover, the remote access to the JTAG chain via Ethernet opens the possibility for remote debugging; the operator is able to printout on the remote screen the internal FPGA signals, real time, and without access physically the JTAG pins. The designed solution is shown in Figure 2.7.
The Digilent's JTAG-SMT2 surface mount programming module together with a couple of analogue switches bridge the COM Express USB bus to the physical JTAG chain. The SMT2 module can be accessed directly from all Xilinx tools, *cse_server* application included. This executable runs in the COM Express and opens a TCP port with a default value of 50001. The local computer running iMPACT or ChipScope tools must be configured via the cable setup dialog box to send the JTAG commands to the remote host by settings the IP address and TCP port field in agreement.
Multiple chips on a electronics card can have their JTAG lines daisy-chained together. In the IOC board the FMC mezzanine card, optionally, can be included in the JTAG chain in agreement with the FMC specifications. If there is not a mezzanine plugged in, or if the mezzanine does not integrate chips requiring reconfiguration and debug operations via JTAG, the FMC connector can be bypassed. It is worth of notice that different devices may use different I/O technologies that comply with different digital logic levels; on the IOC board this issue has been addressed using voltage level translators where necessary.

2.4 Power Distribution

One of the most important consideration in high speed and mixed analog digital board design is the power distribution. Upon completing the definition of the functional blocks to be integrated into the IOC board, we have designed a detailed power distribution scheme. The block diagram in Figure 2.8 shows the main backbones of the IOC board power tree with the related current limits. The IOC board can be powered by any ATX power supply, commonly used in desktop pc, or by an external 12 V DC power source with ±5% tolerance, or it may behave as a Power over Ethernet device. The choice of the power supply to be used is dictated by the final installation in the field and by the overall power consumption that depends a

lot from the COM Express module chosen and the additional commercial modules plugged into the carrier. There are two main 12 V rails, one is reserved to supply the COM Express module while the other feeds the whole carrier power distribution network; all the necessary voltages are generated on-board. The power rails have been grouped in three main modules that guarantee an optimal load distribution.

The power rails defined in the analog section feed the ADCs and DACs and the related signal conditioning circuitry. The supply voltages defined within the FPGA power section are routed exclusively to the Spartan-6 and the FMC mezzanine connector. The FPGA power section deserves special attention since it addresses the Spartan-6 DC and switching characteristics as specified by the manufacturer [19]. When designing multi-functions controller like the IOC board, the FPGA power requirements have to be determined early in the design cycle, even before the logic inside the FPGA has been developed. The worst-case power analysis has

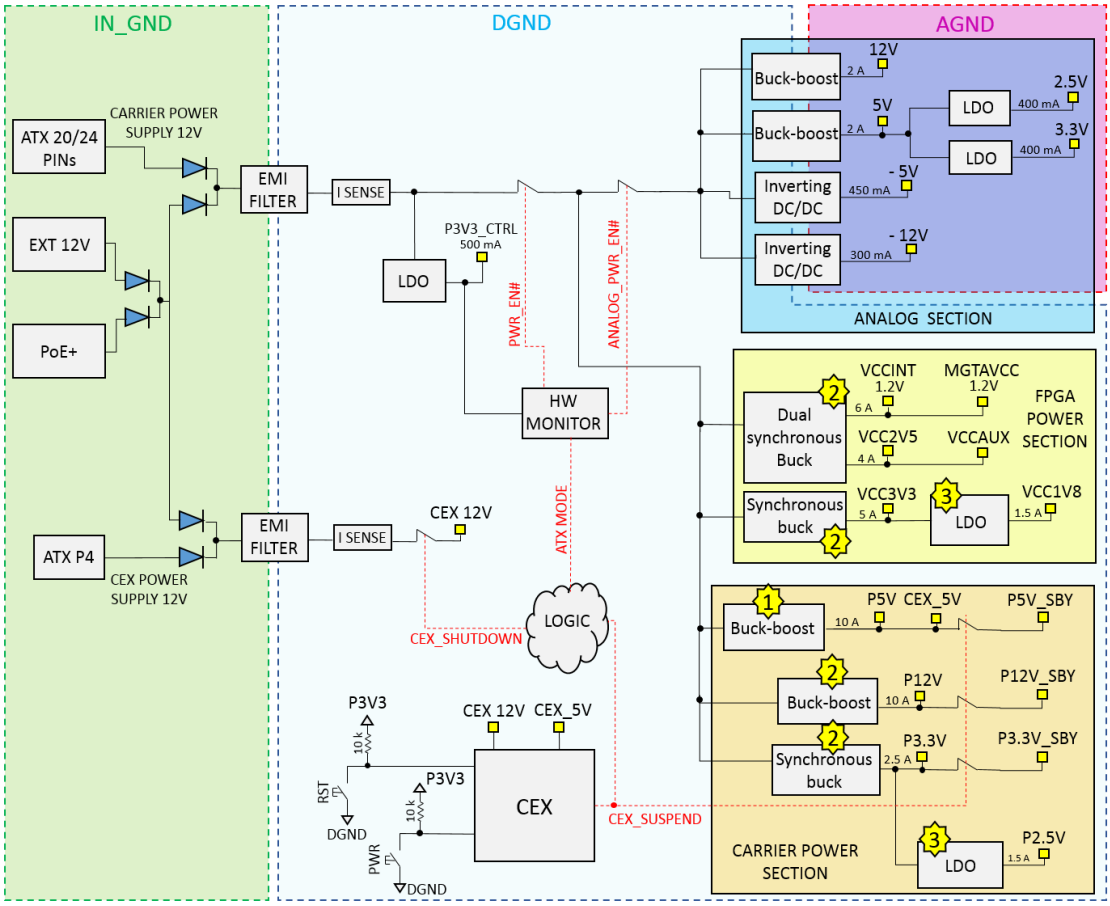

Figure 2.8: IOC board power distribution network.

been performed using the Xilinx's power estimator tools that yield the power and thermal information necessary to choose adequate DC/DC converters [20]. The Spartan-6 devices do not have a required power-on sequence.

All the remaining subsystems are supplied by the carrier board power section that accommodates the *P12V*, *P5V*, *P3.3V* and *P2.5V* power rails that feed those devices that are in working state even when the COM Express is suspended. Additional voltages, *P12V_SBY*, *P5V_SBY* and *P3.3V_SBY* are used to supply the standard pc peripherals that are suspended together with the COM Express module. The IOC board, like a desktop pc, implements the processor suspend functions, the wake on USB activity, the wake on LAN, and the wake on power button press. If the COM Express is not plugged into the IOC board all the carrier functionalities are guaranteed apart those standard pc peripherals fed by the standby rails that are powered off in this case.

To step the input voltage down to the desired output level we have used different DC/DC converter topologies depending on the area, complexity, and maximum power capability and efficiency to be achieved.

The *P12V* and *P5V* voltage regulators are based on the single-inductor buck-boost solution that consists of a single PWM controller that drives the four transistors in the buck, boost or intermediate transition region depending on the input voltage level with respect to the desired output voltage. This topology ensures a regulated output over a wide range of input voltages, and, it yields a high output power, up to several tens of watts; therefore, it has been exploited in the power rails with high load requirements. The same buck-boost topology but with a PWM controller with integrated switches has been exploited for the analog 12 V and 5 V regulators that foresee lower output currents.

Instead, the 3.3 V rails are based on synchronous buck regulators, with integrated power MOSFETs, that feature high efficiency, an excellent setpoint accuracy and low ripple. These buck regulators suit for supplying the FPGA and the main ICs on-board.

The IOC board is also endowed with Low Drop Out (LDO) linear regulators that are a compact and simple solution to generate low voltage and very low noise power rails driving low current and noise sensitive loads.

The output stage of any voltage regulator consists of an LC filtering network in order to reduce high frequency noise and the voltage ripple. Those ICs for which the ripple and high frequency noise is not critical are directly supplied by the main backbones shown in the block diagram, whilst, the most sensitive chips are supplied

by dedicated rails derived by primary voltages using ferrite beads for decoupling and suppressing the high frequency harmonics.
The on-board hardware monitor chip controls the setpoint accuracy of all the regulated voltages, the current consumption, and the temperature information. In case of over-current or over-temperature events, it is able to power-off the COM Express and, in case, the whole carrier. The final user may configure the hardware monitor to leave unpowered the analog section of the board to reduce the default power consumption in those applications where the power consumption and cooling are a concern.
The block diagram provides a first draft of the intended grounding strategy. The EMI filters used to attenuate the common-mode and differential mode noise on the two main 12 V power supplies offer galvanic isolation between the input and output bus; therefore, these EMI filters represent an opportunity to split the *digital logic ground* ($DGND$) from the *input ground* (IN_GND) of the IOC board section that handles the power over Ethernet and the external power sources interface. The *digital logic ground* is the reference terminal for the FPGA power rails and the carrier power section. The *analog ground* ($AGND$) is the reference terminal of the supply that power the analog circuit. The *analog ground* and the *digital logic ground* are tied together close to the analog DC/DC controllers, and, via ground plane bridges that provide an optimal return path to the digital signals connecting the FPGA to the ADC and DAC converters.

2.4.1 ATX Power Control

The possibility to power-on the IOC board using standard computer power supply units represents an important advantage for an embedded pc and facilitates the wiring and installation into commercial industrial pc boxes. The challenge here was to make the ATX power supply hot-swappable and coexisting with PoE+ or any external benchtop DC power supply. In an ATX system [21], the power supply unit status (on/off) is controlled by the computer via the signal $PS\_ON\#$. This signal is tied to ground in the IOC board converting the ATX to an AT scheme where the software has no control of the power supply unit; as soon as the power supply unit is switched on, the carrier is fed with the main 12 V rail via the ATX 24-pin connector, and, this power rail is active even when the COM Express is powered-off or it is not plugged into the carrier. The ATX 24 pin connector provides the 3.3 V, 5 V, -5 V and -12 V alongside the main 12 V. Nevertheless, only the 12 V rail

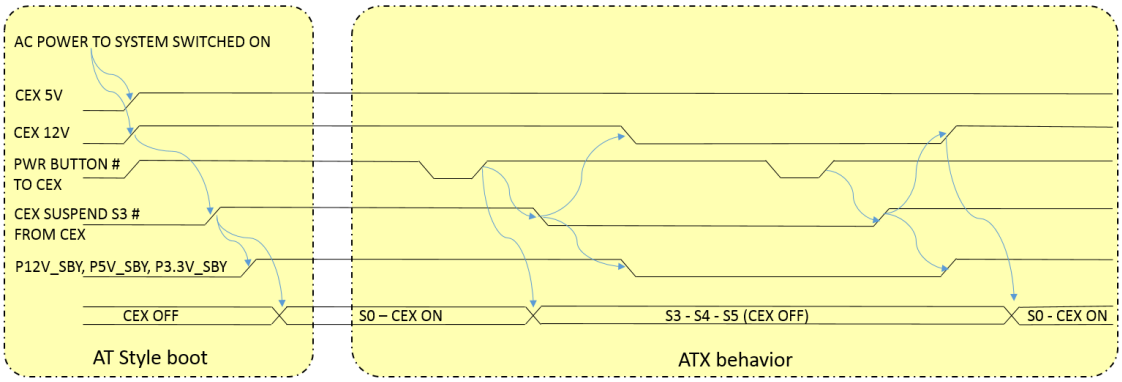

Figure 2.9: AT style power-up boot and the ATX style emulation using the suspend S3 signal.

is used and all the necessary power rails are generated on-board in order to have the three power supply strategies hot-swappable. Despite the ATX power supply unit is converted to the AT power style as required by the carrier design, we are able to emulate an ATX power delivery to the COM Express module thanks to a custom additional logic circuitry integrated downstream of the ATX P4 connector. In the IOC board, like in any ATX system, the operating system controls the shut-down sequence and it has the time to save all the information and to perform the necessary tasks before the CEX 12 V rail is switched off. Figure 2.9 shows the expected COM Express power sequencing diagrams. When the user switches on the AC power to the system we have an AT style boot from the mechanical off state. From here upon, the CEX 5V rail which is present whenever the power supply unit has AC input power, supplies the on-module circuits necessary to feature the suspends state activities, the soft power control, the wake on LAN, the wake on power button press. The suspend S3 signal [22] from the COM Express to the carrier board drives the switches that operate the standby rails and the CEX 12 V power supply, emulating the ATX power scheme with the consequent power efficiency improvement.

2.4.2 Power Over Ethernet

Power over Ethernet implements a way of supplying DC power over the Ethernet cable and allows a single cable to be used for data and power transmission reducing the cable laying cost and complexity [23]. Two of these techniques for transmitting power over Ethernet, referred as *Alternative A* and *Alternative B*, have been standardized by IEEE 802.3 in 2003 and are supported by many commercial Ethernet

switches installed in the SPES control system. The availability in our laboratory of distributed PoE+ injectors persuaded us to design a multi-purpose controller capable of acting as a PoE+ device in low-power applications. The IOC board has been designed in compliance with IEEE 802.3at-2009 power over Ethernet standard, also known as PoE+, which provides up to 25.5 W of DC power delivered to the device using only two out of four pairs in a CAT5 or superior cable.
The input voltage, between 36 V to 57 V, is extracted by the center taps of the primary of the Ethernet transformer and is converted to a regulated 12 V output by the isolated Fly-back converter shown in Figure 2.10. The design is compatible with both *Alternative A* and *Alternative B* foreseen by the standard thanks to the double fully diode bridge. The PoE+ controller negotiates the power class during the time of initial connection and implements additional features like the under voltage lock out, current limitations, thermal shutdown and soft-start.

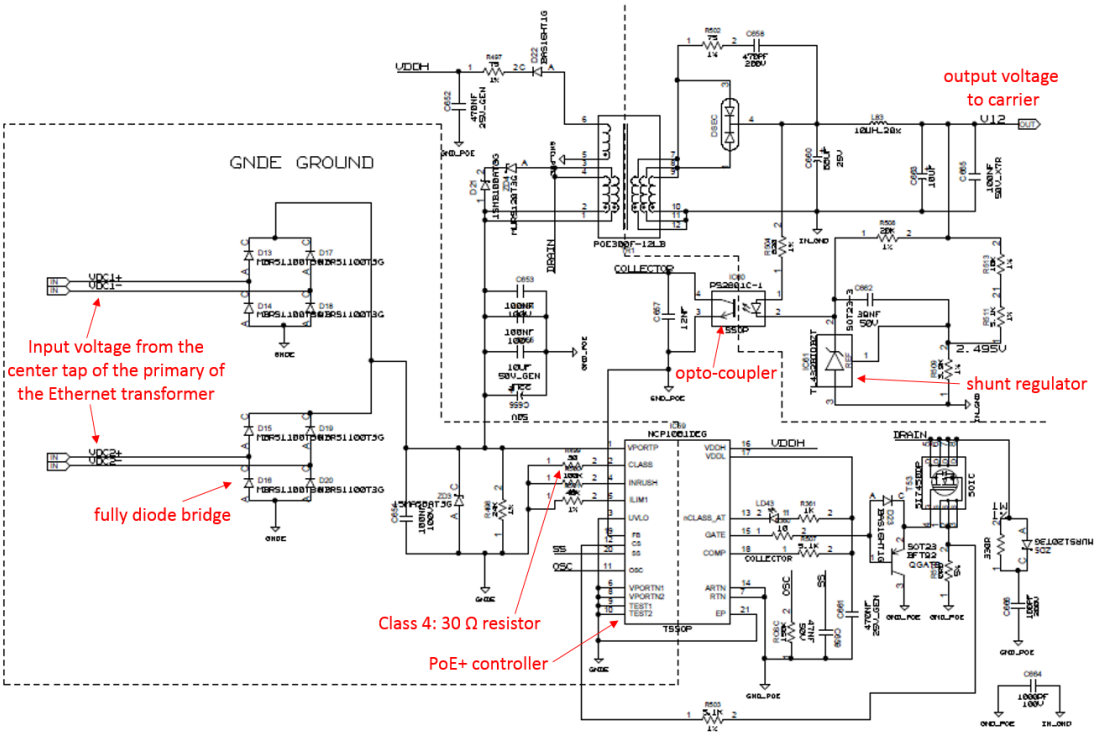

Figure 2.10: PoE+ schematics. Isolated Fly-back converter with extra winding.

Chapter 3

IOC Board Design Process

This chapter resumes the design process followed to turn the IOC board conceptual block diagram into a working prototype.

3.1 Tools and Design Flow

The design flow is resumed in Figure 3.1. The work environment chosen for the IOC board development is the Cadence Allegro toolset. This project manager offers the navigation and organization of design files across all the project phases and combines in a unified environment all the tools needed to develop complex Printed Circuit Board (PCB) designs.
Starting from the functional block diagram and using the Allegro Design Entry HDL tool we captured a detailed electrical schematics that integrates over 2000 components. The capture of the connectivity of the main components required a lot of effort since the IOC board is a new product in our laboratory and we could reuse just a few sections of previous designs. The integrated Cadence Allegro PSpice simulator has been exploited to perform behavioral simulations and analysis necessary to validate subsections of the schematics. Using hierarchical blocks the schematics editing has been boosted and made more modular and intuitive to read. The component selection deserved special care; for any active component, we have verified its availability on the market and in the stock of the main distributors as well as the declared mean time before failure, packaging, price and its voltage compatibility with respect to other components. The accurate management of the library of components was fundamental to capture the design intent correctly the first time avoiding big issues; any new component has been

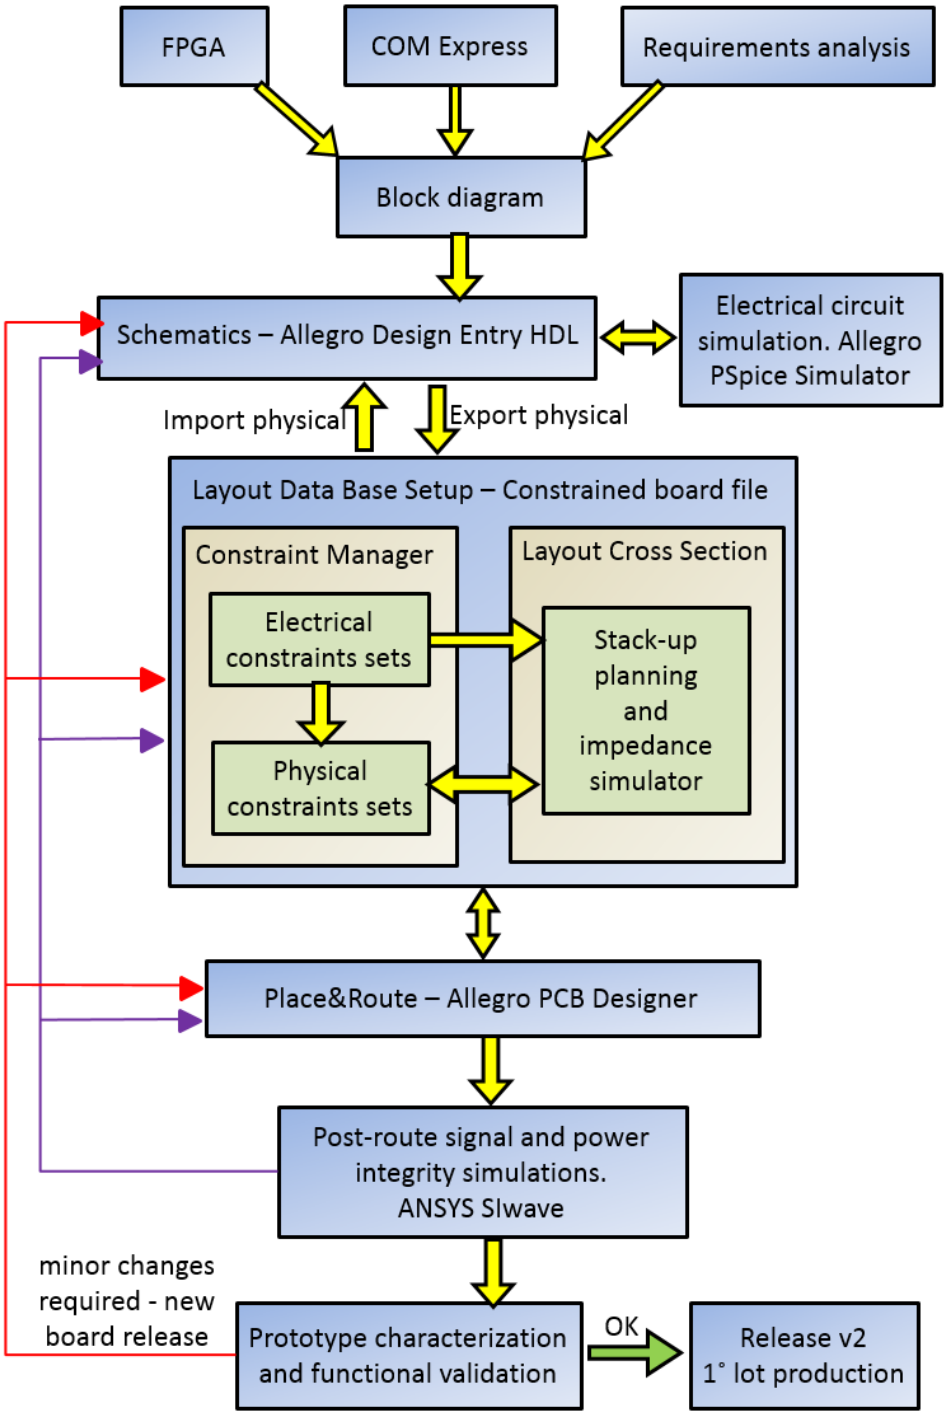

Figure 3.1: IOC board workflow.

attached with a schematics symbol and its main properties (manufacturer, status, the complete part number, value, description, class, package information) and the footprint and pads. The finalized schematics was then transferred to the layout

environment: Allegro PCB Designer Layout. The outcome of the design export to the layout environment is a board file that includes all the relevant PCB information: the complete bill of material, netlist, list of components to be placed and relative footprints, the board stack-up, and the full set of electrical and physical constraints. The layout guidelines of high speed differential transmission lines (e.g. USB 3.0, PCIe, SATA) have been collected in electrical constraint sets within the Allegro Constraint Manager utility. These electrical constraints define the design requirements of the controlled impedance differential pairs and have been turned into physical constraint sets exploiting the Allegro Layout Cross Section tool. The PCB stack-up was planned alongside impedance simulations necessary to tune the microstrip and stripline physical constraints (line width and primary gap) and the stack-up parameters (core and prepreg thickness, copper weight and layer distributions). This iterative process ended up with a PCB cross section proposal that minimizes the impedance mismatch seen by a differential signal crossing two or more layers, and with the controlled impedance signals assigned to a set of physical constraints. This well-honed set of design rules was enriched with information about the board outline and a first draft of the placement of the main I/O connectors imposed by mechanical requirements arising from the front and rear panels of the pc chassis used to install the embedded computer in the field. At this point, we could count on the expertise of a third party company for the placement and routing of the PCB, whilst, we took care of the signal and power integrity analysis, the power and thermal management, manufacturability and testability aspects. The outcome of the place and route is a double sided 14-layer PCB. Signal and power integrity post layout analysis were performed using the ANSYS SIwave tool that, by means of full-wave finite element algorithms, is able to compute the trace characteristics and to evidence discontinuity, reflections and cross coupling issues. The issues highlighted by this post layout analysis have been addressed with minor changes at layout level or, in some cases, at the schematics level with the addition of missing decoupling capacitors or termination resistors. These post layout simulations excluded serious power integrity or signal integrity weaknesses that may lead to malfunctions hard to debug, and validated the proposed board stack-up.
The prototype characterization and functional validation have been carried out alongside the firmware development. The outcome of the test of the first IOC board prototype, v0, was positive and the minor changes required have been addressed with the two subsequent IOC board releases that included changes at the schematics level as well as optimizations to the board layout. All the essential fea-

tures are working fine in the IOC board release v2 that has proved to be a reliable solution during the tests performed in the field. The first production lot has been commissioned and qualified. Twenty IOC boards are ready to be installed in the SPES beam diagnostic system.
Worth of notice is that the custom IOC board has not been designed for commercial purposes, and, will be installed in control applications not subject to the functional safety rules. Therefore, the ElectroMagnetic Compability (EMC) and ElectroMagnetic Interference (EMI) certifications that are the critical part of bringing new electronics products to the market have not been required yet.

3.2 PCB Stackup

To ensure a good level of performance we have designed a multilayer stackup based on the competitive Isola FR408HR high performance laminate dielectric that ensures good thermal performances, reliability and high electrical bandwidth (low loss) [24]; it addresses the requirements of the high speed and controlled impedance signals of the IOC board. The Isola FR408HR is available in different resin content and any material within this family is characterized by specific dielectric constant and loss tangent values. To reduce the production costs, we have chosen only two base materials: one for the core and one for the prepreg. Thereupon, each dielectric layer is obtained by overlapping one or more sheets of a single species material and its thickness is a multiple of the thickness of one of the two base materials. The thickness of any dielectric layer has been calculated alongside the physical constraints set as will be explained later in section 3.3.
The total number of layers required to complete the IOC board layout is mainly determined by the FPGA [25]. The number of signal layers required to break out the signals from the Ball Grid Array (BGA) pads of the FPGA can be estimated using (3.1):

$$Layers = \frac{396}{(26-1)x4} = 4 \tag{3.1}$$

where 396 is the maximum number of user I/O signals available for the FGG676 package, 26 is the number of BGA pins per side, minus 1, times four sides. These 4 layers, in addition to top and bottom layers, are completely dedicated to the signal routing. To optimize the signal and power integrity we foresee the usage of three segmented planes for the power distribution network and five solid ground planes that fill the entire layer. The result is the symmetrical 14-layers stack-up shown in

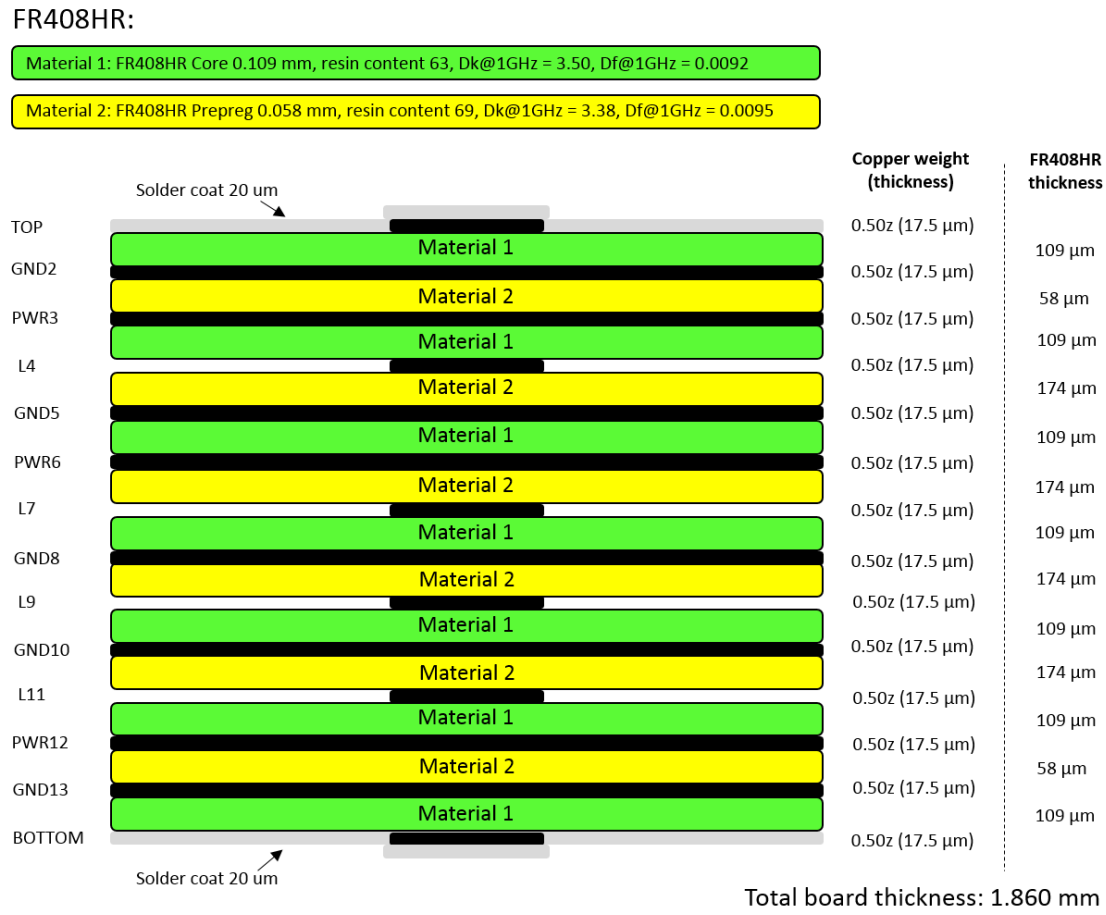

Figure 3.2: IOC board stack-up.

Figure 3.2.

The copper layers distribution follows well defined criteria in order to reduce the electromagnetic emissions and crosstalk that may degrade the electrical performances causing intermittent operations. The total number of copper layers is even and the PCB cross section is symmetrical to improve the impedance symmetry and continuity. Two solid ground planes, GND2 and GND13, face the top and bottom assembly layers and provide the best ground strategy to the active components, reducing the ground bouncing phenomena. Microstrip transmission lines, on top and bottom layers, offer very good impedance continuity and optimal current return paths, therefore, are mainly exploited to route the digital high speed signals and the most sensitive analog signals. The layers PWR3 and PWR12 are reserved for the power distribution network being close enough to the PCB surface. The most demanding component in terms of power integrity is the FPGA whose power supply is delivered through a grid of via and power islands underlying the device

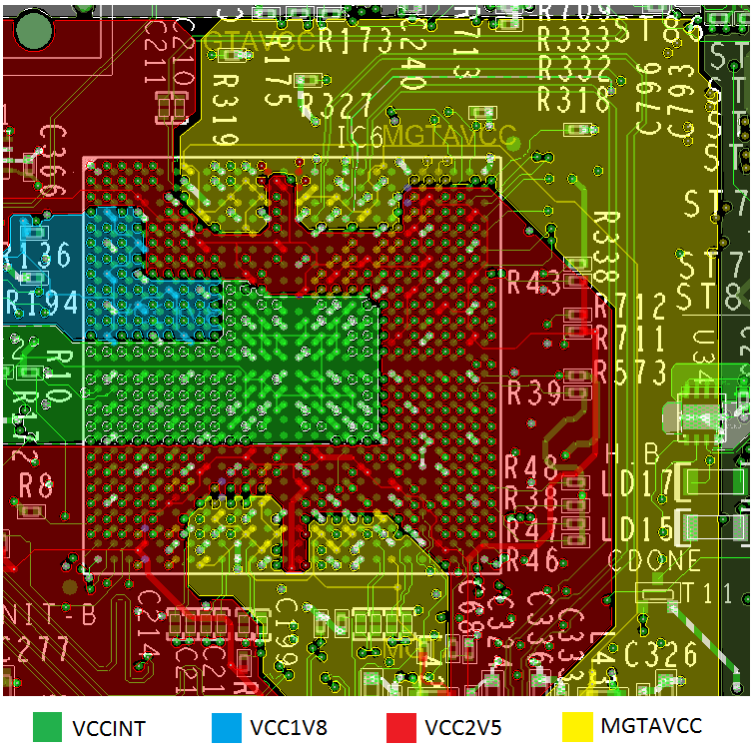

Figure 3.3: PWR3 layer, FPGA power distribution.

area, as shown in Figure 3.3. Any power layer is placed next to a ground layer to create a planar capacitance that aids high frequency decoupling, reduces EMI radiation, and enhances the EMC robustness. To maximize the impedance control and provide a short current return path, the signal layers are tightly coupled (less than 0.174 mm) to an adjacent ground reference plane. This also reduces the switching noise, likelihood produced by the FPGA and digital ICs, coupled with the power planes improving the performances of most sensitive devices like gigabit transceivers and Phase Locked Loops (PLL). Any signal layer is completely isolated between ground and power layers that provide shielding for high speed and sensitive channels. Blind and buried via, back-drilled via, and via on pad are avoided to reduce the fabrication cost of up to 40%. Any via on a high speed digital signal is accompanied by a ground via that provide a low inductance current return path. The PCB is covered by a 20 µm thick solder coating that should be considered when calculating the microstrip impedance.

3.3 Constraint Manager

The electrical constraints considered during the IOC PCB design are the digital high speed layout guidelines: the target single ended and differential impedance and tolerance, the maximum number of via, the minimum/maximum propagation delay, the total etch length and the maximum stub length. Any class of high speed differential signals has been associated to the corresponding electrical constraints set in the Allegro Constraint Manager utility. The constraint manager has been used to define and validate, at each step of the design flow, the layout guidelines of the PCIe, USB Super Speed differential pairs, Gigabit Ethernet connection, SATA, SFP interface and all the FMC differential pairs. In our case, the electrical constraints mapping did not require the building and validation of detailed electrical topology models using the Allegro SigXplorer tool. Indeed, this modeling process is usually required to prototype a new technology or to understand the scenarios of designing custom interconnection solutions, which is not the case of the IOC board. The IOC builds on standard bus topologies whose electrical constraints are defined by the specifications and have been readily mapped in the corresponding electrical constraint sets. These electrical constraints resume the design requirements and have been converted to the physical constraint sets shown in Figure 3.4. Each differential pair has been assigned to the physical constraint set of competence that defines, layer by layer, the trace width, the differential primary gap, the neck and related minimum width and maximum length, and the via enabling. The physical

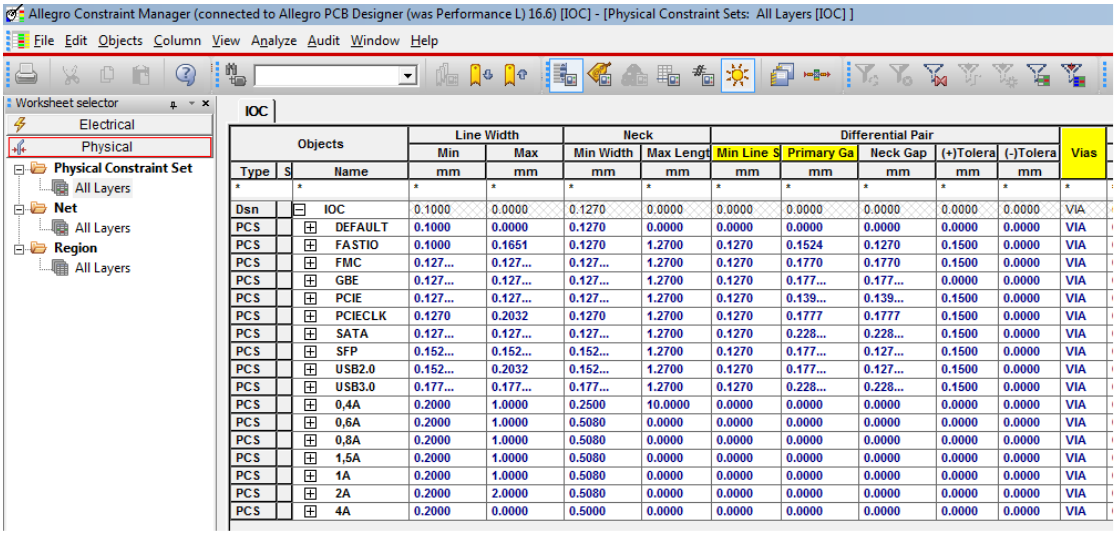

Figure 3.4: Physical constraints mapping to the Allegro constraint manager.

constraints drive the routing of high speed differential nets and are essential when dealing with a multi-layer mixed PCB. An optimal naming convention, followed while defining the basic connectivity in the Allegro Concept HDL tool, facilitates the physical constraint mapping of the design objects considered: net, extended nets, and differential pairs. The routing electrical constraints and physical constraints are computed real time by the cadence and any violation will cause an error to pop-up in the screen reducing, at design time, the probability of signal and power integrity problems. We did not define timing constraints because are computed by the post-layout simulations that were performed in a different work environment.

3.3.1 Impedance Calculator

The stack-up planned has been imported into the Layout Cross Section tool integrated in the Allegro PCB Designer. This is one of the most important phases in the design flow because gives a feedback on the stack-up correctness and on the PCB technology necessary to satisfy the physical requirements (e.g. minimum line width and process tolerances). The cross section table, shown in Figure 3.5, defines an agreement between the system engineer who has designed the electrical

Layout Cross Section

	Subclass Name	Type	Thickness (MM)	Dielectric Constant	Loss Tangent	Negative Artwork	Shield	Width (MM)	Impedance (ohm)	Coupling Type	Spacing (MM)	DiffZ0 (ohm)
1		SURFACE		1	0							
2		DIELECTRIC	0.02	3.3	0							
3	TOP	CONDUCTOR	0.052	4.5	0.035	☐		0.1270	56.688	EDGE	0.1524	89.985
4		DIELECTRIC	0.109	3.5	0.0092							
5	GND2	PLANE	0.017	4.5	0.035	☐	☒					
6		DIELECTRIC	0.058	3.38	0.0095							
7	PWR3	PLANE	0.017	4.5	0.035	☐	☒					
8		DIELECTRIC	0.109	3.5	0.0092							
9	L4	CONDUCTOR	0.017	4.5	0.035	☐		0.1300	49.849	EDGE	0.1300	89.22
10		DIELECTRIC	0.174	3.38	0.0095							
11	GND5	PLANE	0.017	4.5	0.035	☐	☒					
12		DIELECTRIC	0.109	3.5	0.0092							
13	PWR6	PLANE	0.017	4.5	0.035	☐	☒					
14		DIELECTRIC	0.174	3.38	0.0095							
15	L7	CONDUCTOR	0.017	4.5	0.035	☐		0.1270	50.413	EDGE	0.1524	92.501
16		DIELECTRIC	0.109	3.5	0.0092							
17	GND8	PLANE	0.017	4.5	0.035	☐	☒					
18		DIELECTRIC	0.174	3.38	0.0095							
19	L9	CONDUCTOR	0.017	4.5	0.035	☐		0.1270	50.413	EDGE	0.1524	92.501
20		DIELECTRIC	0.109	3.5	0.0092							
21	GND10	PLANE	0.017	4.5	0.035	☐	☒					
22		DIELECTRIC	0.174	3.38	0.0095							
23	L11	CONDUCTOR	0.017	4.5	0.035	☐		0.1300	49.851	EDGE	0.1300	89.21
24		DIELECTRIC	0.109	3.5	0.0092							
25	PWR12	PLANE	0.017	4.5	0.035	☐	☒					
26		DIELECTRIC	0.058	3.38	0.0095							
27	GND13	PLANE	0.017	4.5	0.035	☐	☒					
28		DIELECTRIC	0.109	3.5	0.0092							
29	BOTTOM	CONDUCTOR	0.052	4.5	0.035	☐		0.1270	56.688	EDGE	0.1524	89.985
30		DIELECTRIC	0.02	3.3	0							
31		SURFACE		1	0							

Figure 3.5: IOC layout data base setup.

schematics, the layout designer who is in charge of placing and routing the board, and the manufacturing company that must provide a feedback on the manufacturability and relative cost of the PCB. The layout cross section tool computes accurate transmission line modeling and provides a unified user interface of those parameters that determine the transmission line impedance. The impedance of microstrip and stripline is computed in real time opening the possibility to tune the dielectric layers thickness and to fill the physical constraint sets with the optimal line width and primary gap values in order to satisfy the electrical constraints.

3.4 Layout

The physical and electrical constraints, defined in the previous steps, together with the analysis of the mechanical, thermal, grounding, signal and power integrity aspects, and a pre-placement floor-plan, have driven the external expert during the layout design. In particular, the pre-placement draft opened the possibility to see the best signal flow between functional blocks, and, to correct potential mechanical issues in an early design phase.

3.4.1 Placement

The IOC board outline has been set to 250 mm x 200 mm; these dimensions are similar to those of a commercial motherboard, but, significantly smaller than those of the commercial COM Express evaluation board of reference. The board outline and the arrangement of the connectors in the outer ring have been agreed with mechanical and system engineers. As shown in Figure 3.6, the top layer is almost reserved for the digital signal processing. The COM Express board to board connectors have been instantiated in the middle of the layout together with the mini, compact, and base form factor mounting holes and corresponding shapes, which delimit a constrained region where it is not possible to place tall components. The FPGA, corresponding to the largest BGA footprint in the picture, is also in the middle of the board to reduce the total etch length of constrained signals. The FMC connector and the three PCIe slots have been placed in agreement with the standard requirements. The most sensitive active devices have been placed far from the edge of the board to preserve the controlled impedance and to reduce the chance of electromagnetic interference (EMI). Thermal effects have been considered to optimize the placement: the FPGA and the hotter chips are assembled

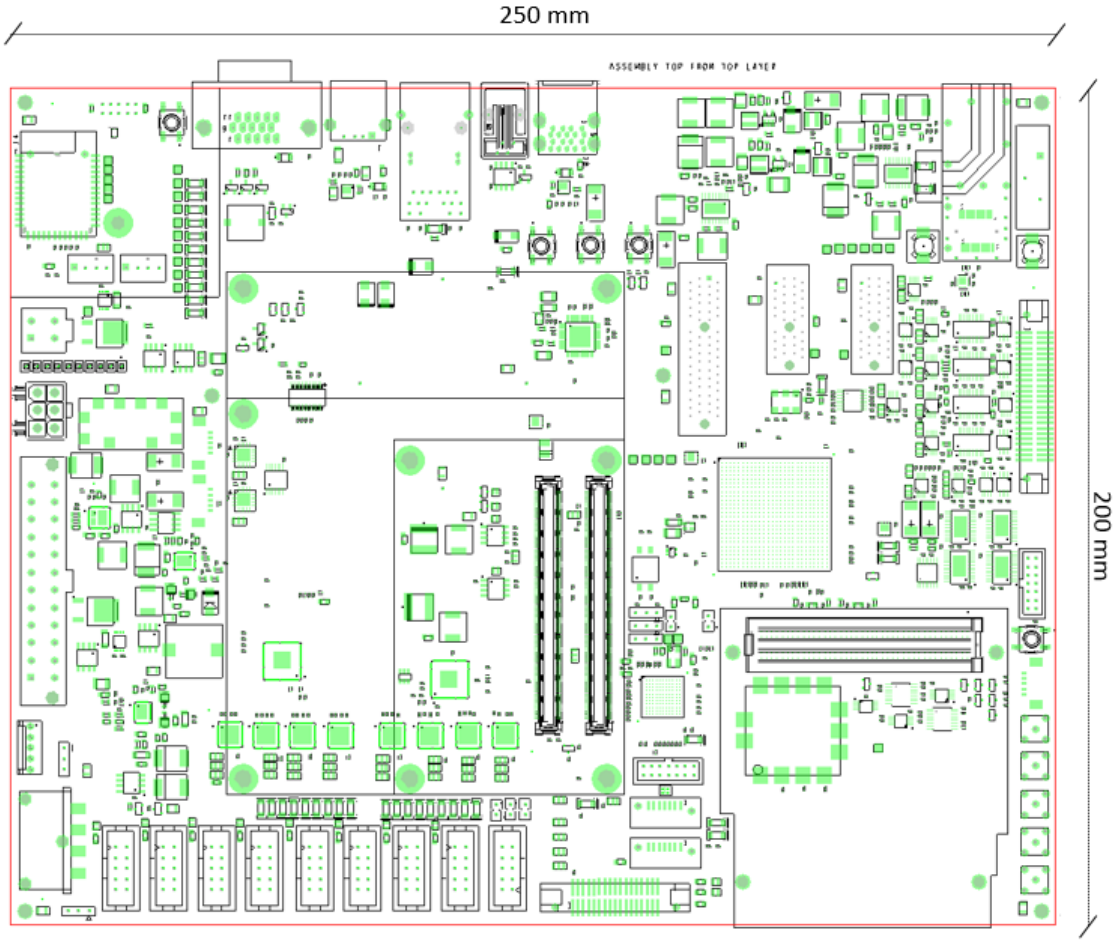

Figure 3.6: IOC placement top.

in the top side of the board to receive an unrestricted airflow provided by two fans installed into the chassis. Functional blocks of circuitry have been organized and placed as groups in order to minimize signal crossing and the spread apart of bypass capacitors and termination resistors.

