THÈSE POUR OBTENIR LE GRADE DE DOCTEUR DE L'UNIVERSITÉ DE MONTPELLIER

En Électronique

École doctorale Information, Structures et Système (I2S)

Unité de recherche IES Montpellier, UMR-CNRS 5214

Conception et Intégration d'une liaison sans fil tolérant aux radiations pour l'Equipment de contrôle dans les accélérateurs de particules du CERN

Design and development of a radiation tolerant wireless link for the control equipment of CERN particle accelerators

Présentée par Antonio Scialdone

Le 17 juin 2025

Sous la direction de Luigi Dilillo, Jérôme Boch et Frédéric Saigne

Devant le jury composé de

[Daniela Dragomirescu, Professeur INSA, Toulouse]	[Présidente du jury, Rapporteure]
[Almudena Lindoso, Professeur, Universidad Carlos III de Madrid]	[Rapporteure]
[Luigi DILILLO, Docteur, Université de Montpellier, IES]	[Directeur de thèse]
[Jérôme BOCH, Professeur, Université de Montpellier, IES]	[Co-encadrant de thèse]
[Frédéric SAIGNE, Professeur, Université de Montpellier, IES]	[Co-encadrant de thèse]
[Salvatore DANZECA, Docteur, CERN]	[Invité]
[Rudy FERRARO, Docteur, CERN]	[Invité]

UNIVERSITÉ DE MONTPELLIER

Abstract

The Large Hadron Collider at CERN requires thousands of instruments to operate reliably: radiation monitors, power converters, quench protections, beam instrumentations, and more. All of them are operated through a wired control network, which guarantees reliable and deterministic communication. Nevertheless, the instruments are distributed all along the accelerator tunnel, requiring more than 400 kilometres of cables. In such a scenario, deployment is expensive and time consuming, requiring months of work for deploying a new system. In fact, the cable trays are currently full, limiting new installations. These challenges will be even more difficult to overcome for future particle accelerators, like the Future Circular Collider envisioned by CERN, which is planned to be 100 kilometres long. The required number of instruments and consequently, the cable length, will increase dramatically. A wired control network for such a scenario is impractical. Wireless communications could bring many advantages: reduced cabling, cheaper infrastructure and faster installations. However, current wireless solutions are inadequate for the CERN accelerator. Electronic systems are subject to radiation effects, which could impact the reliability of the wireless link, making instruments unreachable. Therefore, commercial off the shelf devices must be qualified under radiation before being used. Moreover, current wireless standards do not fit within the requirements of control applications that require very low latency. The objective of this research is to advance the state of the art of wireless networks in the control network of CERN particle accelerator. Firstly, a methodology is proposed to evaluate the effects of radiations on radio frequency front-ends—the core of any wireless system—and their impact on communication links. Then, to address the limitations of current wireless standards, a custom wireless physical layer is being developed at CERN and presented in this work, which leverages the flexibility of Software-Defined Radio. This approach not only highlights the advantages of achieving high radiation tolerance but also demonstrates the significant benefits of a tailored physical layer. By utilizing Software Defined Radios, first a solution is proposed to extend the existing wired infrastructure of CERN control network. Secondly, a custom physical layer capable of overcoming the limitations of current wireless standards is shown. Designed to meet the stringent requirements of control applications, the solution exhibits robust radiation tolerance, achieves more than 99% availability with respect to the total run time of the accelerator, and proves to be a promising foundation for building wireless networks in future particle accelerators.

Keywords

Wireless Networks, Radiation to Electronics effects, Wireless Fieldbuses, Software Defined Radio, Custom Wireless Physical Layer

Résumé

Les communications sans fil sont devenues fondamentales dans notre vie quotidienne. Leur succès fulgurant s'explique par les nombreux avantages qu'elles offrent en connectant des appareils sans les contraintes des câbles. Par conséquent, un intérêt croissant se manifeste pour l'utilisation des technologies sans fil dans divers secteurs au-delà de l'électronique grand public, notamment dans les usines, les installations industrielles, et même dans des environnements extrêmes tels que les centrales nucléaires ou les accélérateurs de particules. Cet intérêt repose sur la possibilité d'améliorer l'efficacité et l'accessibilité des données tout en s'affranchissant des contraintes liées au câblage. Dans des environnements où les connexions filaires traditionnelles sont peu pratiques, les solutions sans fil pourraient offrir une flexibilité et une adaptabilité incomparables. Le Centre Européen de Recherche Nucléaire (CERN) illustre bien ce type de scénario. Au CERN, les physiciens explorent les fondamentaux de l'univers en accélérant des faisceaux de particules dans une chaîne d'accélérateurs. Le Large Hadron Collider (LHC) constitue le dernier accélérateur de cette chaîne, où les faisceaux sont accélérés jusqu'à une énergie de 7 TeV. Ces faisceaux entrent en collision en quatre points d'interaction, produisant des gerbes de particules analysées pour découvrir de nouvelles particules fondamentales. Des milliers d'instruments, tels que les convertisseurs de puissance, les dispositifs de surveillance des radiations et les systèmes cryogéniques, sont nécessaires pour garantir le fonctionnement fiable de l'ensemble du complexe. Actuellement, ces instruments sont surveillés et contrôlés via un système basé sur le réseau FieldBus, un réseau câblé qui permet un contrôle distribué en temps réel. Les FieldBus, introduits dans les années 1980, avaient pour objectif de relier les niveaux supérieurs d'une usine (interfaces utilisateur) aux niveaux inférieurs (capteurs et actionneurs). Depuis, plusieurs normes ont été développées pour répondre aux exigences de différentes applications. Au CERN, plusieurs FieldBus sont utilisés dans l'architecture de contrôle, le plus répandu étant le WorldFIP, qui connecte plus de 10 000 nœuds sur plus de 100 km de câbles.

Cependant, l'installation de câbles dans le tunnel du LHC est une opération coûteuse et nécessite une planification minutieuse. Elle offre peu de flexibilité lorsqu'il faut déplacer des équipements. Plus récemment, de nouvelles installations se sont révélées encore plus complexes, les chemins de câbles ayant presque atteint leur capacité maximale. Par ailleurs, le projet de Future Circular Collider (FCC), proposé par le conseil du CERN comme prochaine génération d'accélérateur de particules, prévoit un tunnel de 100 km de long. Considérant cette longueur et le probable accroissement du nombre d'instruments, les connexions filaires deviennent peu pratiques. En revanche, les solutions sans fil, en cours d'évaluation, pourraient offrir de nombreux avantages : des coûts d'installation réduits grâce à l'absence de câblage et une flexibilité accrue pour l'installation, les mises à niveau et les déplacements d'équipements. Néanmoins, l'adoption des technologies sans fil pose des défis spécifiques. L'un des principaux réside dans la moindre fiabilité des systèmes de communication sans fil, ce qui complique leur

usage pour des applications de contrôle où la fiabilité est cruciale. Les systèmes de communication modernes utilisent généralement des techniques de correction pour atténuer les perturbations naturelles, améliorant le débit et la fiabilité au prix d'une latence accrue. Cependant, dans les applications de contrôle nécessitant un fonctionnement en temps réel, une très faible latence est indispensable. Les protocoles sans fil standards utilisés dans l'électronique grand public ne conviennent donc pas à ces scénarios.