The topography of the bottom side of the IOC board is shown in Figure 3.7. The bottom layer faces to a ground plane (GND12) that shields the most sensitive digital circuits and signals from the radiated and conducted switching noise emitted by the DC/DC controllers, mostly assembled in the bottom side of the PCB.

The components belonging to the analog section have all been placed in the top and bottom right side of the PCB in correspondence of an analog ground segment. The IOC board necessitates of three segmented power planes, each distributing multiple voltages, and of a digital and analog ground separation, operated on GND2 and GND13 layers. The ground split is necessary to maintain a wide input and output

analog dynamic range with low noise in a digital environment. We exercised cau-

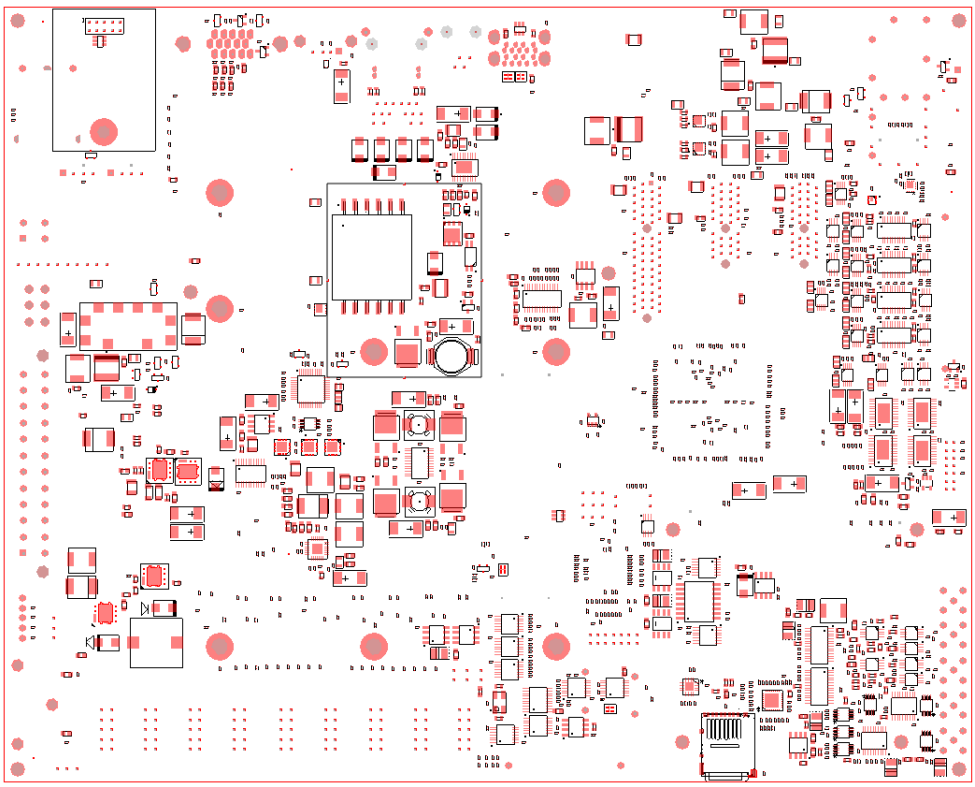

Figure 3.7: IOC placement bottom.

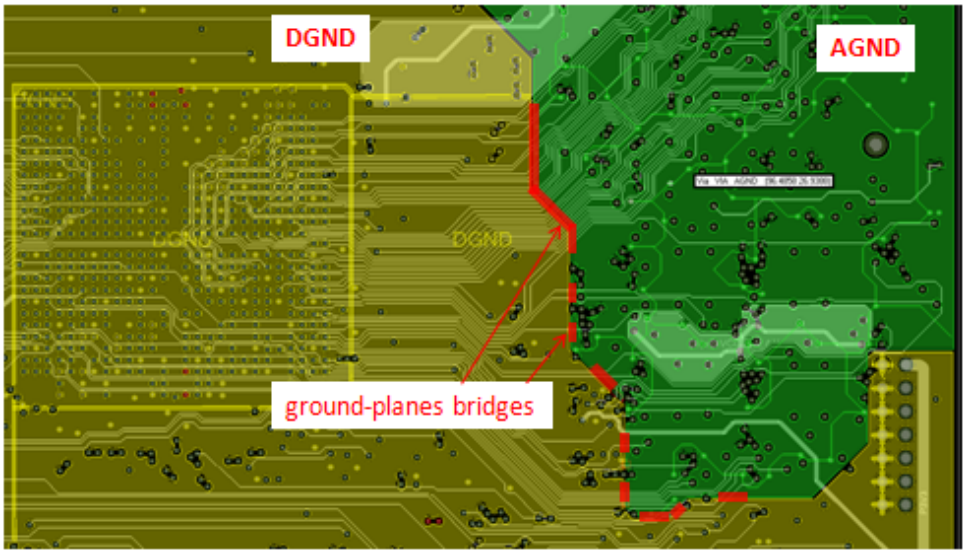

Figure 3.8: Ground plane bridges for traces.

tion placing and routing connected components across the planes splits; high speed

transmission lines should not cross splits in the reference power or ground plane as that will break up the current return path with a consequent impedance mismatch and increment of the losses. The analog and digital grounds are connected at the power supply (close to the DC/DC converter), and, via ground plane bridges for the return currents of the ADCs and DACs digital interfaces. Figure 3.8 highlights the ground bridges between the data converters and the FPGA.

3.4.2 Routing

The IOC board was manually routed using top, bottom and the four signal inner layers L4, L7, L9, and L11. Figure 3.9 shows the IOC completely routed with all the signal layers on. First have been routed the high speed differential impedance pairs, whereupon all the other slow digital signals. The IOC board layout complies with the routing practices normally adopted to reduce the crosstalk problems and to improve the signal and power integrity. The high speed channels acting as

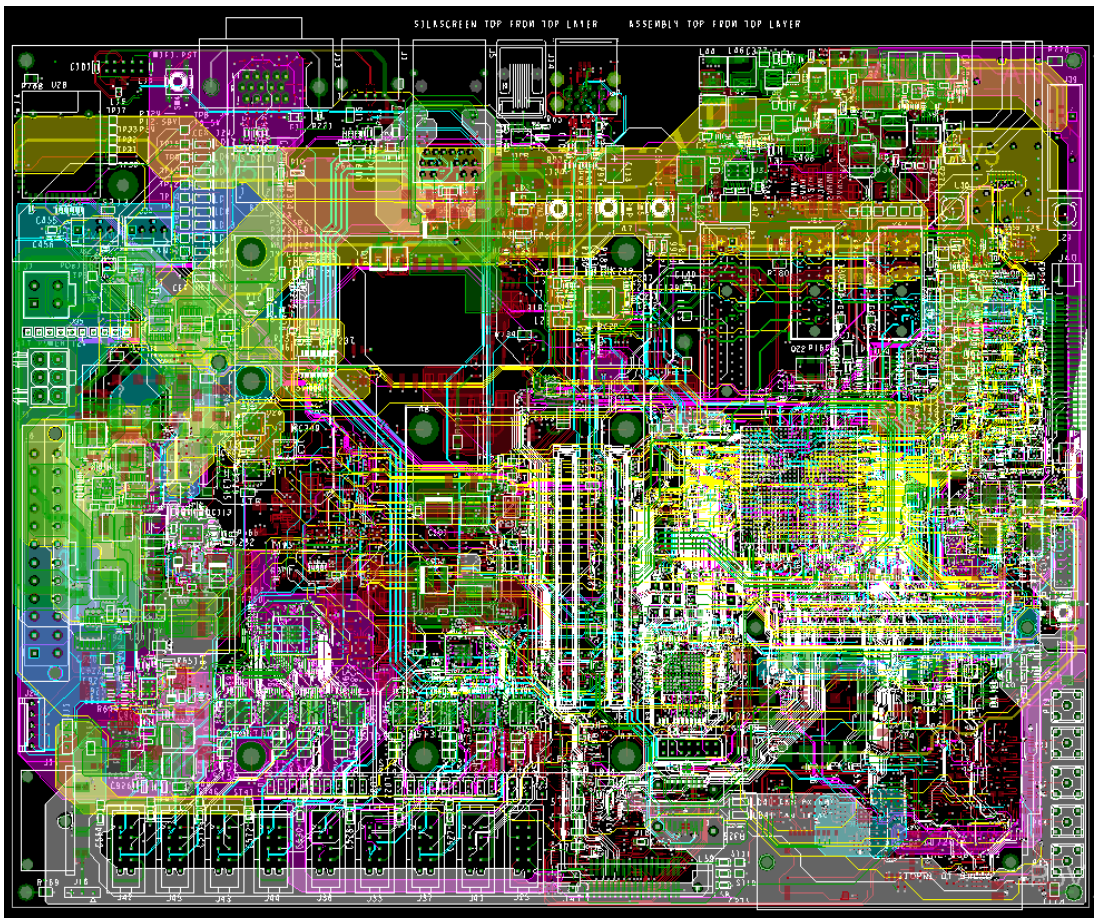

Figure 3.9: IOC board fully routed; layers top, bottom, L4, L7, L9 and L11 on.

transmission lines have been analyzed with post-layout simulations.

3.5 Post Layout Signal and Power Integrity Simulations

The post-layout simulations performed before releasing the production files are essential in order to correct signal and power integrity issues afterward difficult to reproduce during the prototyping phase. We imported the layout into ANSYS SIwave, a specialized design platform for power integrity, signal integrity and EMI analysis of complex PCBs [26]. As resumed in Figure 3.10, we used this platform to validate high speed channels and the power delivery system. Basically, for any digital IC, in order to operate correctly, we must ensure logic signals to be transferred accurately and a proper power delivery. At this purpose, the performed signal integrity simulations validate the transmission lines impedance matching and the crosstalk against the electrical specifications. The transmission line quality is well described by the scattering matrix. Whilst, the power integrity simulations guarantee that the Power Delivery Network (PDN) delivers a sufficiently clean power supply to any IC. A well designed PDN is required to have a very low impedance. Low impedance means low voltage ripple; the voltage ripple is caused by the combination of the switching currents of all ICs to which a power rail provides power. SIwave opens the possibility to measure the power delivery impedance and also to estimate the simultaneous switching noise which manifests as a shift of the supply and ground voltage in response to the simultaneous multiple output drivers switching. In the IOC board, the DC low impedance of the PDN is ensured by the three segmented layers where any supply voltage is delivered through power

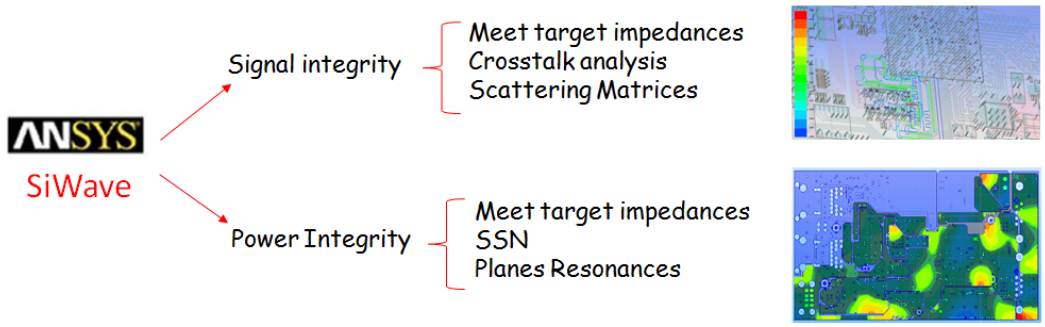

Figure 3.10: Simulation performed using SIwave.

islands. Nevertheless, planes resonances may significantly increase the power delivery impedance at the resonance frequency, degrading the board performances. Hence, we performed several measures of power planes resonances.

3.5.1 Signal Integrity Analysis

Impedance matching is the foundation to get signal integrity. The impedance measurement is the first analysis that has been performed on the high speed channels where the transmission line length is long compared to the foreseen wavelength of the signal. SIwave integrates an impedance analyzer that measures and plots the single-ended or differential impedance of the selected nets versus the propagation delay, as shown in Figure 3.11. Exploiting this signal net analyzer utility is possible to identify any impedance mismatch and its position along the transmission line. The impedance mismatch causes a signal reflection which reduces the noise margin. In the proposed example, the differential characteristic impedance of the net under test is 90 ohm and it is routed partly in the top and partly in the bottom. The inductive effect of the via on the transmission line is well visible and generates a thin impedance discontinuity. The narrow impedance discontinuities are significant only if comparable with the signal wavelength, otherwise negligible. A similar impedance discontinuity is observed whenever a transmission line crosses a reference plane split. Exploiting this tool we validated the microstrip and strip-

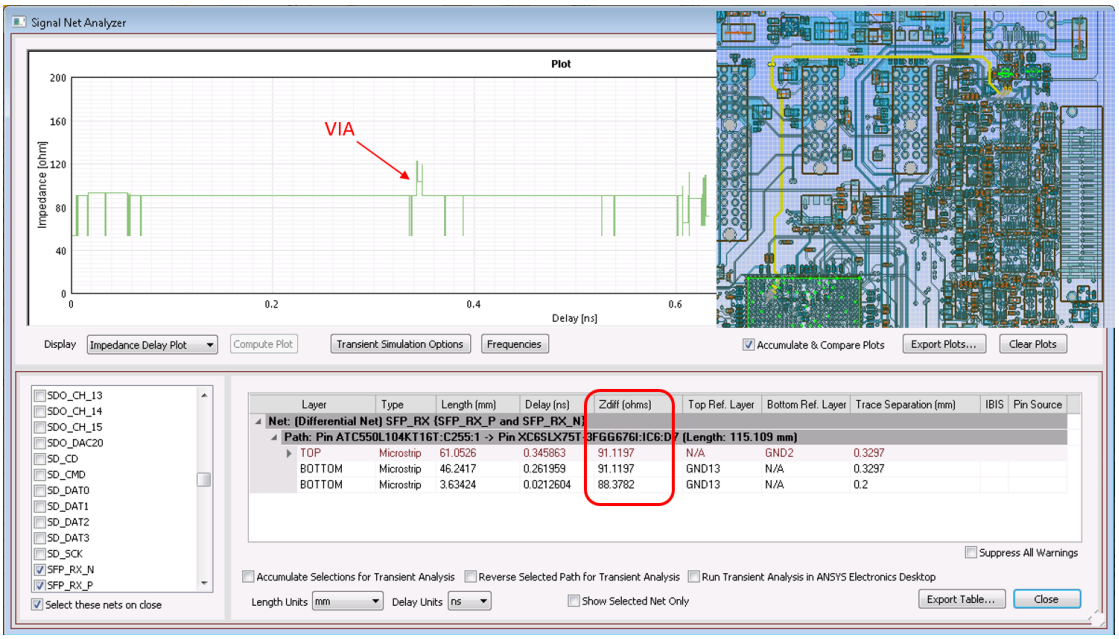

Figure 3.11: Impedance delay plot of the SFP RX pair.

line geometrical parameters of the PCIe, SATA, USB 3.0, FMC, SFP and Gigabit Ethernet links.

The same tool offers the possibility to perform the transient analysis which is useful to understand the behavior of the transmission line in terms of ringing and reflections. The result of the transient analysis applied to the SFP reception pair is shown in Figure 3.12. The source signal chosen for the test is a square wave with a rise time of 35 ps, and, imposing the right termination, the plot does not show ringing or reflection phenomena as a proof of a good impedance matching.

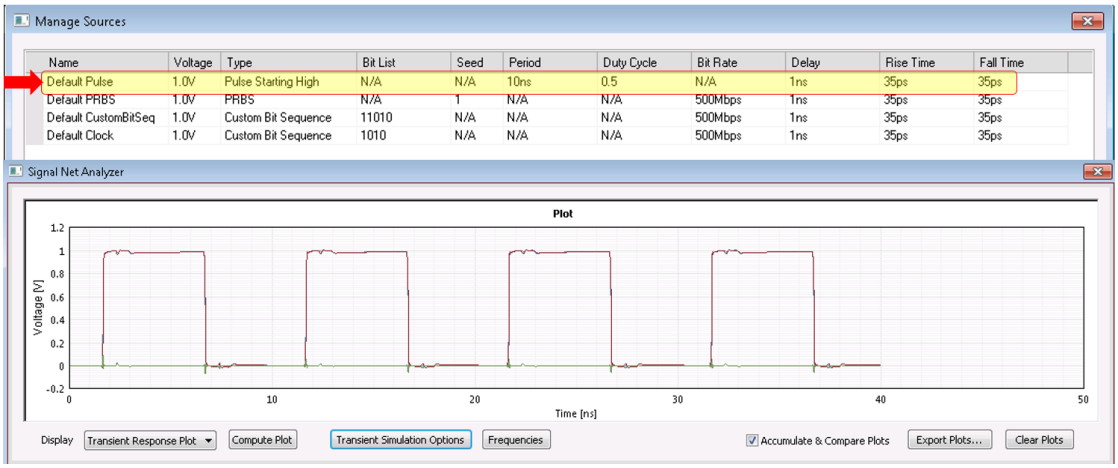

Figure 3.12: Transient analysis of the SFP RX pair.

The SIwave tool opens the possibility to calculate the S-parameters for a complete post-layout analysis. The plot of the S-parameters provides important information about the transmission, reflection, and crosstalk. Combining the measured S-parameter elements with designed-in circuit elements allow the user to quickly estimate the insertion loss of the transmission line. The attenuation, or insertion loss, is one of the signal integrity impairments that mostly impact the end-to-end link performance. Using SIwave, we computed the S-parameters of the USB super speed microstrip. The IOC board complies with the USB 3.0 specification that foresees an overall channel differential insertion loss at the Nyquist frequency of approximately 20 dB, with about 10 dB allocated for the host, 7.5 dB for cable and assembly, and 2.5 dB reserved for the device [27]. From these considerations, the acceptable differential insertion loss of the USB 3.0 microstrips on the custom COM Express carrier board is less than 7.5 dB at 2.5 GHz (7.5 dB allocated for the carrier and 2.5 dB allocated for the CEX module). The S-parameters, of the USB super speed pairs, have been computed with a frequency sweep analysis from the DC up to 10 GHz, placing the ports on the type A USB receptacle and the

COM Express board to board connector pins. The transmission terms are shown in Figure 3.13 and refer to a total etch length equal to 157.48 mm. The attenuation is about 1.8 dB at 2.5 GHz, well within the requirements. From the plot, we can extrapolate the 3 dB bandwidth, which is 4.8 GHz, and compare it with the Isola FR408HR specifications that state a 5.6 GHz 3dB bandwidth for a 150 mm long microstrip line. The computed insertion loss include the contribution of the diode clamp and the via.

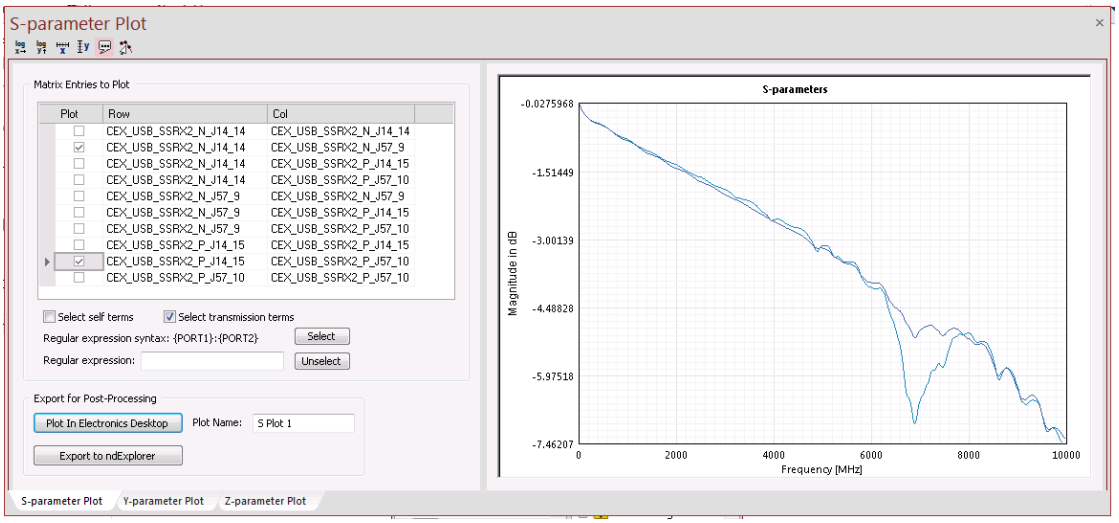

Figure 3.13: Scattering parameters S21: forward transmission terms for the USB super speed RX pair two port networks.

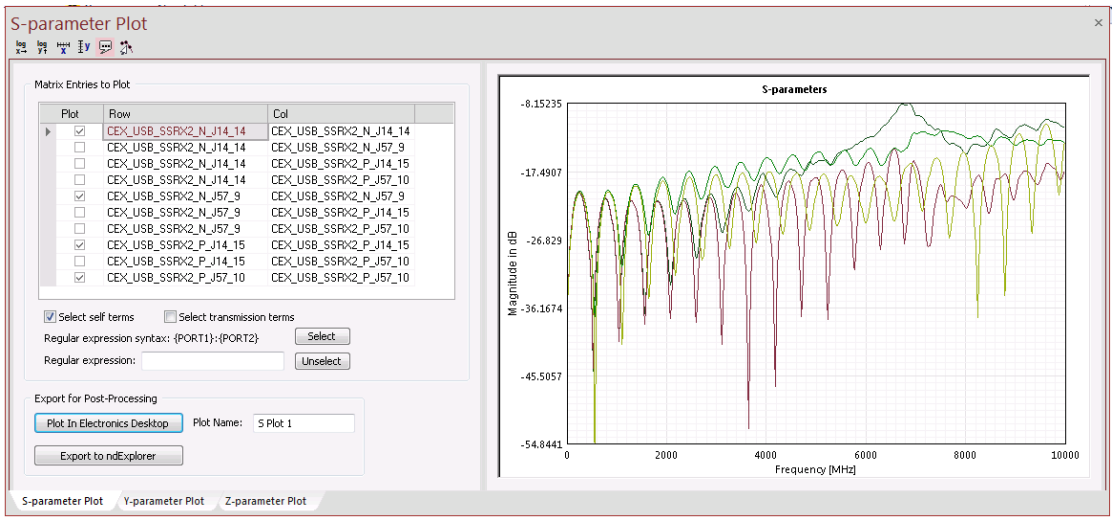

Figure 3.14: USB super speed RX pair return terms.

Plotting the self-terms (S11 and S22 scattering parameters), as shown in Figure

3.14, it is possible to have an idea about the return losses. High values of the return loss coefficients indicate an impedance mismatch between the source and the transmission line with the consequent generation of reflected waves that decrease the signal intensity at the receiver.
The frequency sweep applied to differential net yields also information about intra-pair coupling. The near-end and far-end crosstalk, computed on the USB super speed receiver pair, are shown in Figure 3.15.

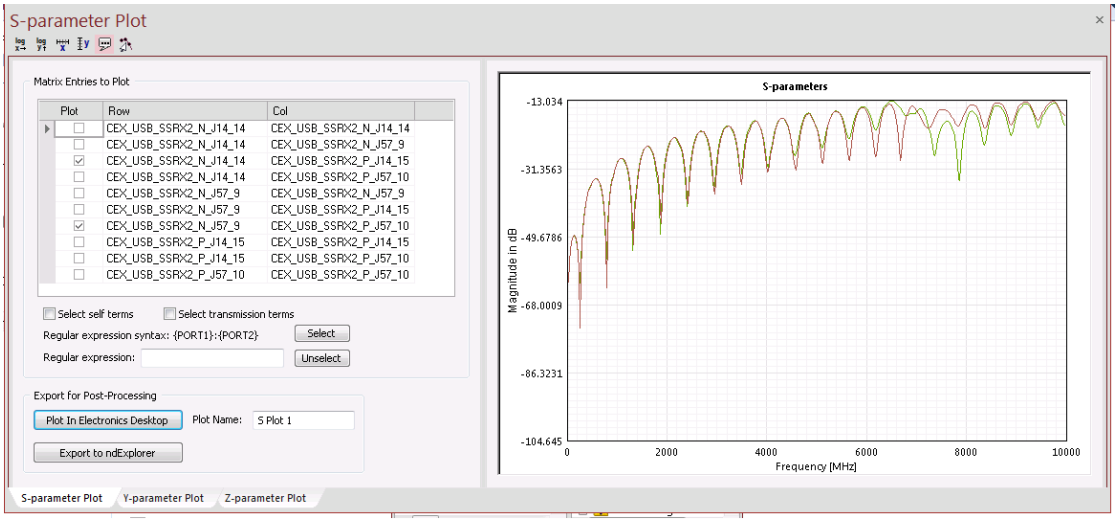

Figure 3.15: USB super speed RX pair coupling terms.

3.5.2 Power Integrity Analysis: Computing Resonant Modes

A stable power distribution network is fundamental to deliver clean power from the source to all ICs, to provide the best return path for signals and to keep the emitted radiation within EMI limits. A PDN with poor stability results in impedance peaks that degrade the flat and low impedance profile with consequent performance reduction of the circuit being powered. The geometrical PCB structure and the derived parasitic elements can form resonant circuits with high Q factors and resonant frequencies in the hundreds of MHz range (within the signaling bandwidth) that reduce the power stability. Moreover, these resonant behaviors affect the signal transmission. The resonances lead the nets to perform differently at the resonant frequency, introducing a harmonic distortion that reduces the signal integrity. The IOC board cross section is composed of three segmented power planes that, together with the ground planes, form cavities of different shapes with different resonance frequencies. In a multi-layer PCB with segmented power planes, like the

IOC, it is necessary the use of specialized tools to model the cavities and the PDN impedance profile, given the complexity of the cavities geometries with irregular shapes. This highlights the importance of performing post-layout power integrity simulations. The SIwave tool opens the possibility to compute the cavity-mode resonances of a multi-layer PCB in the frequency range of interest. Figure 3.16 and Figure 3.17 show the areas of the IOC PCB concerned by natural resonances at 264 MHz and 272 MHz respectively. These images display the transverse dependence of the voltage difference between PWR6 and GND8 planes at the resonant frequency, where a red color indicates higher voltage amplitude. These resonant modes may be a concern since they occur at relatively low frequencies if compared with the bandwidth of many digital signals crossing these cavities on the layer L7, and since they appear in crucial areas of the layout and affect the power rails of important active devices.

Decoupling capacitors have been used to dump the voltage swing of the selected modal resonances [28]. Not only the value of the capacitor but also its position in the board is important to dump the voltage amplitude and move the resonance out of the band of interest, with a much smaller magnitude. Adding two 220 pF

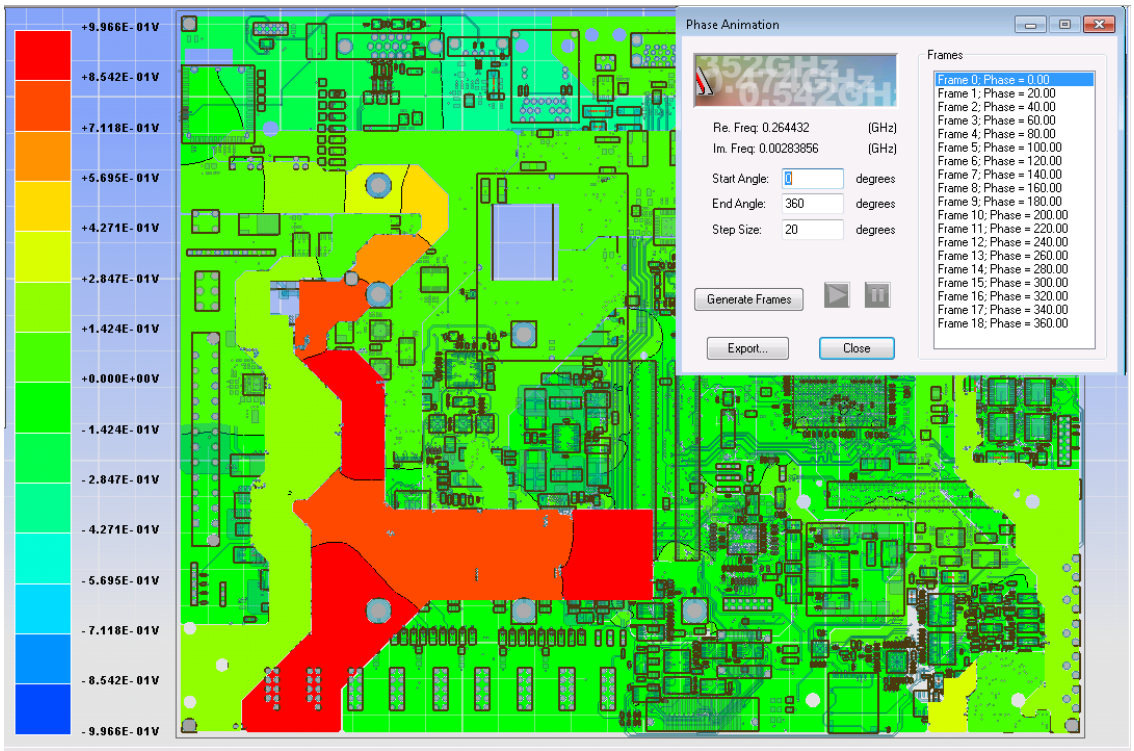

Figure 3.16: Voltage distribution between PWR6 and GND8 induced by the resonant mode 3 that occurs at 264 MHz.

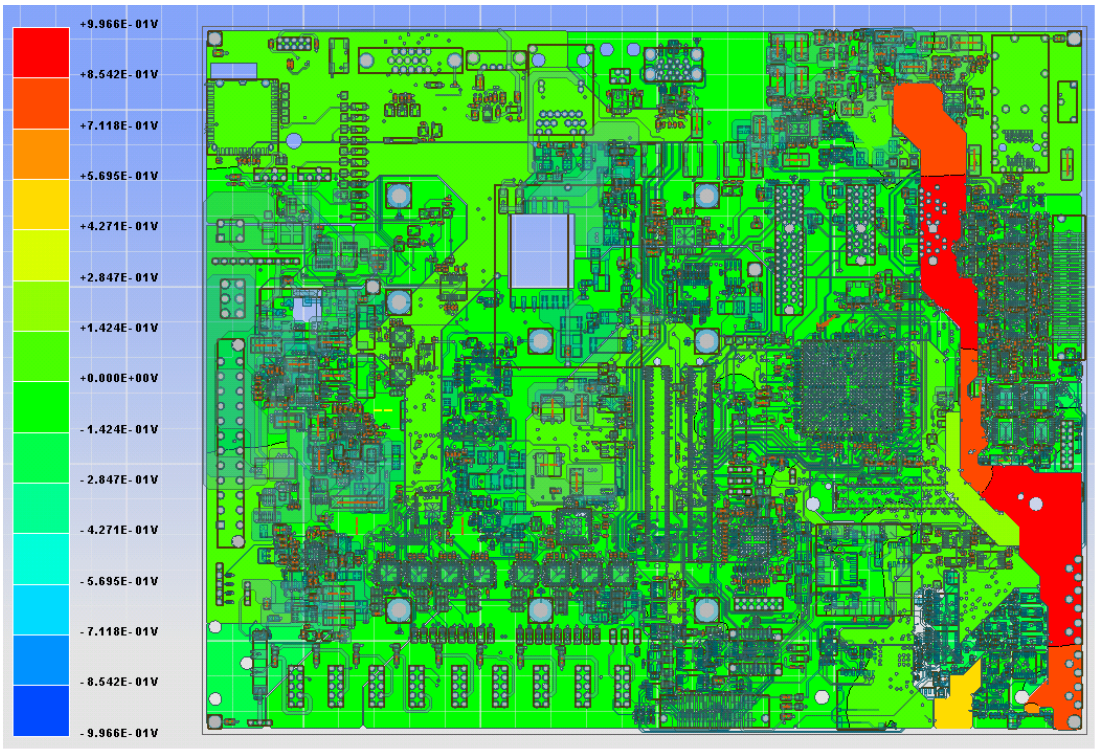

Figure 3.17: Voltage distribution between PWR6 and GND8 induced by the resonant mode 4 that occurs at 272 MHz.

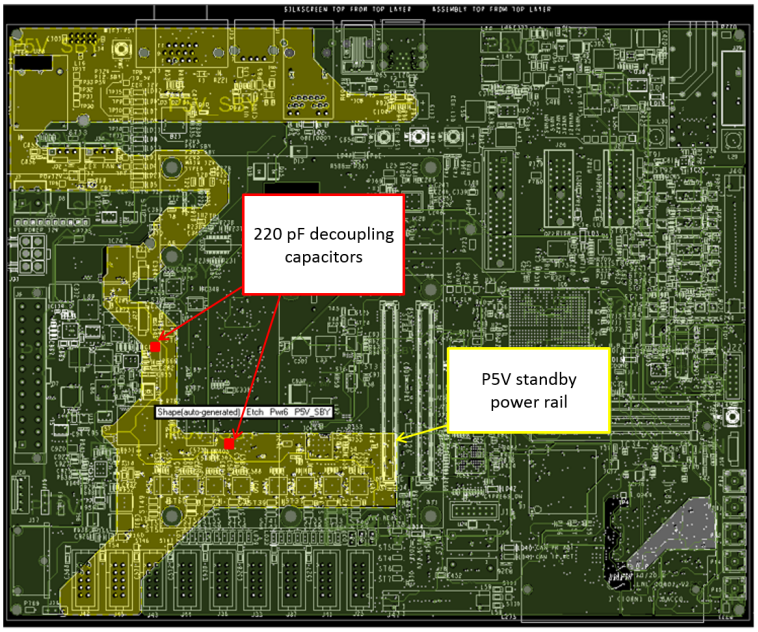

Figure 3.18: Two 220 pF decoupling capacitors across $P5V\_SBY$ and GND.

capacitors across the $P5V\_SBY$ and ground, as shown in Figure 3.18, we could greatly reduce the large impedance variation related to the third resonance mode.

The capacitor value has been calculated assuming the leakage and mounting inductance of approximately 0.5 nH and placing the self-resonant frequency of the decoupling capacitor at a higher frequency, 480 MHz, than the third resonance mode (264 MHz). The impact of the two decoupling capacitors on the $P5V\_SBY$ power rail is evident from the result of the new resonance mode analysis shown in Figure 3.19. As expected the $P5V\_SBY$ power island has been effectively decoupled adding the two capacitors and the related resonant mode, previously at 264 MHz, has been shifted to lower frequencies and with a negligible magnitude; it does not appear anymore between the computed resonant modes . With a similar approach, further resonant modes have been dumped.

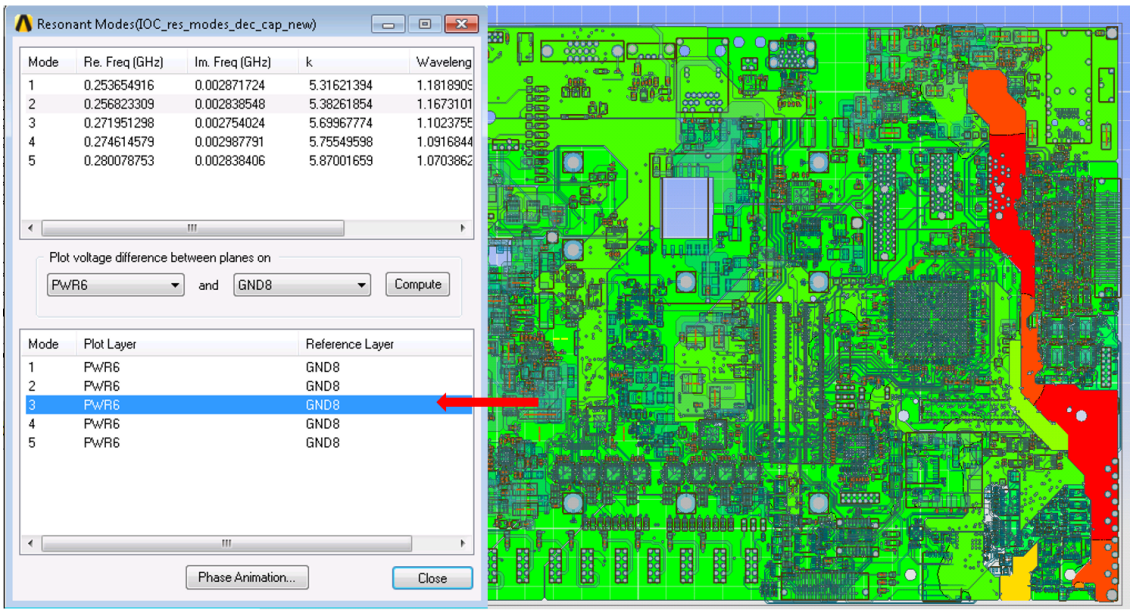

Figure 3.19: New resonant modes with decoupling capacitors.

Chapter 4

Firmware Implementation

4.1 Digital Design Flow

As it has already been discussed, the on-board FPGA provides support for the peripherals not commonly found in a commercial motherboard and will be exploited in control applications requiring fast data acquisition, real time processing capability and high bandwidth communication interfaces. The FPGA is on field programmable and addresses the need of customization required by several applications in the SPES control system. Nevertheless, already during the hardware development, with particular reference to the schematics capture phase, it is important to define the main FPGA functionalities and the integration of the FPGA resources into the embedded computer board. The firmware developed for the FPGA has been coded in VHDL language, which is an IEEE standard since 1987 [29]. The VHDL code has been developed following professional coding style rules intended for writing portable, readable and of quality code and to make finding errors easier [30].
The firmware implementation followed envisages the traditional digital design flow shown in Figure 4.1. The synchronous digital design to be inferred into the FPGA was first implemented using the synthesizable constructs of the VHDL. After the synthesis which generates a netlist for the target FPGA, we performed functional simulations to validate the logical correctness of the Register Transfer Level (RTL) design implemented. This phase avoided logic errors in the code and has eased the debug of the IOC board prototype. A complete timing analysis of the design is available upon successfully complete the automatic design implementation (place and route). The timing report provides information about the constraint coverage statistics and the timing violations that must be fixed to avoid metastability and

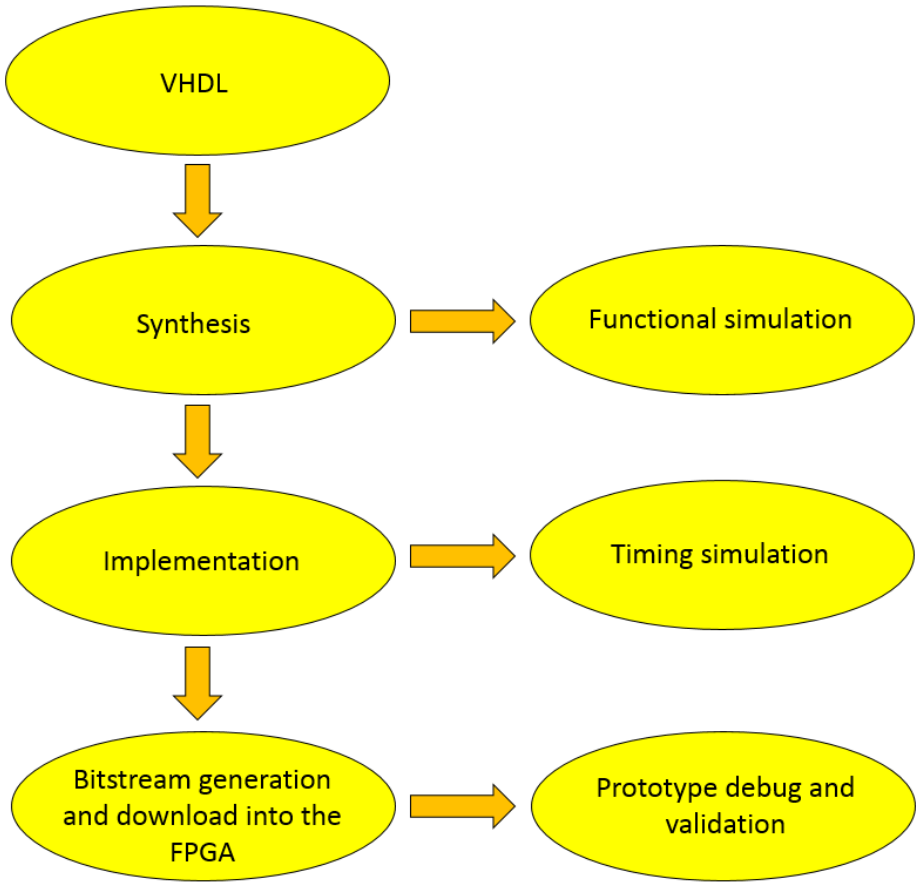

Figure 4.1: Digital design flow.

hardware malfunctions difficult to reproduce and debug. The Xilinx's ISE design suit 14.7, used to implement the firmware, supports all the 6 series devices and integrates into a unified design environment the tools to synthesize, implement and analyze the digital design. It has enabled the possibility to generate post place and route simulation models that include the timing delay information necessary to compute the slack associated with each connection. The last step of the firmware implementation was the test with hardware; the configuration file has been downloaded into an evaluation board, before the arrival of the prototype, and in the IOC board prototype subsequently. The code has been developed following a modular and hierarchical approach. The cores developed for hardware debug purposes during the IOC board validation have been integrated, with minor changes, into the final application top-level entity where they serve dedicated interfaces.

4.2 RTL Design Overview

An overview of the RTL design with the flow of the digital signals and the main data processing modules is given in Figure 4.2. The primary task of the firmware is to provide a local management and transparent mapping of the relevant process variables of the peripherals under control. Each module on the right side of the block diagram controls a peripheral or a data bus rather than an on-board IC and represents a hardware abstraction layer that exposes in a memory space all the relevant information and control bits of the device under control. These cores hide to the end user the complexity of the low level communication protocol; I/O operations often must comply with several time constraints. The RTL design integrates together many digital cores. Some of them are completely custom modules that interface custom peripherals. Others include proprietary or open source cores that

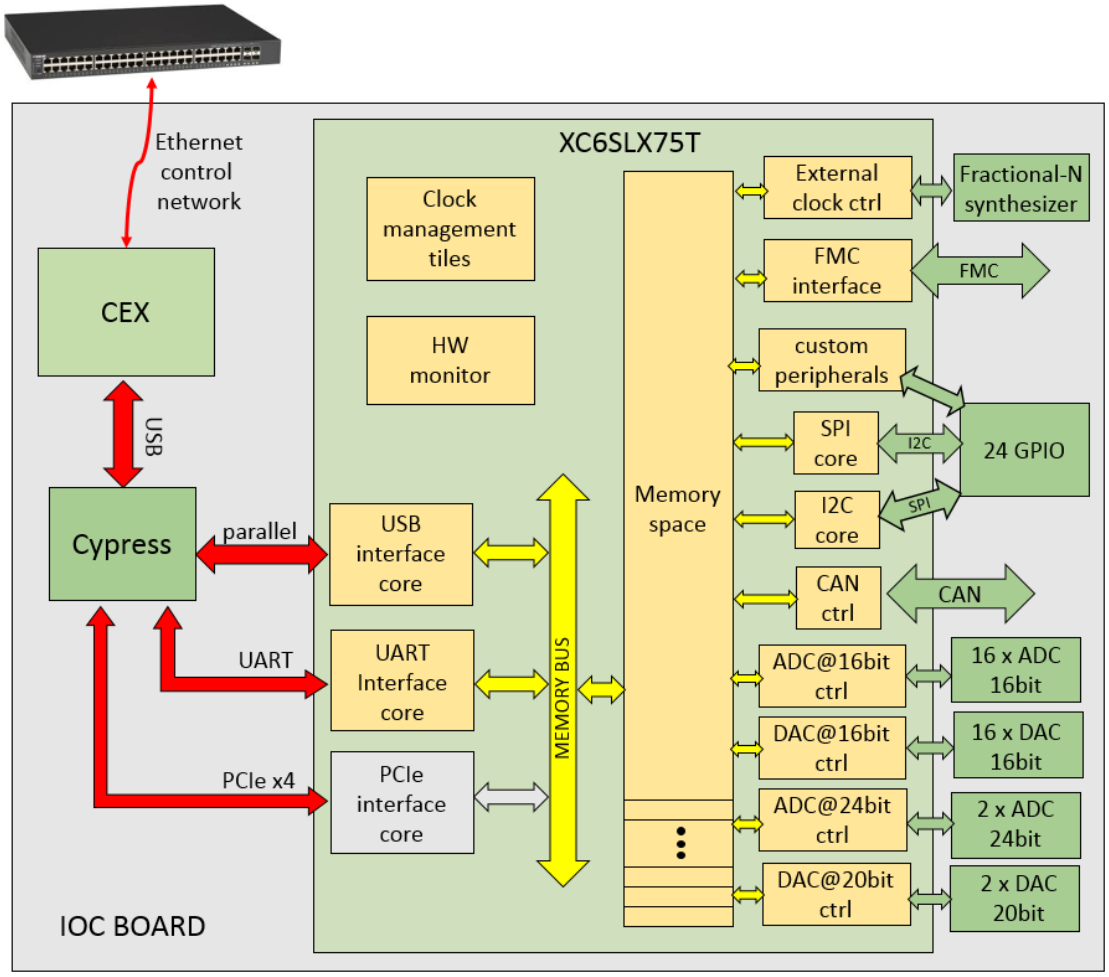

Figure 4.2: RTL design overview.

implement the physical layer of standard communication protocols like the SPI, I^2C or CAN bus, and whose characteristics are parameterized at a higher layer via the VHDL generic construct or the addition of external custom logic that addresses the requirements of a specific device.
The RTL design time has been accelerated exploiting the Xilinx CORE Generator system that provides access to highly parameterized intellectual property cores, among which, sophisticated clock management tiles and Multi-Gb/s transceivers for serial standards such as PCIe and Gigabit Ethernet. The memory space, in the middle of the block diagram, has been conceived as a set of memory blocks each composed of a stack of registers or a dual port RAM or a First Input First Output (FIFO) memory, depending on the specific application requirements. It implements the user defined addressing space of the IOC board and can be remotely accessed run-time for configuration, initialization, to monitor the status of the system and for data acquisition. This addressable space is mapped in the EPICS database loaded in the COM Express module via a USB interface core or a UART interface core. The PCIe endpoint has not been implemented yet, therefore, left in grey in the picture. These interface cores transfer the software read and write operations to the memory bus whose architecture accomplishes with our aim to create a simple and common interface between the master (USB or UART interface core) and the addressable memory slaves. The data width on the memory bus has been fixed to 32-bit and, so far, the sign has been handled using the two's complement format. Nevertheless, a fixed point arithmetic is supported by the interconnection architecture implemented.

4.2.1 Digital Design Performances Boosting Techniques

This brief subsection is devoted to the digital design aspects encountered in synthesizable designs which have been paramount for the success of the FPGA project [31].

- In order to produce a synthesizable code and for easing the debug we avoided complex processes containing many conditional conditions; the process represents an hardware circuit. The pure algorithmic thinking typical of C programming leads to unsuccessful FPGA projects.

- Pipelining to reduce the latency of the largest combinatorial logic blocks making them compatible with the clock frequency.

- Fanout. In modern devices, the delay due to routing is often a significant percentage of the overall latency, if compared with the logic delay upon pipeline. Reducing the fanout and buffering helps to reduce the capacitive loading on the nets and to prevent timing violations.

- Look-ahead techniques can help in boosting the performance of digital filters.

- Synchronous resets. Any core implemented uses the reset signal as a synchronous reset; the rest is synchronized to the clock domains at top-level.

- Complete state coverage in the finite state machines. If for any reason the state machine goes into an illegal state it must be taken to a safe state, usually, to the idle state.

- Synchronization between clock domains to avoid metastability problems. The synchronization of external inputs or of single-bit slow control signals coming from the memory space has been achieved using a synchronization circuit composed of two cascaded flip-flops (double flopping stage). Whilst, to pass multi-bit data from one clock domain to another asynchronous clock domain we adopted dual-port RAMs or FIFOs with independent read and write clocks, both readily available as Xilinx intellectual property cores.

- Synchronization of the high speed serial streams. In the IOC board, the synchronization of the optical link is guaranteed by the Xilinx gigabit transceivers that automatically locks on the input data stream and tracks the phase of the input data in order to shift the recovered clock to the best sampling point minimizing the possibility of having the marginal capturing phenomena. This is fundamental for establishing a reliable bidirectional communication.
  More in general, the synchronization of user defined high speed serial links can be achieved by exploiting a clock data recovery core.

We did not implement redundancy techniques since the IOC board will not be installed in rad-hard environments and is not intended for applications where the safety is a primary concern.
The firmware developed is based on these guidelines and the digital design implementation succeeded without timing violations. The consequent firmware debug on the IOC prototypes has been straightforward.

4.3 Slow Control and Data Acquisition via COM Express

In the SPES distributed control system the most likely configuration of the IOC board is with COM Express and FPGA strictly coupled, with the FPGA that assists the main processor in real time applications and handles the data digitization, processing and provides access to hardware specific devices. In this configuration, the primary task of the firmware is to establish a reliable communication channel to the local processor. This communication channel is exploited for monitoring and slow control operations as well as to readout beam diagnostic data and the information necessary to transport the beam. The relevant status information, the control bits, and the data collected from the sensors on the field are gathered to the EPICS local database loaded into the COM Express through a communication channel that relies on the USB bus or, for debug and applications with low bandwidth requirements, on the UART bus. In mostly control applications at LNL, the worst case scenario in terms of data rate is in the order of few Mbps, well in the range of the USB interface core and the EPICS channel access architecture. High speed data acquisition and processing from instruments can be eventually handled within the FPGA, with the processed and reduced data presented to the USB interface core instead of the raw data.

4.3.1 UART Interface Core

The developed UART interface core, whose block diagram is shown in Figure 4.3, supports the bidirectional transmission of American Standard Code for Information Interchange (ASCII) characters between the COM Express and the FPGA. Upon having instantiated the UART core inside the FPGA, the user via the command prompt of the COM Express module can open the serial port connected to the Spartan-6 and can exchange messages with the hardware. The message consists of a string of bytes, where each byte represents an ASCII character. The input parser inside the UART interface core analyzes the string of ASCII symbols received and decodes the message in agreement with a predefined command table, see Table 4.1, that defines the protocol used to perform single read and single write operations to the memory space. The master in the communication is the COM Express that initiates any data transaction with the hardware. The control handshaking mechanism is implemented at the protocol level; any read or write request forwarded

to the FPGA by the software is followed by an acknowledge or error message, in agreement with Table 4.2.

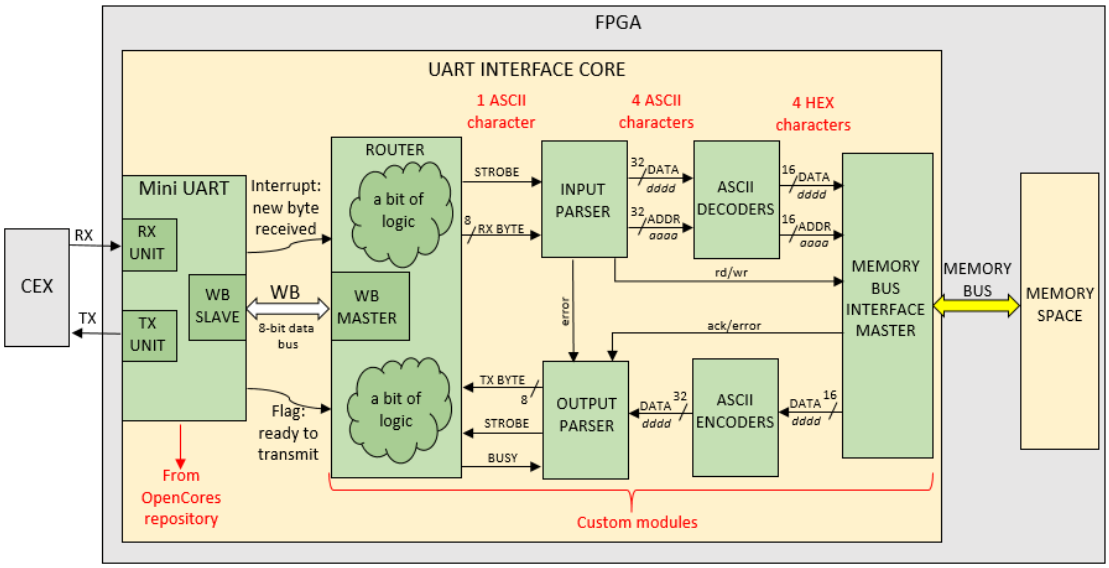

Figure 4.3: UART interface core overview.

Command	**ASCII characters**
Write the word *dddd* to the hex address *aaaa*.	$!aaaa : dddd \setminus n$
Read the word at address *aaaa*. The data field can be left to 0.	$?aaaa : 0000 \setminus n$

Table 4.1: Input parser command table.

Command	**ASCII characters**
The word *dddd* has been written successfully.	$\$dddd \setminus n$
Write command failed	$\#ffff \setminus n$
Read command returns the data *dddd*.	$\$dddd \setminus n$
Read command failed	$\#ffff \setminus n$

Table 4.2: Output parser command table.

The UART physical layer is provided by the Mini UART core downloaded from the OpenCores repository [32]. It features a WishBone (WB) [33] slave interface in 8-bit data bus and it does not implement transmit and receive FIFOs. Any time a new character is received, an interrupt is sent to the router core which starts a WB read transaction and bridges the received character to the input parser. Inside the

input parser, the received chars are gathered in an 11-byte deep buffer. As soon as a string of ASCII chars matches with one of the lines of the input parser command table, the address and data fields are converted in hexadecimal characters and the memory bus master starts the corresponding read or write operation on the memory bus. The memory bus is a simplified version of the WB system on chip architecture. The read and write operations return with an acknowledge or an error that is used by the output parser to build the reply message. The UART interface core supports the standard serial baud rates, selectable at synthesis time, and can address 2^{16} words. It has been developed during the hardware debug to perform the first functional tests on the IOC board prototype when the USB interface core was still under design. It is available for low power and low bandwidth applications, or as a training for newcomers students.

4.3.2 USB Interface Core

The primary communication channel between the COM Express and the FPGA is the Cypress FX3 super speed USB controller which, in our configuration, virtualizes a 16-bits parallel bus on the host USB 2.0 bus. In the processor, a user-space application programming interface provides the user with all the necessary calls to perform single read, single write, block read and block write operations to the hardware. The Cypress FX3 IC, loaded with a custom configuration firmware at power-up, transfers these read and write operations to the FPGA where the USB interface core acts as a parallel port endpoint. The signaling flow is shown in Figure 4.4. The USB interface core supports two data transfer types:

- memory: single read and write operations on the memory bus. 32-bit data width. We exploit the memory port data transfer to read and write the stack of configuration and status registers; slow control and monitoring operations.

- stream: block read and write operations similar to the Direct Memory Access (DMA) data transfer. The output stream data transfer type has been exploited to readout streams of raw data digitized with the on-board ADCs. Conversely, the input stream can be used in conjunction with DACs to emulate a signal generator. The maximum data transfer block size is configurable before synthesis in agreement with the depth of the FIFO instantiated into the memory space.

The communication between host and hardware takes place through the exchange

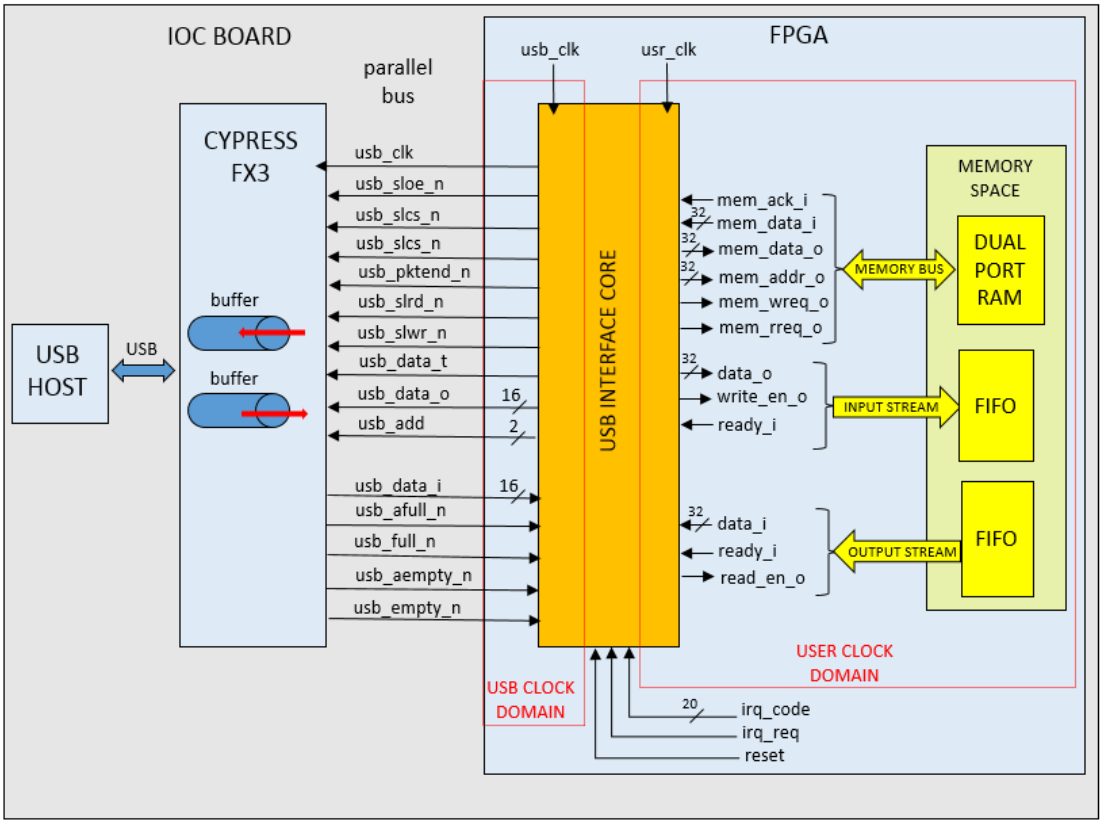

Figure 4.4: USB interface core. In this example we have only one memory port, one input stream port, and one output stream port. Multiple ports configurations are allowed.

of frames containing the address, data and header fields. The header specifies the source ID, destination port, direction, type of interface, data size, error codes. The USB interface core decodes/encodes these packets and routes the information to/from the correct stream or memory port. The USB interface core supports multiple memory ports and multiple I/O stream ports. When generated with multiple memory ports, the address lines are shared among ports and the active memory bus is selected using the read and write request vectors. Therefore, each memory bus maintains a 32-bit addressing space. In our firmware release, any memory slave has its dedicated memory or stream port in a point-to-point connection.

Any data transaction is started by the USB host, however, hardware interrupts are also supported by the core. For example, the interrupt routine may be used as a back-pressure signal indicating that the internal FIFO memory is almost full. The interrupt request is explicated by a 20-bit interrupt code that selects the correct interrupt service to be executed.

The core integrates two separate clocks, one for the parallel bus and a user clock

for the memory space and the user logic. The synchronization between these two clock domains is handled inside the USB interface core using FIFOs. The readout capability, in terms of data rate, is limited by the host and by the EPICS architecture; clocking the parallel bus with a 40 MHz clock, the hardware saturates the USB 2.0 bandwidth during block transfers. The USB interface core addresses the bandwidth demand of almost all the IOC applications in the SPES control system.

4.3.3 Example of Analog Input Data Acquisition and Control

Figure 4.5 describes a typical firmware implementation for the 16-bit analog input data acquisition. A custom core, in yellow in the picture, handles the data conversion and the data acquisition via a serial interface. This core implements the hardware interface and provides a good level of hardware abstraction to the end user. All the significant control variables and the raw data are exposed to the memory space, and, the data digitization process is conceived as a sequence of simple read and write operations by the end user. In order to digitize an analogue voltage signal fed into the analogue input channel, the user has to set, in the configuration registers, the desired acquisition mode (continuous acquisition or a predefined number of samples), the sampling rate, and the start/stop acquisition. All samples will then made available to the software in a FIFO memory. The FIFO can be mapped in the memory port or in the output stream port of the USB interface core, depending on the application. E.g., in continuous acquisition mode at hundreds of ksps, the user may want to continuously readout the raw data using the output stream port. Whilst, if the user wishes to monitor slowly varying signals, or just getting a bunch of consecutive samples, the FIFO can be mapped in the memory port, the same as the configuration space like in the picture or, eventually, a dedicated one. The low-level communication with the ADC device and the complete digital signal flow is handled by a finite state machine, and, therefore, completely transparent for the software. It is a hardware task to meet all the time specifications and to know the details of the physical layer communication model.

4.4 Gigabit Ethernet 1000BASE-X Implementation

The SFP module directly coupled with the Spartan-6 gigabit transceiver opens the possibility to implement an optical Gigabit Ethernet connection. This intercon-

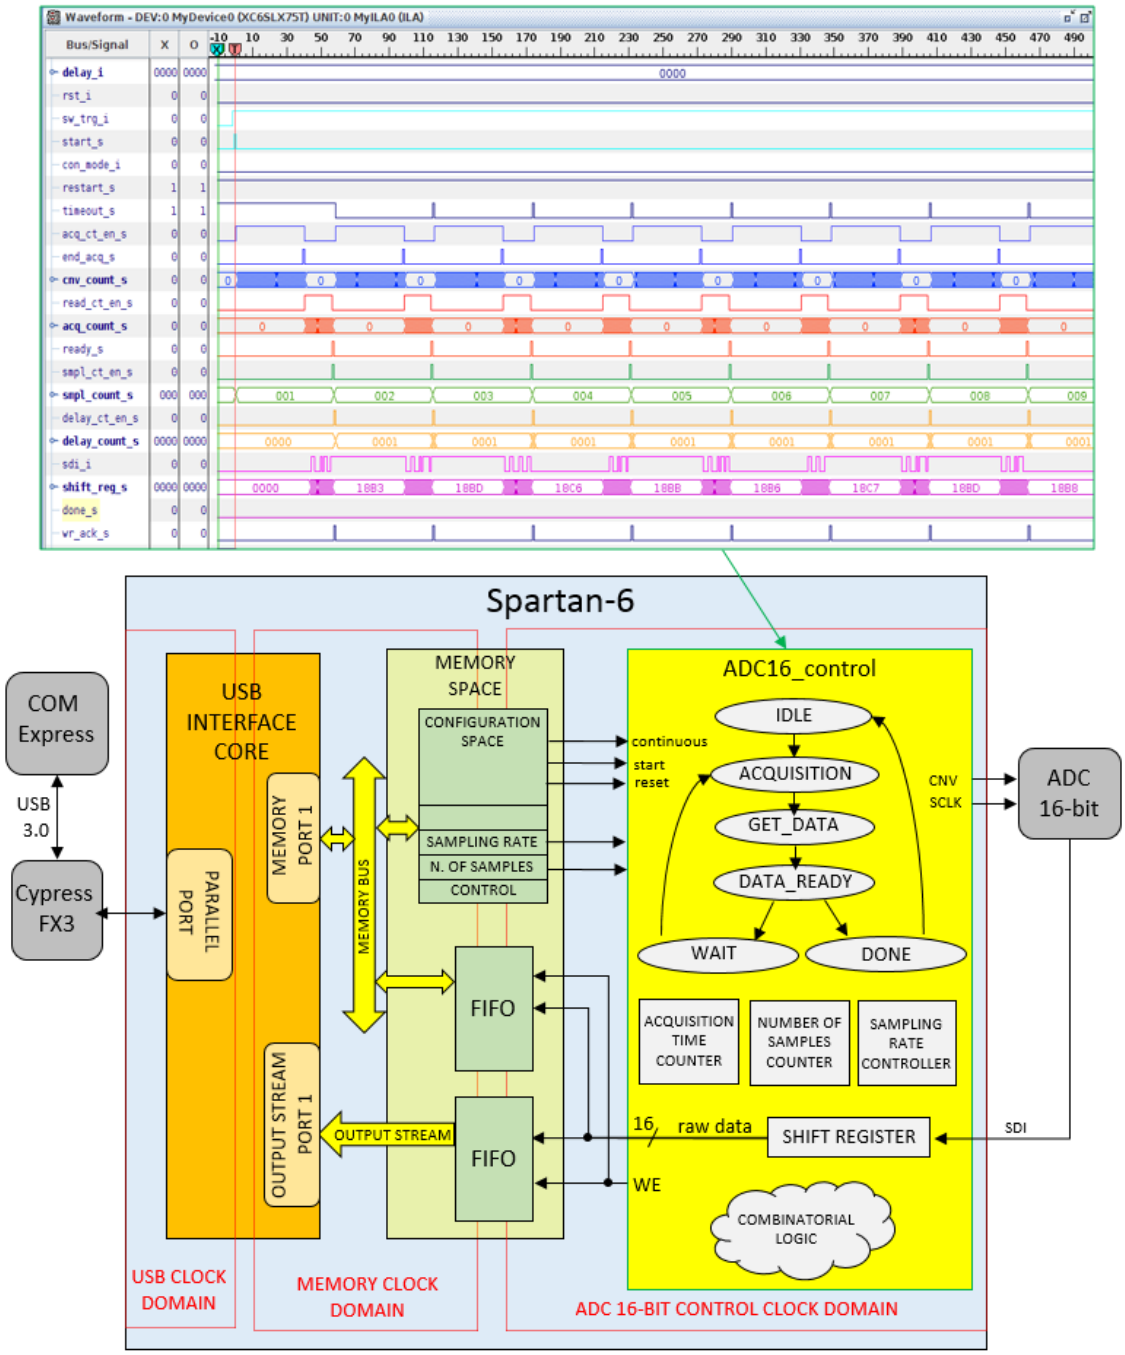

Figure 4.5: ADC 16-bit control and data acquisition via the USB interface core.

nects solution provides the IOC board with a modular interface that can be easily adapted to several optic and copper networking standards by plugging the proper SFP transceiver. The implementation of an optical Ethernet channel is a highly appreciated feature to be exploited in those applications where the FPGA handles high speed data acquisition and readout tasks in a standalone manner (without

COM Express), and, in centralized control systems.
The 1000BASE-SX is a fiber optic standard for transmitting Ethernet frames at a rate of the Gigabit per second. The electro-optical conversion is handled by the SFP transceiver that, in our case, operates in the near-infrared wavelength (850 nm) with a range capability between 220 m and 550 m. As shown in Figure 4.6, the Ethernet physical layer is built around the Spartan-6 Gigabit Transceiver Port (GTP) that handles the data serialization/deserialization and the 8b/10b encoding [34]. The physical layer is completed by the Xilinx PCS/PMA intellectual property core that communicates with the MAC via the Gigabit Media Independent Interface (GMII). The MAC implements the data link layer foreseen by the Open System Interconnection (OSI) model of the telecommunication network, and it has been instantiated in the design as a parameterizable logic core that complies with

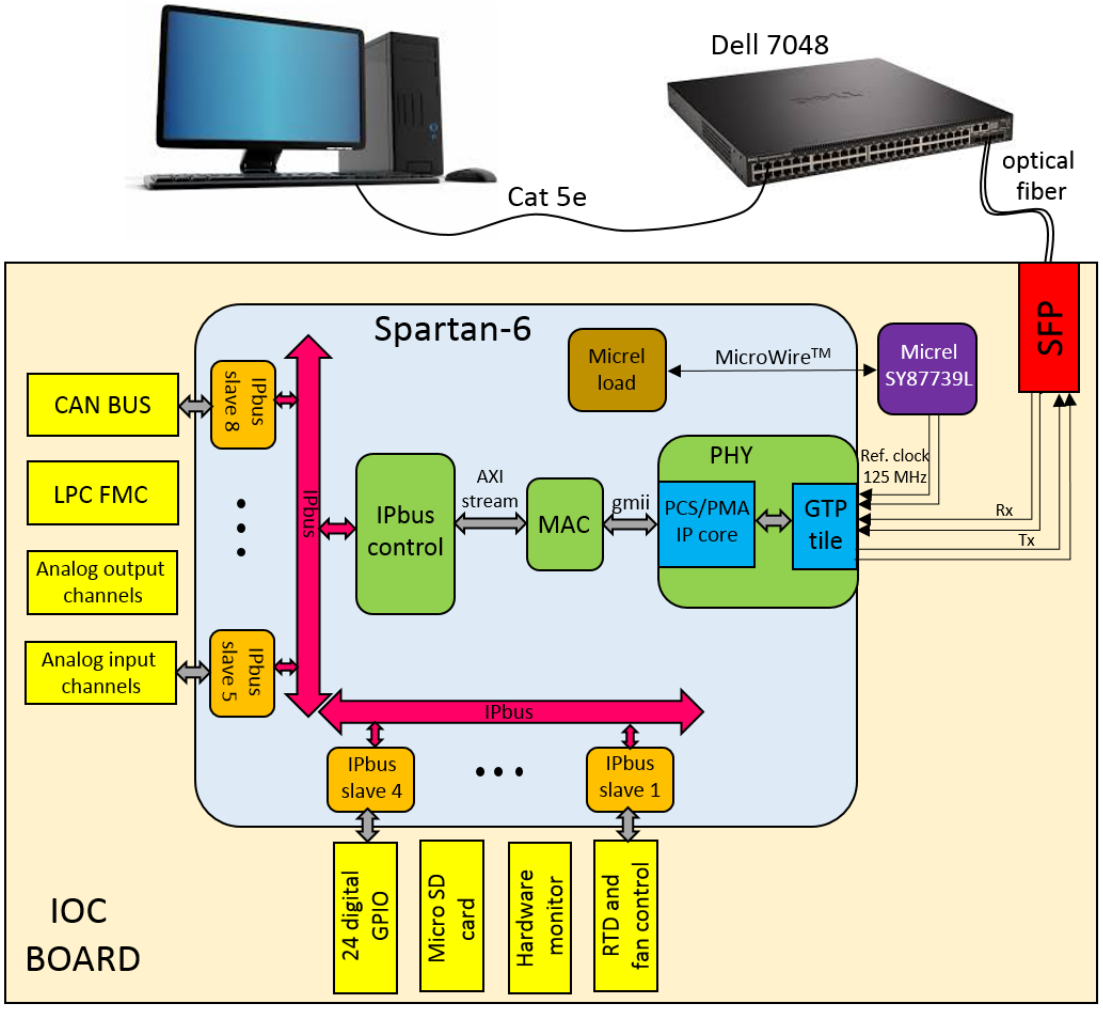

Figure 4.6: IOC board control via IPbus.

the IEEE 802.3 specifications; this enables the user to implements a variety of Ethernet-based designs, from low cost 10/100/1000 Mbps to 2.5 gigabit Ethernet ports [35]. The main tasks performed by the data link layer are framing, synchronization, error detection, flow control. Despite the physical layer and data link layer operation complexity is hidden to the end user thanks to the Xilinx's off-the-shelf logic cores, special care must be taken during their instance in a hierarchical design since the power-up sequence, the reset and clocking strategies may affect the proper operation of such cores. On top of the data link layer, in the OSI model, there are the network layer and the transport layer. The network layer implements the Internet protocol that specifies addressing, format and routing information. There are two standard protocols available: the *transmission control protocol* (TCP) that is a highly reliable protocol that guarantees data delivery, and, the *user datagram protocol* (UDP) that is less reliable but simple and fast protocol ideal for streaming sequences in audio/video applications. The transport layer specifies the mechanisms for flow control, error detection and error recovery. Both the internet protocol and transport protocol are handled by the IPbus control core provided by CERN [36] [37]. The IPbus is an IP-based protocol for controlling hardware devices. The IPbus software client applications generate IPbus transactions (read and write requests) transported to the hardware target device via the UDP transport protocol. On the client side, a hardware access library provides the user with a Python/C++ application programming interface for starting IPbus read, write and read-modify-write operations. Whereas, within the FPGA, the IPbus control server provides a UDP endpoint and bridges the UDP packets on the IPbus system-on-chip interconnection architecture. This hardware bus foresees 32-bit data transfer and 32-bit word addressing. Each instance of the IPbus control core responds to a unique IP address hard-coded at synthesis time, or, dynamically assigned with the RARP protocol in order to have a single firmware for multiple hardware platforms. The IPbus control core behaves as IPbus master and initiates the transfer cycles with the addressed slaves upon the reception of a UDP packet containing one or more IPbus transaction requests. The IPbus slave can be a simple block of memory or a complex synchronous logic block controlling a device linked to the FPGA. The IPbus provides a flexible and uniform interface between slave modules and guarantees a unified protocol to access the slaves internal registers containing relevant data and control variables. The implementation of an Ethernet end-point inside the FPGA exposes all the IOC board resources directly connected to the Spartan-6 on the accelerator control network, without the need of the COM Express module.

All the analog and digital input and output channels, the CAN bus, the FMC mezzanine board are remotely accessible for those applications requiring a centralized control architecture.

The Gigabit Ethernet transmission over optical fiber (1000BASE-X) is defined by the IEEE 802.3z standard which includes various technologies for transmitting Ethernet frames at a rate of 1 Gbps. The design implemented and described in this section refers to the 1000BASE-SX standard which foresees data transmission over multi-mode fiber using 770 to 860 nm wavelength up to 550 m distance. The SFP transceiver is the part of the Gigabit Ethernet physical layer that defines the communication media used to transmit the Ethernet frames. The IOC board, running the RTL design described above, complies also with the 1000BASE-LX and 1000BASE-CX standards by plugging the proper SFP module. The 1000BASE-LX standard is intended for multi-mode fiber using 1,270 to 1,355 nm wavelength for moderate distance transmission (up to 550 m) using multi-mode fibers, or, for long distance (up to 5 km) transmission over a single-mode fiber. The 1000BASE-CX standard is intended for transmissions over balanced shielded copper cables and foresees a maximum communication distance of 25 meters. The short transmission distance is due to the very high signal transmission rate and the high attenuation of copper cables. The 1000BASE-T has succeeded it for general copper wiring applications; the usage of 4 twisted pairs and the diverse data channel encoding strategy have extended the transmissions range up to 100 m via copper cables.

Chapter 5

IOC Board Characterization

5.1 Power Consumption

Upon arrival of the first IOC board prototype, the validation process started powering the board and checking the voltage at the output of the voltage regulators in different load conditions. The main 12 V can be supplied by a standard ATX module, by a benchtop power supply, and, the IOC board is PoE+ compatible. These three power strategies are hot-swappable and may coexist. The power consumption is a key parameter for the IOC board and should be limited due to the requirement of being remotely powered via PoE+ in several applications. The power consumption fluctuates depending on the CPU usage and on the logic designed within the FPGA. The FPGA power consumption is strongly affected by the percentage of resources utilization, the number and rate of I/O pins used, the clock frequency of synchronous circuits and the toggle percentage of internal logic. Table 5.1 resumes the measured power consumption of the IOC board in different configurations. This power analysis does not include the contribute of commercial module in PCIe or FMC form factor eventually plugged in the carrier. The observed power consumption with all the power rails enabled, the Spartan 6 loaded with a low activity hardware validation firmware release, and a compact COM Express module with an Intel Atom processor E3845 @ 1.91 GHz is about 22 W. Increasing the CPU usage and the Spartan-6 power demand, the observed IOC board power consumption has reached peaks of about 31 W. These values are slightly out of range for the PoE+ standard. Indeed, considering the efficiency of the PoE+ flyback converter equal to 85%, then, temporary power requests above 26 W would force the PoE+ controller chip in overcurrent protection mode causing

undesired shuts down.
Shutting down the COM Express module the power consumption reduces to about 12 W. For those applications not requiring the processor and that exploit only the FPGA plus the analog input/output channels the target power consumption ranges from 10 W to 15 W. The PoE+ is a feasible solution in these use cases.
Programming the onboard power manager chip, the overall power consumption can be further reduced of about 4 W disabling the analog section.

IOC board operating conditions	**Power consumption**
Full board power consumption: all power rails enabled, COM Express running, FPGA configured.	21 - 31 [W]
Carrier board power consumption without COM Express module	10 - 15 [W]
Carrier board power consumption without COM Express module and analog section disabled	6 - 11 [W]

Table 5.1: IOC board power consumption measurements.

5.2 Thermal Distribution Analysis

The PCB layout density makes the heat management critical and of crucial importance; excessive heating of the PCB causes performance degradation and reliability reduction. The high temperature accelerates failures mechanisms both at die level and PCB level, and this highlights the importance of performing power distribution simulations already at the design phase: simulations allow the correction of macroscopic errors and facilitate the integration of cooling solution at design time. Nevertheless, in complex PCBs that integrate over 2000 components, like the IOC board, a lot of potential issues may arise during the PCB manufacturing and assembly. These tiny faults in PCB components might manifest as bigger failures during the first power-up of the board or, worse, as malfunctions very difficult to debug. The solution often is a thermal imaging. Thermal imaging cameras are getting higher resolutions and open the possibility to perform real time heat distribution analysis and help the diagnostic of potential overheating problems in the PCB in an early prototyping stage. Figure 5.1 shows the temperature profile of the IOC board during nominal operating conditions. The heat distribution in the IOC board is clearly visible and thanks to the good resolution of the infrared camera it is possible to identify the hottest components and areas of the PCB. The

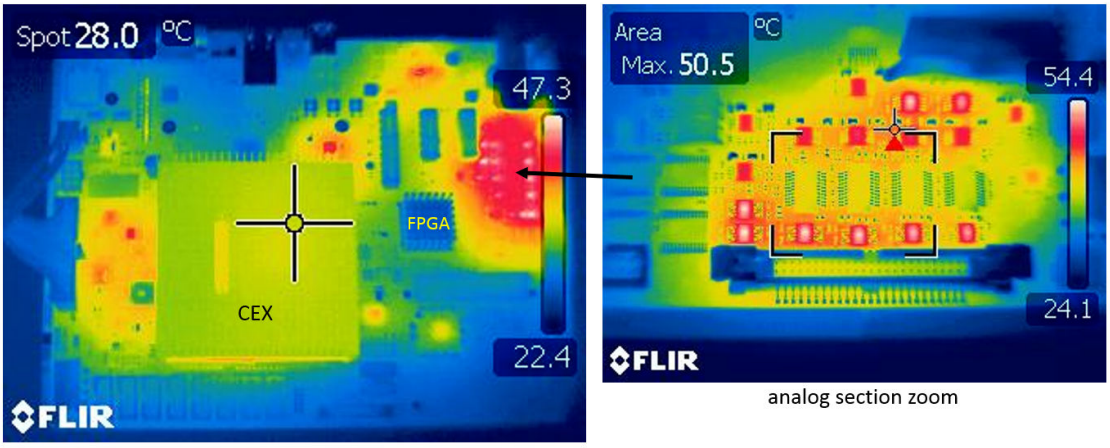

Figure 5.1: IOC board thermal image.

FPGA and the COM Express module are not a concern, indeed, even increasing the processing rate, exploiting simple solutions such as a cooling fan and heat sinks these components reach the thermal equilibrium with temperatures below 45 °C. The analog section is the warmest area of the PCB, as expected. A zoom on the analog section shows the high density of this area and allow the identification of criticalities: it turns out that the 16 operational amplifiers are the hottest chips. These low noise high speed amplifiers have a quiescent power of about 100 mW and are available in a small MSOP 8-pin package without the exposed pad. The

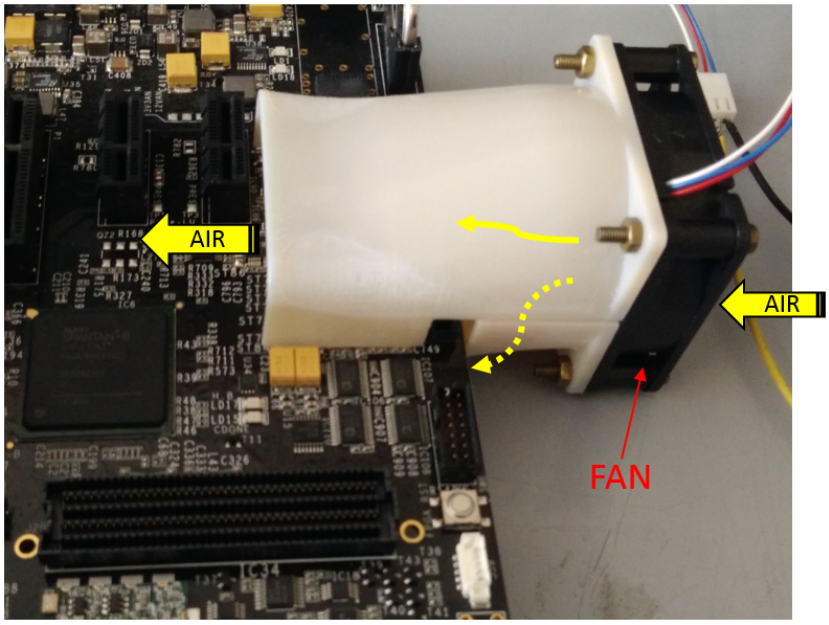

Figure 5.2: PCB analog section custom cooling solution.

chip is supplied via the ±12 V rails and has no connection to the ground planes, therefore, the conduction heat transfer is limited. The on-board hardware monitor can be programmed to disable all the analog power rails to reduce the overall power consumption and heating of the IOC board in those applications that do not require analog input and outputs channels. For those applications requiring the input/output analog channels, a custom cooling solution based on a custom shaped heatsink plus a dedicated fan conveyor system has been developed, and, it is shown in Figure 5.2. A thermal image is not only a feedback of the correctness of the PCB design but is also a good diagnostic tool. Often tiny faults cause hot spots on the PCB that can be easily identify by a thermal image. Figure 5.3 is an example of two defective IOC boards. On the left a failing component was loading excessively the 3.3 V power rail thus forcing the DC/DC controller thermal shutdown; on the right an improper solder bridging easily detected with the thermal image and afterword diagnosed through visual inspection.

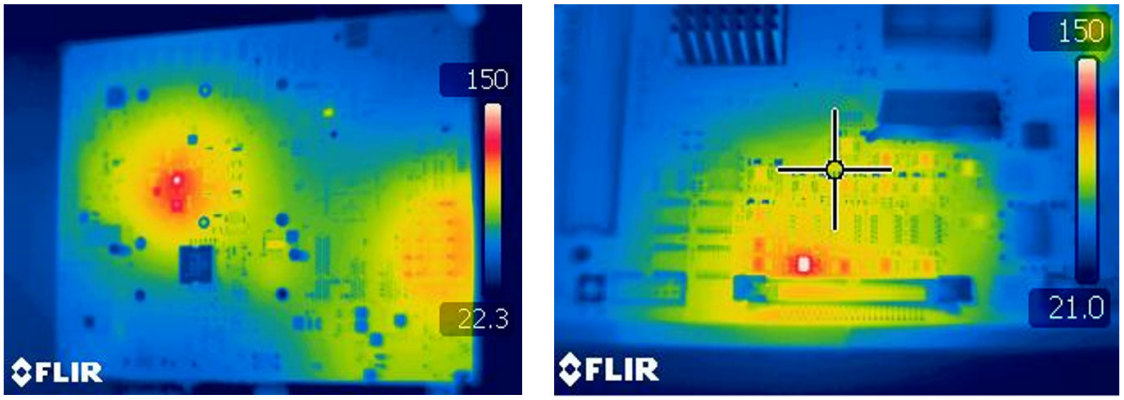

Figure 5.3: Tiny faults very often generate hot spots highlighted by the thermal image.

5.3 Data Digitization Accuracy

The Effective Number of Bit (ENOB) is often used as an indicator of the overall performance of the analog input channels [38]. The dual differential ADCs assembled on the IOC board feature a theoretical 16-bit resolution with a maximum throughput of 1 Msps. The ADC manufacturer specification reports an ENOB of about 15.1 with a reference voltage of 4 V (our case). The ENOB of the 16 channels has been computed from the baseline noise measurements given the absence, in our Lab, of a sine wave generator having noise level at least four times less than

the level of accuracy required for the ADC random noise measurement. We configured the board to record long traces from the input data streams by setting the ADCs in continuous sampling mode with a chosen sampling rate of 690 ksps. The baselines were recorded channel by channel grounding the corresponding analog input on the connector; the measurements account for the overall accuracy of the acquisition system inclusive of the frontend analog electronics and other circuitry that feeds the analog input channel. Whereupon, the baseline data were transferred to a server running Matlab® [39] where the ENOB was computed following two different procedures.