Des efforts sont néanmoins déployés pour répondre aux exigences de la quatrième révolution industrielle, également appelée Industrie 4.0. Cette révolution envisage des usines intelligentes, intégrant des technologies internet à faible latence et haute fiabilité pour permettre des applications nécessitant des temps de réaction extrêmement courts. L'échange d'informations avec les capteurs et actionneurs de manière fiable et rapide est au cœur de cette transformation. Pour répondre à cette demande, les bus de terrain sans fil ont émergé comme une solution standardisée, conçue pour offrir une fiabilité ultra-élevée dans les environnements industriels. Cependant, ces solutions ciblent principalement les applications de surveillance, avec une latence typiquement de l'ordre de quelques millisecondes, et non le contrôle en temps réel qui exige une latence inférieure à une milliseconde.

Toutefois, l'adoption de nouvelles technologies dans un environnement comme le CERN est loin d'être immédiate. Les équipements électroniques dans le LHC sont exposés aux effets des radiations, qui peuvent perturber leur fonctionnement ou entraîner des défaillances. Par exemple, un appareil sans fil pourrait subir des erreurs de communication ou des pertes de signal, rendant les instruments inaccessibles au système de contrôle. Dans les cas extrêmes, les radiations peuvent provoquer des défaillances destructrices. Bien que des techniques spécifiques permettent de concevoir des équipements résistants aux radiations, le CERN privilégie le développement de systèmes basés exclusivement sur des composants standards (COTS) pour des raisons de coûts. Ces composants sont donc largement testés dans le cadre du programme d'assurance de résistance aux radiations du CERN. Toutefois, les systèmes COTS intégrant des fonctionnalités sans fil, comme le Wi-Fi ou la 5G, sont intrinsèquement complexes, tant sur le plan matériel que logiciel. Cette complexité rend difficile l'évaluation des effets des radiations sur ces systèmes et la proposition de stratégies d'atténuation efficaces.

Dans cette optique, l'objectif principal de cette thèse est d'explorer l'intégration des technologies sans fil dans les accélérateurs de particules actuels et futurs du CERN, tout en relevant les défis mentionnés. Cela nécessite plusieurs étapes : a) analyser le système de contrôle actuel du CERN pour identifier les exigences des applications et les contraintes physiques ; b) développer une solution sans fil pour résoudre les problèmes de câblage à court terme ; c) évaluer l'impact des radiations sur la liaison sans fil et proposer des techniques d'atténuation pour améliorer sa fiabilité ; d) utiliser ces résultats pour concevoir une liaison sans fil résistante aux radiations, servant de base pour des systèmes futurs dans des environnements irradiés.

L'infrastructure actuelle du CERN impose différentes contraintes sur les technologies sans fil qui pourraient être adoptées. Premièrement, compte tenu de la complexité de l'accélérateur et de la longévité des machines, une solution fixe a été écartée au profit d'une solution basée sur la Software-Defined Radio (SDR). Le concept de SDR consiste à implémenter la plupart des fonctionnalités sans fil dans la logique programmable, tout en minimisant le matériel fixe, comme la partie Radio Fréquence (RF) et la conversion analogique/numérique. Théoriquement, cela permet d'implémenter n'importe quel protocole sans fil simplement en reconfigurant la logique programmable. Dans un contexte comme celui du CERN, cela

représente un grand avantage, car changer de protocole devient beaucoup plus facile et moins coûteux. Deuxièmement, les équipements des accélérateurs, qui sont exposés à des radiations, sont fabriqués sur mesure par différents groupes au CERN. Ainsi, redessiner ces systèmes pour les rendre compatibles avec une norme sans fil n'est pas une option, car cela serait trop coûteux et chronophage. Cette contrainte a nécessité le développement d'une solution ad hoc pouvant être ajoutée à l'infrastructure actuelle du Fieldbus. L'évaluation a mis en évidence certains inconvénients, qui pourraient être surmontés en développant une solution sans fil spécifiquement adaptée à l'application. Par conséquent, une solution sans fil personnalisée a été développée. Il s'agit d'un lien sans fil basé sur le SDR, construit à partir de composants COTS tolérants aux radiations, testés sous irradiation dans le cadre de cette thèse de doctorat. Le système développé implémente toutes les fonctionnalités de base pour transmettre et recevoir des données brutes entre deux systèmes de manière sans fil, avec une latence inférieure à la milliseconde. Son développement représente une étape importante vers l'utilisation des technologies sans fil pour les applications de contrôle dans les accélérateurs de particules. En effet, il peut servir de plateforme de base pour implémenter de futurs protocoles de communication sans fil, avec l'avantage principal de fournir déjà des caractéristiques de tolérance aux radiations. En outre, étant basé sur le SDR, il peut être modifié en fonction des besoins pour répondre à différentes exigences, telles qu'une latence plus faible. Ces activités sont décrites tout au long du manuscrit, qui est organisé comme suit.

Le chapitre 1 décrit le complexe du CERN, en mettant particulièrement l'accent sur le LHC et son environnement radiatif auquel les systèmes électroniques sont soumis. Les mécanismes fondamentaux des effets des radiations sur les électroniques y sont également expliqués.

Le chapitre 2 présente les bases des réseaux de communication pour les systèmes de contrôle en décrivant les exigences principales qu'un réseau de contrôle doit remplir et ses caractéristiques, tout en soulignant la différence avec les réseaux traditionnels utilisés à domicile ou au bureau. La deuxième partie du chapitre décrit l'infrastructure de contrôle du LHC et les exigences des applications, qui sont fondamentales pour concevoir une solution sans fil. En plus des exigences standard d'un réseau de contrôle, celui destiné au LHC doit être tolérant aux radiations, étant donné qu'une partie de l'infrastructure de contrôle est exposée à des zones irradiées.

Le chapitre 3 explore la caractérisation radiative d'un circuit intégré radiofréquence, qui constitue le composant central de tout système sans fil. Une méthodologie pour étudier les effets des radiations sur ces dispositifs est discutée en prenant l'émetteur-récepteur RF AD9361 comme exemple. Les résultats des tests de radiation sont présentés, ainsi que l'impact des effets des radiations sur un éventuel système de communication basé sur ce composant. Grâce à sa tolérance aux radiations, ce composant a été sélectionné comme cœur des solutions sans fil pour le système de contrôle du LHC développées dans le reste du manuscrit.

Le chapitre 4 présente une recherche bibliographique sur les tentatives passées visant à étendre un FiledBus avec des technologies sans fil, en montrant les défis rencontrés par les projets précédents avec le même objectif. Ensuite, sur la base des résultats de cette recherche, une liaison sans fil personnalisée et tolérante aux radiations est discutée, et son développement est décrit dans la seconde moitié du chapitre. La solution proposée permet d'étendre l'infrastructure actuelle du CERN avec des fonctionnalités sans fil, tout en restant complètement transparente pour l'architecture de contrôle du CERN. La mise en œuvre matérielle a été testée sur du matériel réel pour démontrer son principe de fonctionnement. En

outre, les résultats des évaluations en laboratoire sont présentés. Le chapitre se termine par une discussion sur les avantages et les inconvénients de l'utilisation d'une telle solution pour l'architecture du CERN.