5.3.1 ENOB Calculation From SNR

This approach exploits the quantization noise model [40], and replaces the maximum error of an ideal converter with the rms noise value extrapolated from the experimental data. The baseline values, properly sorted into bins, have been fitted using a single term Gaussian distribution model, as shown in Figure 5.4. Assuming a Gaussian distribution also for the noise, the standard deviation σ of the histogram represents the ADC input noise in terms of Least Significant Bits (LSB) rms. The theoretical signal-to-noise ratio in dB can now be calculated assuming a full-scale input sine wave as described by (5.1):

$$SNR = 20 \log_{10} \frac{rms\_full\_scale\_input}{rms\_ADC\_baseline\_noise} = 20 \log_{10} \frac{2^N/(2\sqrt{2})}{\sigma} \qquad (5.1)$$

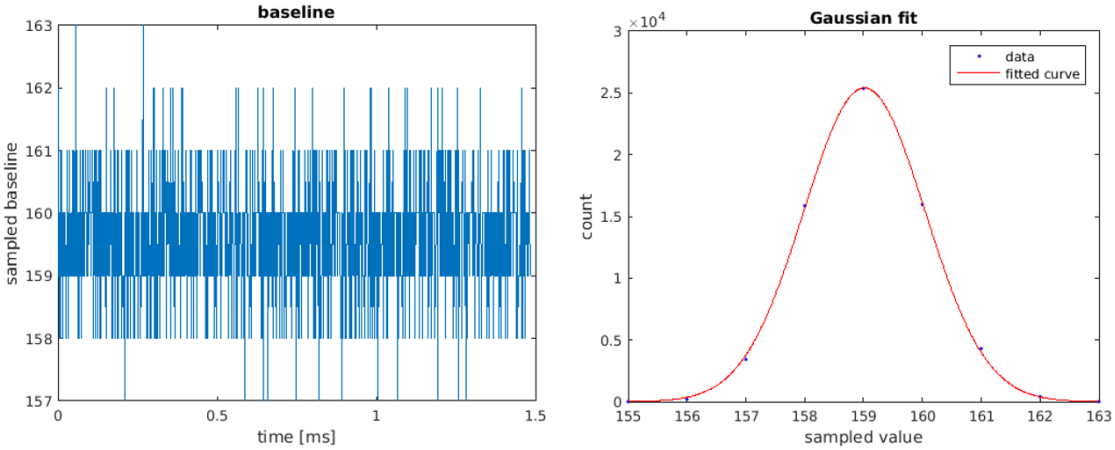

Figure 5.4: Digitized trace for the ADC baseline noise analysis. On the right the histogram bin counts interpolated with a Gaussian distribution.

where N is the declared ADC number of bits. The ENOB can be calculated from the theoretical Signal to Noise Ratio (SNR) of an ideal N-bit ADC (5.2):

$$ENOB = \frac{SNR - 1.76}{6.02} \tag{5.2}$$

The Matlab code is reported in Appendix B.

5.3.2 ENOB Calculation From SINAD

A more accurate method for quantifying the ADC dynamic performance is the SIgnal to Noise And Distortion ratio (SINAD) [41]. The spectral output of the Fast Fourier Transform (FFT) applied to the baseline is shown in Figure 5.5. The sampling rate (fs) was set to 690 ksps and the total frequency range covered by the FFT analysis is from dc to fs/2. The frequency resolution is obtained dividing the sampling rate by the number of samples collected (65536 samples in our case). As visible from the picture, an ideal full-scale sine-wave has been added to the sampled noise floor by software and the FFT output is used to measure the amplitude of the different harmonics and noise components of the digitized data. This analysis accounts for possible distortion products generated by the circuitry surrounding the ADC and coupled with the digitized analog input, and by the Differential Non

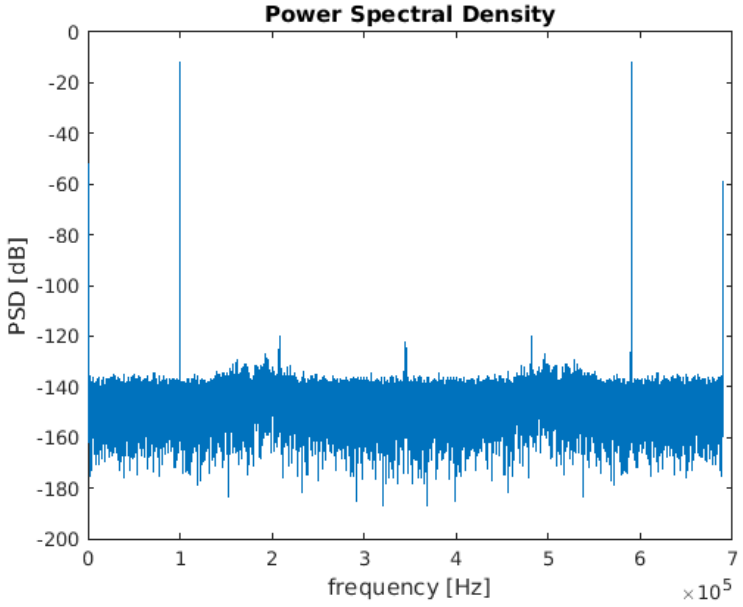

Figure 5.5: ADC baseline power spectral density.

Linearity (DNL) of the converter itself. SINAD value is returned by the Matlab function *sinad* and is a good indication of the overall ADC performance because it includes all components which make up noise and distortion. SINAD can be converted to ENOB using the relationship (5.3):

$$ENOB = \frac{SINAD - 1.76}{6.02} \tag{5.3}$$

5.3.3 IOC Board ADCs Performances

ENOB has been computed for each channel using both methods and the results are resumed in Figure 5.6. As expected, the SNR slightly overcomes the SINAD of all channels and there is a good correlation between the two methods described above. All channels, according to the manufacturer, are at least 14 bits accurate and this proves the quality of the hardware and validates the analog and digital grounds strategy adopted. To meet the mechanical constraints, it was impossible to assemble on the IOC board 16 coaxial connectors for the analog inputs, and, as many for the analog outputs. Therefore the overall 32 analog input/output channels have been routed to a 50 positions header connected to a front-panel via

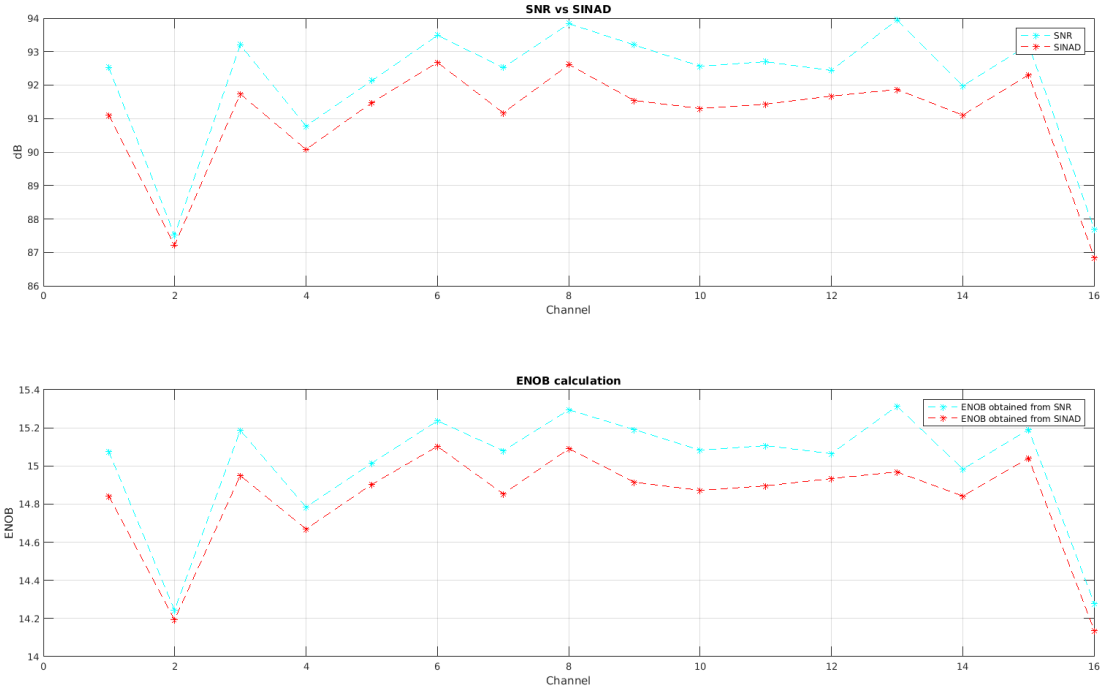

Figure 5.6: SNR, SINAD and ENOB calculations for the 16 analog input channels. The baseline noise has been measured grounding the analog inputs on the connector.

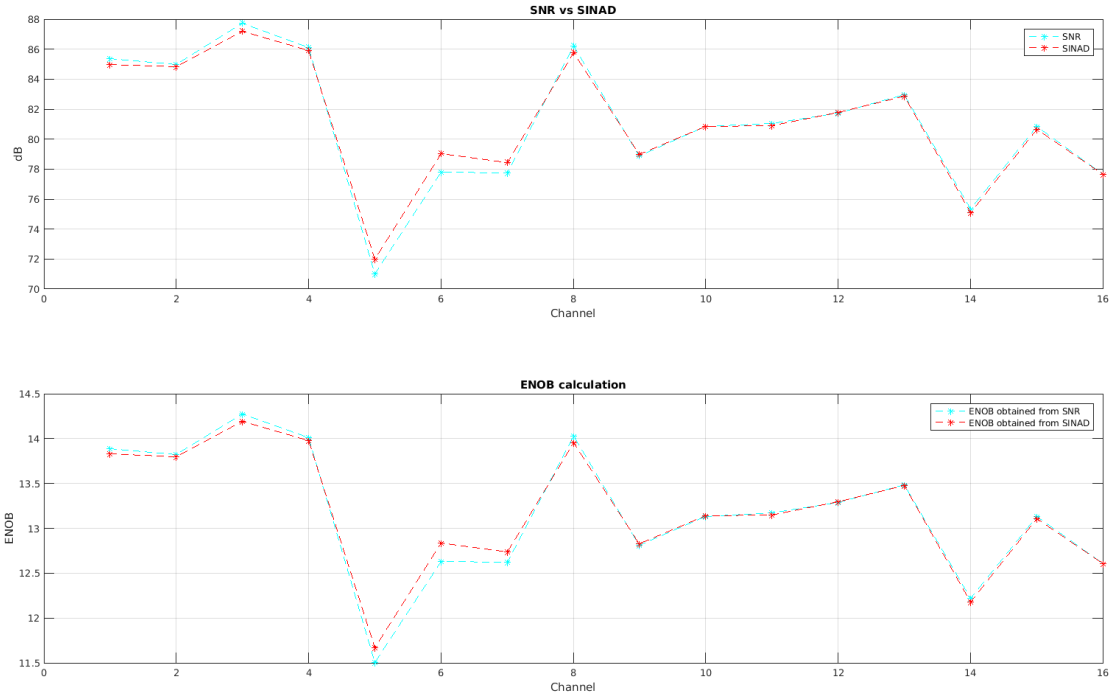

Figure 5.7: SNR, SINAD and ENOB calculations for the 16 analog input channels. The baseline noise has been measured grounding the analog inputs channels after about 20 cm of flat cable.

a flat cable. This solution reduces the overall accuracy of the acquisition system to about 13 bits, as shown in Figure 5.7. The increased noisiness of the digitized data leads to a higher variance of the ENOB after consecutive baseline measurements of the same channel. SINAD and SNR tend to converge to similar values since the noise dominates the harmonic distortion contribution. Nevertheless, interpolated noisy data may lead to non-accurate SNR evaluation if the data-set is not big enough; therefore, it is preferable to refer to the ENOB value derived from SINAD. The measurement accuracy of low-bandwidth parameters, such as a temperature, pressure, beam current, may be increased by oversampling and then decimation within the FPGA [42],[43]. Assuming that the dominant noise is modeled as white noise, for any additional bit of effective resolution the signal must be oversampled by a factor of four as shown in equation (5.4):

$$f_{os} = 4^{w} \times f_s \tag{5.4}$$

where f_{os} is the oversampling frequency, w is the number of additional bits of resolution desired, and f_s is the original sampling frequency requirement imposed by the Nyquist Theorem. Once the oversampling is accomplished, the extra 4^w

samples, gathered and buffered into the FPGA during the sampling period, are accumulated (by adding the 4^w consecutive samples) and then decimated to yield a more accurate sample. The decimation consists in an arithmetic average (dividing by 4^w or right shifting of 2^w bits) and then truncation of the result by using the upper 16 bits or to retain one or more extra bits to extend the ADC resolution beyond the theoretical 16-bits limit.

5.4 ADC Offset and Gain Calibration

The effective number of bit is a good indicator of the ADC performances and takes into account the ADC error caused by quantization and the system noise coupled with the ADC. But the accuracy of the measurements is affected also by non-ideal behaviors of the ADC itself like the offset and gain errors. These variations must be characterized for any ADC assembled in the IOC board and compensated since we are possibly dealing with high gain settings and large conversion range ($\pm 10V$) that may yield an uncompensated error above 20 LSB. Offset and gain errors are compensated in software in accordance with equation (5.5):

$$voltage[V] = value \times m \times \gamma + \alpha \tag{5.5}$$

where *value* is the ADC output value in LSB, m is the ideal gain, γ is the gain correction factor and α is the offset correction.
The calibration parameters γ and α are stored in the EEPROM memory on-board and have been calculated, channel by channel, using the following calibration procedure [44].

1. During the calibration process it is necessary to digitize DC voltages with an accuracy of about 1 LSB, therefore, we need to extend the ENOB to the theoretical number of bits, 16. The FPGA must be configured to handle the oversampling and decimation technique.

2. Compute the ideal ADC gain: $m = 10/2^{15} = 0.00030518[V/bit]$

3. Set the output voltage value in the software as: $voltage = value \times m$

4. Force the analog input channel to 0 V using a stable and precise voltage source. The voltage source resolution requirement is to be better than 1 LSB (0.3 mV).

5. Measure and store the output voltage. Let's call it A.

6. Increase the input voltage until the output voltage changes. Store the output as B.

7. Store the input voltage as ACTUAL.

8. Compute the input voltage required for the ideal ADC to change its output from A to B and store this value as PERFECT: $PERFECT = B - m/2$

9. Compute the offset error (input voltage to apply to get an output 0): $\alpha = ACTUAL - PERFECT$

10. Set the output voltage value in the software as: $voltage = value \times m + \alpha$

11. Set again the input voltage to 0 V and verify that the measured voltage is now 0 V.

12. Set the input voltage (y) to a DC voltage close to the positive (or negative) end point scale and measure the voltage. Store it as y_{fs}

13. Divide the ideal output (y) by the measured output to get the gain correction factor: $\gamma = y/y_{fs}$

14. Set your software to compensate for the gain and offset errors using equation (5.5).

The calibration results are shown in Table 5.2 and validate the proposed procedure. Upon correcting offset and gain errors the actual ADC curve may still slightly deviate from the ideal input-output curve due to non-linearity errors. These are below 2 LSBs by manufacturer specification, therefore negligible for our typical applications. Offset and gain errors specified by the ADC manufacturer have typical values of the order of few LSBs, but, these errors are compounded by the analog

DC input voltage	**voltage measured**
$-9.0002V$	$-8.99925701874V$
$2.5e^{-05}V$	$6.862034388e^{-05}V$
$8.9997V$	$8.99969822909V$

Table 5.2: ADCs 16-bit calibration results. On the left the DC input voltage, on the right the voltage measured.

frontend circuitry that must be included in the offset and gain error calibration procedure.
In the IOC board, the ADCs are driven by an analog frontend composed of a low noise amplifier in a follower configuration, and a differential instrumentation amplifier with a selectable gain that performs the single-end to differential conversion. As measured in practice, the overall offset and gain errors rise to tens of LSBs if we do not account for this frontend circuitry during the calibration procedure. Moreover, dealing with single-ended analog inputs, the operating conditions, the cabling and grounding strategy affect deeply the offset and gain errors and, ideally, the calibration procedure should be periodically implemented at runtime by software.

5.5 DAC 16-bit Functionality Test

The IOC board features 16 independent analog output channels implemented by 8 dual, 16-bit, DACs. A dedicated core running within the FPGA controls the digital to analog conversion via a serial interface that operates at 30 MHz. The DACs feature a power down mode that is selectable via internal registers and in power down mode the outputs are grounded. The DACs assembled in the IOC board operate from dual supplies ($\pm 12V$) to enable the possibility to select via software the output ranges: $5V$, $\pm 5V$, $10V$, $\pm 10V$, $10.8V$, $\pm 10.8V$. Figure 5.8 shows the

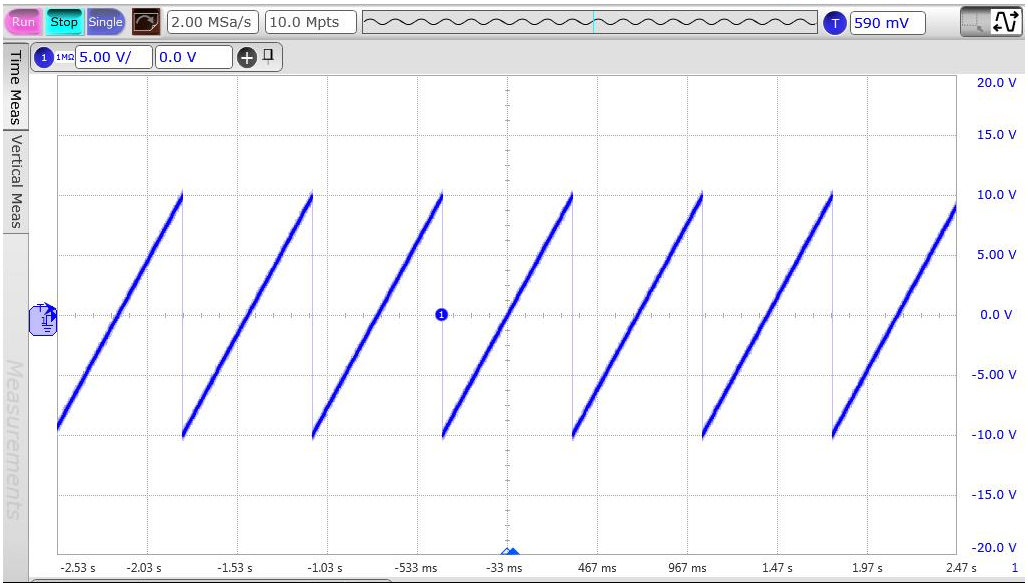

Figure 5.8: DAC $\pm 10V$ output sweep test.

outcome of the functional tests. The output sweep is generated in response to a digital stimulus consisting of a counter increment. The DAC output update is controlled by a DAC control core within the FPGA that receives the incremented count value from a Python script and shifts this value inside the DAC data input shift register. The new digital data is then loaded to the analog output upon reception of the load command via the corresponding bit in the control register.
To perform a more advanced characterization based on non-linearity tests and total harmonic distortion analysis, a sine output sweep response should be captured with a 20-bit resolution high speed digitizer, currently not available in our laboratory. The measured power and ground integrity, together with the fact that the DAC outputs directly drive the analog output connector without any extra circuitry in between, make us confident that there won't be significant performance degradation, at PCB level, with respect to the manufacturer specifications. Moreover, the need for DACs with an effective resolution higher than 12 bits is currently restricted to few applications already covered by legacy commercial solutions, therefore the DACs characterization in terms of non-linearity and total harmonic distortion is postponed as a future development.

5.6 FPGA Mezzanine Card Interface Validation

The FMC connector bridges the mezzanine card front panel I/O to the FPGA processing module and must comply with the electro-mechanical specifications foreseen by the FMC standard. The design and the connectivity correctness of the FMC interface have been verified with a commercial floating ammeter, shown in Figure 5.9. This low pin count mezzanine card implements four high resolution bipolar current meters. Each channel has two software selectable measuring ranges of $\pm 1mA$ and $\pm$ 1 µA [45]. The data conversion and acquisition is controlled by a logic core implemented into the FPGA: $FMC\_control$. This logic core interfaces the mezzanine card via serial bus plus additional control signals assigned between the FMC connector user-defined I/O pins. The 12 V and 3.3 V power rails are supplied by the carrier to the mezzanine via the dedicated power pins as specified by the FMC standard. A finite state machine controls the digital interface and guarantees that all the timing parameters such as the minimum acquisition time or the minimum time between two consecutive conversions is respected. The data conversion starts upon reception of the start acquisition signal from the user, and it ends on the falling edge of the busy indicator flag from the mezzanine to carrier.

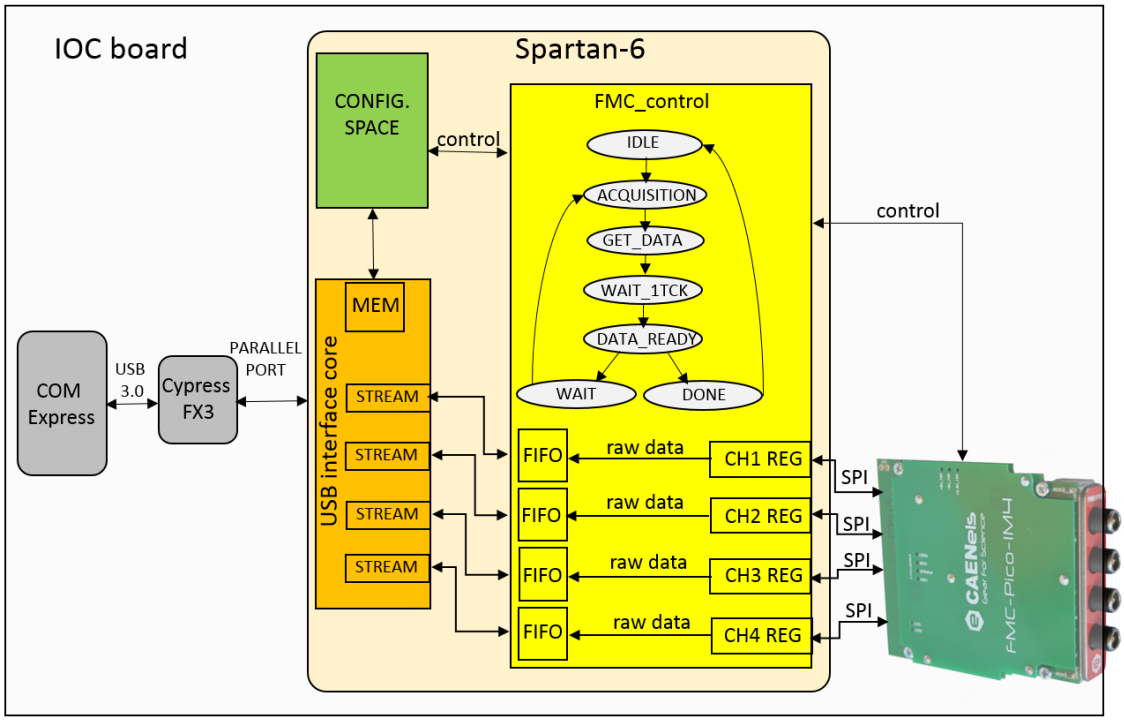

Figure 5.9: 4-channel, 20-bit, 1 Msps, FMC pico-ampere meter readout scheme.

The conversion phase is followed by the acquisition phase during which the digitized data is received from the serial bus and registered inside a shift register for the serial to parallel conversion. The FIFO memory performs the cross-clock domain synchronization and decouples the data sampling rate, tunable at runtime by software, from the raw data streams readout. The *WAIT* state delays the start of the new conversion of a software-predefined number of clock periods to match the desired sampling rate. The *FMC_control* core provides support for two working modes: continuous and fixed number of samples. If the end user specifies a number of samples different from 0, then, upon receiving the start acquisition command, the finite state machine cycles the predefined number of times before going back to the *IDLE* state, passing through the *DONE* state that rises an interrupt to the software, signaling that the requested data are ready. In continuous sampling mode (number of samples set to 0) the finite state machine continues to cycle indefinitely, until the user stops the acquisition. The FIFO depth is 1024 samples and whenever the FIFO is full, the FMC control core rises an interrupt to software. To avoid data loss there are more sophisticated techniques concerning asynchronous FIFO designs whose implementation is beyond the debug purpose, therefore this task is left as a future development [46].

By measuring the current values injected into the mezzanine analog inputs with a precise current source, it was possible to validate the functionality of the FPGA to mezzanine card interface. Noise and bandwidth measurements were also checked against the mezzanine manufacturer electrical specifications.
Noise tests have been performed by collecting the digitized trace of the baseline (as explained in section 5.3), but with the mezzanine input channels floating. Figure 5.10 shows the baseline digital trace sampled at about 570 ksps and the corresponding power spectral density.
The SNR and ENOB have been computed with the Matlab tool for both the mezzanine measurement ranges. A value of 95.85 dB has been achieved for the mA range, slightly lower than the specification of 100 dB, and 94.80 dB for the µA range, outperforming the specification of 90 dB.
The bandwidth of the mezzanine card can be measured by observing the amplitude of the digitized trace as the input sine wave frequency is swept from low to high frequency. The amplitude of the input signal is held constant and close to the full-scale value while the frequency is increased. The frequency sweep analysis is shown in Figure 5.11. The bandwidth at 3 dB attenuation is greater than 10 kHz for all the 4 input channels, in agreement with the mezzanine specification.
The low pin count FMC connector provides 68 user-defined single-ended signals (or 34 user-defined differential pairs) and one serial transceiver pair with a dedicated clock for signaling in the Gbps range or above. The pico-ampere meter mezzanine does not use all these resources, therefore, it is not the ideal test bench for the FMC interface; but, this commercial mezzanine will be used in several applications in the SPES control system and therefore it was important to validate the FMC

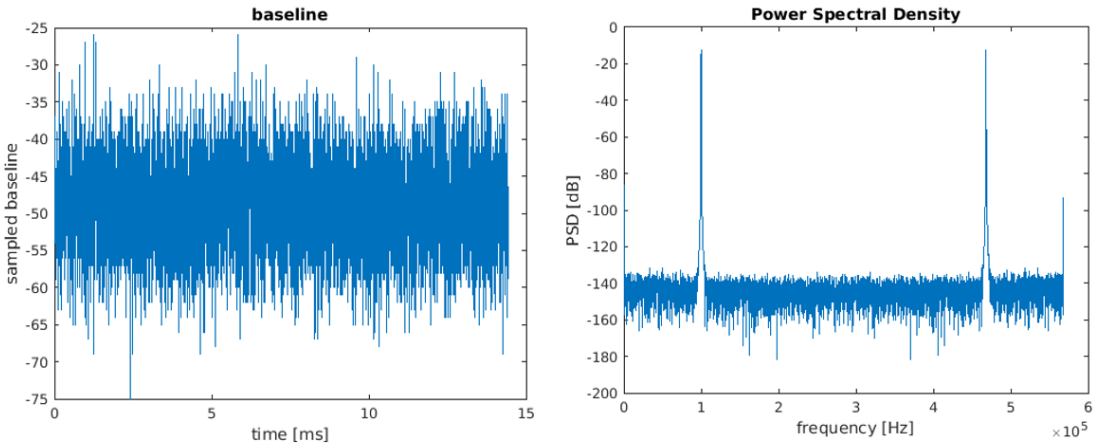

Figure 5.10: FMC mezzanine channel 1 baseline and power spectral density.

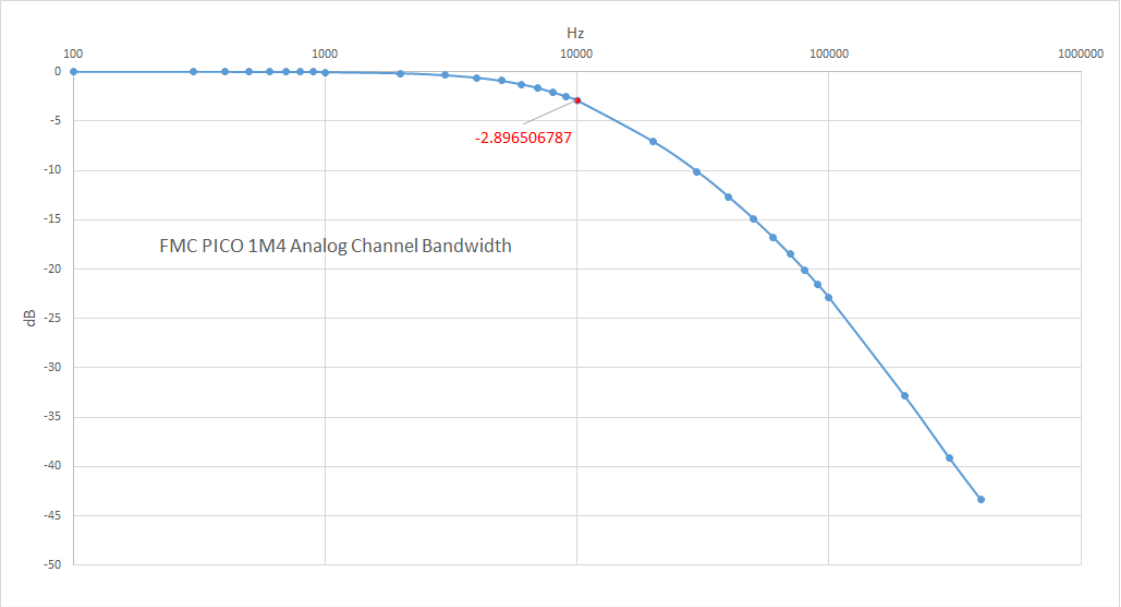

Figure 5.11: FMC mezzanine channel 1 bandwidth (-3 dB).

interface functionality using this mezzanine.

5.7 Signal and Power Integrity Measurements

5.7.1 Clock Jitter Measurements

Assessment of clock signals quality is one of the most important steps when dealing with fast analog to digital converters and high speed digital channels. Jitter on clock signals, defined as the timing variations of a set of clock edges from their ideal position, is typically caused by thermal noise and power ripple, but also by cross-talk and interference coupled from nearby circuits. Clock jitter leads to performance degradation and therefore the characterization of the main system clocks is a very good indicator of the quality of the designed hardware. This paragraph describes the jitter measurements performed on the reference clock of the Spartan-6 transceiver. The GTP tile provides Current Mode Logic (CML) serial drivers and buffers to implement high speed serial link up to 3.125 Gbps. The IOC board exploits one of these transceivers to build a high bandwidth and low latency communication channel via the on-board SFP module. In a typical application, the SFP may host a hot-pluggable optical module transceiver to implement an optical Ethernet physical layer. The GTP tile requires a low jitter reference clock that is provided by the Micrel SY87739L frequency synthesizer. This clock is routed

to the FPGA dedicated reference clock input pin pairs as shown in Figure 5.12.

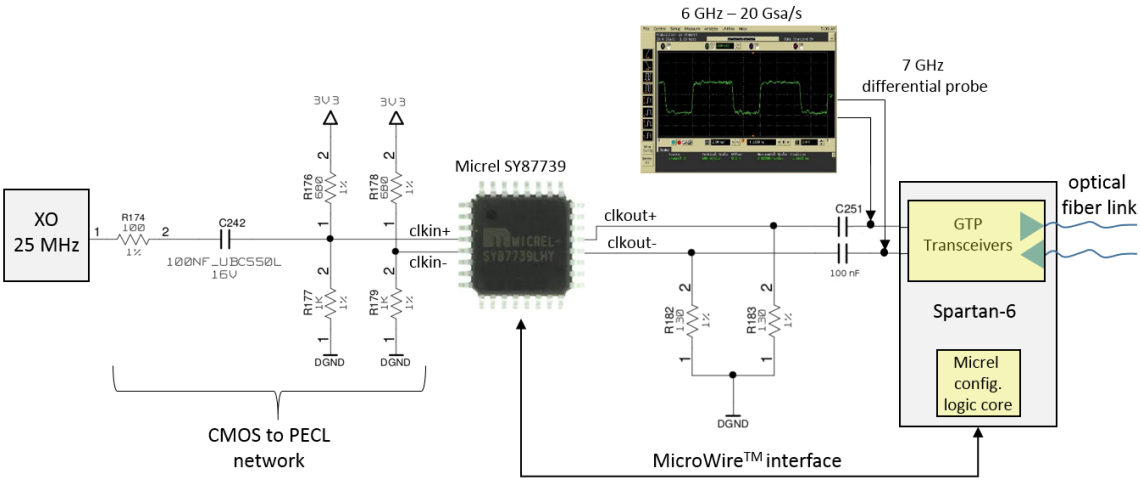

Figure 5.12: Clock jitter measurement setup.

0000	Q_P	Q_{P-1}	P	000	D	N	M
0	11	3	27	0	5	14	15
0000	01011	00011	1010	000	00101	101	111

Table 5.3: Fractional-N synthesizer programming bits to get a 125 MHz output from a 25 MHz input.

The SY87739L generates any differential Pseudo Emitter-Coupled Logic (PECL) output frequency from 10 MHz up to 729 MHz upon being properly configured via the MicroWire™ interface. The optical Ethernet application requires a reference frequency of 125 MHz obtained using the 32-bit programming string reported in Table 5.3. The output frequency can be deduced from (5.6) where f_{ref} is the 25 MHz clock fed to the synthesizer.

$$f_{out} = \frac{\frac{N}{M}[P - \frac{Q_{P-1}}{Q_{P-1}+Q_P}]f_{ref}}{D} \tag{5.6}$$

The generated clock is probed downstream of the two AC coupling capacitors using an Agilent Infiniium 20 Gsps wide bandwidth oscilloscope (6 GHz) and a 7 GHz differential probe. The choice of the oscilloscope and the clock jitter measurement setup deserve special care to achieve accurate measurements. The signal bandwidth can be approximated by the signal rise/fall time as expressed by the general equation (5.7).

$$bandwidth_{3-dB} = \frac{0.338}{min(risetime, falltime)} \tag{5.7}$$

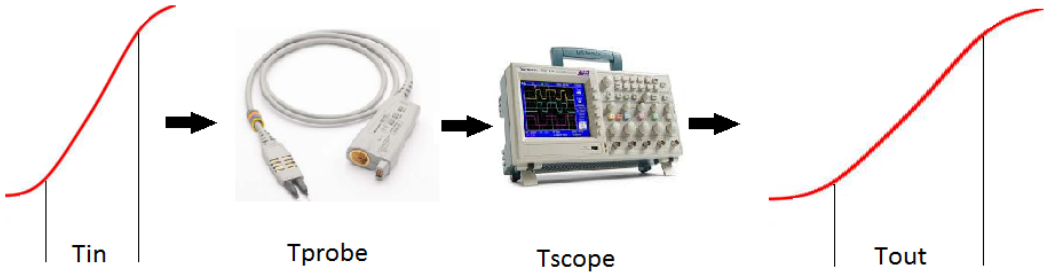

Figure 5.13: Measured signal bandwidth.

The equivalent bandwidth of the probe plus the oscilloscope should be at least 3 times greater than the signal bandwidth in order to avoid poor jitter measurements. As a matter of facts, a low bandwidth probe behaves like a capacitive load and degrades the signal rise time. The rise time of the measured signal depends on both the probe bandwidth and the oscilloscope bandwidth and is well expressed by (5.8) with reference to Figure 5.13.

$$T_{out} = \sqrt{T_{in}^2 + T_{probe}^2 + T_{scope}^2} \tag{5.8}$$

In our case, it turns out that a 6 GHz oscilloscope with a 7 GHz probe is equivalent to an ideal 4.5 GHz instrument and the minimum rise time of the observed signal is about 75 ps. Considering the differential clock under test, the measured rise time is about 200 ps; thus, using (5.8) one can estimate the error introduced by the measurement setup on the rising time that, in our case, is about 8%; it is an acceptable percentage.
Moreover, the quantization noise due to the vertical gain setting as well as the low sampling rate affect the jitter measurement. During the proposed measurements the vertical gain of the scope has been adjusted to make the signal span the full height of the display, and, the maximum sampling rate of 20 Msps was selected.

Period Jitter

The period jitter is defined as the deviation in cycle time of a signal with respect to the ideal period that can be reasonably approximated with its average value. Since the clock period jitter is random and is described with a Gaussian distribution, it can be completely expressed in terms of root mean square (rms) value [47]. The period jitter is particularly useful to compute time margins in high speed digital systems. The measurements have been performed following two methods.

Method 1:

- Change the horizontal time scale of the scope to capture one full clock cycle.
- Configure the scope to measure the clock period (rise edge to rise edge).
- Turn on the histogram of the measurement to record the mean and standard deviation of the period realizations.
- Wait until 10000 (JEDEC standard) or more clock periods have been recorded (indicated as khits in some scopes) before stopping the acquisition. Read the standard deviation on the screen. See Figure 5.14.

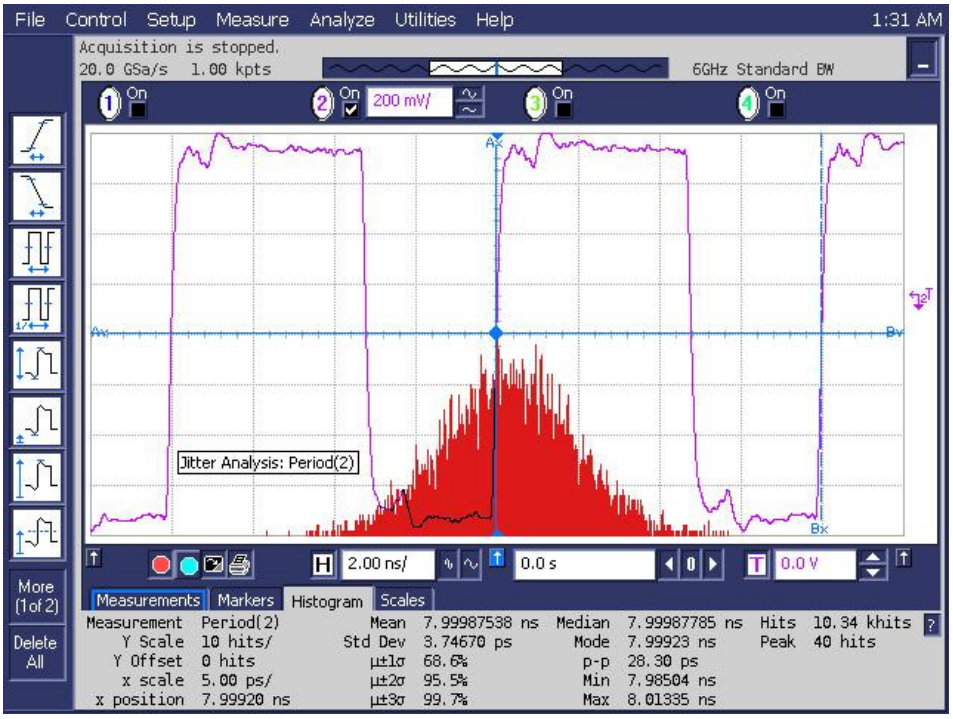

Figure 5.14: Period jitter method 1. Standard deviation: 3.75 ps.

Method 2:

- Change the horizontal time scale of the scope to capture 10000 (JEDEC standard) or more clock cycles.
- Configure the scope to measure the clock period (rise edge to rise edge). The scope will measure the period of all the captured clock cycles that appear in the screen.

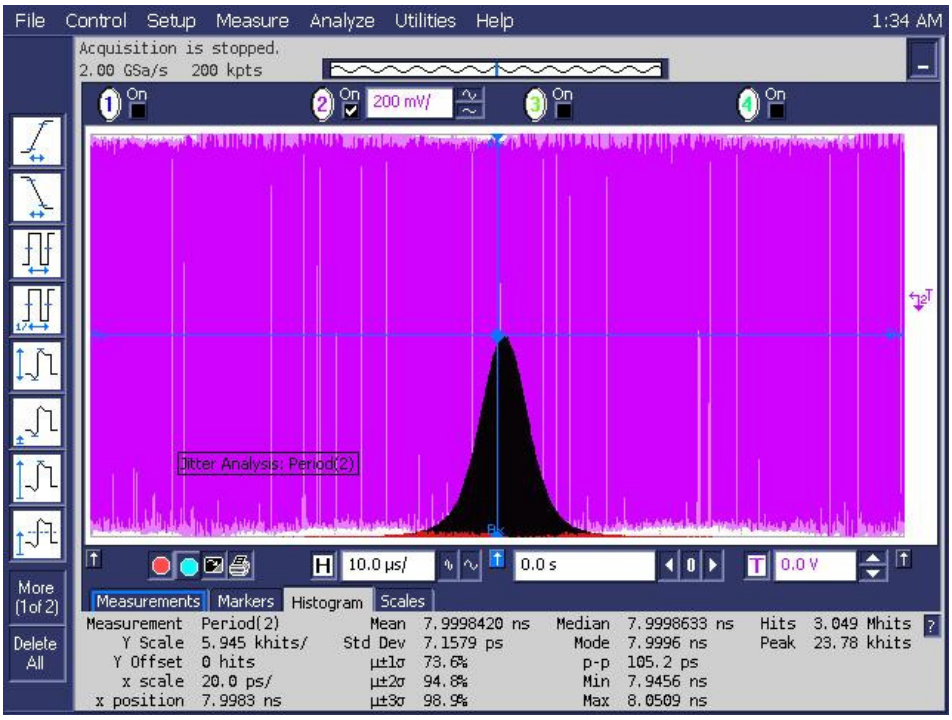

Figure 5.15: Period jitter method 2. Standard deviation: 7.16 ps.

- Turn on the measurement histogram and read the standard deviation on the screen. See Figure 5.15.

Method 2 is more accurate because the statistical analysis is based on a larger data set as indicated by the Hits number of the last picture. Increasing the data set beyond 3 Mhits, the rms jitter does not change significantly, whilst the peak-to-peak value steadily increases because of the nature of the Gaussian distribution. In the digital world, this means that it is impossible to completely exclude timing violations in synchronous circuits and therefore a digital design must somehow accept and manage a minimal percentage of error.

Cycle-to-Cycle Jitter

Another short-term jitter definition is the cycle-to-cycle jitter which is defined as the difference, in period time, between 2 consecutive clock cycles. In this case there is no reference to an ideal clock period. The cycle-to-cycle measurement is shown in Figure 5.16. The number of measurements used to build the histogram must be at least 10000. The cycle-to-cycle jitter measurement has been repeated changing the horizontal scale of the scope, increasing the number of hits, and the result does not change significantly since this short-term jitter is dominated by high-

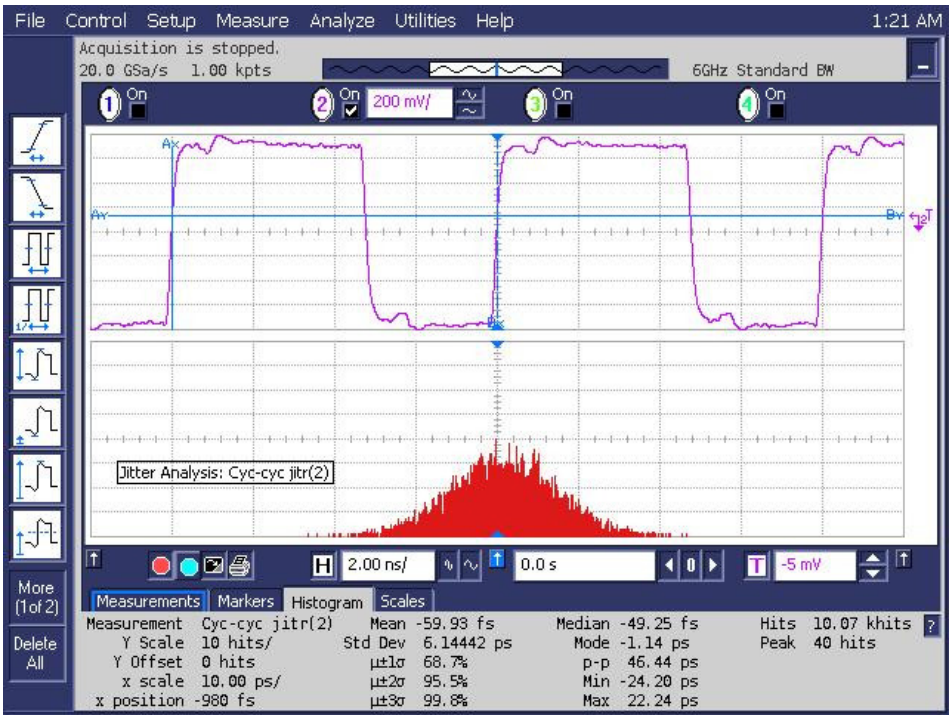

Figure 5.16: Cycle-to-cycle jitter. Standard deviation: 6.14 ps.

frequency noise. The measured rms jitter of 6.14 ps, as expected, exceeds slightly the specification value which is 4.8 ps at 125 MHz. This slight difference, hereby called noise contribution to jitter and indicated with σ_{noise}, can be quantified using equation (5.9):

$$\sigma_{noise} = \sqrt{\sigma_M^2 - \sigma_{die}^2} = \sqrt{(6.14ps)^2 - (4.8ps)^2} = 3.8ps \tag{5.9}$$

where σ_M is the measured jitter, σ_{die} is the expected jitter and σ_{noise} is the total noise contribution to the jitter, induced by the hardware plus the oscilloscope. The oscilloscope contribution to σ_{noise} is relevant and mainly induced by the sampling clock jitter (intrinsic jitter), by the limited bandwidth and by the quantization noise stated before. Therefore, from this analysis, we can conclude that the contribution to the total jitter of the IOC board is below 3.8 ps; this is a good design quality indicator.

Long-Term Jitter

The long-term jitter measures the change of the clock's output transition from its ideal position over several consecutive cycles. Figure 5.17 refers to the N-cycle jitter measurement performed over 1000 (8 µs) clock cycles. It is also referred as

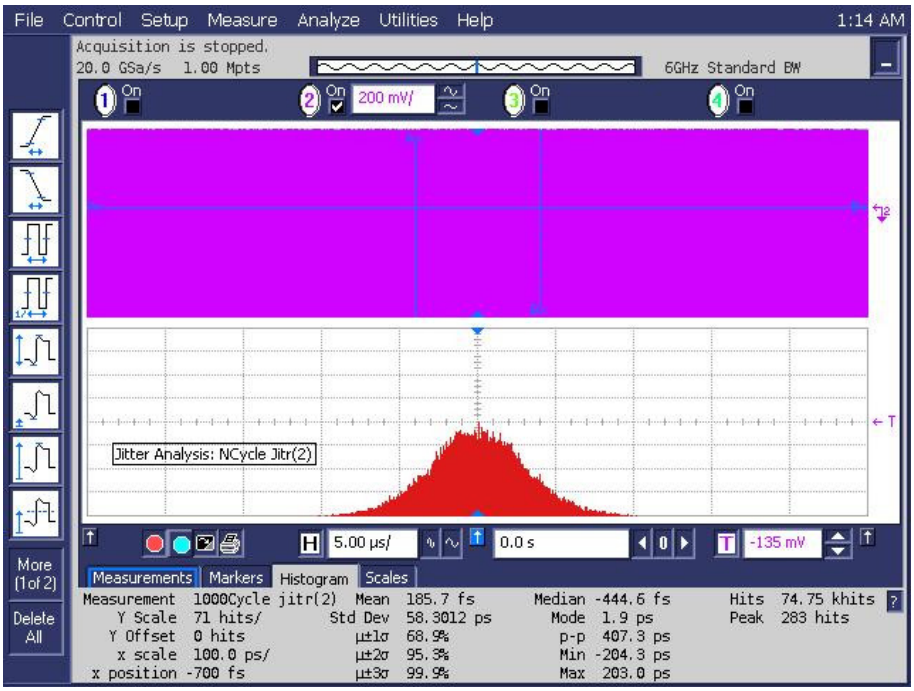

Figure 5.17: 1000-cycle jitter. Standard deviation: 58.30 ps.

accumulated jitter because it represents the integration of the jitter on a continuous stream of clock cycles; therefore, its rms value measured is higher: 58 ps in our case. This long-term jitter is sensitive to low-frequency phase noise of the clock signal. It is particularly significant for application requiring synchronization of timing events like sub-nanosecond timing systems which are widely used in nuclear physics experiments.

5.7.2 Time Domain Reflectometry Measurements

Time-domain reflectometry technology represents a convenient method to check the characteristic impedance of the transmission lines. Measuring the electrical properties of transmission lines is crucial for high speed digital signals and to validate the PCB manufacturing process. The instruments usually used to measure the electrical characteristics of the transmission lines are the Time Domain Reflectometer (TDR) and the Vector Network Analyzer (VNA). The TDR propagates a very fast rising edge waveform through the transmission line and measures the reflected waveform, yielding the instantaneous impedance information. VNA operates in the frequency domain; it generates a sine waveform and measures the reflected amplitude and phase at each frequency value to extract the reflection parameters (scattering matrix).

Higher data rates mean also tight intra-pair and inter-pair skew requirements. The intra-pair skew can be extracted by the differential TDR measure while comparing the TDR response of all the differential pairs of a data bus, it is possible to estimate the inter-pair skew. TDR measures represent an effective way to get a feedback on the most important interconnections in the prototyping phase of the project. Performing TDR measurements on a prototype it is possible to quantify runtime all the reflections caused by parasitic elements, termination mismatches or by the effective loading of ICs or cards hosted by the PCB under test.

USB 3.0 SuperSpeed Differential Pairs Electrical Characterization

Figure 5.18 shows an example of TDR analysis applied to the USB 3.0 super speed differential pairs. The differential characteristic impedance for the USB 3.0 super speed pairs is recommended to be within $90\,\Omega \pm 7\,\Omega$. The impedance trend along the transmission line describes how the waves interact with the PCB [48]. Let's consider the measurement setup shown in the picture, where the signal propagates from a region with a calibrated impedance Z_1 equal to $50\,\Omega$ (represented by the instrument and the probes) to the device under test (our PCB) represented by the impedance Z_2 to be measured. The ratio of the reflected voltage measured by

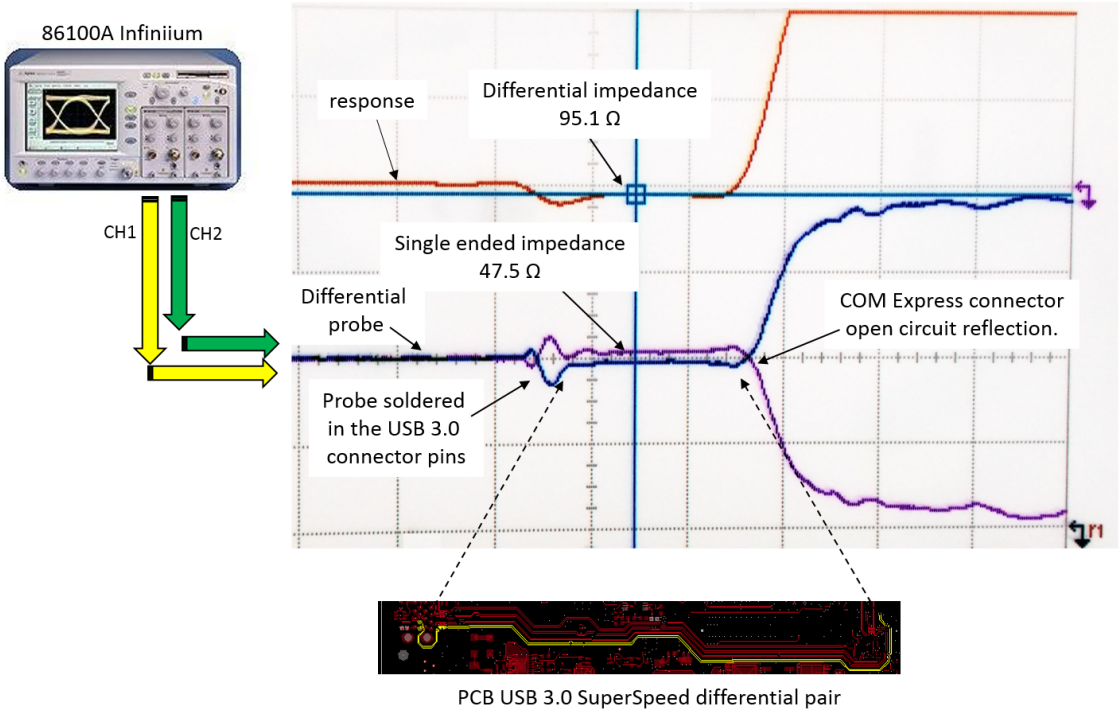

Figure 5.18: TDR technique for USB 3.0 SuperSpeed differential pair characterization.

TDR to the incident voltage, indicated as reflection coefficient ρ, is related to the two impedance by equation (5.10).

$$\rho = \frac{V_{reflected}}{V_{incident}} = \frac{Z_2 - Z_1}{Z_2 + Z_1} \tag{5.10}$$

From (5.10) it is clear that any voltage reflection and its time delay, accurately detected by the TDR, represent an impedance discontinuity in the device under test. In the proposed measurement, the only two discontinuity points observed are represented by the probe soldering on the USB 3.0 connector pins and by the COM Express module plugged into the carrier that does not implement the super speed transceiver. The measure performed shows no impedance discontinuity on the PCB traces and the differential impedance represented by the orange response is within the acceptance limits everywhere from the connector to the COM Express module. The analysis was performed on all the USB 3.0 super speed pairs.

Measuring Parasitic Elements Using TDR

Time domain reflectometry is also an effective way to measure impedance mismatches induced on the transmission lines by parasitic capacitance and inductance of devices and structures of the system containing the transmission lines under test [49]. It is for example the case of the capacitance seen at the input pad of the FPGA die. Figure 5.19 shows the single ended TDR analysis applied to a user defined I/O pin of the FMC connector. The single ended stimulus is applied to a I/O pin of the FMC connector linked to the FPGA. The signal reflection at

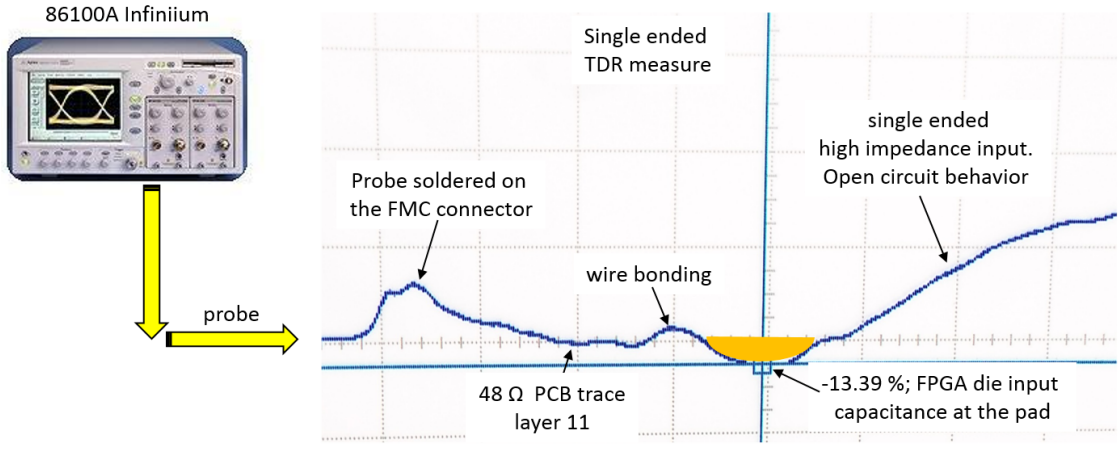

Figure 5.19: TDR single ended measure. FPGA excess die input capacitance.

the FPGA input pad is well visible. The sign of the reflection, in agreement with (5.10), tells us the inductive or capacitive nature of the impedance discontinuity. The signal entering the FPGA is subject to an inductive reflection first, probably due to the internal package wire bonding architecture, and then, the FPGA die input capacitance causes a negative reflection. Using TDR technology, the total capacitance at the input pad of the FPGA can be determined by integrating the normalized reflected waveform, orange area in the picture. This capacitance must be taken into account when dealing with fast rise/fall time signals (e.g. a rise time less than 200 ps). Moreover, the TDR analysis highlights the status of the internal differential termination of a chip. Figure 5.20 shows a LVDS transmission line not terminated on the left, and the corresponding pair with the 100 Ω internal terminator activated on the right.

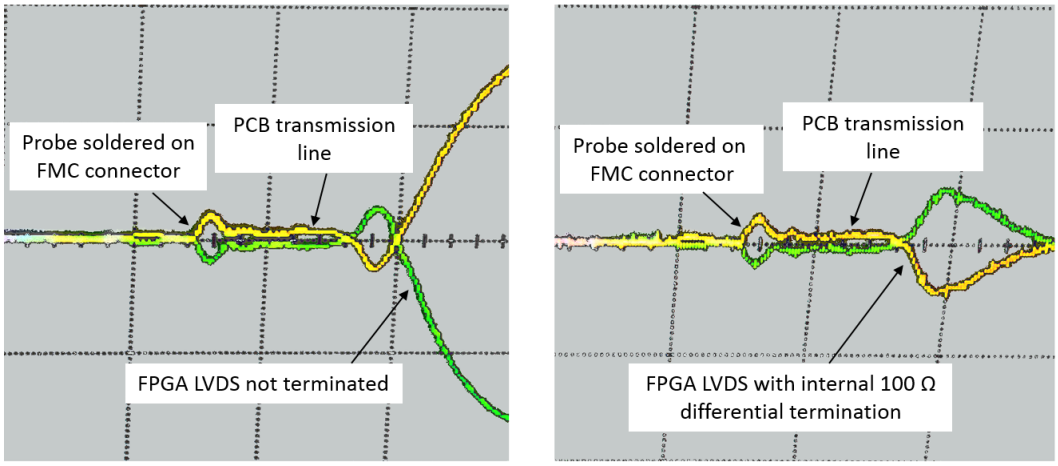

Figure 5.20: FPGA LVDS differential termination.

5.7.3 Power Integrity Measurements

Ripple and noise on DC power rails are a major source of clock jitter and signal integrity issues. Therefore, the analysis of the quality of the on-board DC supplies represents one of the most important phases during the hardware validation. In the IOC board, power integrity is of crucial importance since it may compromise the performance of the analog section of the board or degrade the data rate of the fast digital interfaces. This paragraph reports the analysis of the Spartan-6 core power supply: V_{CCINT}. This approach can be extended to measure the quality of the remaining DC rails. In literature, the DC supply is sometimes represented as the victim of simultaneous transitions on fast digital I/O pins of an IC, or

the aggressor: power supply noise and power supply induced jitter are the main contributors to clock and data jitter. In our analysis, we tend to consider the power supply as the aggressor of the digital system assuming an optimal bypass network composed of different bypass levels that supply transient currents to the system in response to rapid I/O commutation. The challenge is to measure the DC supply without adding noise, without loading the target power rail, with an high bandwidth to detect high speed transients and noise content above 1 GHz, and, it is very important to maximize the offset range of the measurement system. This is achieved using a specialized $1:1$ active power integrity probe [50]. Figure 5.21 shows the voltage ripple measurement at nominal load condition. The voltage regulator provides a steady-state DC value of 1.214 V, within the allowable range of 1.2 V $\pm$ 5% as specified by the FPGA manufacturer. This power rail is generated by a synchronous buck converter whose switching action cause a low ripple on top of the DC value. The output voltage ripple is commonly considered a key figure of merit of the design and it has been analyzed enabling the histogram feature of the oscilloscope. The horizontal time scale has been changed to 100 µs per division to include in the measurements the high frequency voltage spikes generated by the switching voltage regulator and by probe grounding noise pickup.

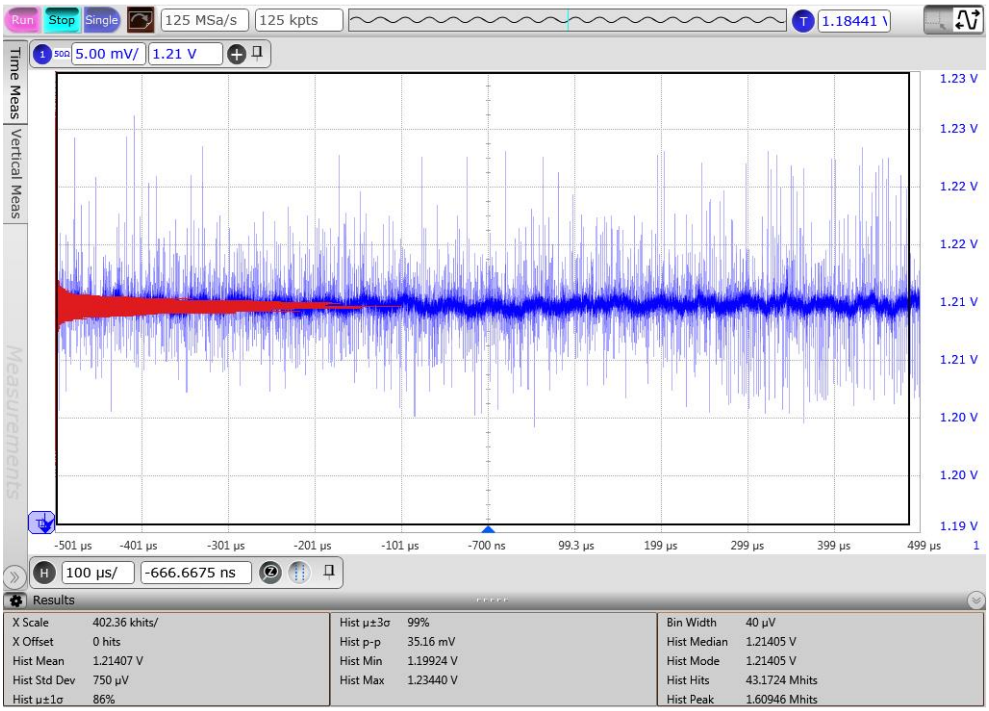

Figure 5.21: Spartan-6 1.2 V core power supply ripple measurement.

The peak-to-peak ripple value measured is 35 mV, about 3% of the target DC value, therefore acceptable. The standard deviation is 750 µV. Looking at the power rail in the frequency domain provides an useful feedback to improve the hardware design. The FFT analysis of the DC trace allows to quickly identify coupled aggressors. Often the unwanted harmonics match with the switching frequency of the voltage regulator itself or neighboring DC/DC converters, but, in our case, the FFT analysis of the V_{CCINT} power rail does not show significant harmonics.
An analysis of the ramp time minimum and maximum requirements, according to the specifications of the FPGA manufacturer, has also been performed. The V_{CCINT} power rail requirement, in terms of ramp time, is between 0.20 to 50 ms. The switching regulator chosen has an internal 1.2 ms voltage-servo soft start, as confirmed by Figure 5.22, that meets this requirement.
The voltage ripple analysis has been performed on all the power supplies of the IOC controller board, and, as a rule of acceptance, all the peak-to-peak values must fall within 5% of their corresponding target DC voltage if not otherwise specified by IC manufacturers.

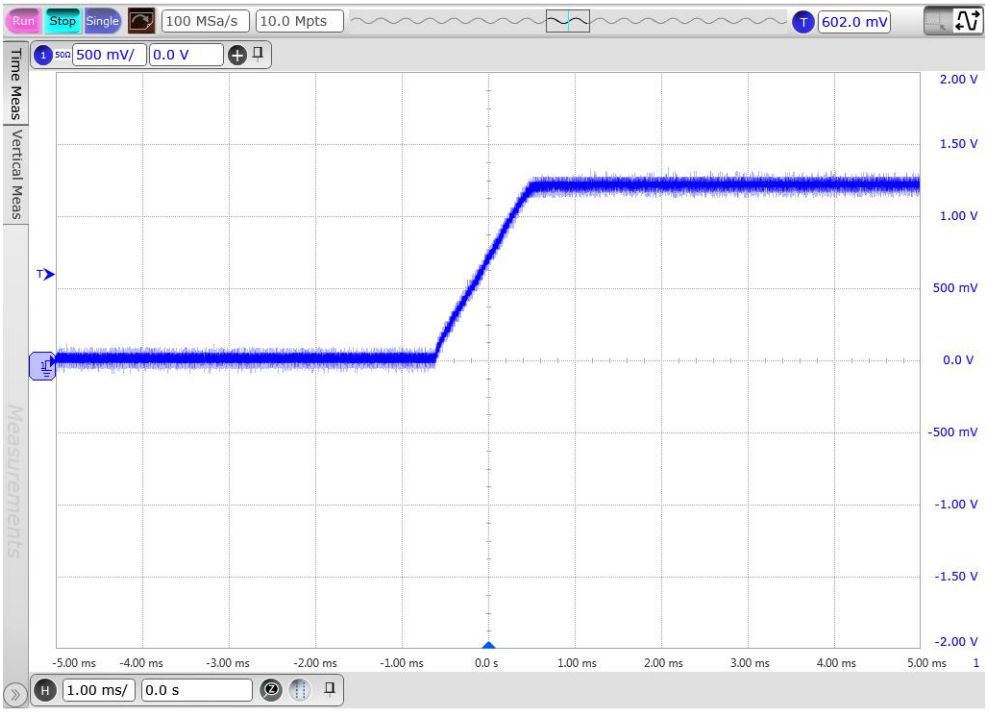

Figure 5.22: Spartan-6 1.2 V core power supply soft start.

5.8 SFP Data Communication Link

The SFP connector connected with the Spartan-6 GTP transceiver opens the possibility to implement high speed data links between two FPGA boards or an optical Gigabit Ethernet connection. Dealing with high speed data links, signal and power integrity are of crucial importance since may partially o completely compromise the reliability of the communication channel. The implementation of an optical Ethernet channel, as described in section 4.4, is the ideal test bench to qualify the serial link. Figure 5.23 shows the 1000BASE-SX test setup and the IOC board without COM Express plugged in. The first test performed to check the connection integrity was the ping utility that exploits the Internet Control Message Protocol (ICMP) echo function. The ping report showed 0% of packets loss thus excluding any hardware malfunctions. A more deep test, involving the IPbus master and slave interconnection, has been performed writing to and reading from an IPBUS slave consisting of a stack of 64 registers. The Python script developed, in sequence, selects randomly one out of the 64 registers, generates a random integer between 0 to 1000, writes the integer into the selected register, reads it back from the hardware and compares the received data against the transmitted one. These steps were repeated thousands of times and no inequalities were detected. The test never failed and always exited with a positive message. Two transceiver status LEDs are also available for an optical inspection as shown in Figure 5.24. The green LED on indicates that all the clock signals are within specification and both the SFP module and GTP receiver are synchronized with incoming data. The red light on indicates a faulty condition of the Ethernet physical layer detected via the loss of synchronization flag available at the SFP connector, or by the GTP transceiver. Usually, the red light on is caused by a lack of power integrity, or a reference clock out of specification, or a faulty physical transmission link.
The absence of Ethernet packets loss is a good indicator of the hardware functionality. However, we performed a more complete analysis inclusive of information about the channel reliability and transmission quality. The first qualitative analysis of the high speed optical channel, hereby presented, has been carried out measuring the eye diagram [51]. An eye diagram is a representation in the time domain of overlapping data bit periods that allows the combined evaluation of the channel noise and intersymbol interference. Figure 5.25 shows the eye diagrams catch at different points along the gigabit Ethernet transmission and reception paths. Eye diagrams represent an intuitive way to estimate the serial link performance; the

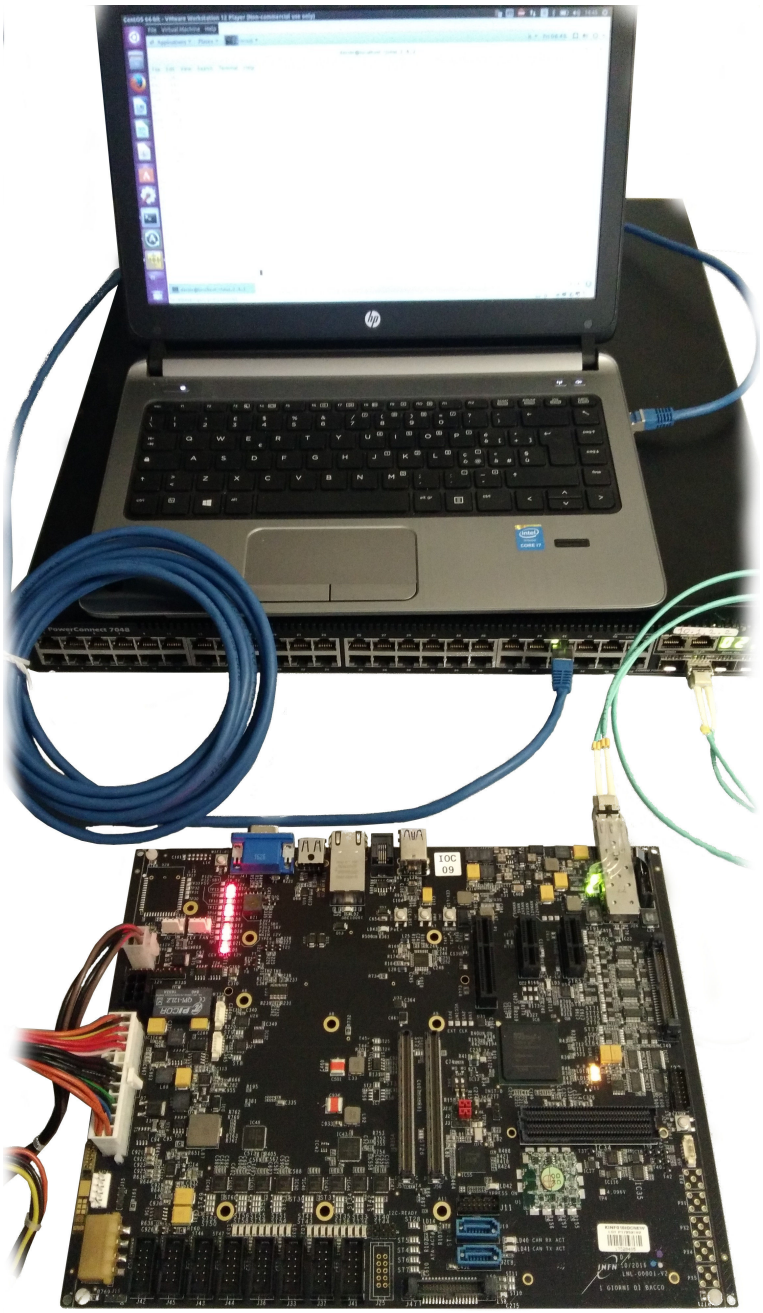

Figure 5.23: IOC board 1000BASE-SX test setup.

eye diagrams measured are all wide open as a confirm of the absence of malformed Ethernet packets observation during the tests. The eye opening in the horizontal time axis, referred to as eye width, gives important information about the signal jitter and intersymbol interference. Whilst, the vertical opening, usually referred to as eye height, is affected by noise, crosstalk and impedance mismatch that causes

```
770 = 770
541 = 541
438 = 438
370 = 370
166 = 166
819 = 819
613 = 613
850 = 850
211 = 211
5 = 5
292 = 292
442 = 442
164 = 164
774 = 774
489 = 489
669 = 669
214 = 214
494 = 494
323 = 323
425 = 425
99 = 99
301 = 301
1000 BASEX TEST PASSED
[davide@localhost uhal_2_4_2]$
```

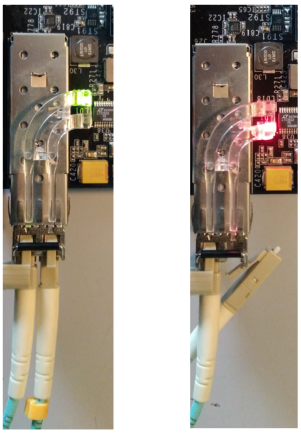

Figure 5.24: IPbus test script and LEDs status indicator.

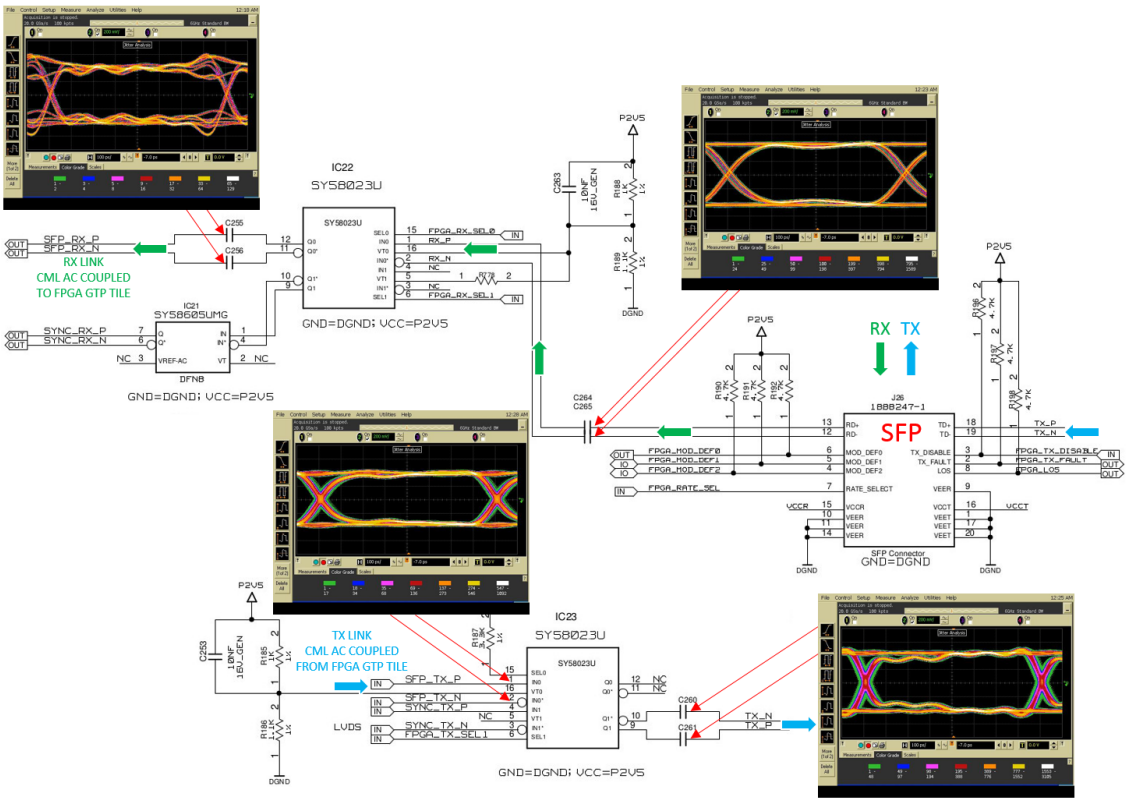

Figure 5.25: 1.25 Gbps eye diagram.

reflection and ringing. The bit error rate which tells us how many bits arrive at their destination in error yields a quantitative information about the serial link performance and is strongly affected by the eye closure that precludes the possibility

to determine whether a received data bit is a 0 or 1.

5.8.1 Bit Error Rate Test at 3.125 Gbps

The SFP interface provides support for several standard protocols running at data rates higher than the Gigabit. This interface can be exploited in a point to point connection to implement a serial data streaming between two intelligent nodes in the control network of the SPES facility. In this case, it is interesting to evaluate the maximum data rate achievable with the IOC board. The maximum performanace of the GTP transceiver is 3.125 Gbps. Xilinx® makes available a customizable Integrated Bit Error Ratio Tester (IBERT) core design for the Spartan-6 transceiver that suits the evaluation of the IOC serial interface [52] [53]. The IBERT core can be generated for the desired line rate. Figure 5.26 shows the test setup. All main GTP parameters are monitored and accessible via the ChipScope™ console shown in Figure 5.27.
The off-the-shelf logic core includes a pattern generator and a pattern receiver. The pattern generator presents a variety of Pseudorandom Binary Sequence (PRBS)

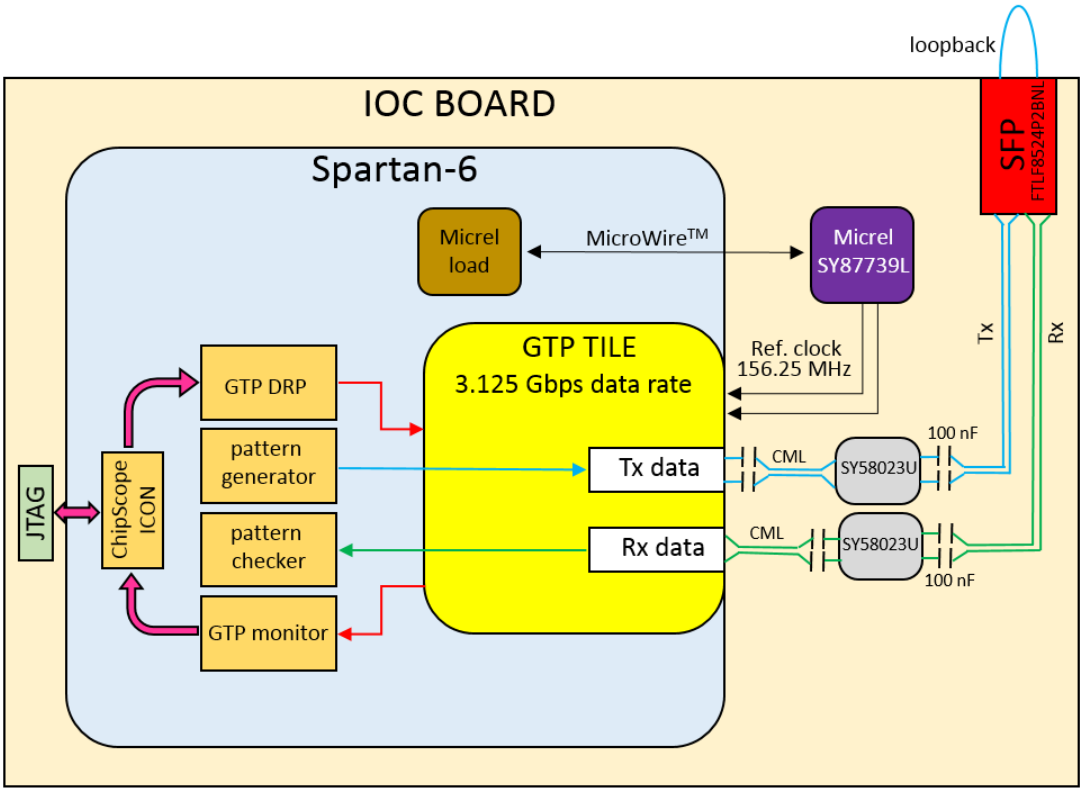

Figure 5.26: Integrated bit error ratio test setup.

patterns, runtime user selectable, that can be continuously sent over the channel. The fiber optic loopback redirects the PRBS patterns to the error detector that checks incoming patterns against the generated ones. The IBERT core has been generated to provide a line rate of 3.125 Gbps to stress the channel at the maximum sustainable data rate. As visible from the console, after a run of about 16 minutes with a PRBS 31, the receiver bit error count was still zero. This brilliant result is again confirmed by the eye diagrams shown in Figure 5.28. The user may

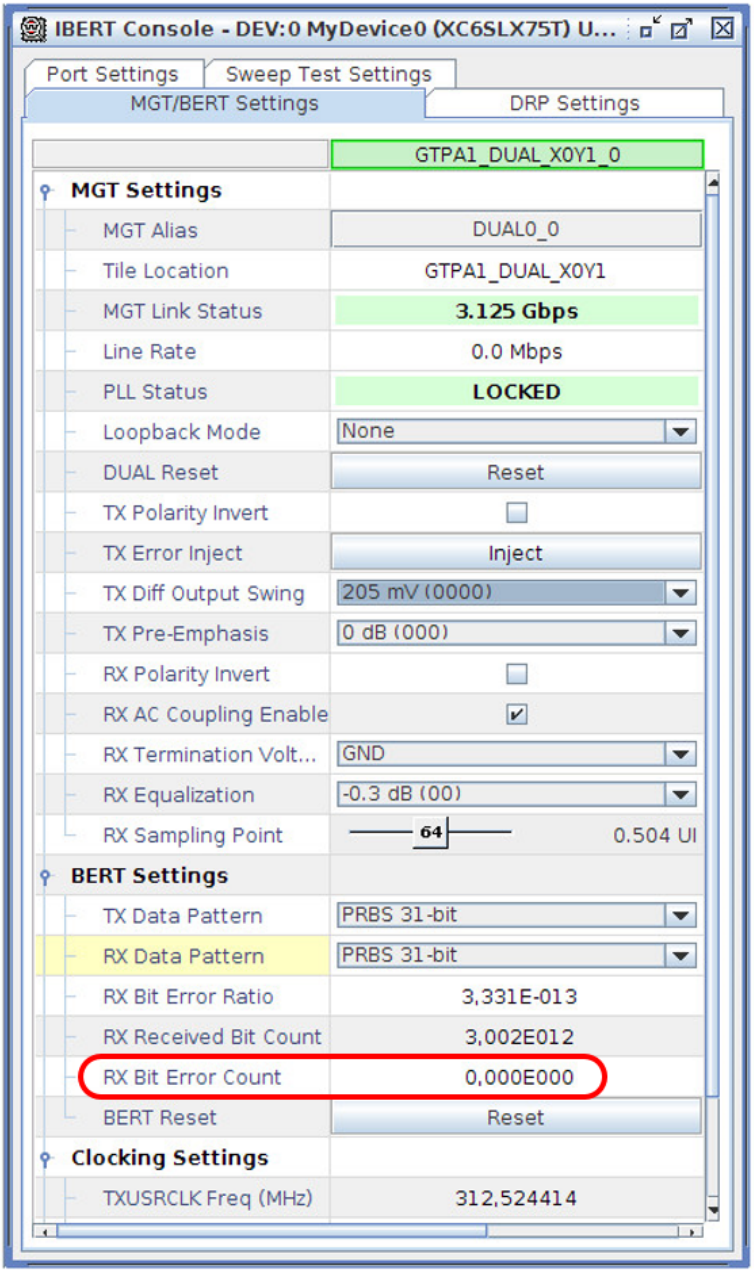

Figure 5.27: IBERT console; bit error rate test result.

perform the sweep test via the IBERT console, consisting in a 2D full scan of the horizontal and vertical offset sampling points in order to generate the eye diagram of incoming data. We measured the eye diagrams with the external oscilloscope, at the same positions as before along the transmit and receive paths. Comparing the eye diagrams in figure 5.28 with those in figure 5.25 we can notice that there is no significant degradation of signal quality by increasing the bit rate. The explanation for this is that the rise and fall times of the 3.125 Gbps link are almost equal to those at 1.25 Gbps. The signal jitter does not scale with the data rate and it starts to be well visible in the eye diagram. The component with the worse performance is the FPGA's GTP transceiver as expected. This is highlighted by the higher amount of eye closure of the top left eye diagram, likely due to parasitic effects introduced by the FPGA I/O pad as shown before. Whilst the cross point switches and the SFP module present perfect terminations and an almost no reflection phenomena.

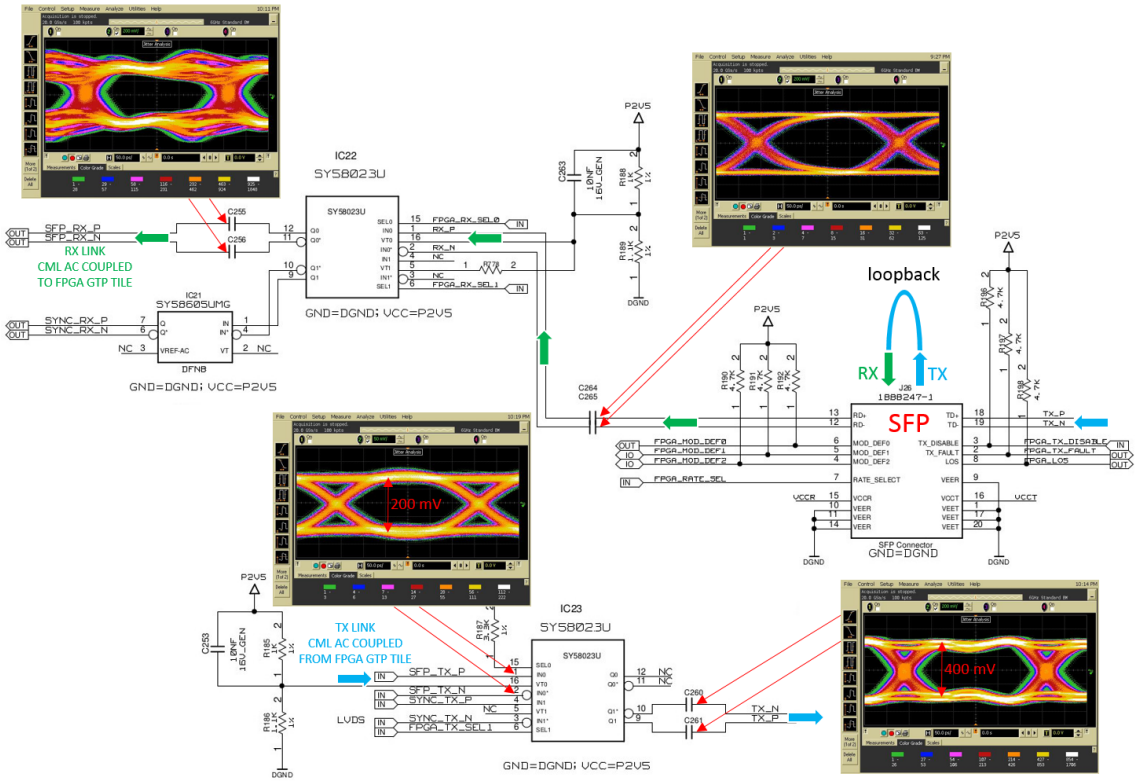

Figure 5.28: 3.125 Gbps data rate eye diagram.

5.9 IOC Project Status

The IOC project roadmap and the complete status overview is given in Table 5.4. The debug of all the standard pc peripherals has been straightforward thanks to

Date	Event
October 2014	IOC project starts.
January 2015	IOC block diagram and main features specification written.
May 2015	IOC v0 schematics design released. Board stack-up defined. Order sent to a company for the layout design.
August 2015	IOC v0 layout completed.
September 2015	IOC v0 review and signal integrity simulations finalized. Production files released. Two prototypes ordered.
November 2015	Two IOC v0 boards fully assembled received at LNL. Starting the debug.
October 2016	All standard desktop pc features ok. FPGA configurable via JTAG connector and via the COM Express. All power rails are within specifications. PoE+ ok, ADCs 16-bit ok, DACs 16-bit ok, ADCs 24-bit ok, DACs 20-bit ok, 8 RS232/RS422 interfaces ok, Cypress fx3 USB to parallel port bridge ok. Starting to work on the IOC v1 revision.
January 2017	IOC v1 production files released. Minor changes on passive components, fixed a CMOS oscillator footprint, changed the serial port headers pinout, SPI connector removed, added 24 digital GPIO on a 34 positions header for flat cable, removed the audio jack (space reasons), CAN bus DSUB 9 connector replaced with a 10 contacts header for flat cable, fixed the PCIe x1 connector footprint, removed Cypress watch dog, optimized the layout and silkscreen layers.
April 2017	IOC v1 prototype received at LNL. Starting the v1 board validation.
October 2017	Fixed an issue on the FMC connector. IOC v2 production files released. Ordered 19 IOC v2 boards.
December 2017	19 IOC v2 boards received at LNL. Starting the IOC v2 boards qualification and integration in the SPES beam diagnostics data readout system.
June 2018	The IOC board is the core of the SPES diagnostics data readout system and the results of the tests performed under real beam conditions prove that the IOC board outperforms the legacy VME technology and is a reliable solution to standardize the control of several appliances in the SPES control system.