Pour surmonter les limites imposées par l'architecture du FieldBus, une nouvelle solution spécifiquement adaptée aux exigences des applications du LHC a été proposée et explorée dans le chapitre 5. La première partie discute de la caractérisation radiative des blocs de construction nécessaires pour construire un système sans fil complet, étant donné que la tolérance aux radiations est d'une importance primordiale pour l'interface de contrôle. Suite à la sélection des composants, la plateforme sans fil complète est présentée, et sa caractérisation radiative est discutée. Enfin, la fiabilité de ce système est analysée dans le contexte du système de contrôle du LHC.

Contents

LIST OF ACRONYMS **11**

INTRODUCTION **16**

1 RADIATION EFFECTS ON ELECTRONICS AND IMPACT ON CERN PARTICLE ACCELERATORS **20**

1.1 BASIC MECHANISMS OF RADIATION EFFECTS 21
1.1.1 Total Ionizing Dose *22*
1.1.2 Displacement Damage (DD) *24*
1.1.3 Single Event Effects *25*
1.1.4 Soft errors mitigation techniques *28*
1.2 LHC LAYOUT AND RADIATION ENVIRONMENT 29
1.2.1 Layout *29*
1.2.2 Operating principle *31*
1.2.3 Radiation effects impact on LHC operation *32*
1.2.4 Radiation sources *33*
1.2.5 Current and future radiation levels *34*
1.3 RADIATION TEST FACILITIES 36
1.3.1 PSI Proton Irradiation Facility *37*
1.3.2 TRIUMF Neutron Facility *38*
1.3.3 CERN Co60 Facility *39*
1.3.4 CERN CHARM facility *39*
1.4 CHAPTER SUMMARY 41

2 CONTROL NETWORKS IN THE LHC ACCELERATOR **42**

2.1 COMMUNICATION NETWORKS FOR CONTROL SYSTEM 43
2.1.1 Industrial Control Network *44*
2.1.2 Wired control network architecture *47*
2.1.3 Fieldbus *47*
2.2 THE LHC CONTROL SYSTEM 48
2.2.1 Architecture *48*
2.2.2 WorldFIP fieldbus for slow control of LHC equipment *50*
2.2.3 LHC equipment requirements *51*
2.2.4 Summary requirements *53*
2.3 CHAPTER SUMMARY 54

3 STUDY OF RADIATION EFFECTS ON WIRELESS COMMUNICATION **55**

3.1 BACKGROUND ON WIRELESS COMMUNICATION 56
3.1.1 Source encoding *57*
3.1.2 Channel encoding *57*
3.1.3 Digital modulation *57*
3.1.4 Receiver synchronization *63*
3.2 RADIATION TESTING METHODOLOGIES FOR RF INTEGRATED CIRCUITS 70
3.2.1 Radio Frequency Front-End *70*
3.2.2 Modern wireless architecture: Software-defined Radio *71*
3.2.3 Review on radiation effects on RF circuits *73*
3.2.4 Using standard metrics to evaluate radiation effects on RFIC *75*

3.3 Exploring radiation effects on a RFIC 78
3.3.1 Device under test *78*
3.3.2 Application under test *79*
3.3.3 Radiation test setup *82*
3.3.4 Radiation test results *83*
3.3.5 Discussion *88*
3.4 Chapter summary 91

4 A RADIATION-TOLERANT WIRELESS EXTENSION FOR THE WORLDFIP LHC CONTROL NETWORK 92

4.1 Integration of wireless technologies into fieldbus systems 93
4.1.1 Challenges *93*
4.1.2 Past attempts *94*
4.1.3 Summary *97*
4.2 Extending the WorldFIP LHC control network using SDR 98
4.2.1 The WorldFIP protocol *98*
4.2.2 Design choices *101*
4.2.3 Repeater extension *103*
4.2.4 Gateway extension *111*
4.2.5 Discussion *112*
4.3 Implementation and characterization of a radiation tolerant repeater extension 113
4.3.1 CRaTeBo: A radiation tolerant testing platform for FPGA-based systems *113*
4.3.2 Repeater extension: component selection and platform description *116*
4.3.3 Neutron radiation testing *117*
4.3.4 Proton radiation testing *121*
4.4 Discussion 124
4.4.1 Link availability and lifetime across the LHC *126*
4.5 Chapter summary 127

5 DESIGN AND DEVELOPMENT OF A WIRELESS PHYSICAL LAYER FOR CONTROL APPLICATIONS 128

5.1 State-of-the-art wireless control networks in radiation environments 129
5.2 Design and evaluation of a custom wireless PHY for LHC control applications 132
5.2.1 Hardware platform *132*
5.2.2 Gateware implementation *133*
5.2.3 Software description *135*
5.2.4 Platform working principle *136*
5.2.5 Performance characterization of the wireless PHY *136*
5.3 Radiation characterization of the wireless physical layer 137
5.3.1 Test methodology *138*
5.3.2 Test results *140*
5.3.3 Results analysis and discussion *142*
5.4 Chapter summary 145

CONCLUSIONS 146

LIST OF FIGURES 151

LIST OF TABLES 155

BIBLIOGRAPHY 157

List of Acronyms

ADC Analog-to-Digital Converter

AGC Automatic Gain Control

AM Amplitude Modulation

API Application Programming Interface

ASIC Application Specific Integrated Circuit

BER Bit Error Rates

BLM Beam Loss Monitors

B-PAM Binary Pulse Amplitude Modulation

CAN Controlled Area Network

CCC CERN Control Center

CCR CERN Control Room

COTS Commercial off the Shelf

CRC Cyclic-Redundancy-Check

CSMA Carrier Sense Multiple Access

DAC Digital-to-Analog Converter

DD Displacement Damage

DPPM Differential Pulse Position Modulation

DS Dispersion Suppressor

DSSS Direct Sequence Spread Spectrum

DUT Device Under Test

DWC Duplication with Comparison

ECC Error Correction Codes

ESA European Space Agency

EVM Error Vector Magnitude

FCC Future Circular Collider

FCS Frame Check Sequence

FECs Front-End Computers

FGC Function Control/Generator

FES Frame End Sequence

FIP Factory Instrumentation Protocol

FIR Finite Impulse Response

FM Frequency Modulation

FPGA Field-Programmable Gate Array

FSS Frame Start Sequence

GEFE GBT-based Expandable Front End

GPIO General-Purpose Input/Output

GUI Graphical User Interface

HeH High Energy Hadron

HL-LHC High-Luminosity LHC

HPC-FMC High-Pin Count FPGA Mezzanine Card

ICN Industrial Control Network

IEL Ionizing Energy Loss

IIoT Industrial Internet of Things

IF Intermediate Frequency

IP Internet Protocol

IoT Internet of Things

IR Insertion Regions

ISA International Society of Automation

LAN Local Area Network

LET Linear Energy Transfer

LHC Large Hadron Collider

LLC Logical Link Control

LNA Low Noise Amplifier

LOs Local Oscillators

LSS Long Straight Section

LUT Look Up Table

MAC Medium Access Control

MM Muller and Mueller

MOS Metal-Oxide Semiconductor

MRAM Magneto-Resistive Random-Access Memory

NCS Networked Control Systems

NIEL Non Ionizing Energy Loss

NR New Radio

OFDM Orthogonal Frequency-Division Multiplexing

OHWR Open Hardware Repository

PHY Physical Layer

PIF Proton Irradiation Facility

PL Programmable Logic

PLL Phase-Locked Loop

POR Power-On-Reset

PPM Pulse Position Modulation

PS Processing System

PSK Phase Shift Keying

QAM Quadrature Amplitude Modulation

QPS Quadrature Phase Shift Keying

QPSK Quadrature Phase Shift Keying

QPS Quench Protection System

RadMon Radiation Monitoring

RHA Radiation Hardness Assurance

RF Radio Frequency

RFICs Radio Frequency Integrated Circuits

RFSoC Radio Frequency System-on-Chip

RSSI Receive Signal Strength Indicator

SBC Single Board Computer

SCA Slow Control Adapter

SDR Software Defined Radio

SEE Single Event Effects

SEB Single Event Burnout

SEFI Single Event Functional Interrupt

SEGR Single Event Gate Rapture

SEL Single Event Latchup

SET Single Event Transient

SEU Single Event Upset

SiGe Silicon-Germanium

SOM System on Module

SPI Serial Peripheral Interface

SPS Super Proton Synchrotron

SRRC Square Root-Raised Cosine

TID Total Ionising Dose

TMR Triple Modular Redundancy

TN Technical Network

TNF TRIUMF Neutron Facility

ThN Thermal Neutron

TDMA Time Division Multiple Access

VCO Voltage Controlled Oscillator

WAN Wide Area Network

ZC Zero-Crossing

Introduction

Wireless communications have become fundamental in our daily lives. Their unprecedent success is due to the myriad benefits of connecting devices without the constraints of cables. As a result, there is a growing interest in leveraging wireless technology across various sectors beyond consumer electronics, including industrial plants, factories, and even in extreme environments such as nuclear power plants or particle accelerators. The interest stems from the potential to enhance efficiency and data accessibility without cables constraints. In environments where traditional wired connections may be impractical, wireless solutions could offer unparalleled flexibility and adaptability.

The European Organization for Nuclear Research (CERN) represents one of these scenarios. At CERN, physicists study the fundamentals of our universe by accelerating two particle beams through an accelerator chain. The Large Hadron Collider (LHC) is the final accelerator of the complex, where the particle beams are accelerated up to 7 TeV. The two beams are then collided at four interaction points and the resulting particle shower is analyzed to find fundamental particles. Thousands of instruments, i.e. power converters, radiation monitoring devices, cryogenics and more, are necessary to ensure a reliable operation of the whole accelerator complex. Currently, they are monitored and controlled through a control system based on fieldbus, a wired network that enables real-time and distributed control. Fieldbuses were first introduced in the 1980 with the objective of connecting the higher-level of a factory, i.e. the operator and the user interface, with the lowest level, i.e. sensors and actuators. Since then, multiple standards have been developed to cope with the requirements of the different applications. At CERN, multiple fieldbuses are employed in the control architecture, with WorldFIP being the most adopted. It connects more than 10’000 nodes, spreading over 100 kms of cables.

However, deploying cables in the LHC tunnel is an activity that requires careful planning, it is expensive, and it does not allow much flexibility when devices must be relocated. More recently, new installations have become even more difficult because the cable trays have almost reached their maximum capacity. In addition to the current challenges, the Future Circular Collider (FCC) proposed by the CERN council as the next generation of particle accelerator, is foreseen to be 100 kms long. Considering such a length and that the number of instruments will likely increase, wired connections will be very unpractical. On the other hand, wireless solutions could bring multiple benefits to this scenario and are therefore under evaluation. The installation costs will drastically reduce due to little or no cabling, and the flexibility of wireless devices will make installation, upgrade, and relocation much easier.

Nevertheless, the adoption of wireless technologies also introduces challenges that must be carefully addressed. One of the main concerns lies in the lower reliability of the wireless channel, which complicates the adoption of such solutions in control applications where reliability is of paramount importance. Modern communication systems typically employ correction techniques to mitigate natural impairments, enhancing data rate and reliability at the cost of latency. However, in control applications requiring real-time equipment control, a very low latency is necessary. Thus, standard wireless protocols used in consumer electronics for data-centric applications prove inadequate for such scenarios.

Despite that, efforts are underway to meet the demands of the fourth industrial revolution currently in progress, also known as Industry 4.0. This revolution envisions smart factories equipped with internet technologies characterized by low latency and high reliability, seamlessly integrated into production facilities and operations to enable a series of applications that require very low reaction times. Clearly, the core of such a revolution is the ability to exchange information with sensors and actuators in a timely and reliable manner. In response to growing demand, wireless fieldbuses have emerged as a solution, standardized in recent years with a focus on delivering ultra-high reliability within industrial environments. However, they are primarily targeting monitoring applications rather than real-time control due to their inherent latency, typically in the range of milliseconds. Another popular approach is the integration of Internet of Things (IoT) devices into factory operations for data collection and slow control, also referred to as Industrial IoT (IIoT). Yet, these solutions are not designed for real-time equipment control, which demands sub-millisecond latency. To address this need, ongoing advancements such as the 5G New Radio (NR) standard have emerged, specifically targeting low-latency applications with sub-millisecond requirements.

However, while the adoption of new technologies can be immediate in other contexts, it is not so at CERN. The electronics surrounding the LHC is subject to radiation effects, which may impact the device functionalities and performance. For instance, a wireless device may experience communication errors or loss of communications, which results in the instruments unavailability from the point of view of the control system. In the worst case, radiations may even induce destructive events. Whilst specific techniques exist to design electronics to work in radiation environments, at CERN the cost constraint impose the development of systems based exclusively on Commercial Off the Shelf (COTS) components. Therefore, COTS are extensively tested under radiation as part of the CERN Radiation Hardness Assurance (RHA) developed in the past years. Nevertheless, COTS systems implementing wireless functionalities, i.e. Wi-Fi and 5G, are inherently complex. The intricacy arises from software and hardware complexities and the scarcity of adequate test points, which make it challenging to assess radiation effects on communication and propose effective mitigation strategies for ensuring system reliability.

In this perspective, the main objective of this doctoral work is to investigate the opportunity to leverage wireless technologies for current and future CERN particles accelerators, while addressing the challenges described above. Several activities are necessary to achieve this goal: a) analyze the current CERN control system to identify the application requirements and the

physical constraints; b) develop a wireless solution that can fit within the current system to overcome the short-term cabling issue; c) evaluate the impact of radiation effects on the wireless link and propose mitigation techniques to improve its reliability; d) use the findings to build a radiation-tolerant wireless link that can serve as starting point for future wireless systems for control interfaces in radiation environments.