Table 5.4: IOC project roadmap and state-of-art.

the COM approach; the IOC board behaves like the motherboard of a desktop pc and the end user can plug standard peripherals (e.g. mouse, keyboard, Ethernet, monitor) and boot a standard Linux distribution from a solid state drive disk. The three PCIe slots were validated using a 1-port Gigabit Ethernet network card in PCIe form factor. During the speed test the Gbps band was almost saturated and this in principle excludes any hardware problem on the PCIe links. The debug of the most important application-specific interfaces, linked to the FPGA, has been carried out alongside the firmware development. The first VHDL code released was a collection of modules developed with the purpose of debugging the IOC revision v0 board shown in Figure 5.29.

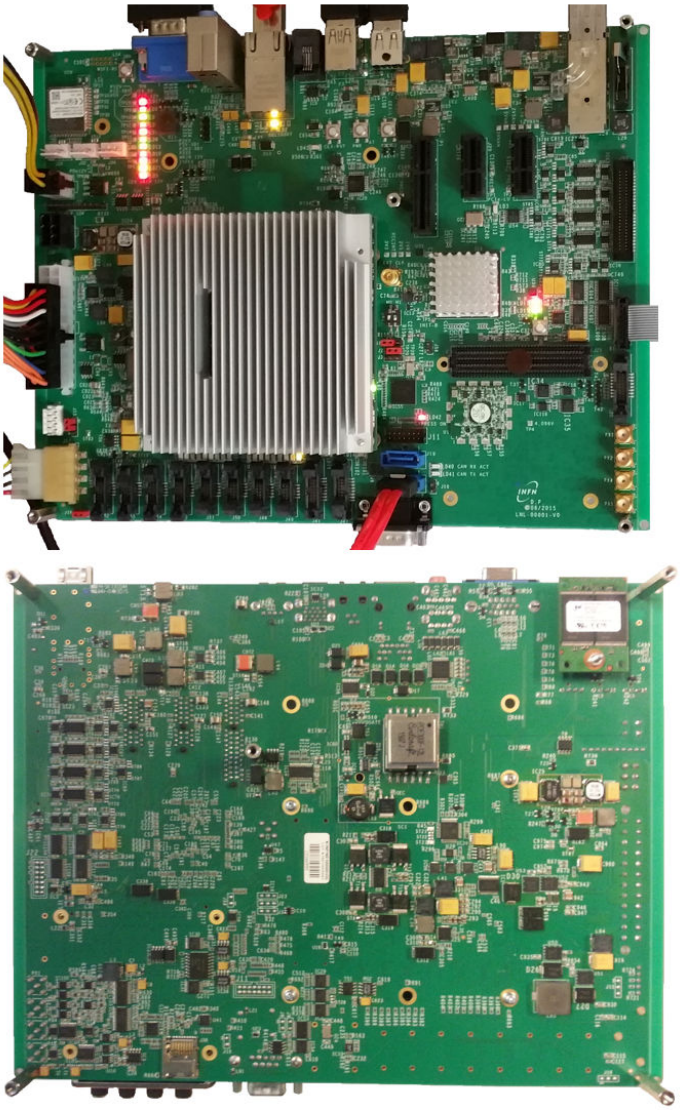

Figure 5.29: IOC v0 prototype top and bottom views.

During this phase, the goal was to debug only the essential features in order to converge to a first production release eligible to satisfy those applications requiring the IOC board in a very short time. The outcome of the v0 prototypes was positive since the main IOC board features were working fine. The subsequent release v1 fixed several minor issues found during the debug phase and introduced a pinout optimization where necessary. The v2 revision was necessary to fix a mechanical issue on the FMC connector. The IOC board revision v2 is shown in Figure 5.30.

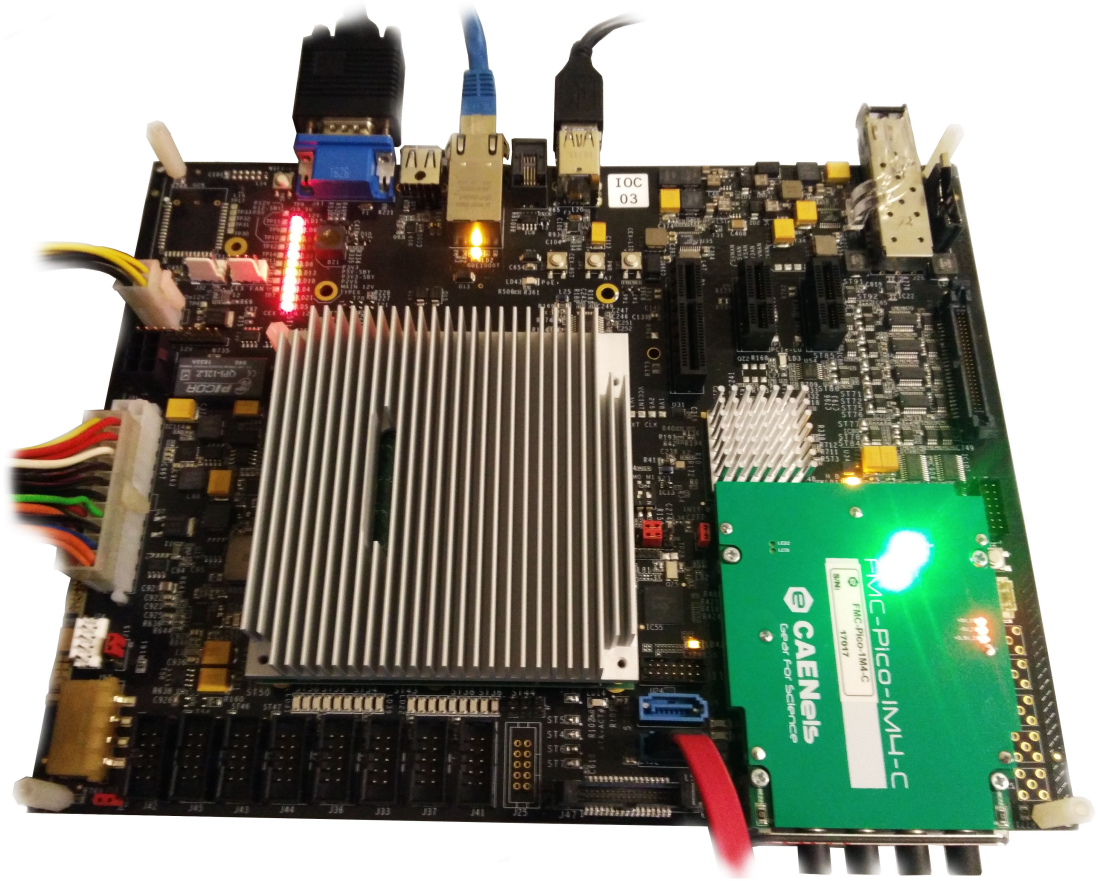

Figure 5.30: IOC v2 production board.

Nineteen boards have been produced and represent the first IOC board production lot that will be devoted to the new beam diagnostic system, as described in the next chapter. The block diagram in Figure 5.31 gives a complete overview of the features tested and working on the IOC v2 production release. In green are the supported features, while in gray are those features left as ongoing work or as a future development, as listed below.

- There are a few features that have never been tested so far. This is the case

of the CAN bus, one wire, micro-SD card slot, the two resistance to digital converters, the total application time-of-use recorder, the EEPROM memory attached to the COM Express, the Wifi module, the fan speed controller linked to the FPGA.

- The two 24-bit resolution analog input channels and the two 20-bit resolution analog output channels were working fine on the v0 prototype but have not been assembled on the v2 boards (cost saving reasons). It would be interesting to assemble a full IOC v2 board and characterize these high-resolution analog input channels in terms of ENOB. The result could then be compared with the 16-bit input channels performance.

- The development of a VHDL core that acts as an I^2C master and periodically reads the on-board digital temperature sensors and modulates the FAN speed in agreement. This is important to increase the hardware reliability in view of the installation of the IOC board in the field.

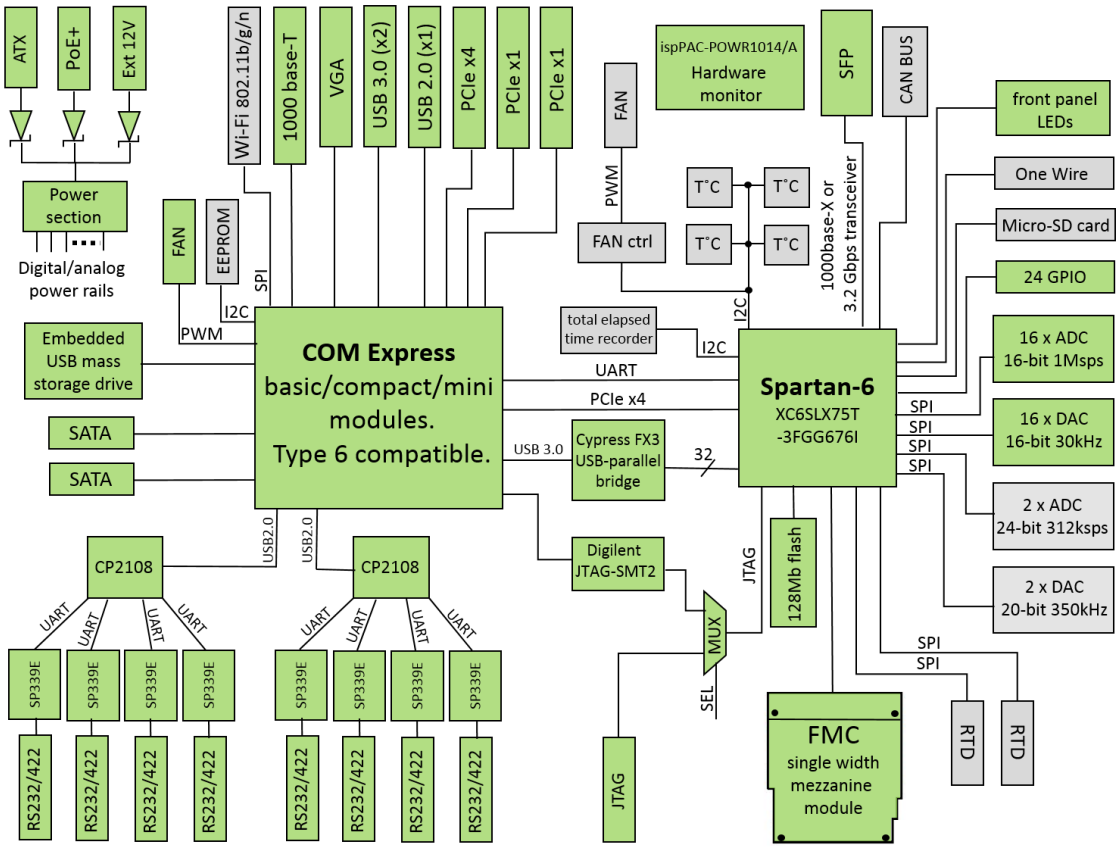

Figure 5.31: IOC v2 production board validation progress.

- The low pin count FMC LVDS pairs and the FMC serial transceiver link performances should be characterized with a loopback mezzanine card together with a PRBS pattern generator and checker.

- The on-board ispPAC hardware monitor has been used to switch on/off the analog power supply, but, it embeds a programmable PLD and timers, that together with the analog and digital inputs may be exploited to implement soft or hard protection shutdown, rise alerts to the FPGA, monitor all the power supply rails, and force the FPGA reconfiguration from the attached EEPROM whenever the FPGA core power supply drops temporarily below a predefined threshold.

- The installation of the IOC board in the SPES control system does not require EMC and EMI emission and immunity certifications, nevertheless, performing these tests in a specialized laboratory would be a good feedback on the project robustness. The outcome of these tests may highlight important improvements for future IOC board releases.

- The execution of accelerated aging tests, in particular, thermal cycling, it would be important to understand the IOC board reliability.

Chapter 6

Beam Diagnostic Data Readout System

The ongoing upgrade campaign required to accommodate the SPES facility demands a new beam diagnostic system that outperforms the legacy VME based current and beam position monitor. The diagnostic is an essential component of the beam transport system without which the operator would blindly attempt to drive the beam from the source to the experimental halls. It provides a feedback while setting all the relevant parameters of the RF subsystem and of the magnetic and electrostatic lenses. The main parameters to be measured and that characterize the beam are the intensity, the transverse centroid, and the transverse beam profile. Moreover, the beam diagnostic system should help the understanding of the beam behavior under normal and abnormal conditions. It opens the possibility to monitor and correct the beam losses along the accelerator pipe increasing the transport efficiency. The diagnostic data readout represents a typical use case of the IOC board designed. The IOC, at the core of the new diagnostic readout system, collects, processes and bridges the diagnostic information to the SPES control network for remote visualization and control. The test performed on the field demonstrate that the new electronics leads to a performance improvement and extends the current measurement resolution to a few pA, addressing the SPES requirements.

6.1 Beam Diagnostic Instruments

SPES foresees the acceleration of about 42 different single species (among which Cs, Xe, Rb, Ba, Br, Sn) neutron-rich ion beams generated in the target-ion source chamber where the primary proton beam is converted into a radioactive ion beam through a multi-foil UCx target. Regardless of the accelerated species, the beam diagnostic system should help the understanding of the beam behavior under normal and abnormal conditions, and it must cope with a wide range of RIB energies and intensities moving along the beam pipe from the target-ion source chamber to the experimental halls, passing through the superconductive linear post-accelerator ALPI. The diagnostic boxes and their readout electronics installed in ALPI must be upgraded to cope with the low-intensity RIBs foreseen by SPES and, at the same time, ALPI must be backward compatible with the Tandem [54] and Piave [55] sources that inject in the superconducting linear accelerator single-charged stable heavy ion beams with intensities up to mA. The extensive experience in the design of diagnostic instruments for stable ion beams accrued at LNL, suggests the

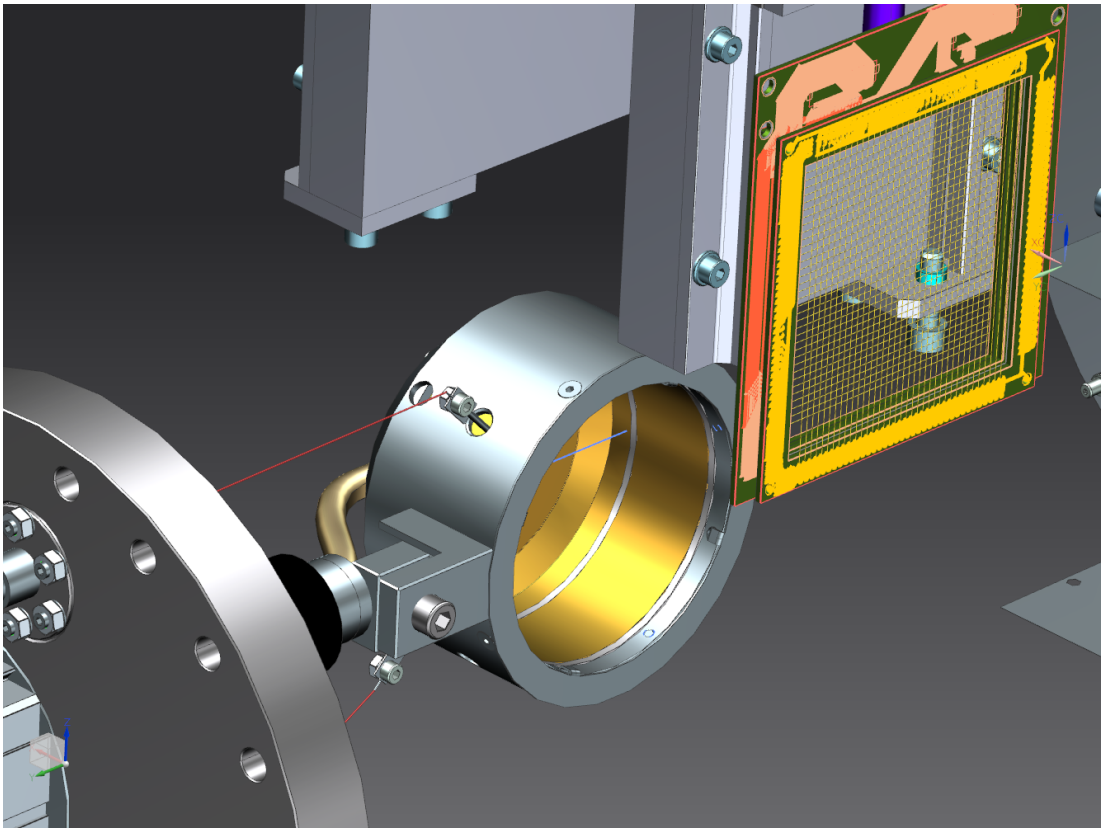

Figure 6.1: Diagnostic box: Faraday cup and beam profile monitor. Picture courtesy of Ing. Claudio Fanin.

use of diagnostic instruments and readout electronics with a wide range of capability, and/or realized with customizable parameters to bring all the beam transport lines at nominal conditions regardless of the accelerated species and intensity. In particular, the goal of the upgrade campaign concerning the diagnostic electronics in ALPI, hereby treated, is to extend the current sensitivity to tens of pA for transporting the low-intensity RIB.
Figure 6.1 gives an internal view of the diagnostic boxes installed in ALPI, entirely designed at LNL. The diagnostic box is a vacuum chamber that accommodates a Faraday cup and a Beam Profile Monitor (BPM), both highlighted in yellow. The BPM and the Faraday cup have independent precise motion control positioning systems based on stepper motors or pneumatic actuators.

6.1.1 Beam Profile Monitor

The beam profile and the beam position in the vacuum chamber are influenced by the transversal and longitudinal beam dynamics, and, therefore, beam profile measurements are important to check the transverse beam shape and the transverse matching between different part of the accelerator. Depending on the beam size, beam energy and currents, a variety of profile monitors exist. Legacy diagnostic systems installed at LNL confirm the possibility to intercept the beam with a grid of wires to yield the 2D projection of beams with sizes in the cm or sub-cm range. The grid is composed by forty horizontal and forty vertical thin wires (50 µm diameter), made of gold plated tungsten. The beam crossing the mesh produces a percentage of charge on the wires, proportional to the beam intensity in that section of the vacuum chamber. The grid is produced in 250 µm, 500 µm, 1 mm and 2 mm spacing to intercept beams of sizes from 1 cm to 8 cm. Alongside profile grids currently installed in ALPI [56], new BPMs are under study and design to address the mechanical and vacuum requirements introduced by the new SPES diagnostic boxes.
The new BPM, shown in Figure 6.2, consists of a carrier board where it is possible to plug a top and bottom boards that hold the mesh of wires. Each wire has one extremity grounded through an high-value resistor that discharges the wire when left floating, preserving the readout electronics from dangerous electrical discharges during the readout. These PCB have been manufactured in Rogers RO3003™ that suits high-vacuum applications [57]. The vacuum achieved with the new PCBs inside the diagnostic box is in the order of 10^{-8} mbar in agreement with the SPES

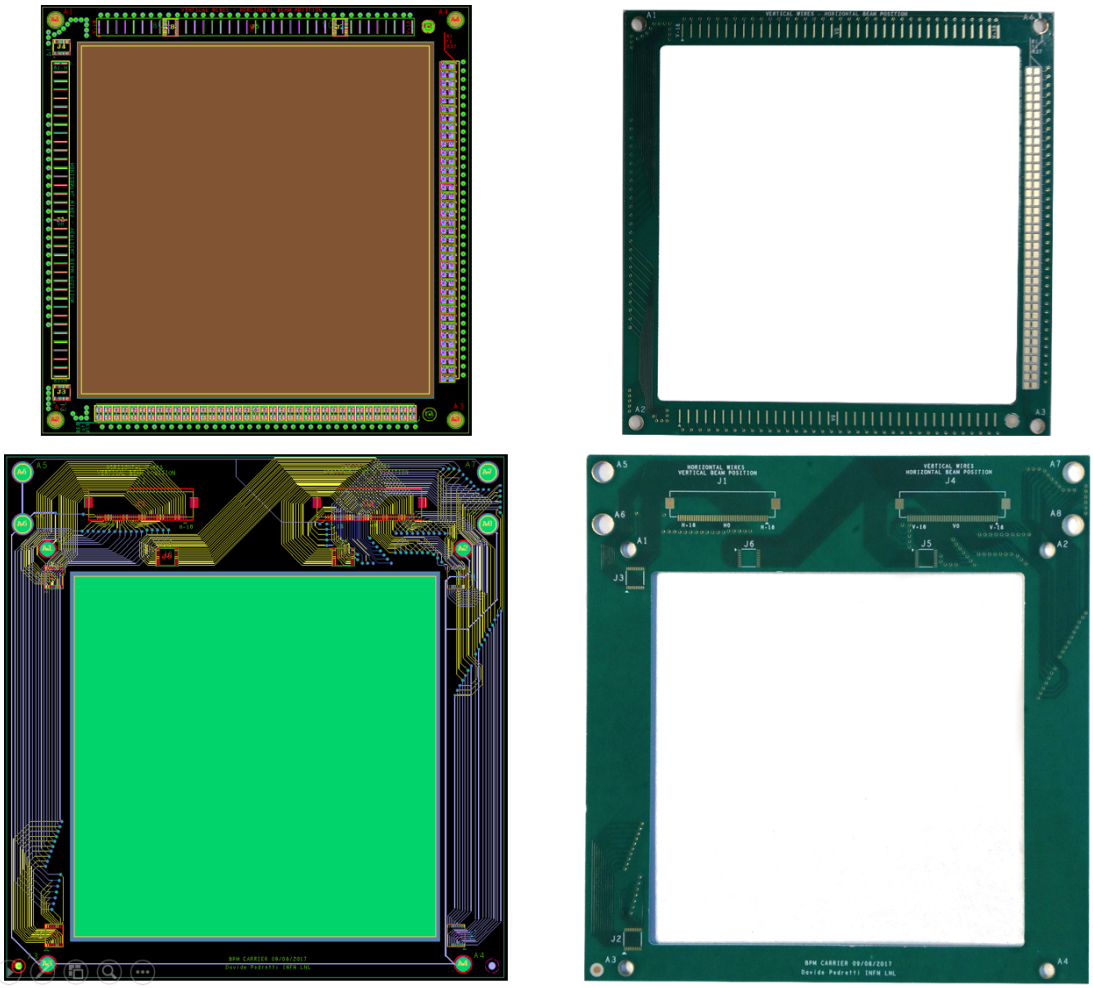

Figure 6.2: The new beam profile monitor layout and prototype.

requirements, and an order of magnitude better compared with the previous FR4 PCBs.

The significant information is the average charge trapped in the wire, and, at this purpose, profile grids are able to characterize DC beams as well as bunched beams like the one observed in ALPI. The beam sensitivity is dependent on the readout electronics and on the wire diameter. So far using legacy VME technology, the minimum beam intensity measured, with stable ion beams, was about hundreds of pA per wire. The expectation is to extend the sensitivity to few pA per wire with the new readout system that will be based on the IOC board.

The main concern using profile grids is represented by the long exposure to high-intensity and high-energy beams that may damage the grid. For this purpose, it is interesting to understand the working conditions of the wire intercepting the beam and, in particular, to estimate the maximum temperature reached by the wire.

Wire Temperature Profile

The methods of heat transfer considered in the high vacuum chamber are the radiation and conduction. The radiation heat transfer is described by the *black body* Stefan-Boltzmann law:

$$E_n = \sigma_b T^4 \tag{6.1}$$

where E_n is the energy radiated per unit surface area per unit time (W/m^2) by a black body, σ_b is the Stefan-Boltzmann constant and T is the absolute temperature. A body that does not absorb all incident radiation also emits less energy compared with a black body. The correction factor to be introduced is the emissivity ϵ, minor than 1 for any *gray body*:

$$E_n = \epsilon\sigma_b T^4 \tag{6.2}$$

Starting from (6.2), the net radiation heat transfer of a *gray body* is given by:

$$E_n = \epsilon\sigma_b(T^4 - T_0^4) \tag{6.3}$$

where T_0 is the ambient temperature inside the vacuum chamber that, with a good approximation, can be fixed to 300 K.
The thermal conduction transports the heat inside the wire and according to Fourier's law the heat flow $[\frac{W}{m^2}]$ is found by the relation:

$$q = -k\frac{dT}{dx} \tag{6.4}$$

where k is the thermal conductivity and $\frac{dT}{dx}$ is the temperature gradient in the direction of the flow.
The linear approximation of (6.4) is the Newton's law of cooling that gives the temperature drop in relation to the heat transfer rate Q per unit of time and the thermal resistance for an unitary cross section and unit of length:

$$\Delta T = QR_{th} \tag{6.5}$$

In analogy with the electric capacitance, the thermal capacitance C_{th} determine the temperature variation per unit time per unit volume caused by a net heat transfer rate P:

$$\Delta T = \frac{P}{C_{th}} \tag{6.6}$$

In our case study, the heat source is represented by the beam whose power density distribution can be well approximated using a standard Gaussian profile with the beam emittance expressed by the variance σ_{beam}:

$$p(x,y) = \frac{P}{2\pi\sigma_{beam}} e^{-\frac{x^2+y^2}{2\sigma^2_{beam}}} \tag{6.7}$$

where P is the total beam power and p(x,y) the power density distribution.
The temperature distribution along the wire and the heat transfer problem can be solved using the finite element method as shown in Figure 6.3. We consider a golden tungsten wire with a diameter of 50 µm and a length of about 10 cm. The wire is soldered to the PCB at the two extremities were the temperature is kept at 300 K (the initial temperature value of the wire). Now let's divide the wire in a collection of N finite elements of length 1 mm corresponding to an array of temperature values T_x. For each element, equations (6.3), (6.5), (6.6) hold and describe the system. The power transferred by the beam to the wire is distributed on the N elements following a standard Gaussian distribution as assumed before. At time t_0 we assume that the beam power is switched on and the temperature transient starts. The transient is divided into time steps lasting dt seconds, during which the temperature of each element is assumed to be constant and the finite element method solves the system of equations governing each element. At any iteration (see python code in Appendix A) the new temperature of each element is calculated using (6.6). The net heat power transferred to the finite element in a dt time is the balance of the power transferred to the element by the beam, the conducted power exchanged with neighboring cells, positive if incoming, and

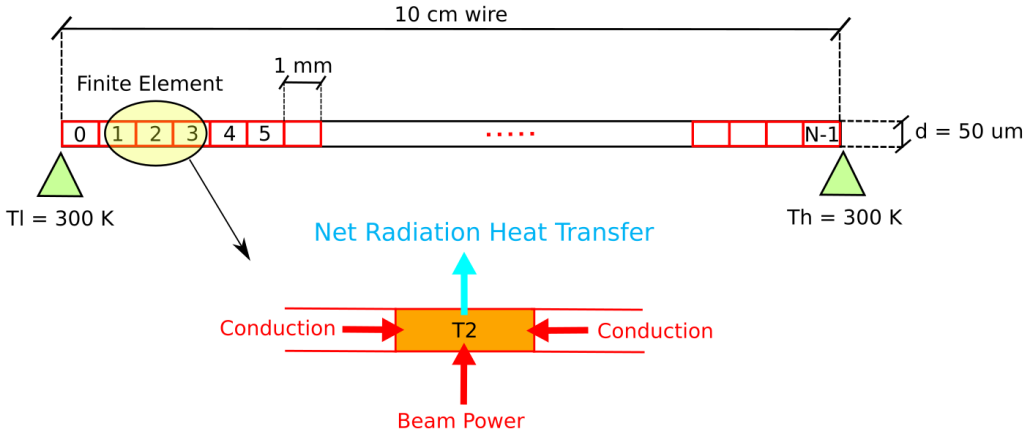

Figure 6.3: Finite element heat fluxes.

the power loss by radiation. When a steady-state condition is reached the net heat transfer to the finite element is zero and the temperature is invariant among iterations. The emissivity coefficient was set to 0.1 (gold emissivity, worst case). The first simulation, shown in Figure 6.4, was performed assuming a wire current of 10 nA and a beam energy of 20 MeV with a charge state 1^+. The equivalent beam power on the wire is 0.2 W. The temperature in the middle part of the wire reaches 1678 K and all the power is dissipated by the radiation flow as predictable given the low wire section that increases the thermal resistance. These high temperature justify wires made of tungsten. The conduction heat transfer shapes the temperature profile and, at the wire extremities, the conducting contribution is almost null; the primary heat dissipation mechanism is the radiation, as expected. The thermal resistivity of tungsten increases roughly linearly in the temperature range of interest; this effect has been considered in the simulation as visible in the conduction heat transfer subplot.

The maximum wire temperature is influenced by the beam emittance. A beam more focused, as shown in Figure 6.5, increases the maximum wire temperature to 2252 K. One can repeat the simulation with a higher beam current intercepted

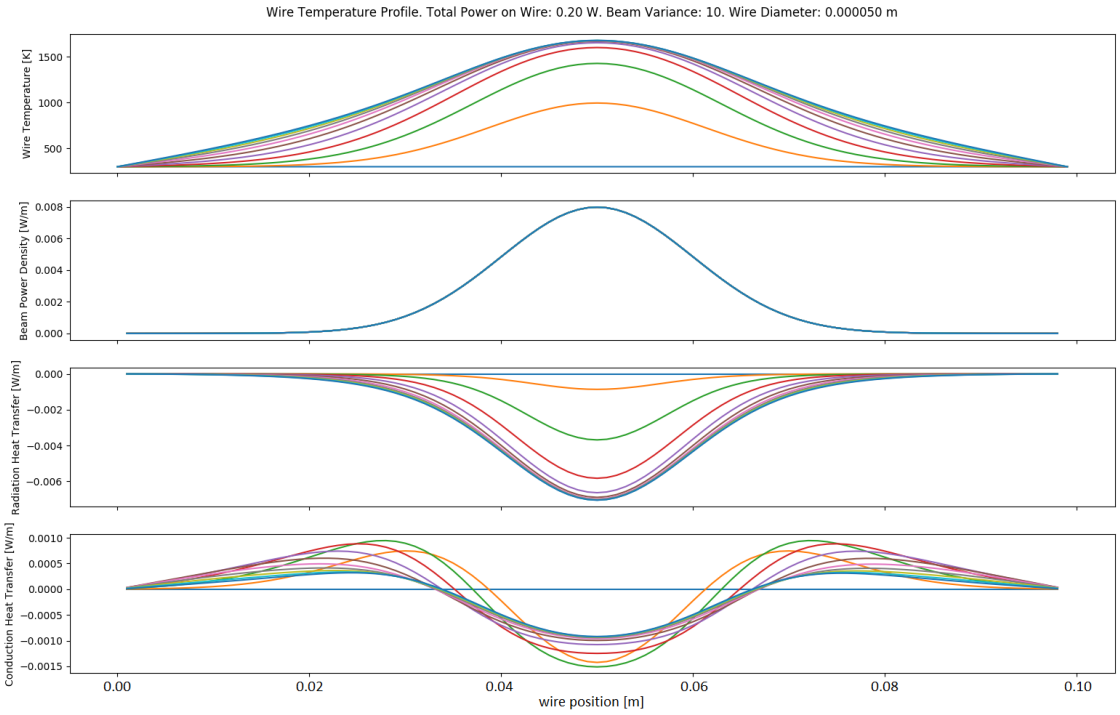

Figure 6.4: Wire temperature profile; beam energy 20 MeV, beam current intercepted by the wire 10 nA, beam power distribution variance 10, wire diameter 50 µm.

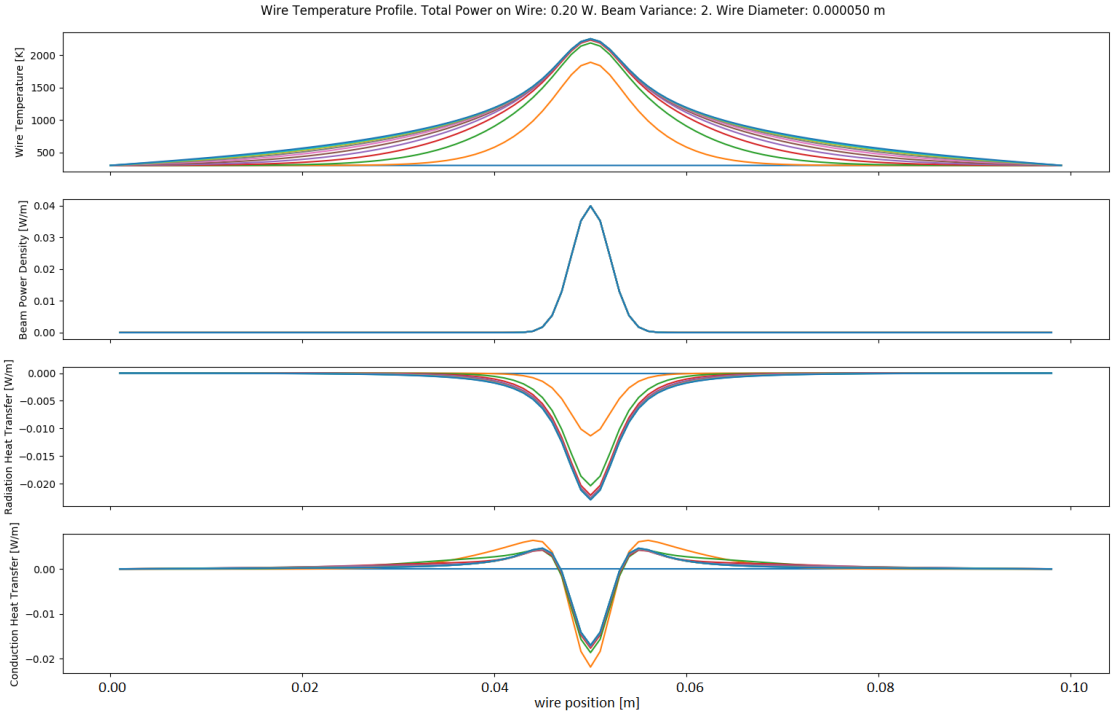

Figure 6.5: Wire temperature profile; beam energy 20 MeV, beam current intercepted by the wire 10 nA, beam power distribution variance 2, wire diameter 50 um

by the wire or changing the wire diameter. With a 50 µm wire diameter and hundreds of nA per wire, at 20 MeV, the maximum temperature exceeds the tungsten melting point. Higher diameters result in higher radiation and conduction flows thus decreasing the wire temperature, but as a con, reduce the grid transparency. The transparency of the BPMs installed in ALPI ranges from 64% up to 90%. With the proposed method, that avoids more complex analysis to calculate the maximum wire temperature from the total beam power [58], it is possible to deduce that the profile grids designed at LNL are able to intercept beams with power up to a few Watts. Transporting stable ion beams, the worst case scenario is represented by the last diagnostic station in ALPI where rarely have been observed beam with powers greater than 1 W. This analysis allow us to conclude that the diagnostic readout electronics must cope with currents from tens of pA up to tens of nA per wire. When transporting low intensities RIBs the wire temperature is not a concern, whilst conversely, the main problem is to extend the sensitivity to very low-intensity currents, down to the fA range. This is accomplished replacing the grid with the Micro-Channel Plate technology as will be described later in section 6.4.2.

6.1.2 Faraday Cup

As well documented in literature [59], several destructive and nondestructive methods have been studied to measure the beam intensity. Nondestructive methods are of particular interest because allow the monitoring of online beam intensity fluctuation. Those techniques are usually complex and rarely cope with sub-nA beams, therefore, their usage is limited in the SPES complex. When transporting RIBs in the pA or sub-pA range, low current measurement techniques are more suitable. Between low current measurements, scintillation counters are one of the most interesting and extend the resolution to the fA range. These particle detectors require a frontend electronics based on fast digitizers and are ideal for high energy beams (hundreds of MeV or higher). Dealing with low energy beams, the Faraday Cup is the simplest and cost-effective solution for the measurements of currents from pA to mA; therefore, this is the technique used to monitor the RIB intensity in SPES. The number of ions trapped in the cup per unit of time, n, is proportional to the measured current I divided by the elementary charge e, assuming each ion with a single state charge.

$$n = \frac{I}{e} \tag{6.8}$$

The new diagnostic readout electronics converts and digitizes this current information in a dynamic range 2 pA to 1 mA. Figure 6.6 shows the Faraday cup prototype designed at LNL. The cooling system is based on a heat pipe connected to the cup with a Shapal™ tablet [60]. This ceramic guarantees high thermal conductivity and an excellent electrical insulation at the same time. The heat pipe terminate with fins cooled by the fan. The emission of secondary electrons from the copper surface struck by incident ions represents the main source of error that can be reduced with an electron suppressor ring kept at negative voltage. The electric field generated by this ring prevents the escape of secondary electrons. Considering a positive particle beam, the escape of secondary electrons would increase the beam current measured, as shown in Figure 6.7.

To quantify the contribute of secondary electrons on the current measurements, the voltage on the suppressor ring was gradually changed with the beam current kept at 110 nA during the whole test. The vertical error bars in figure represent the error introduced by secondary electron. The voltage to supply the secondary electron suppressor in order to measure true values is strongly influenced by the Faraday cup radius and shape. E.g., for the prototype under test this voltage is well above the usual value of 500 V, emphasizing the need of a geometry improvement.

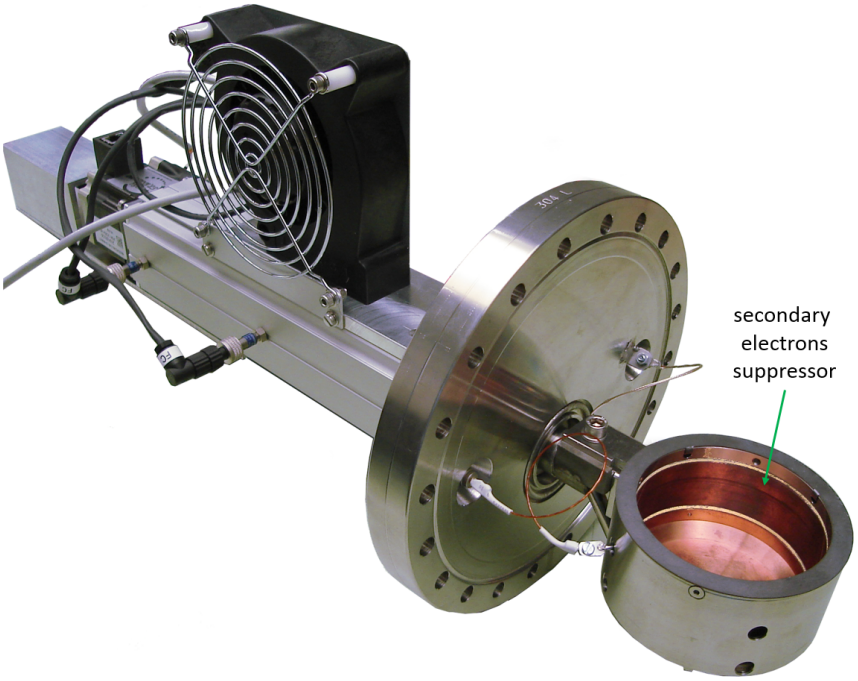

Figure 6.6: Faraday cup prototype.

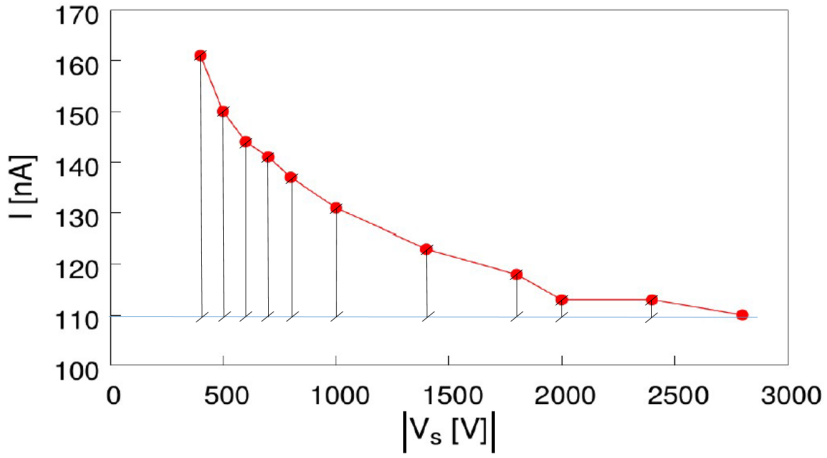

Figure 6.7: Faraday cup measured current vs secondary electrons suppressor voltage.

6.2 Diagnostics Readout Architecture

The sensitivity to low intensity beam is imposed by the quality of the readout electronics. The IOC board will be the core of the new beam diagnostic data

acquisition and signal processing system, where it will replace the legacy VME technology. A block diagram of the new readout electronics is shown in Figure

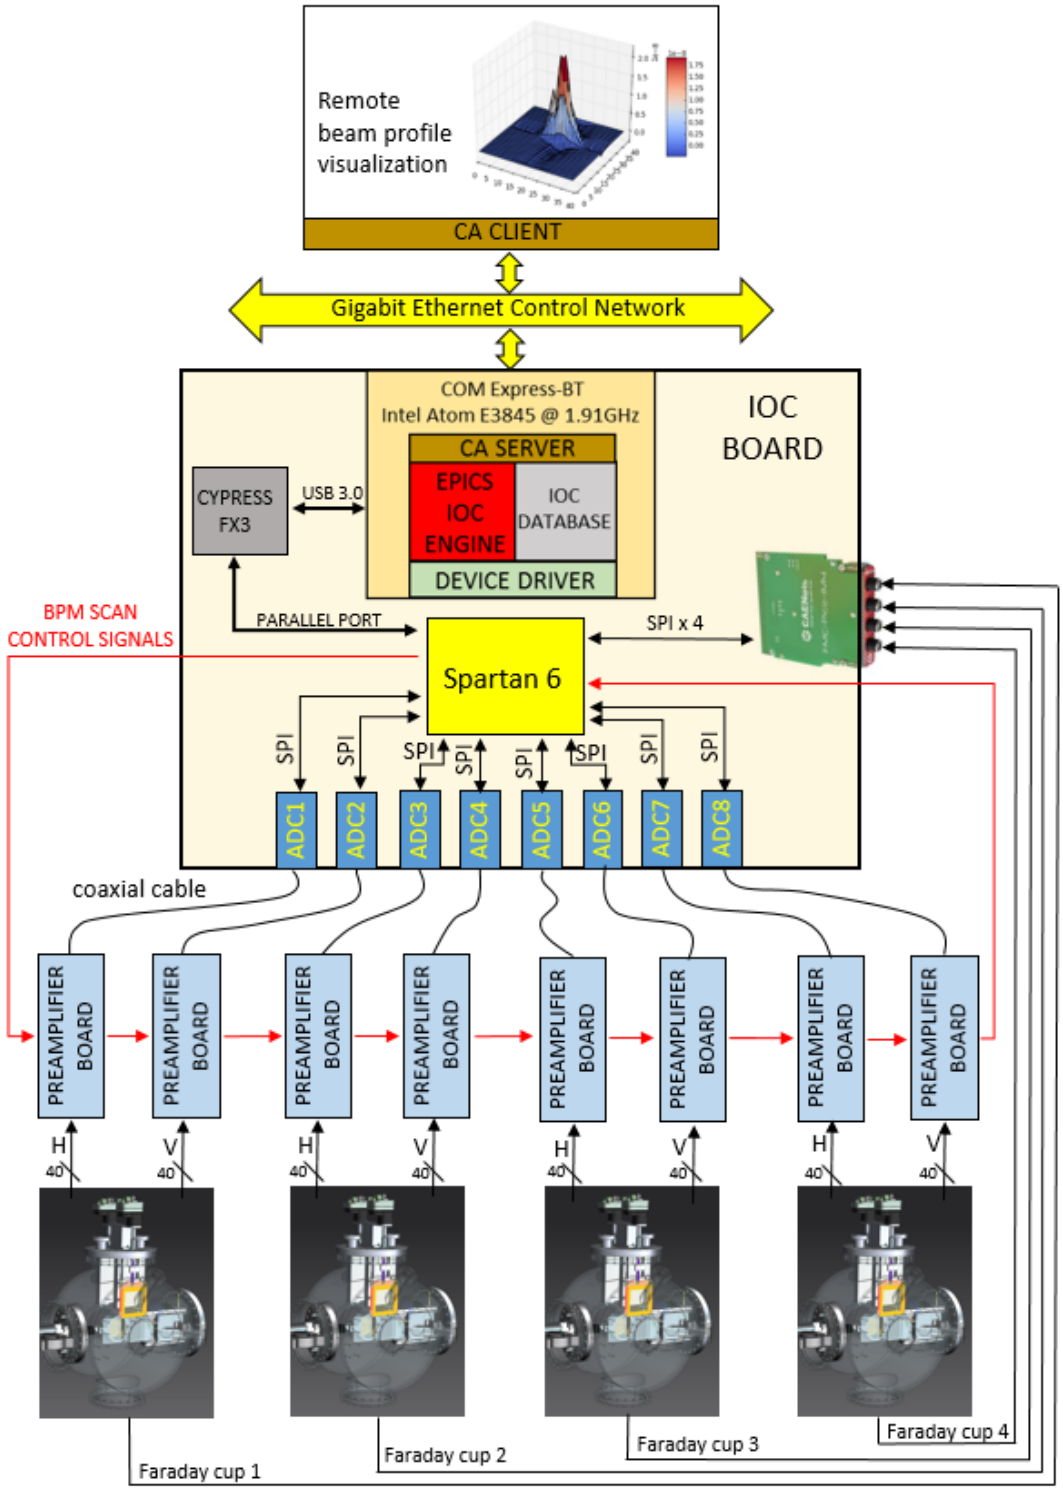

Figure 6.8: Block diagram of the diagnostic data readout chain.