The current CERN infrastructure imposes different constraints on the wireless technologies that could be adopted. Firstly, considering the complexity of the accelerator complex and longevity of the machine, a fixed solution was discarded in favor of a solution based on Software-Defined Radio (SDR). The SDR concept consists of implementing most of the wireless functionalities into programmable logic, keeping at the minimum the fixed hardware, such as the Radio Frequency (RF) section and analog/digital conversion. Therefore, theoretically it allows to implement any kind of wireless protocol by reconfiguring the programmable logic only. In a context like CERN, this is of great advantage because changing protocol would be much easier and less expensive. Secondly, the accelerator equipment, which is exposed to radiation, is custom made by the different groups at CERN. Therefore, redesigning those systems to be compliant with a wireless standard is not an option because it would be too much expensive and time consuming. This constraint required the development of an ad-hoc solution that could be added on top of the current fieldbus infrastructure. The evaluation highlighted some drawbacks, which could be overcome by developing a wireless solution specifically tailored to the application. Therefore, a custom wireless solution has been developed. It is a wireless link based on the SDR paradigm, and it is built using radiation-tolerant COTS, which have been tested under radiation as part of this doctoral thesis. The device implements all the basic functionalities to transmit/receive raw data between two systems wirelessly, with sub-millisecond latency. Its development represents an important step towards the usage of wireless technologies for control applications in particle accelerators. In fact, it can serve as a basic platform to implement future communication protocol for wireless control, with the main advantage of already providing radiation-tolerance characteristics. In addition, being based on SDR, it can be modified depending on the need to accommodate different requirements, such as lower latency. These activities are described throughout the manuscript, which is organized with the following structure.

Chapter 1 describes the CERN complex, with a particular emphasis on the LHC and its radiation environment to which electronics is subject. Moreover, the basic mechanisms of radiation effects to electronics are explained.

Chapter 2 provides the fundamentals of communication network for control systems, the Fieldbus, describing the main requirements that a control network should fulfill and its characteristics, emphasizing the difference with traditional networks used in home/office. The second part of the chapter describes the LHC control infrastructure, and the application requirements, fundamentals for the design of the wireless solution. In addition to the standard requirements of a control network, the one to be used in the LHC must be radiation tolerant, given that part of the control infrastructure is exposed to radiation areas. Hence, Chapter 3 delves into the radiation characterization of a Radio Frequency Integrated Circuit, which is the core component of any wireless system. The chapter discusses a methodology to study radiation-induced effects on such devices, taking the AD9361 RF Transceiver as an example.

Radiation test results are shown, and the impact of radiation effects on a possible communication system based on this component are discussed. Given its radiation-tolerance, the component was selected as the core of the wireless solutions for the LHC control system developed in the rest of the manuscript.
Chapter 4 presents literature research on the past attempts that have been made to extend a fieldbus with wireless technologies, showing the challenges faced by previous projects with the same objective. Then, based on the outcome of the research, a custom radiation-tolerant wireless link is discussed, and its development is described in the second half of the chapter. The proposed solution allows to extend current CERN infrastructure with wireless features, being completely transparent for the CERN control architecture. The hardware implementation has been tested used actual hardware, to demonstrate its working principle. Moreover, results from lab evaluation are shown. The chapter concludes with a discussion of the pros and cons of using such a solution for the CERN architecture.
To overcome the limitations imposed by the fieldbus architecture, a new solution specifically tailored to the LHC applications requirements was proposed, and it is explored in Chapter 5. The first part discusses the radiation characterization of the basic building blocks necessary to build a full wireless system, given that radiation-tolerance is of paramount importance for the control interface. Following the components selection, the complete wireless platform is presented, and its radiation characterization is discussed. Eventually, the reliability of such system is analyzed in the context of the LHC control system.

1 Radiation Effects on Electronics and impact on CERN particle accelerators

Much of this research delves into building a wireless system that can survive in a radiation environment, particularly the one of CERN particle accelerators, emphasizing the importance of understanding how electronic is influenced by radiation effects.

Particles interact with matter in various ways, releasing energy upon impact until they are either fully absorbed or exit. This energy release can occur through ionizing or non-ionizing processes, both of which affect electronic components. Understanding these effects is crucial for comprehending how electronic components respond to radiation exposure, assessing their reliability, and devising techniques to mitigate radiation-induced effects. In fact, electronics installed in CERN particle accelerators is subject to radiation effects, hence it must be properly qualified before being installed. As it will be discussed, radiation-induced failures negatively impact the accelerator operations, reducing the effective operational time, and the amount of useful data that can be collected. Part of the CERN control architecture, which is fundamental to ensure a reliable operation of the accelerators, is exposed to radiation. Nonetheless, commercial wireless devices are not made to survive in radiation environments, thus they cannot be used for the LHC control architecture. Hence, one of the main challenges addressed in this research is the realization of a radiation-tolerant wireless system.

The first part of this chapter discusses the basic mechanisms of radiation effects on electronics and outlines the metrics used to quantify these effects. Later, an overview of the LHC structure and its operational principles is provided, alongside the varied radiation environment found within different areas of the tunnel. These fundamentals are crucial for comprehending the challenges and the design considerations involved in the design of any radiation-tolerant electronic system tailored for the LHC, including the wireless solutions explored in this work.

1.1 Basic mechanisms of radiation effects

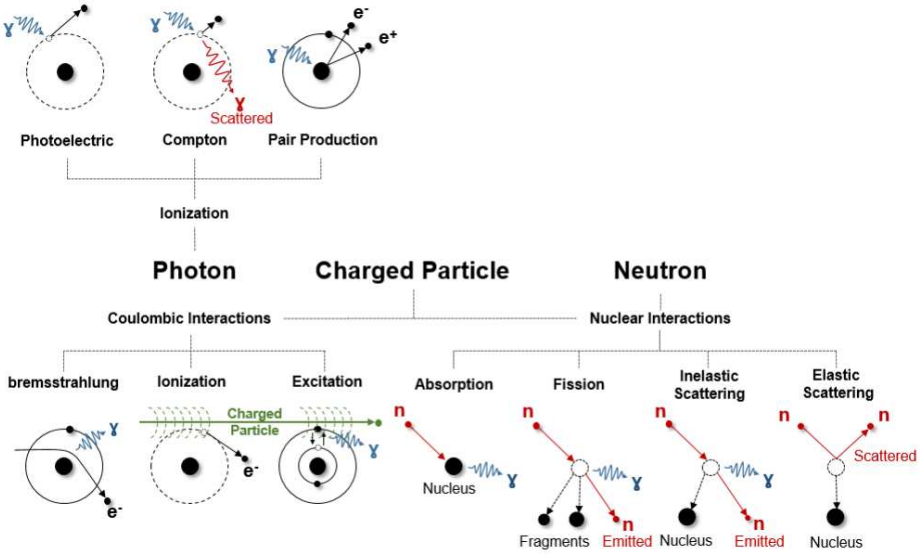

Figure 1 - Particle interactions [1]

Exposure to radiation can alter the properties of materials due to the interactions between incoming particles and the material itself. Illustrated in Fig. 1, different processes can be initiated by the impact of particles depending on their characteristics and energy.
When photons interact with atoms, they can initiate ionization through three primary mechanisms. Firstly, in the photoelectric effect, a photon transfers all its energy to an electron, ejecting it from the atom. Secondly, in the Compton effect, a photon with greater energy than the electron's binding energy scatters away, giving some energy to the electron, which is then ejected. Thirdly, when a high-energy photon interacts near an atom's nucleus, it can create an electron/positron pair, ionizing the matter, known as Pair production.

Charged particles, on the other hand, interact through Coulombic and nuclear interactions. Neutrons, being electrically neutral, solely undergo nuclear interactions. Coulombic interactions include Bremsstrahlung, Ionization, and Excitation. Bremsstrahlung occurs when a charged particle is deflected by the nucleus's electrical force, loosing kinetic energy through photon emission. The Bremsstrahlung mainly affects lightly charged particles like electrons. Ionization happens when charged particles pass near an orbital electron with enough energy to eject it, resulting in indirect ionization as the ejected electron ionizes matter. Excitation occurs when a particle lacks sufficient energy to eject an electron but can excite it to move between inner shells, generating photons. Excitation and Ionization mainly concern heavily charged particles.

Neutron instead can interact through nuclear interactions which can be of four types. In neutron absorption, thermal neutrons (neutrons with an energy of approximately 25 meV) are absorbed by the nucleus, causing excitation. Upon de-excitation, one or more photons are emitted. In elastic scattering, after the neutron collides with the nucleus, both the neutron and the nucleus are scattered, with the neutron losing energy. If the incident neutron is absorbed by the recoiling nucleus, it is called inelastic scattering. In this case, it emits a neutron with lower kinetic energy that exits in an excited state and emits one or more photons. Lastly, fission reactions occur when a heavy nucleus absorbs a neutron, subsequently splitting into smaller fragments known as fission fragments.

These interactions can be initiated not only by incident particles but also by secondary particles generated in previous interactions. For example, a photon from excitation can trigger electron generation, leading to indirect ionization. However, primary interactions, thus direct ionization, predominate due to the higher energy of incident particles compared to secondary ones.

In both direct and indirect ionization, after the interaction, particles lose energy until they are stopped. The stopping power S can be defined as a function of both the ionizing and nuclear interactions normalized to the material density, as

$$S = \frac{1}{\rho} \cdot \left(\frac{dE}{dx}\Big|_{IEL} + \frac{dE}{dx}\Big|_{NIEL} \right)$$

Where ρ is the mass density of the material expressed in g/cm^3, whereas dE/dx is the rate of Ionizing Energy Loss (IEL) and Non-Ionizing Energy Loss (NIEL) along the path.
When a particle moves through a material, its energy loss is not uniform. Instead, this loss varies with the particle's speed. Slower particles have a higher likelihood of interacting with electrons. As a result, as the particle decelerates, the probability of these interactions rises. Consequently, the rate of energy loss increases along the particle's path until it reaches a maximum point, known as the Bragg Peak, after which it decreases to zero.

For energetic charged particles the IEL dominates, whereas neutrons are the main responsible for NIEL. Both IEL and NIEL are responsible for three types of effects on electronic devices: Single Event Effects (SEE), Total Ionising Dose (TID), which are mainly caused by IEL, and Displacement Damage (DD), which is caused by NIEL, and describes the rate of energy loss due to the atomic displacement as a particle traverses the material.

TID and DD are both types of cumulative effects, occurring continuously as long as the device is exposed to radiation. These effects can lead to the gradual degradation of the device, such as an increase in current consumption. In contrast, SEEs are caused by a single ionized particle striking the device's sensitive region, which alters its behaviour.

1.1.1 Total Ionizing Dose

The energy lost by charged particles is absorbed by the material, and this is referred to as the Total Ionizing Dose (TID). The unit of measurement for TID is the Gray (Gy), which represents one joule of energy absorbed per kilogram of matter. It can also be expressed in rad, where 1 Gy is equivalent to 100 rad. The energy deposited can be expressed by the Linear Energy Transfer (LET), and represents the amount of energy lost per unit of length, which can be defined as:

$$LET = \frac{dE}{dx}\Big|_{IEL}$$

The TID deposited by a certain particle can be calculated using its *LET* and the **Fluence** $\Phi(E)$ $[\mathrm{p} \cdot \mathrm{cm}^{-2}]$, which corresponds to the number of particles hitting the device over a specific period. Fluence is obtained by integrating the **Flux** ϕ $[\mathrm{p} \cdot \mathrm{cm}^{-2} \cdot \mathrm{s}^{-1}]$ over time, where the flux is the rate at which particles strike the material. Therefore, considering the LET and fluence of a charged particle, the TID can be expressed as:

$$TID = K_{Gy} \frac{LET(E)}{\rho} \Phi(E)$$

with $K_{Gy} = 1.6 \cdot 10^{-10}$ $[\mathrm{Gy} \cdot \mathrm{g} \cdot \mathrm{MeV}^{-1}]$ is the unit conversion parameter, and $\Phi(E)$ is the particle fluence.

Therefore, one can evaluate the TID deposited by a charged particle as a function of its energy. Considering an environment mixed with different charged particles, the TID will be the sum of the different contributions produced by each particle according to their LET and Fluence:

$$TID = K_{Gy} \sum_{p} \int \frac{LET(p,E)}{\rho} \Phi(p,E)\, dE$$

Once the TID is deposited in the material, it induces long-term effects on the electronic device. The degradation is caused by the ionization of the insulating material used in integrated circuits, such as silicon. When the transistor is exposed to radiation, electron-hole (e-h) pairs are generated in the oxide layer, leading to charge buildup and subsequent device degradation [2]. For example, in a Metal-Oxide Semiconductor (MOS) based device, the radiation-induced charge generation can be split in four processes, as illustrated in Figure 2.

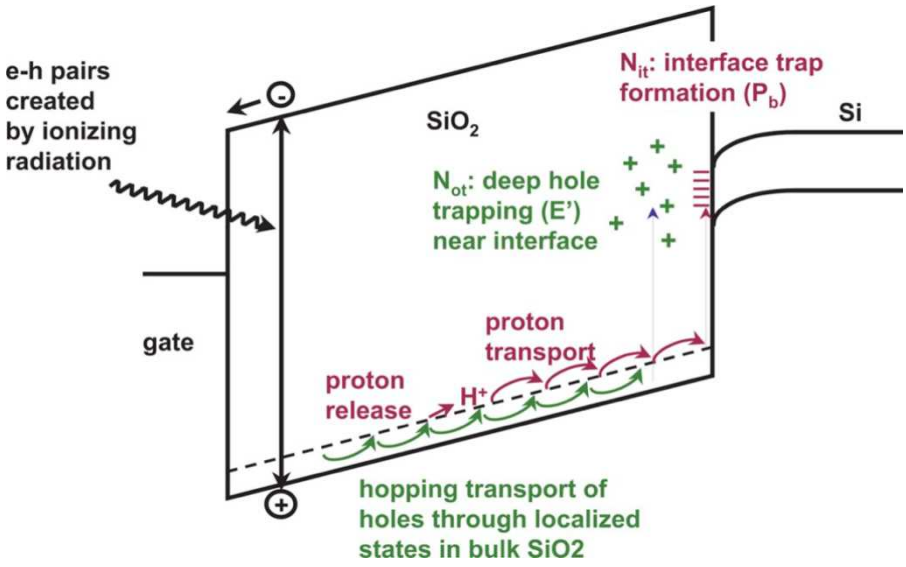

Figure 2 Radiation-induced charge generation in a MOS [3]

In a MOS device with a positive gate bias, an electric field is created along the gate oxide and an inversion layer is formed at the Silicon (SI) – SiO_2 interface. When the device is exposed to radiation, depending on the energy of the striking particle, first e-h pairs are generated within the Si and dioxide dielectric layer, with a rate reported in Table 1.