6.8. Each IOC controller serves up to 4 diagnostic stations, each endowed with a Faraday cup and a BPM. The BPM considered are mostly grids, but, this readout architecture, without any variation, is able to readout beam profilers based on the micro-channel plate technology. Horizontal and vertical beam profile information are first elaborated by analogue frontend boards, shown in Figure 6.9, that convert the charge deposited by the beam on the wires in a voltage signal and perform the 40 to 1 channel multiplexing in a time division technique. The TransImpedance Amplifiers (TIA) are built around a very low bias current JFET amplifier [61]. Indeed, the bias current is the main source of error in this configuration. Whereas the voltage offset of the amplifier stage is negligible in our measurements. The TIA stage designed guarantees also a very low input impedance at low frequency, and, the feedback network is digitally selectable by the user depending on the current to be measured. The bandwidth is fixed to 10 Hz for all feedback configurations, while the gain value available are 10^6, 10^7, 10^8 and 10^9 depending on the beam current sensitivity to be achieved. The outputs of the 40 TIA stages are multiplexed to a single voltage signal digitized by the ADC on-board the IOC card. The preamplifier boards are connected in daisy chain topology of up to 8 boards; the grids scan is simultaneous and controlled via digital I/O signals (red arrows in the picture). These single-ended digital I/O are fed into an RS422 transceiver

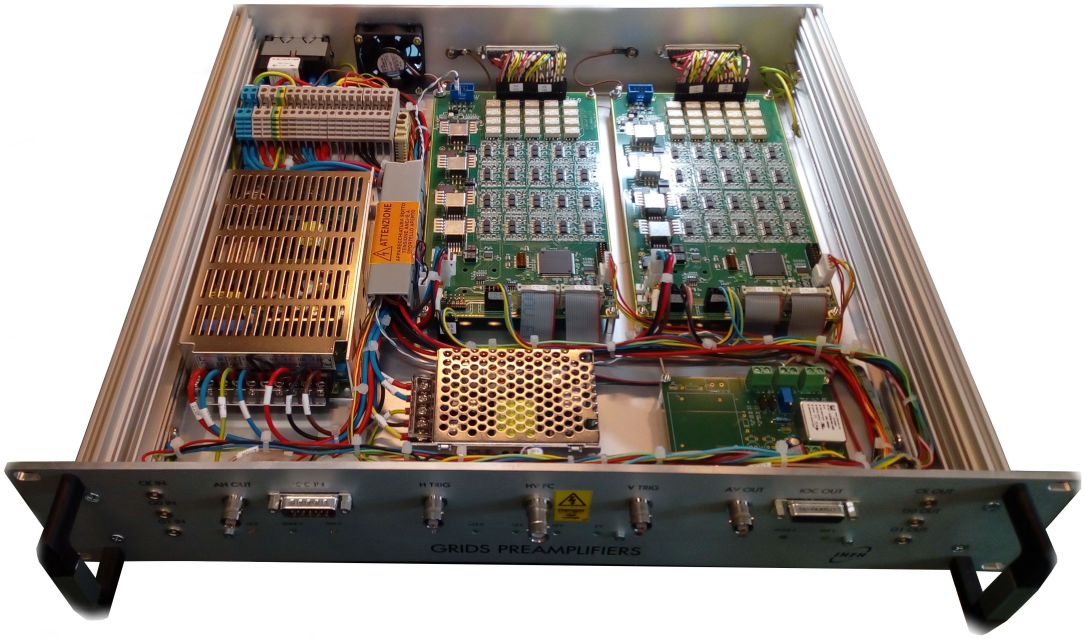

Figure 6.9: Grid horizontal and vertical channels preamplifier boards.

to suit a differential data transmission over twisted pairs. The beam current information generated by the Faraday cups are fed to the commercial 4-channels FMC picoampere meter. This mezzanine card integrates analog front-end and the digitization stage and interfaces with the FPGA through serial SPI interfaces. The readout of diagnostics data is an example of an application where the Spartan6 performs time-critical operations, like the BPM scan, freeing the operating system from real-time tasks. The FPGA controls the data conversion, data acquisition, data buffering and first digital processing and transfers the diagnostic data to the processor. The COM Express runs and EPICS input output controller application that implements a database and channel access server sharing the diagnostic data on the control network for remote visualization. The COM Express module chosen for this application is the compact Express-BT which exploits an Intel Atom E3845 @ 1.91 GHz processor. The operating system installed is a standard Linux distribution. The network connectivity is ensured by the Gigabit Ethernet port available on the COM Express module. Remote FPGA debug tools and reconfiguration, temperature and power monitor, fan speed control, LCD display control are all additional features supported by the IOC card that will be exploited in the field in this application.

6.2.1 Diagnostic Data Digital Processing

The frontend computer performing the diagnostic data acquisition must run several parallel tasks, some of them having time-critical requirements. For this reason, past solutions, based on the VME technology, have foreseen the usage of VxWorks real-time operating system in the VME CPU board [62]. Whilst, in the new readout system, these time-critical tasks are conveniently implemented in hardware inside the FPGA. The VHDL code developed handles the data acquisition (grids scan) and opens the possibility to operate digital filters fundamental to reach a current resolution in the pA range. An overview of the digital data flow and signal processing is given in Figure 6.10. The primary task of the firmware is to expose to the software all the relevant process variables of the peripherals under control together with the beam current and profile information. The Cypress fx3 chip bridges software read and write requests to the FPGA. These read and write operations are decoded by the *USB interface core*. In the diagnostic application there is no need to readout blocks of raw data, therefore, the I/O stream ports of the USB core are not exploited. Whilst, only two memory ports are used. One is reserved for the

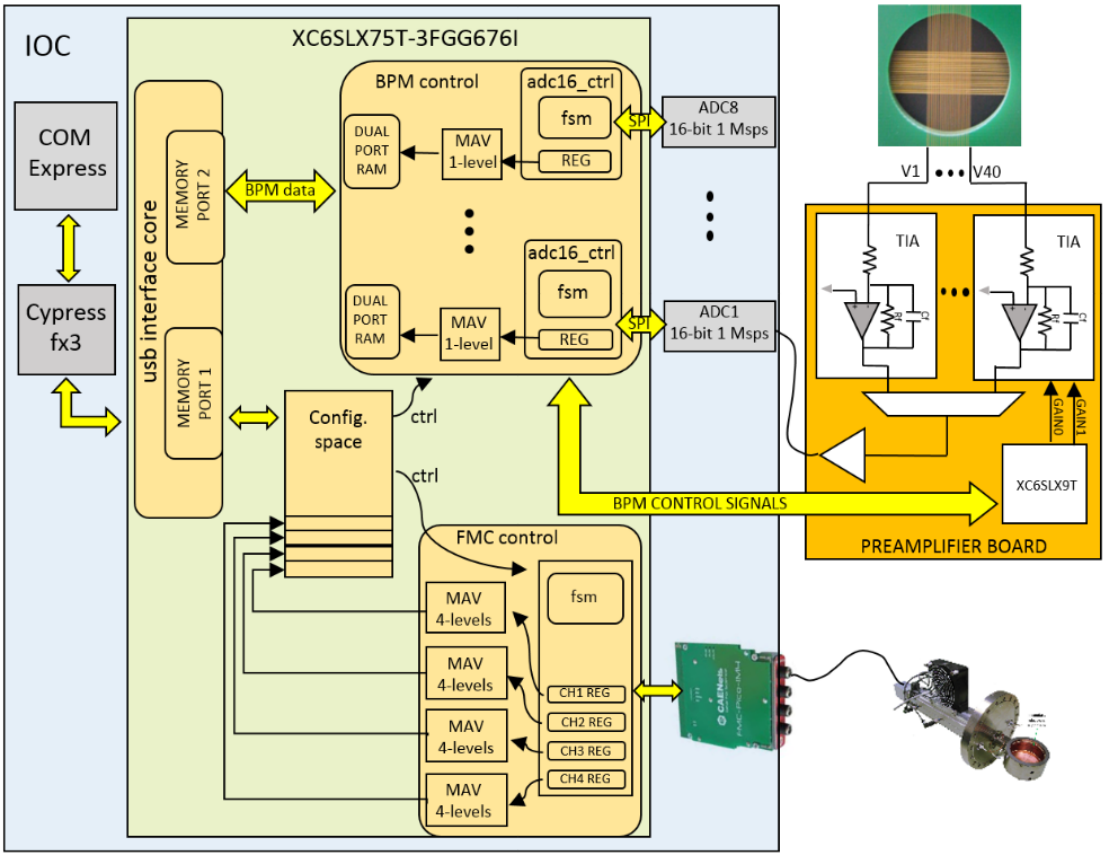

Figure 6.10: Diagnostic data readout RTL design.

configuration space that consists of a stack of registers involved in the diagnostic readout system as resumed in Table 6.1.

The BPM control register and FMC control register reported in tables 6.2 and 6.3 respectively, allow the end user to control the profile grids and Faraday cups scan. Four registers in the configuration space store the beam current information (read out from the Faraday cups), instead, the beam current intercepted by the 40 hori-

Address	Register	Description
0x53	`FMC_CTRL_REG`	FMC control register
0x54	`FMC_DATA1`	Faraday cup 1 current
0x55	`FMC_DATA2`	Faraday cup 2 current
0x56	`FMC_DATA3`	Faraday cup 3 current
0x57	`FMC_DATA4`	Faraday cup 4 current
0x58	`BPM_CTRL_REG`	BPM control register
0x5b	`BPM_DELAY`	BPM clock rising edge to start acquisition delay

Table 6.1: Diagnostic registers.

Bit	Description
0	BPM start. 1 = start acquisition. 0 stop acquisition.
1	BPM gain 1
2	BPM gain 2
3	BPM continuous acquisition mode
4	BPM selection. 1 = positive beam. 0 = negative beam.
5	BPM test mode
31	BPM reset.

Table 6.2: BPM control register.

Bit	Description
0-9	Number of samples. To be set to 0 for continuous acquisition
10	1 to start acquisition. 0 to stop acquisition
11	FMC CH1 range. 1 = 1 mA. 0 = 1 µA
12-14	FMC CH1 MAV number of samples per stage
15	FMC CH2 range. 1 = 1 mA. 0 = 1 µA
16-18	FMC CH2 MAV number of samples per stage
19	FMC CH3 range. 1 = 1 mA. 0 = 1 µA
20-22	FMC CH3 MAV number of samples per stage
23	FMC CH4 range. 1 = 1 mA. 0 = 1 µA
24-26	FMC CH4 MAV number of samples per stage
31	FMC reset

Table 6.3: FMC control register.

zontal and 40 vertical wires of each BPM is stored in 8 consecutive banks of dual port RAM, each consisting of 40 registers. The beam current and profile information are periodically refreshed by the *BPM control* and *FMC control* cores upon setting the corresponding start acquisition flag to logic state '1'. The mezzanine card implements 4 analog input channels. The conversion is simultaneous and is controlled by the *FMC control* core via a rising edge on the *CNV* digital signal. The conversion time, indicated by the mezzanine to the carrier with a *BUSY* flag signal, is followed by the acquisition phase during which the *FMC control* core readout the sample from each channel. The digital interface is based on a SPI bus. During the acquisition time, data must be registered inside the *FMC control* core using the return clock signal from the mezzanine to the carrier. Setting the number of samples field to 0 in the FMC control register, the continuous acquisition mode is selected, and, upon starting the conversion, the raw data are continuously gathered into the 4-level Moving Average filter (MAV) [63]. The DC beam current measurement, output of the MAV filter, is stored in the corresponding register

in the configuration space. Here, the current value is refreshed with a latency proportional to the filter depth given by (6.9).

$$MAVdepth = (samples\_per\_stage)^{number\_of\_stages} \tag{6.9}$$

The number of stages is hard-coded and it is equal to 4, while, the number of samples per stage is user selectable via the configuration space and is a power of 2 in the range 8 to 128. In a similar way, the *BPM control* core implements 8 SPI master devices that interface with the ADCs and control the current conversion and acquisition phases via a finite state machine. The hardware takes care of this low level protocols guaranteeing a good level of hardware abstraction to the software. The complexity of the *BPM control* core is not limited to the ADCs interface, but, it also generates the control signals that govern the BPM scan, as shown in Figure 6.11. The gain signals select the feedback network of the TIA stage in the preamplifier boards and accordingly the DC current gain. When in test mode the input of TIA is fed with a known reference current and this mode is exploited during the initial calibration phase. The calibration procedure aims to correct the offset and gain error of the whole acquisition chain (TIA stages plus ADC). The calibration parameters can be saved on-board in the EEPROM memory and loaded at power-up in the configuration space. A dedicated signal is used to select the beam polarity. These slow control signals are directly routed from the configuration space to the IOC board digital outputs. The scan of the 40 horizontal or vertical wires is synchronized with the *BPM clock* signal which alternates 40 x 200 µs pulses to pauses lasting 1.6 ms. The generation of such clock signal was a concern in the previous diagnostic data readout architecture based on the VME technology and led to the usage of a real-time operating system. Each *BPM clock* period schedules the charge acquisition from the corresponding vertical and horizontal wire in all the grids. To reduce the equivalent input noise of the frontend electronics, during each clock pulse 64 consecutive samples are gathered

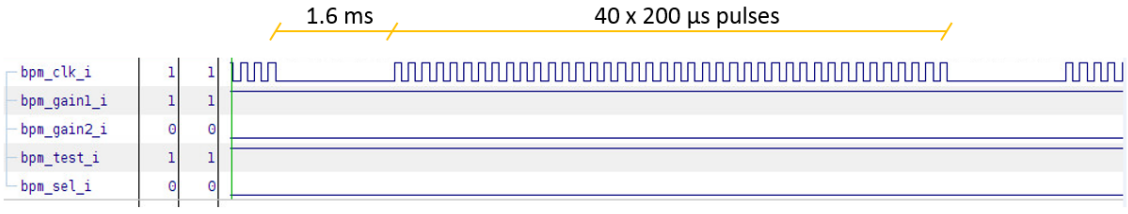

Figure 6.11: BPM scan digital control signals.

inside a single level MAV and the average value output by the filter is stored in the corresponding register in the dual port memory.

6.3 Current Sensitivity Tests

Before the installation in the field for measurements under real beam conditions, a full diagnostic readout chain has been assembled and fully characterized in the Lab. The purpose of these preliminary tests was to release a first version of the firmware and to understand the expected current sensitivity. The current sensitivity characterization setup is shown in Figure 6.12 together with the preamplifier box and the IOC box assembled. Leaving the 40 TIA inputs floating (thus virtually grounded) and executing a scan it is possible to have a rough measure about the noise introduced by the electronics (cables, connectors, and the grounding strategy may increase the noise on field) which is about 5 pA, deduced by the variance of the measurements relative to consecutive scans of the vertical or horizontal wires. This value can be further reduced considering, for each channel, the average realization of 10 seconds of acquisition. Figure 6.13 is the average result of a BPM acquisition lasting 10 seconds with a scan rate fixed to 10 Hz and the analog inputs

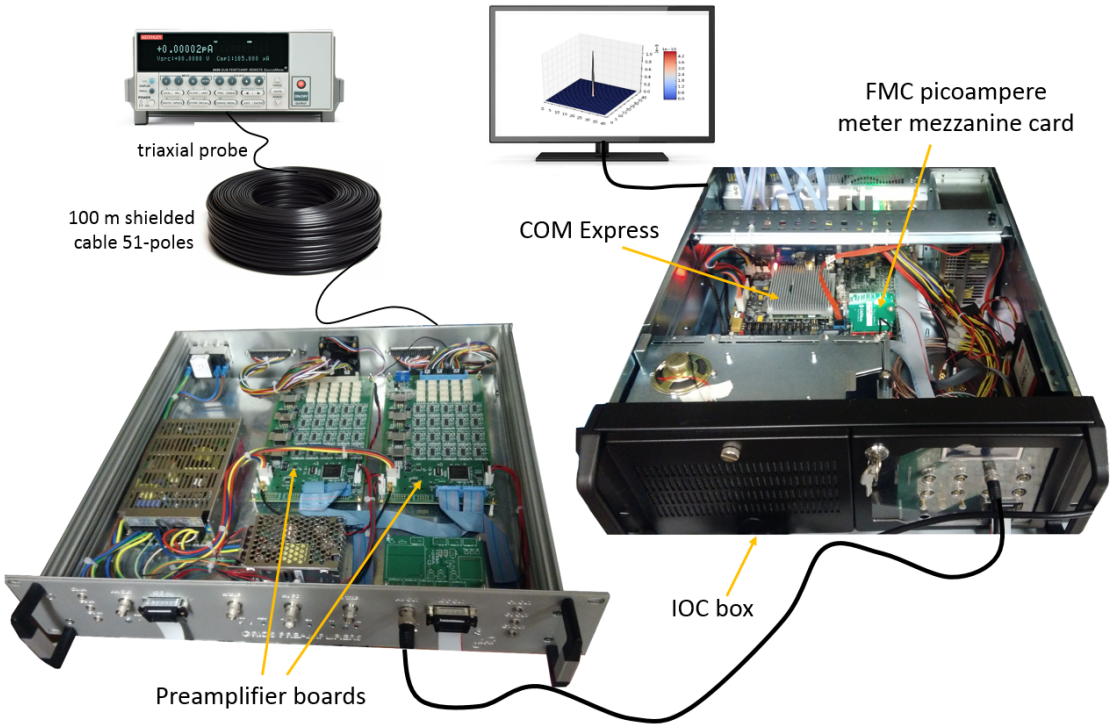

Figure 6.12: Current sensitivity characterization setup.

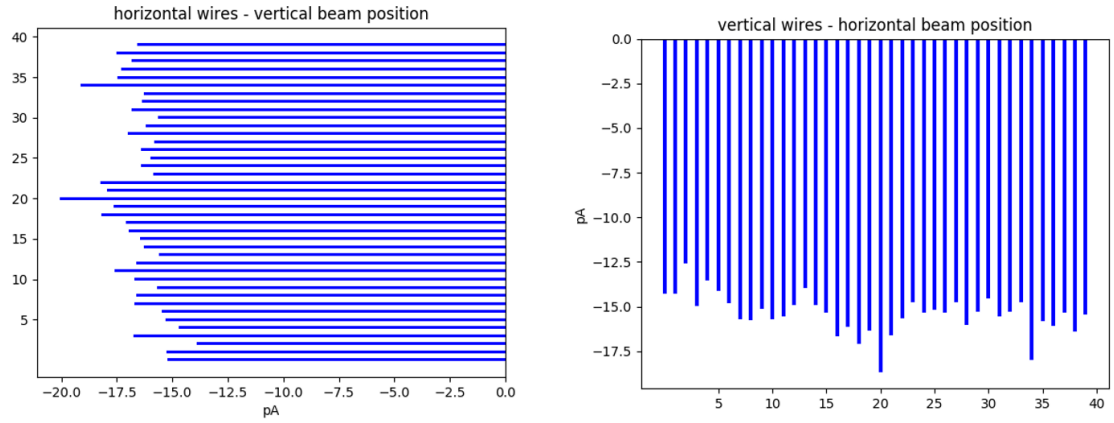

Figure 6.13: BPM offset measured with the inputs virtually grounded by the input stage of the TIA.

floating. Ideally, the average current measured on the wires should be zero, but, it is well visible how all the measurements are affected by a negative offset. The offset fluctuates significantly with temperature and time, therefore a precise calibration procedure makes sense only if performed periodically and automatically with the hardware support. To validate and characterize the current sensitivity of the hardware developed, the offset measurements have been stored in a file and offline subtracted to the current measurements. The profile grid has been emulated with a Keithley 6430 sub-femtoampere high stability current source. The stable DC current has been injected to one of the 40 analog inputs via a 100 m shielded cable to reproduce the on field installation. At this point, making use of a Python script, it was possible to control the acquisition and save the measured currents in a file. The test was repeated decreasing the injected current in one of the forty wires and any acquisition consists of 100 scans (10 seconds at 10 Hz rate). The vertical and horizontal profiles can be reconstructed from the 10 seconds averaged vectors. Figure 6.14, Figure 6.15, and Figure 6.16 validate the proposed BPM readout method and confirm the good current sensitivity achievable with the new electronics.

Current as low as 5 pA are well in the range of the new diagnostic system and these tests demonstrate that the preamplifier boxes can be placed tens of meters away from the beamline without degrading the measurements. In this way, the electronics is preserved from radiation and is accessible for maintenance during the beam time. The total beam current information is provided by the Faraday cup and this information is digitized by the mezzanine card that has two measuring ranges of ± 1 mA and ± 1 µA selectable by the end user via the configuration space. This

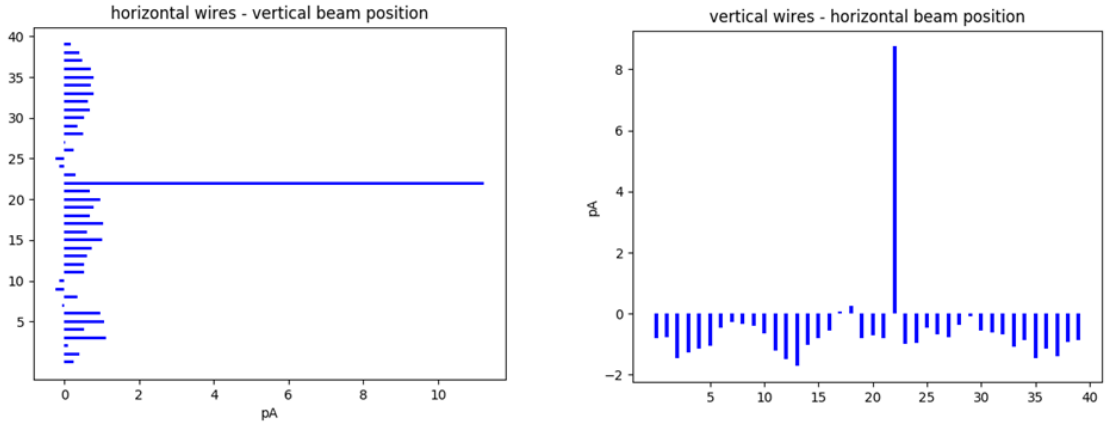

Figure 6.14: BPM acquisition with the injected current set to 10 pA.

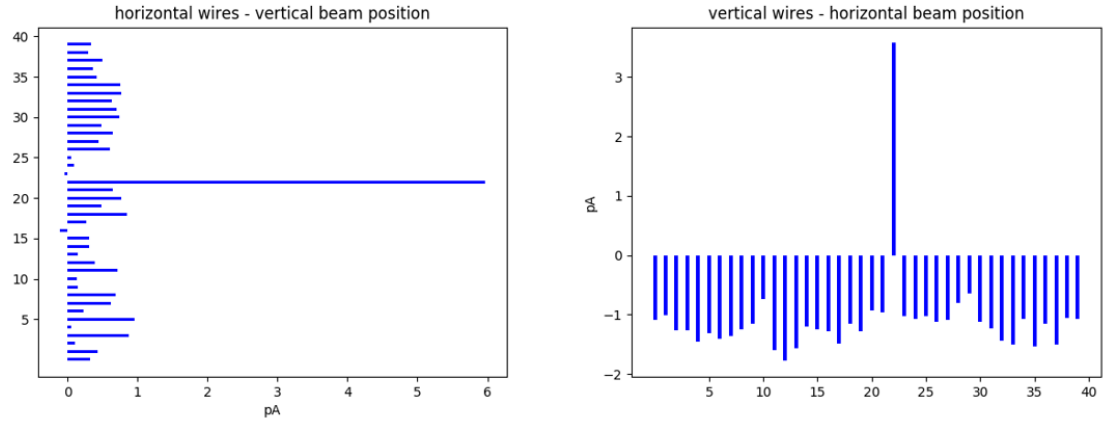

Figure 6.15: BPM scan with the injected current set to 5 pA.

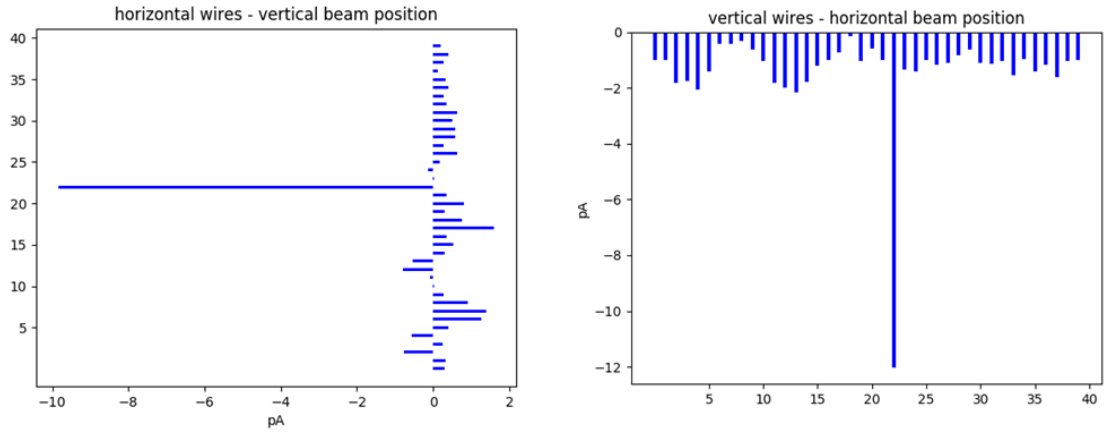

Figure 6.16: BPM scan with the injected current set to minus 10 pA.

commercial card has been chosen because integrates 20-bits resolution ADCs that, when in µA range, guarantee an ideal current resolution of 1.9 pA. The ENOB is

then limited by all the noise and distortion introduced by the electronics (cables, mezzanine carrier board and by the ADC itself). However, since the interesting value is the DC beam current, the digital implementation of the MAV filter has the effect of a reduction of the input equivalent noise of the ADCs, maximizing the current resolution as confirmed by the tests performed.

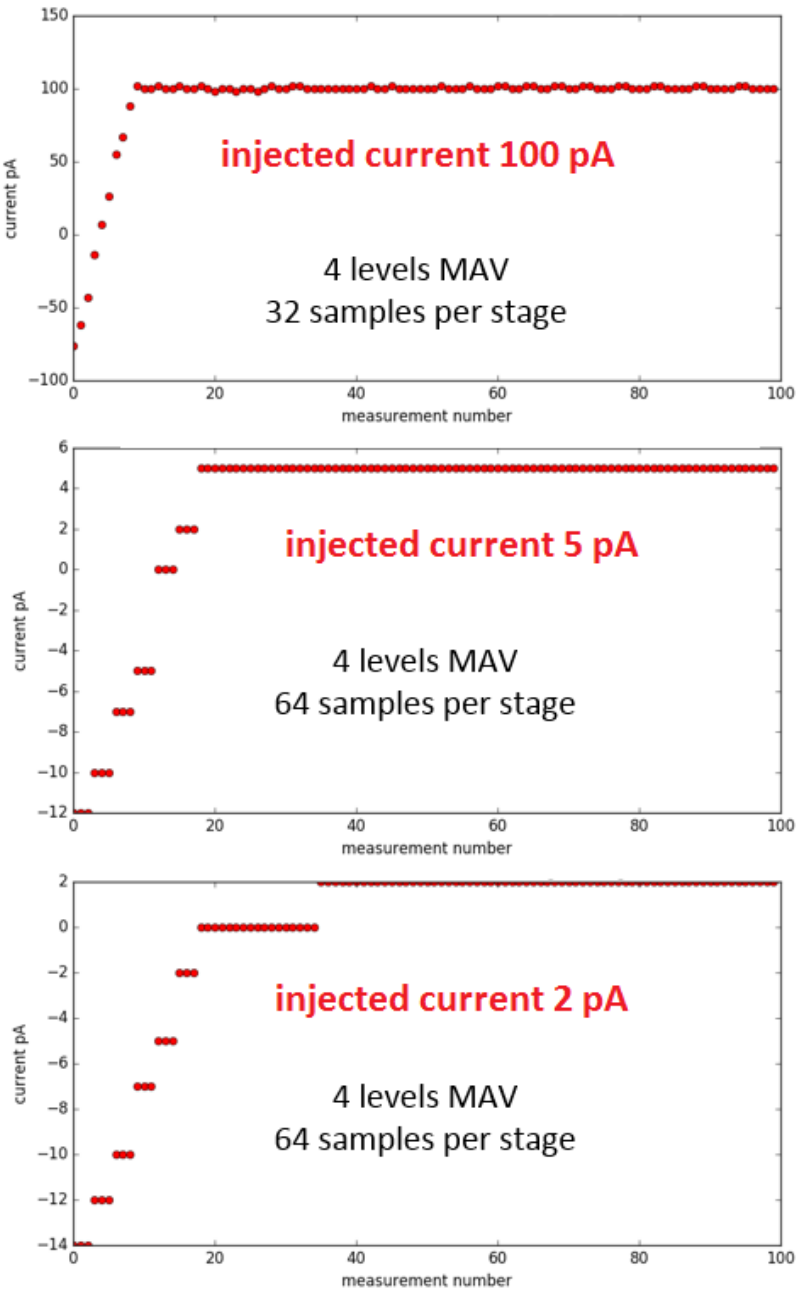

Figure 6.17: Faraday cup readout sensitivity test. The current measurements have been repeated with different injected currents.

The precise current source has been used to inject currents in the pA range to a coaxial cable connected to the mezzanine picoampere meter hosted by the IOC card. Also in this case all settings have been done running a Python script in the COM Express module. The mezzanine was operated in continuous sampling mode with a sampling rate of 864 Ksps. Samples are gathered into the 4-level MAV filter that outputs the average current value, refreshed with a latency given by (6.9). The current measurements have been repeated with injected currents of 100 pA, 5pA and 2pA. For each current value a run of the Python script executed 100 read operations, represented by red dots in Figure 6.17. To successfully measure 2 pA it was necessary to decrease the ADC bandwidth to the sub Hz range (64 samples per stage). This bandwidth reduction does not represent a problem while transporting the beam; the primary requirement of the new electronics was the DC current range extension to few pA.

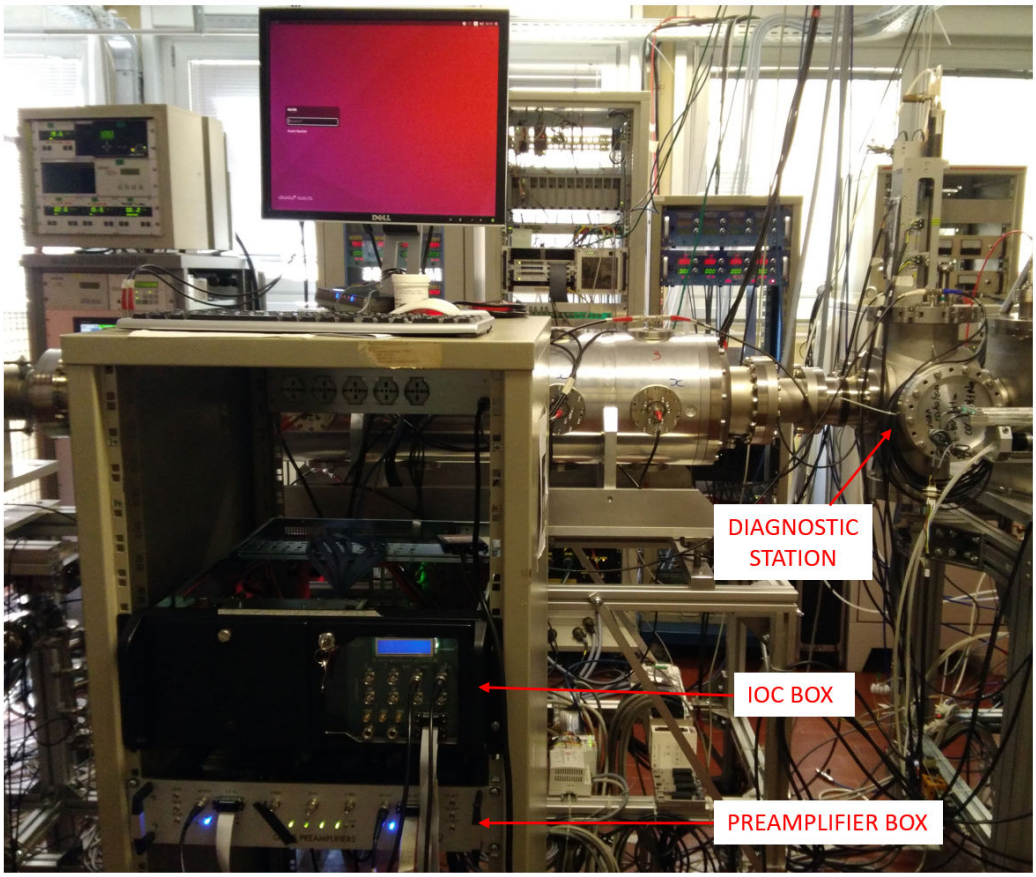

Figure 6.18: Diagnostic rack in the SPES frontend laboratory.

6.4 Real Beam Measurements

The positive outcome of the characterization tests described above, candidate the proposed diagnostic readout electronics as a cost effective solution to replace legacy VME crates in the ALPI's diagnostic subsystem, and to be integrated as SPES diagnostic frontend computers along the whole RIB transport line. The first beam profile measurements have been carried out in the SPES frontend laboratory as shown in Figure 6.18. The SPES frontend laboratory emulates the target-ion source chamber and the exotic particles extraction and selection. The laboratory is endowed with a low energy ion source and with all the elements necessary to transport the beam, including a diagnostic station. The beamline chamber is always accessible during the operation given the very low energy beams delivered. Therefore, this Lab is the ideal test bench for testing the new diagnostic devices developed at LNL, giving the opportunity to measure the beam profile with the new signal processing system [64]. Figure 6.19 shows the outcome of a BPM scan

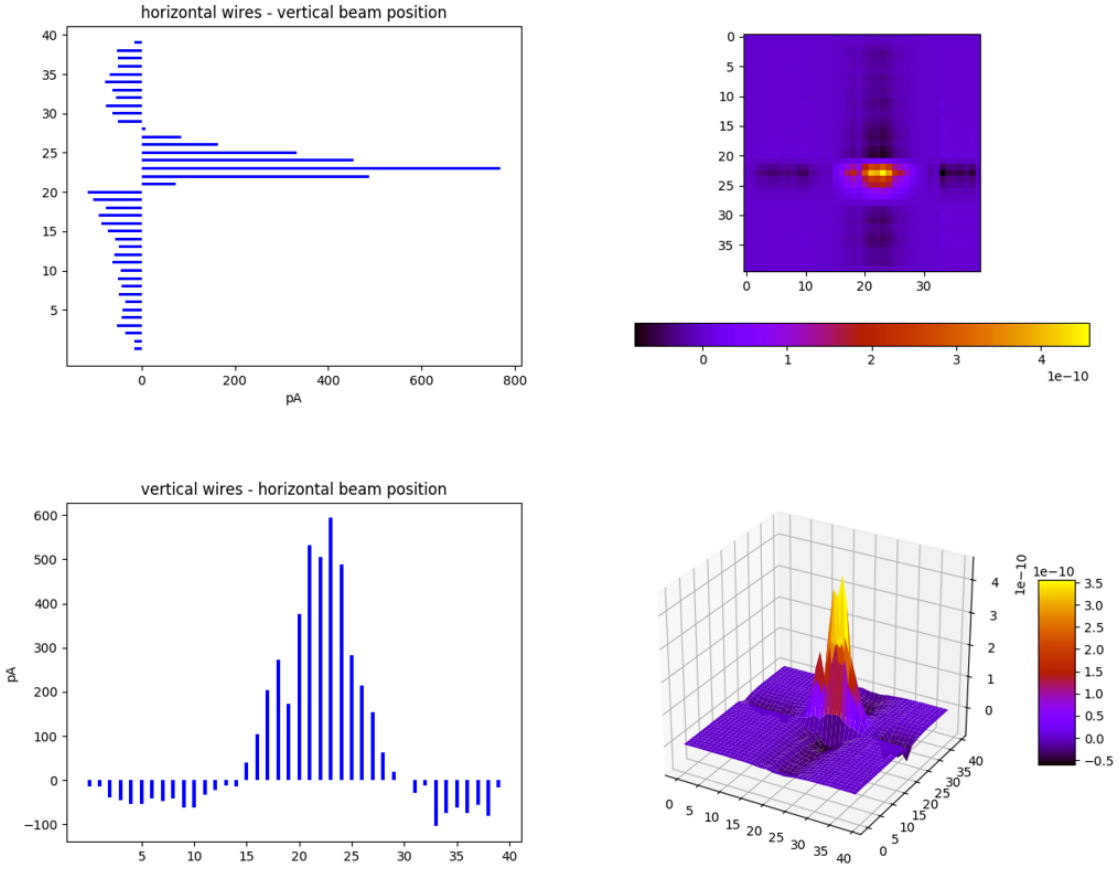

Figure 6.19: SPES frontend laboratory, 25 keV, 1.6 nA, N_2^+ beam profile visualization.

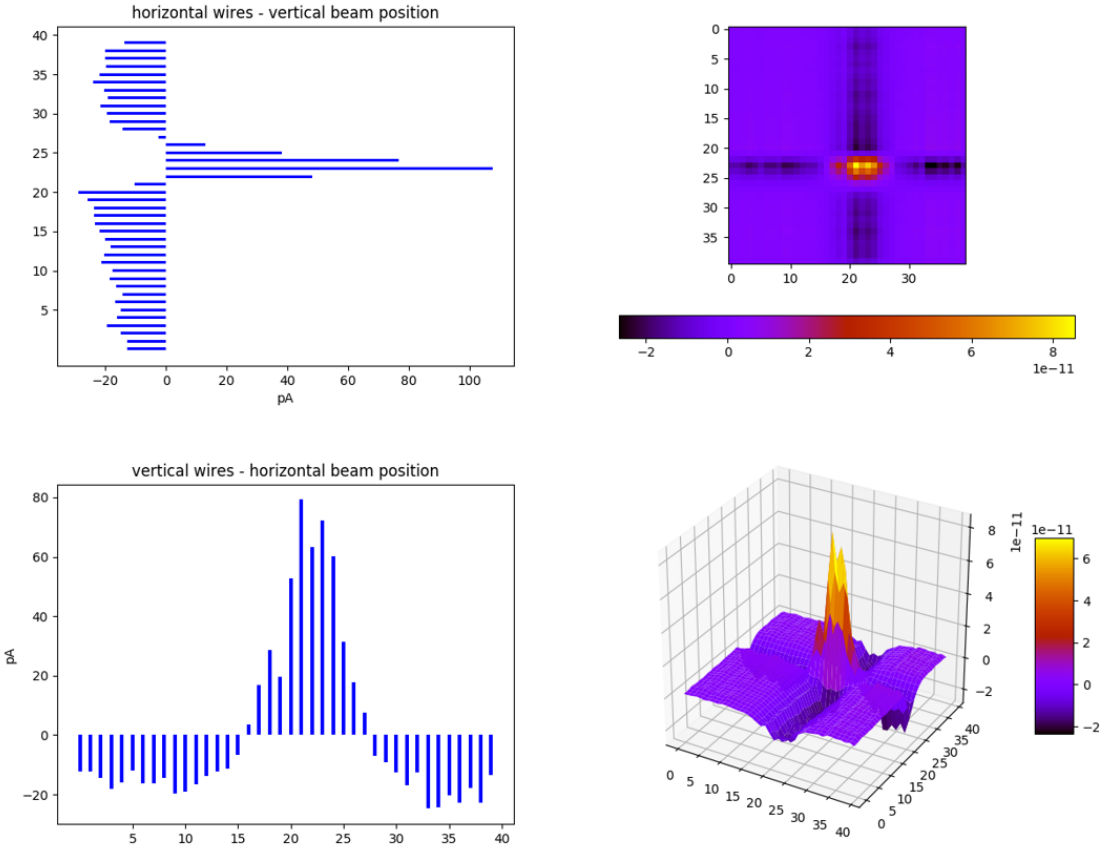

Figure 6.20: SPES frontend lab, 25 keV, 250 pA, N_2^+ beam profile visualization. © 2018 IEEE.

lasting 10 seconds at 10 Hz with a profile grid intercepting a 25 keV N_2^+ beam with intensity 1.6 nA. The BPM installed at the SPES frontend is a bit different from the grid designed for ALPI and consists of 4 cm mesh made of 400 µm steel wires spaced 1 mm. The transparency of this grid is about 36 %. The emission of secondary electrons, well visible in the proposed 2D beam reconstruction, increases the positive charge of those wires struck by the primary positive ion beam to be measured, and, it generates an excess of negative charge on the neighboring wires not intercepting the beam. Therefore, the equivalent effect of secondary electrons is that of a multiplication factor of the primary beam that increases the imaging contrast and the current sensitivity. The 2D and the 3D beam surface reconstructions are the outcomes of the multiplication of the 10 seconds vertical and horizontal average vectors. The plots were generated using Matplotlib, a Python plotting library. The BPM scan was repeated with lower intensity beams. Figure 6.20 shows the beam profile plots achieved with a beam current of 250 pA. As

expected tens of pA per wire is well in the capability of the new electronics that outperforms the old VME based diagnostic system that, instead, was not able to reconstruct the transverse profile of the 250 pA beam.
The Faraday cup was connected to the picoampere meter mezzanine card, as described previously, in order to monitor the beam intensity during the tests. The contribution of secondary electrons, escaping the Faraday cup, on the beam current measurements can be deduced switching on and off the secondary electrons suppressor as shown in Figure 6.21. The secondary electrons contribution on the measured current can be quantified in a factor of about 4.6. The motion of the BPM and the Faraday cup are completely independent and it is possible to measure the beam current with or without the grid inserted in the beam line in front of the Faraday cup. In this way, we could estimate the grid transparency that turned out to be about 33 %, in good agreement with the 36 % calculated given the grid geometry information.