Table 1 - e-h pair generation rate

Material	**Ionization Energy [eV]**	**Pair generation rate [e-h pairs · $Gy^{-1} \cdot cm^{-3}$]**

Silicon (Si)	3.6	4×10^{15}
Silicon Dioxide (SiO2)	18	8.2×10^{14}

After the creation of e-h pairs, electrons move toward the gate, while holes migrate toward the Si/SiO2 interface. Some electrons recombine with holes, while others escape, contributing to the charge yield. The holes that escape recombination move towards the interface by hopping through localized states in the SiO2, with some becoming trapped near the interface, forming a positive oxide-trap charge. During this hopping process, protons are released and move towards the interface, where they can interact with the interface traps. Depending on their positions, traps are divided into oxide, border, and interface traps. The accumulation of these traps may have different impact on the MOS performance [3].

This accumulation of oxide-trap and interface-trap charge causes a drift in the threshold voltage of the device as:

$$\Delta V_{TH} = -q\frac{1}{C_{ox}}\Delta N_{ot} \pm -q\frac{1}{C_{ox}}\Delta N_{it} \quad 1$$

Where C_{ox} represents the oxide capacitance, and N_{ot} and N_{it} denote the oxide-trap and interface-trap respectively. Depending on the type of MOSFET, the sign of the second term can be negative (for p-channel) or positive (for n-channel). According to Eq. 1, in the case of a P-MOS transistor, the accumulation of positive charge necessitates a higher gate voltage to form an inversion channel, thus requiring a higher voltage to turn it on. Conversely, for the NMOS transistor, it becomes easier to generate an inversion layer, hence requiring a lower voltage. Eventually, the threshold voltage may degrade to the extent that the transistor remains consistently on (in the case of NMOS) or off (for PMOS). Additionally, other effects of TID on MOS devices include a decrease in carrier mobility and transconductance, noise degradation and leakage current increase. Hence, it will impact the behaviour of the system using it. Therefore, radiation qualification is essential to assess the impact of TID degradation on the device.

1.1.2 Displacement Damage (DD)

DD can be caused by both charged particles and neutrons through Coulombic and nuclear interactions. If the striking particle has enough energy, it can displace an atom from its original position, creating a defect in the lattice. The first atom displaced is called the Primary Knock-on Atom (PKA). There are four types of defects that can be created:

- Vacancy: when an atom is absent from its position in the lattice.
- Divacancy: when two adjacent positions are missing atoms.
- Interstitial: when an atom occupies an incorrect position in the lattice.
- Frenkel pair: when there is a vacancy and an interstitial close to each other.

Generating a Frenkel pair requires a minimum energy of 21 MeV [4].

If the PKA has sufficient energy, it can create additional defects in the lattice. Low-energy PKA (<1 keV) can create Frenkel pairs or small lattice defects. Intermediate energy PKA (2-12 keV) can form clusters of small defects, while higher energy PKA can cause a cascade of defects [5].

The energy of the PKA depends on the energy of the charged particles impacting the material. Charged particles with energies under 10 MeV can create low-energy PKA in silicon. Electrons, due to their low mass, only produce low-energy PKA. However, heavy charged particles with energies between 10 MeV and 20 MeV can generate intermediate energy PKA, and those with energies above 20 MeV can produce high-energy PKA. Neutrons can generate PKA at very low energies, such as thermal neutrons at 25 meV, through nuclear scattering.

The energy of the PKA, along with the NIEL of the incoming particle, determines the number of Frenkel pairs generated. NIEL depends on the type of particle, its energy, and the material it impacts.

The DD dose can be computed by dividing the particle's NIEL by the material density, as follows:

$$DDD = \frac{NIEL(E)}{\rho}\Phi(\mathrm{E})$$

DD effects are typically reported as a function of the equivalent fluence of a 1 MeV neutron that would produce the same effect as the particle of interest. This quantity, known as the DD Equivalent Fluence (DDEF), is computed as:

$$DDEF = \Phi_{1MeV\ neq.} = \kappa_{n_{1MeV}}(E) \cdot \Phi(\mathrm{E}) = \frac{NIEL(E)}{NIEL(n_{1MeV})} \cdot \Phi(\mathrm{E})$$

Here, $\kappa_{n_{1MeV}}$ is the hardness parameter, obtained by dividing the NIEL of the particle by the NIEL of the 1 MeV neutron. The overall DDEF in a radiation environment can be obtained by summing the contributions of all particles:

$$DDEF = \sum_{p} \int k(p,E) \cdot \Phi(p,E) \cdot dE$$

1.1.3 Single Event Effects

A particle striking the sensitive volume of a device may induce a so-called Single Event Effect (SEE). The type of effect produced depends on the technology and the function of the device. However, two macro-categories can be identified: destructive [6] and non-destructive [7]. The probability of a radiation-induced event to occur is defined as cross-section σ, which can be retrieved experimentally as

$$\sigma = \frac{n}{\Phi}$$

where Φ $[p \cdot cm^{-2}]$ is the total particle fluence and n is the number of events observed.

Destructive event

Destructive events damage permanently a device. These are the Single Event Latchup (SEL), Single Event Burnout (SEB) and the Single Event Gate Rapture (SEGR).

A latch-up is a potentially destructive event that occurs when a low-resistance path forms between the power supply and ground of a device. This can be triggered in any parasitic

structure, such as the p-n-p-n structure in CMOS technology. In this structure, an increase in the PNP transistor's collector current results in an increase in the NPN transistor's base current. This, in turn, boosts the NPN transistor's collector current, which further increases the PNP transistor's base current. Therefore, if a particle strikes the device and activates either the NPN or PNP transistor, latch-up is triggered. The positive feedback loop will then continue to increase the current until it is destroyed.

The SEB instead occurs when the currents originating from charge collection cause a drop in voltage that activates a parasitic bipolar transistor inside the device. If the particle strikes when the MOSFET is subject to a large drain bias, it can induce carrier multiplication, leading to a runaway effect. This creates a short circuit that allows current from the power supply to destroy the device.

Concerning SEGR, the electron-hole pairs generated by the ion strike are separated by the applied bias. For instance, in an N-channel MOSFET, electrons flow toward the drain while holes move toward the gate along the path of the ion strike. This charge collection causes a short-circuit between the drain and the Si-SiO2 interface, resulting in a drop in drain voltage across the gate oxide. This voltage drop increases the oxide field, and if it becomes sufficiently large, it may lead to oxide breakdown and, consequently, short-circuiting of the gate to the substrate [8].