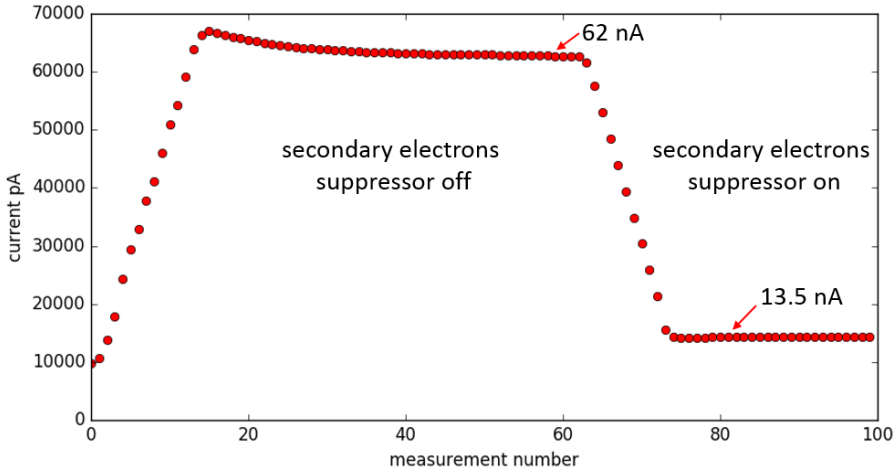

Figure 6.21: Beam current measure. Secondary electrons suppressor switching.

6.4.1 Integration of the New Diagnostic System in ALPI

All the diagnostic boxes in ALPI (and more in general in the SPES facility) must be integrated into the control system network that is based on the EPICS channel access server and client concept. EPICS framework ensures a uniform graphical user interface to operate and control all sub-systems at LNL. The upgrade of the diagnostic software using the EPICS toolset has already been implemented for the old VME based diagnostic readout system [65] and has easily been ported to the new hardware. The EPICS input output controller application runs in the COM

Express module and loads a simple database that provides minimal data processing since the signals from BPM and Faraday cup are real time processed within the FPGA. The IOC server application pulls data from the FPGA and publishes on the Ethernet control network all the relevant information in the form of Process Variables (PV). These PVs are accessible to the clients faced on the control network via the Client Access (CA) protocol. The database is the higher layer of the EPICS server application, where the PVs are defined and made available to clients. At the lowest level, EPICS foresees a device support layer that defines all modules that implements the communication with hardware devices, the Cypress fx3 chip in our case. The device support layer is built around the asynPortDriver, a C++ class that provides uniform and standard methods to handle the communication between the database and the user space or kernel space device drivers.
The diagnostic station DI9 of the ALPI accelerator is already endowed with the new diagnostic readout system. Figure 6.22 refers to a 7.7 Mev, 110 nA, ^{15}N beam transport performed with the new electronics installed in DI9. The BPM used to intercept the beam was a standard grid made of tungsten wires with a diameter of 50 µm and a wire spacing of 500 µm. The transverse profiles present outliers caused by the long exposure of the grid to the beam, and by the consequent high temperature reached by the wires.

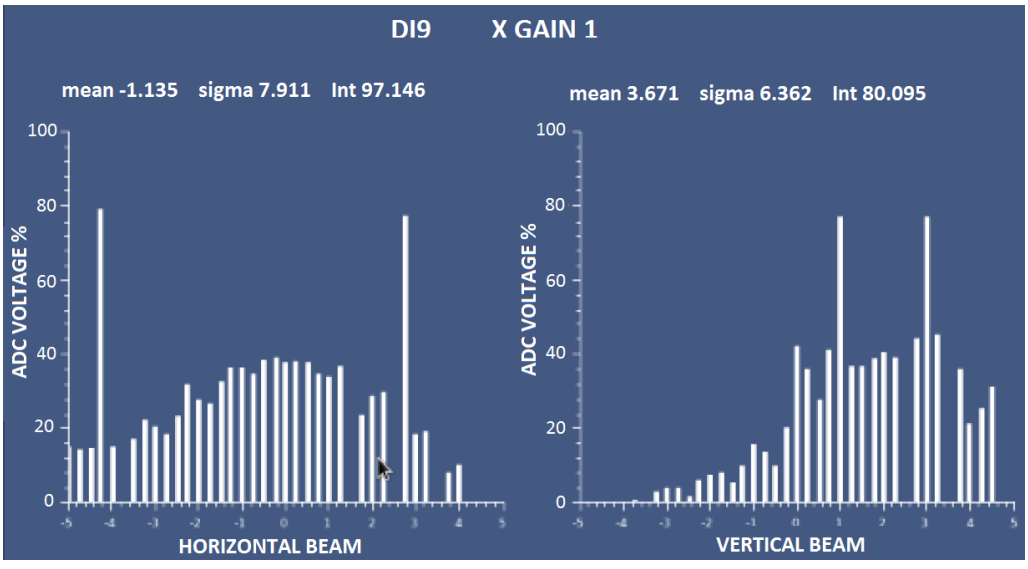

Figure 6.22: ALPI diagnostic station DI9. 7.7 Mev, 110 nA, ^{15}N beam profile. © 2018 IEEE.

6.4.2 Micro-Channel Plate Detectors

Profile grids are a suitable solution to measure the two-dimensional profile of low energy beams with intensities ranging from hundreds of nA down to tens of pA. Yet, with reference to the SPES layout shown in Figure 1.3, the RIB currents foreseen between the target-ion source chamber and the charge breeder may have intensities as low as 10^4 ions per seconds (few femtoamperes). For this purpose, at the Legnaro laboratories, an important research activity related to the design and test of new beam profile monitors based on the Micro-Channel Plate (MCP) technology is being carried out. The MCP consists of a two-dimensional array of thin glass channels. The ions entering a channel struck the channel wall and yield secondary electrons that are accelerated through the channel by the electric field generated by the voltage applied across the MCP. These secondary electrons, in turn, strike the channel surface and produce more secondary electrons as an avalanche multiplier. This cascade process leads to a multiplication factor of about 10^4. The MCP is installed in front of a position sensitive anode matrix that collects the secondary electrons output from the glass channels. This PCB, whose layout is shown in Figure 6.23, consists in a matrix of copper pads connected each other

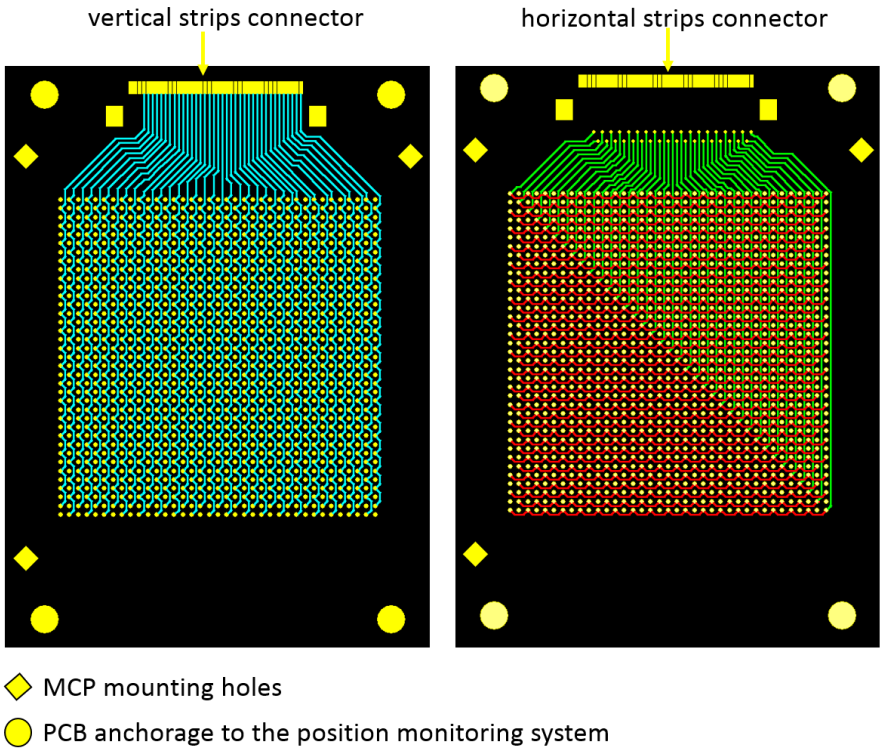

Figure 6.23: MCP position sensitive anode collector matrix.

in a way to form a grid of 40 x 40 strips, similar to the profile grid.
The expectation is to be able to detect very low intensity beams. Indeed, the multiplication factor of the MCP technology is in the order of 10^4. Since the sensitivity of the readout electronics extends to 10 pA our expectation is to be able to characterize beams of tens of fA. The great advantage compared with other low intensity profile monitors is that the MCP exploits exactly the same acquisition system and digital data processing developed for the profile grids. This is very important to achieve hardware and software standardization of the beam diagnostic readout system of the whole SPES facility. The new data acquisition and signal processing system perfectly cope with the anode collector matrix and has been used to test the MCP detector in electrostatic-mirror configuration in the SPES frontend laboratory [66]. Figure 6.24 shows the horizontal beam profile measured with the new electronics and the MCP technology with an estimated beam intensity of about 115 fA.

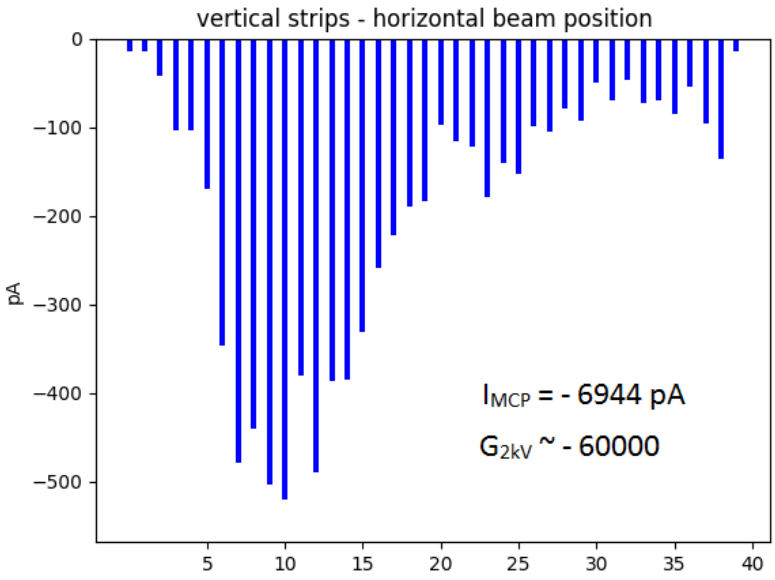

Figure 6.24: Plot of a 25 keV, 115 fA, N_2^+ horizontal beam profile measured with the MCP and a bias voltage of - 2kV.

Currents below 1 pA cannot be measured by standard beam instrumentation available in our lab, therefore, an extrapolation procedure was necessary to estimate the beam current during the tests. The gain of the MCP (multiplication factors), that should be independent on the beam current, at a given bias voltage applied to the MCP, has been extrapolated from beam current measurements performed

in the pA range where the Faraday cup and the IOC board yield accurate current measurements. This gain is defined as the ratio between the total current measured by the MCP I_{MCP} (obtained by subtracting the integrated noise floor current to the integrated current of the 40 strips in the picture) and the real beam current measured by the Faraday cup I_{Fc}. The multiplication factor extrapolated at -2 kV is about $-6x10^4$. Dividing the I_{MCP} by this gain we could estimate the beam current; 115 fA in the proposed example. More tests were performed decreasing the power supply of the ion source and thus decreasing the beam intensity. Figure 6.25 shows the horizontal profile of a 25 keV, N_2^+ beam with an estimated intensity of 10 fA. The beam profile is still visible even if the baseline starts to be quite high if compared with the beam lobe. Further increasing the bias voltage above 2 kV, in order to increase the gain, does not necessarily improve the sensibility because higher electric fields may lead to partial distortion of the measured beam profile.

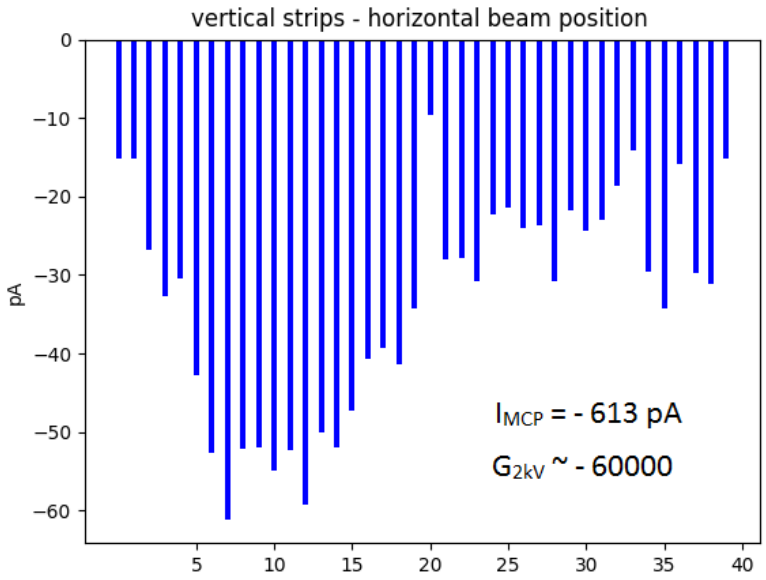

Figure 6.25: Plot of a 25 keV, 10 fA, N_2^+ horizontal beam profile measured with the MCP and a bias voltage of - 2kV.

Chapter 7

Conclusions

7.1 PhD Project Outcome

The research activities concerning electronics and computer science are an essential base to develop the control system of the SPES accelerators complex, where innovative technologies are exploited to embed the control of a wide variety of instruments. In this context, during the PhD studies, I successfully finalized the design and development of a multi-purpose Input Output Controller board capable of embedding the control of almost any equipment involved in the automatic beam transport system of the SPES project. The IOC board is a COM Express carrier board designed to host any COM Express type 6 module and is equipped with a leading semiconductor industry product like the FPGA that assists the main processor in performing real time tasks, data acquisition, and digital data processing. This embedded computer bridges the gap between custom developments and commercial personal computers. The end user can exploit the custom controller board as a general purpose pc and benefit from the I/O features set of a common motherboard. The VGA, USB, SATA, PCIe, and the 10/100/1000 Gigabit Ethernet interfaces are readily available in an off-the-shelf module. The IOC board is equipped with eight serial ports; each of the eight serial links ends on a multi-protocol transceiver that supports RS-232, RS-485, and RS-422 serial standards. WiFi and optical Ethernet are on-board connectivity solutions for those applications requiring galvanic isolation. The carrier can be supplied using the standard ATX power source, and, accommodates the PoE+ as a power solution for low-power applications. The possibility of booting a standard Linux distribution together with the $x86\_64$ architecture makes the software development and

portability from legacy technologies straightforward. The Spartan-6 XC6SLX75T device provides support for those application specific features and peripherals not commonly supported by a standard computer architecture, such as the general purpose digital and analog I/O channels, the optical 3.125 Gbps link, the CAN bus, and the standard FMC interface exploited to extend the board functionalities using commercial modules.

My contribution to the IOC design has been extensive and meaningful in all the project phases. During the first part of the work I finalized the IOC board architecture and, subsequently, I integrated the main features into a detailed electrical schematics. Then, using professional tools, I transferred the electrical schematics to the PCB layout environment. During this phase, I focused on manufacturability aspects, costs, and performances. I planned the PCB stack-up and I collected the electrical and physical constraints that represent our agreement with the PCB manufacturer and the external company in charge of routing the board. I took care of the post-layout signal and power integrity analysis necessary to validate the high-speed transmission lines and the power delivery system before releasing the production files. These post-layout simulations are of crucial importance in order to increase the board performances and the long term reliability, reducing the probability of impedance mismatch, crosstalk, and plane-cavity resonance issues. At the arrival of the first IOC board prototype, I focused on the hardware characterization. The debug of all the standard peripherals has been straightforward thanks to the COM concept. The functional validation of the most important application specific interfaces, linked to the FPGA, has been carried out alongside the development of a collection of independent and synthesizable VHDL cores. These cores will be readily integrated into future firmware application releases. I completed the hardware validation with accurate performance measurements. Some indicators of the overall performances, described in this document, are the thermal distribution, the digitization accuracy, the jitter on the main clock signals, the voltage ripple, and the eye diagrams of the high speed serial links. The minor changes and issues highlighted during the IOC board prototyping phase have been addressed with the two subsequent revisions that fell under my responsibility. I have dedicated the last months of the PhD period to the qualification of the first production lot of twenty IOC boards and their integration in the SPES beam diagnostic data readout and signal processing system.

The design, prototyping, and validation of the multi-purpose controller was supported by exhaustive training activities, provided by the University of Padova,

INFN and CERN, that offered me the opportunity to gain advanced skills in the field of analog and digital mixed PCB design, digital design and FPGA programming, power and signal integrity analysis, real time systems, computer architecture, Linux operating system, networking, and parallel computing. All these competencies have been enriched with a complete overview of the leading semiconductor products available on the market.
The installation in the field of the IOC board represents a great personal reward that crowns these years of busy time during which I turned what was just an idea in 2014, into the working embedded computer shown in Figure 7.1.
This embedded computer integrates into a compact carrier board the advances in computer technology and brings homogeneity to the hardware layer of the front-end computers of the different control subsystems at LNL where it will replace the legacy VME technology so far exploited to achieve standardization and to distribute intelligence and computational power. The advantages introduced by the IOC board are substantial in terms of costs, power consumption, maintainability,

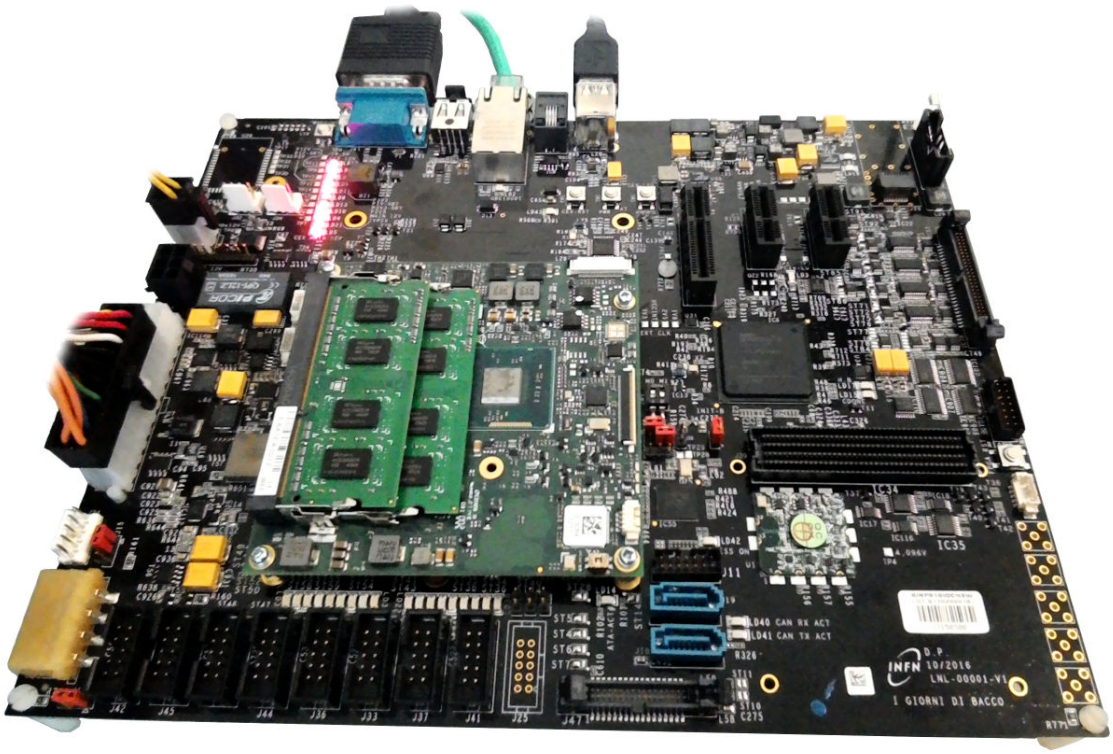

Figure 7.1: PhD work outcome: the Input Output Controller for real time distributed tasks in particle accelerators. This COM Express Type 6 carrier board behaves like a general purpose pc which is able to fulfill several application specific and time critical features normally encountered in a physical facility thanks to the on board FPGA.

performances, flexibility, real time data processing, fast response to input stimuli from the devices under control, and software development which is independent of the underlying hardware layer.
The outcome of this PhD project is an essential base in the perspective of an automatic and high quality radioactive ion beam transport system necessary to boost the research activity at LNL. The radioactive ion beams foreseen by SPES project can be used to carry out important research activities in the field of nuclear physics, nuclear astrophysics, and nuclear medicine.
Concerning the nuclear physics, the study of atomic nucleus structure of unstable nuclides is expected to provide new information on the nuclear forces and nuclear structure. At this purpose, the SPES project will produce single-specie beams of neutron-rich nuclei with mass in the range 80 to 160 with a final energy on target up to 20 MeV/u. In a second phase of the project, SPES will allow also the production of lighter ions or proton-rich nuclei.
Moreover, the production of neutron-rich and proton-rich nuclei far from stability, foreseen by SPES, is essential for addressing open questions about the nuclear physics underlying stellar nucleosynthesis. The goal of nuclear astrophysics is the study of the structure and composition of the universe and SPES may help the understanding of nuclear reactions involved in astrophysical processes at the equilibrium and in explosive environments like supernova explosions.
The primary interdisciplinary application of the SPES facility is the production of radiopharmaceuticals of interest in nuclear medicine. Nuclear medicine uses the radioactive decay of the unstable isotopes in the diagnosis and the treatment of diseases. Some promising radionuclides, already known in nuclear medicine, are ^{64}Cu, ^{89}Sr, ^{131}I, ^{125}I, ^{133}Xe, ^{90}Y.

7.2 IOC Board Applications

With reference to the frontend computer layer of the standard control system architecture introduced in section 1.4, the IOC board has been conceived to act as a local intelligent node in a distributed control system and to interface a wide variety of instruments normally encountered in a particle accelerator: electromagnetic lenses power supplies, charge breeder, beam profile monitors, Faraday cups, Allison scanners, RF controller boards, vacuum pumps, etc. In addition, the IOC board features Megahertz bandwidth closed loop control and the ability to reduce data into the FPGA. The frontend computer and part of the signal conditioning cir-

cuitry are integrated into a single carrier board that moves the sequential controls as close as possible to the physical I/O to achieve the highest performance.
The SPES project is the ideal test bench to develop the applications of the custom I/O controller. In particular, the transport of the neutron-rich ion beams (Cs, Xe, Rb, Ba, Br, Sn) foreseen by SPES requires an upgrade of the VME based current and beam position monitor system to cope with a wide variety of beam intensities and energies. The IOC board is the core of the new diagnostics data readout system and the results of the tests performed under real beam conditions prove that the new hardware extends the current sensitivity to the pA range outperforming the legacy VME technology. The SPES diagnostic system is an IOC target application that enhances the embedded computer capabilities in terms of data acquisition and real time control tasks supported by the on-board FPGA. The integration of the IOC board in the SPES control network has been straightforward thanks to the IOC architecture which facilitates the software development and portability from the legacy technology.
Another target application of the IOC board is the beam focalization. In conjunction with the beam diagnostics, a large number of electromagnetic lenses installed along the beam pipe confine and drive the low intensity radioactive beam from the ion-source chamber to the ALPI superconducting accelerator. Each focalization device needs a high voltage power supply controlled via a serial port to modulate the intensity of the electromagnetic fields that bend the trajectory of the particles inside the vacuum pipe. In this application, the IOC board will replace its precursor: the Intel Edison carrier board. One IOC serves up to eight serial peripherals in a point-to-point communication topology. Nevertheless, in RS-422 mode, each serial port may drive a multi-drop bus extending the control range up to 10 devices per serial port. The IOC board fits also the control of steerers and quadrupoles that rely on a power supply system whose output voltage and current are controlled, in a closed loop, using a couple of analog input and output channels, plus additional digital signals to switch on and off the power.
The implementation of the control of the radioactive ion beam extraction from the target-ion source chamber is similar to the beam focalization system and may exploit the IOC board.
To mention others applications, at LNL, the IOC board addresses the control requirements of the Charge Breeder and the ECR stable ion source system. The peculiarity of this last application is that the instrumentation and the frontend computer are installed inside a high voltage platform isolated from ground; the

communication with the IOC board will be done using an optical fiber. In this case, the COM Express would not be plugged into the carrier board since the optical Ethernet implementation based on the UDP protocol and the data acquisition are handled by the FPGA.
In the field of system monitoring and data visualization, the IOC board architecture fits perfectly with the concept of real time data visualization and control using modern web technologies; embedded web servers running on the COM Express module can be controlled from any browser using basic HTML technologies.
The IOC board is not intended for critical applications requiring a high safety integrity level that would be difficult to achieve with custom hardware developments and with the EPICS concept. In SPES, the hardware standardization of all those controls not subject to the functional safety concept that offers the necessary risk reduction required to achieve safety for the equipment and persons will be addressed by the IOC board. Whilst, the machine protection system and the personal protection system that rely heavily on functional safety are addressed with certificated PLCs, capable of reaching a safety integrity level of 3, according to the international standard for electrical, electronic and programmable electronic safety-related systems: IEC61508.
The cost of a single IOC board fully assembled and inclusive of a COM Express-BT module with the Intel Atom E3845 @ 1.91 GHz processor is negligible compared to legacy VME crates and modules, considering a small production lot of 20 boards.
At the moment, the IOC board is an internal INFN project, non disclosed. We are evaluating the possibility to release the project under the CERN open hardware license.

Bibliography

[1] *The INFN LNL website.* Last accessed: September 2018.
Available: http://www.lnl.infn.it/index.php/en/

[2] *The SPES project website.* Last accessed: September 2018.
Available: https://web.infn.it/spes/

[3] A. Monetti, A. Andreghetto et al, *The RIB production target for the SPES project.* Springer Berlin Heidelberg, Eur. Phys. J. A (2015) 51: 128. https://doi.org/10.1140/epja/i2015-15128-6.

[4] A. Monetti, *The mechanical design and the realization of the new Wien Filter.* Ph.D. thesis, Design and development of the target-ion source system for the SPES project, chapter 4.6, Padova, 2017.

[5] A. Galatà et al, *Status of the SPES-charge breeder (SPES-CB) and its beam line at INFN-LNL.* Nuclear Instruments and Methods in Physics Research, Section B: Beam Interactions with Materials and Atoms, Vol. 376, June 2016. DOI: 10.1016/j.nimb.2015.12.031

[6] M.E. Thuot, L.R. Dalesio, *Control System Architecture: The Standard and Non-Standard Models.* IEEE, Proceedings of International Conference on Particle Accelerators, May 1993, Washington USA, DOI: 10.1109/PAC.1993.309138.

[7] D. Bortolato, S. Pavinato et al, *Upgrade of the LLRF Control System at LNL.* JACoW Proceedings of ICALEPCS2017, Barcelona, 2017, DOI: 10.18429/JACoW-ICALEPCS2017-TUPHA117.

[8] *The EPICS website.* Last accessed: September 2018.
Available: https://epics.anl.gov/

[9] D. Pedretti et al, *Custom Hardware Platform Based on Intel Edison Module.* JACoW Proceedings of ICALEPCS2015, Melbourne, 2015, DOI: 10.18429/JACoW-ICALEPCS2015-WEM307

[10] *The Yocto Project website.* Last accessed: September 2018. Available: https://www.yoctoproject.org/

[11] PICMG, *COM Express overview website.* Last accessed: September 2018. Available: https://www.picmg.org/openstandards/com-express/

[12] PCI-SIG, *PCI Express Base Specification.* Revision 3.1, October 2014.

[13] ICE Systems, *Proxmox Virtual Environment.* Datasheet, 2016, Vienna, Austria.

[14] PICMG, *COM Express Carrier Design Guide.* Rev. 2.0, December 2013.

[15] BOSCH, *CAN Specification.* Version 2.0, September 1991.

[16] ANSI/VITA, *American National Standard for FPGA Mezzanine Card (FMC) Standard.* ANSI/VITA 57.1-2008.

[17] Xilinx®,*iMPACT User Guide.* Last accessed: June 2018. Available: http://web.ee.nchu.edu.tw/ cpfan/FY92a-baseband/iMPACT-guide.pdf

[18] Xilinx®, *Platform Cable USB II.* DS593 (v1.5), June 2015.

[19] Xilinx®, *Spartan-6 FPGA Data Sheet: DC and Switching Characteristics.* DS162 (v3.1.1), January 2015.

[20] Xilinx®, *Xilinx Power Estimator User Guide.* UG440 (v2017.2), June 2017.

[21] Intel, *ATX Specification.* Version 2.03, December 1998.

[22] ADLINK Technology Inc., *cExpress-BT User's Manual.* revision 0.20 preliminary, May 20, 2014.

[23] Microsemi, *Understanding 802.3at.* PoE Plus Standard Increases Available Power, June 2011.

[24] Isola, *FR408HR High Performance Laminate and Prepreg.* Data Sheet, Rev. B, April 2018. Available: https://www.isola-group.com/wp-content/uploads/data-sheets/fr408hr.pdf

[25] Xilinx®, *Recommended Design Rules and Strategies for BGA Devices.* User Guide, UG1099 (v1.0), March 2016.

[26] ANSYS™, *ANSYS SIwave website.* Last accessed: September 2018. Available: https://www.ansys.com/products/electronics/ansys-siwave

[27] Hewlett-Packard Company, Intel Corporation, Microsoft Corporation, Renesas Corporation, ST-Ericsson, Texas Instruments, *Universal Serial Bus 3.1 Specification.* Revision 1.0, July 26, 2013.

[28] J. N. Tripathi, *Selection and Placement of Decoupling Capacitors in High Speed Systems.* IEEE Electromagnetic Compatibility Magazine, vol. 2, October 2013, DOI: 10.1109/MEMC.2013.6714703.

[29] A. Rushton, *VHDL for Logic Synthesis.* Third Edition, April 2011, ISBN-13: 978-0470688472.

[30] P. Loschmidt et al, *Guidelines for VHDL Coding.* 30-09-2010.

[31] J. Serrano, *Digital Signal Processing Using Field Programmable Gate Arrays.* Proceedings of BIW08, Beam Instrumentation Workshop, Tahoe City, California, May 2008.

[32] OpenCores, *OpenCores website.* Last accessed: September 2018. Available: https://opencores.org/?do=svn

[33] OpenCores, *WISHBONE System-on-Chip (SoC) Interconnection Architecture for Portable IP Cores.* revision B4, 2010. Available: https://cdn.opencores.org/downloads/wbspec_b4.pdf

[34] Xilinx®, *Spartan-6 FPGA GTP Transceivers.* Advance Product Specification, UG386, Rev 2.2, April 2010.

[35] Xilinx,®, *LogiCORE IP Tri-Mode Ethernet MAC v4.5.* User Guide, UG138, March 2011.

[36] CERN, *Welcome to the IPbus documentation!.* Last accessed: September 2018. Available: https://ipbus.web.cern.ch/ipbus/doc/user/html/

[37] C. Ghabrous Larrea et al, *IPbus: a flexible Ethernet-based control system for xTCA hardware.* JINST 10 (2015) no.02, C02019. DOI: 10.1088/1748-0221/10/02/C02019.

[38] IEEE standards association, *IEEE Standard for Terminology and Test Methods for Analog-to-Digital Converters.* IEEE Std 1241™-2010, Revision of IEEE Std 1241-2000, DOI: 10.1109/IEEESTD.2011.5692956.

[39] MathWorks®, *Matlab website.* Last accessed September 2018. Available: https://it.mathworks.com/products/matlab.html

[40] Analog Devices, *Taking the Mystery out of the Infamous Formula, "SNR = 6.02N + 1.76dB," and Why You Should Care.* MT-001, TUTORIAL, Rev A, 2009.

[41] Analog Devices, *Understand SINAD, ENOB, SNR, THD + N, and SFDR so You Don't Get Lost in the Noise Floor.* MT-003, TUTORIAL, Rev A, 2009.

[42] Silicon Labs, *Improving ADC Resolution by Oversampling and Averaging.* AN118, Rev. 1.3 July 2013.

[43] Actel®, *Improving ADC Results Through Oversampling and Post-Processing of data.* January 2007. Available: https://www.microsemi.com/document-portal/doc_view/131569-improving-adc-results-white-paper.

[44] Atmel®, *AVR120: Characterization and Calibration of the ADC on an AVR.* Application note, AVR 8-bit Microcontrollers , Rev. 2559E, August 2016.

[45] CAENels, *4-channel 20 bit 1 MSPS FMC Floating Ammeter.* User's Manual, Rev. 1.2, July 2015.

[46] C. E. Cummings, *Simulation and Synthesis Techniques for Asynchronous FIFO Design.* SNUG, San Jose 2002, Rev 1.2

[47] SiTime™, *Clock Jitter Definitions and Measurement Methods.* SiT-AN10007 Rev 1.2, January 2014.

[48] E. Bogatin, *TDR and VNA Techniques for PCB Characterization.* GigaTest Labs, Sunnyvale, CA.
Available: https://www.keysight.com/upload/cmc_upload/All/GTL91.pdf.

[49] D. J. Dascher, *Measuring Parasitic Capacitance and Inductance Using TDR.* Article 11, Hewlett-Packard Journal, August 1996.

[50] Keysight Technologies™, *N7020A Power Rail Probe for Power Integrity Measurements, N8846A Power Integrity Analysis Application.* Data Sheet.

[51] ON Semiconductor®, *Understanding Data Eye Diagram Methodology for Analyzing High Speed Digital Signals.* Application note, AND9075/D, Rev. 1, June 2015.

[52] Xilinx®, *ChipScope Integrated Bit Error Ratio Test (IBERT) for Spartan-6 GTP.* Product Specification, DS782, Rev 2.02a, October 2011.

[53] Xilinx®, *ChipScope Pro Tutorial: Using IBERT with ChipScope.* UG811, Rev 14.5, March 2013.

[54] Laboratori Nazionali di Legnaro, *Ion Beams Available at the LNL XTU Tandem.* UG811, Servizio Sorgenti e Iniettori home page, updates on December 2014, Available: https://www.lnl.infn.it/ ssi/Ion_Beams.htm

[55] A. Lombardi et al, *The New Positive Ion Injector Piave at LNL.* UG811, Proceedings of Particle Accelerator Conference, 1997, DOI: 10.1109/PAC.1997.749951

[56] J. Vasquez, M. Poggi et al, *Beam Diagnostics for the SPES Project.* Annual Report, Legnaro, Italy, 2013.

[57] Rogers corporation, *RO3000® Series Circuit Materials.* Available: https://www.rogerscorp.com/documents/722/acs/RO3000-Laminate-Data-Sheet-RO3003-RO3006-RO3010.pdf

[58] A. Monetti, *Design and development of the target-ion source system for the SPES project.* PhD thesis, University of Padova, Italy, 2016.

[59] P. Forck, *Lecture Notes on Beam Instrumentation and Diagnostics.* Joint University Accelerator School, Archamps, January - March 2011.

[60] Goodfellow, *Goodfellow website.* Available: http://www.goodfellow-ceramics.com/products/ceramics/shapal-hi-m/characteristics/

[61] Analog Devices, *Low Power, Rail-to-Rail Output, Precision JFET Amplifier.* Data Sheet, Rev. F.

[62] M. Bellato et al, *The Beam Diagnostics System of the ALPI Post-Accelerator.* Proceedings of the 3rd EPAC, Berlin, 1992.

[63] Analog Devices, *Moving Average Filters.* dsp-book, chapter 15.

[64] D. Pedretti et al, *The New Beam Diagnostic Data Readout and Signal Processing System at LNL.* IEEE Transactions on Nuclear Science 10.1109/TNS.2018.2808220, February 2018.

[65] M. Giacchini, G. Bassato et al, *Upgrade of Beam Diagnostics System of Alpi-Piave Accelerator's Complex at LNL.* Proceedings of PCaPAC2014, Karlsruhe, Germany, 2014.

[66] F. Galtarossa, M. Poggi et al, *A Beam Profile Monitor for SPES Low-intensity Radioactive Beams.* Annual Report, 2017, Legnaro.

Appendix A

Python Code - Wire Temperature Profile

```
#!/usr/bin/python
# -*- coding: utf-8 -*-
# Authors: Damiano Bortolato, Davide Pedretti, Claudio Fanin
# Wire Temperature Profile Characterization

from pylab import *
import matplotlib.pyplot as plt
import numpy as np

# tungsten
D_w = 19250
Cs_w = 130
rhot_w = 1./174 # 300 K

sigmab=5.67e-8 # Stefan Boltzmann constant
epsilon=0.1 # gold emissivity

# wire geometry
d = 50e-6 # m
S = pi*d**2/4 # m2
L = .1

# material
Cs_mat = Cs_w
D_mat = D_w
rhot_mat = rhot_w

```

```
28 # finite element
29 dx = 1e-3 # m
30 dV = dx*S  # m3
31 N=int(float(L)/dx) # number of finite elements
32 Ct_dx = Cs_mat*dV*D_mat # Thermal Capacity
33 R_0_dx = array([rhot_w*dx/S]*(N-2)) # Thermal Resistance vector
34
35 # Beam
36 E = 20e6 # beam energy
37 Iw = 10e-9 # current per wire
38 var_beam = 2 # beam variance
39
40 Pgx = array([0.0]*(N-2)) # beam power density distribution
41 for i in range(N-2):
42     Pgx[i] = ((Iw*E)/(var_beam*sqrt(2*pi)))* np.exp(-np.power(i-(N-2)
       /2, 2.)/(2*np.power(var_beam, 2.)))
43
44 # Initial conditions and constraints
45 Th = 300
46 Tl = 300
47 Tx = (N-array(range(N),dtype=float))*(Th-Tl)/(N-1)+Tl # temperature
       vector initialization
48 T0 = array([300]*(N-2))
49
50 # simulation
51 dt = .00001 # time interval
52 St = 5 # total simulation time
53 Ni = int(St/dt) # number of iterations
54
55 print "total power on the wire %.2f W" % sum(Pgx)
56
57 # plot settings
58 x=arange(N)*dx
59 f, (ax1,ax2,ax3,ax4) = plt.subplots(4, sharex=True, sharey=False)
60 f.suptitle("Wire Temperature Profile. Total Power on Wire: %.2f W.
       Beam Variance: %d. Wire Diameter: %.6f m" % (sum(Pgx), int(
       var_beam),d))
61 xlabel("wire position [m]")
62 ax1.set_ylabel("Wire Temperature [K]")
63 ax2.set_ylabel("Beam Power Density [W/m]")
64 ax3.set_ylabel("Radiation Heat Transfer [W/m]")
65 ax4.set_ylabel("Conduction Heat Transfer [W/m]")
66
```

```
67 # iterations
68 for i in range(Ni):
69     Pirr = (Tx[1:-1]**4-T0**4)*sigmab*epsilon*dx*d*pi # element power
       loss by radiation
70     R_v_dx = R_0_dx+(Tx[1:-1] - T0)*2.3529e-6*dx/S # updating the
       element thermal resistance vector assuming a linear dependance on
       T
71     Pcon = (Tx[2:]+Tx[:-2]-2*Tx[1:-1])/R_v_dx # conduction power
       transfer between neighboring elements
72     Px = Pcon - Pirr + Pgx # element net power
73     DT = Px*dt/Ct_dx
74     Tx[1:-1] += DT # element new temperature
75     if i % (Ni/10) == 0:
76         print dt*i, max(abs(DT))
77         ax1.plot(x, Tx)
78         ax2.plot(x[1:-1],Pgx)
79         ax3.plot(x[1:-1],-Pirr)
80         ax4.plot(x[1:-1],Pcon)
81
82 print "Total power radiated at steady-state: %.2f W" % sum(Pirr)
83
84 plt.setp([a.get_xticklabels() for a in f.axes[:-1]], visible=False)
85
86 print "Tmax %.1f @pos %.2f" % (max(Tx), argwhere(Tx==max(Tx))[0]*dx)
87
88 show()
```

Appendix B

Matlab Code - ENOB Calculation

ENOB Calculation from SNR

```
close all
clear all

%% Open file, import signed integers and baseline plot
fileID = importdata('ENOB_CH4.dat');
% Find the last two's power in data lenght
N = 2^(floor(log2(length(fileID))));
baseline = fileID(1:N);
time = (1:N)*1.45e-3;
plot(time,baseline);
xlabel('time [ms]'), ylabel('sampled baseline');
title('baseline');

%% Data histogram
bin = [];
[N,edges] = histcounts(baseline);
for j = 1:(length(edges)-1)
    bin(j) = round((edges(j)+edges(j+1))/2);
end
scatter(bin,N,'filled');

%% Gaussian fit
f = fit(bin.',N.','gauss1');
plot(f,bin,N);
xlabel('sampled value'), ylabel('count');
```

```
title('Gaussian fit');
param=coeffvalues(f);
average=param(2);
ADC_rms_noise=param(3)/sqrt(2);

%% SNR
nbits = 16;
SNR=20*log10(((2^(nbits))/(2*sqrt(2)))/ADC_rms_noise);

%% ENOB
ENOB = (SNR-1.76)/6.02;
```

ENOB Calculation from SINAD

```
close all
clear all

%% Open file and import signed integers
fileID = importdata('ENOB_CH4.dat');
% Find the last two's power in data lenght
N = 2^(floor(log2(length(fileID))));
noise = fileID(1:N);

%% ADC Parameters
Fs = 690e3; % Sampling Frequency
fullScale = 32767; % Full scale value
% Calculate values for frequency axes
freq = linspace(0,Fs,N)';

%% Sinusoidal tone
T = 1/Fs;
F = 100e3;
A = fullScale*0.99;
t = ((0:N-1)*T)';
tone = A*sin(2*pi*F*t);
%% Adding quantization noise
% We add white noise to the ideal sinusoidal tone.
% The coefficient 0.29 is calculated to obtain a sine tone
% with an ENOB of 16. This is equal to the theoretical
% quantization noise of a 16-bits ADC.
```

```
x = tone + (0.29*randn(1,N))';

%% Adding analog system noise (baseline noise)
signal = x + noise;

%% Normalize data
normData = signal / fullScale;

%% Apply Window
wData = normData.* hanning(N);

%% Calculate bilateral PSD
PSD = abs((1/N) * fft(wData)).^2;

%% Calculate SINAD
SINAD = sinad(wData,Fs);
ENOB = (SINAD - 1.76) / 6.02;

plot(freq,10*log10(PSD));
xlabel('frequency [Hz]'), ylabel('PSD [dB]');
title('Power Spectral Density');
display(SINAD);
display(ENOB);
hold on;
```