Non-destructive event

This category groups those events that do not destroy the device when they occur, rather they alter the device functionalities. They are also called soft errors [9]. The basics mechanisms inducing soft errors are described in Figure 3. When a charged particles strikes through the device material, it generates e-h pairs, as shown in Fig. 3a. In a reverse-biased p-n junctions, a high electric field exists across the depletion region. This electric field collects the charge generated by the particle as it passes through, and this collected charge creates a transient current/voltage. Due to the conductive nature of the path formed by the charged particles and the separation of charges by the depletion region, the electric field in the vicinity of the particle's path can be disrupted. This disruption spreads the electric field of the junction beyond its immediate area and deep into the substrate (Fig. 3b). Consequently, the collection of charges within the junction increases, enlarging the range of the electric field. Charges that are deposited some distance away from the initial junction can still be effectively collected due to the efficient drift process. This phenomenon significantly contributes to the enhanced charge collection within the junction. Drift is then followed by a diffusion process, Fig. 3c. In this context, both processes contribute to the movement and collection of charges within the device material.

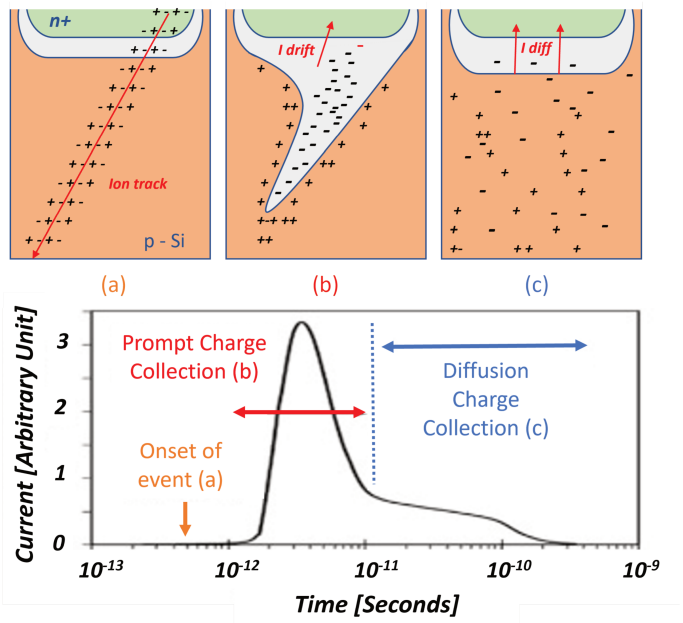

Figure 3 - Charge generation and collection in a reverse-biased junction [10]

Charge collection may cause a Single Event Effect, which can be divided into Single Event Upset (SEU), Single Event Transient (SET) and Single Event Functional Interrupt (SEFI).

A SEU occurs when the energy of the striking particle is high enough to flip the state of a logic element, for instance a memory cell. The process can be explained looking at the typical structure of a CMOS memory cell, shown in Figure 4.

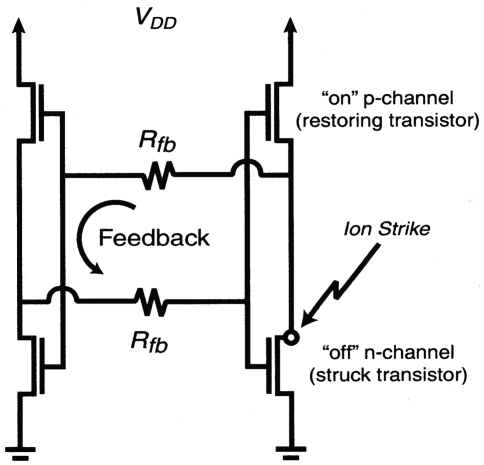

Figure 4 - Ion strike on a memory cell

The memory cell comprises two CMOS inverters, storing either a '0' or a '1' based on the activated transistor pair. The reverse-biased drain junction of the inactive transistors serves as a sensitive location, where energetic particles induce a transient current. Illustrated in Fig. 4, when an energetic particle strikes, the 'n-channel' transistor is affected, generating a transient current. The 'on' p-channel restoring transistor attempts to counterbalance this current, but with limited drive capability. Consequently, the current flow induces a voltage drop at the p-channel transistor's drain, which combined with the transient current may reverse the transistor state. Ultimately, the inverter toggles, leading to a change in the stored value as the memory cell transitions from '0' to '1'.

During a SET, when a particle strike induces a current or voltage transient, the consequences vary depending on the circuit. In digital devices, this transient can result in a glitch in combinatorial logic, which might be later captured by a memory element, leading to storage of an incorrect value. The impact of this depends on the specific application. However, if the transient induced by radiation is brief, it might not be captured at all, resulting in no visible effect. Notably, the likelihood of capturing such events increases with higher clock speeds.

A SEFI arises because of the aforementioned effects, causing the device to malfunction and fail to perform its intended function, because stuck in a wrong logic state. As a result, it necessitates restoration through either a power cycle or a reset procedure.

1.1.4 Soft errors mitigation techniques

Soft errors can be rectified through specific operations. For instance, if a storage element undergoes a change in value, it can be rewritten to restore normal functionality, resulting in only temporary failure. Despite not causing permanent damage to the device, soft errors can still have detrimental effects on the overall application. For instance, a system may make incorrect decisions if the value being considered are wrong because of a SEU. Consequently, various techniques are employed to mitigate these effects.

The most straightforward yet costly approach involves employing redundancy schemes. This method entails creating multiple replicas of the same logic to execute identical operations. One option is to duplicate the function, as in the case of Duplication with Comparison (DWC), or triplicate it, as in Triple Modular Redundancy (TMR). With DWC, the system compares outputs to detect SEUs, whereas TMR utilizes a majority voter to select the correct output based on a majority decision. These strategies can be implemented at various stages of the design process, depending on available resources. For instance, they may be implemented at the system level by replicating the entire system, which represents the most expensive solution. Alternatively, they could be selectively applied to the most critical parts of the design after identifying them. For example, in an Field-Programmable Gate Array (FPGA) design, DWC or TMR can be applied at various levels of granularity [11]. A designer might replicate each individual register, each block of the design, or the entire FPGA design. Alternatively, multiple FPGAs could be used, with additional external circuitry to vote on the output from each FPGA. Similarly, for processors and microcontrollers, these redundancy techniques can be implemented at different levels within the software [12].

In some applications, extensive use of memory is crucial for storing and reading data, making the correctness of memory content essential for reliable operation. An example could be an application code executed by a processor, which if corrupted would cause the execution to hang. Or, for SRAM-based FPGAs, the configuration memory could be corrupted, causing a functional failure. In these occasions, to counter SEUs, in such cases, various techniques can be employed. One approach is using Error Correction Codes (ECC), which can detect and correct errors by means of additional parity bits when the memory is accessed. However, since ECC do not work in all circumstances, especially if the parity bits get corrupted, additional techniques like scrubbing are often used [13]. Scrubbing involves periodically reading the memory, comparing its content to an original reference, and rewriting it if any discrepancies are found. These combined methods help ensure the integrity of memory content, maintaining the reliability of the application.

Regarding Single Event Transients (SETs), they do not always have an observable effect on the circuit as their impact depends on factors like the width and amplitude of the transient.

Various techniques have been proposed to mitigate these effects by filtering the transient through additional logic [14]. For example, in an FPGA, if an SET occurs in the combinatorial logic, it only creates a visible effect if captured by a storage element. This capture depends on the circuit's operating frequency and the duration of the radiation-induced SET. To prevent SETs from causing issues, guard gate circuits can be intentionally added before the combinatorial logic. These guard gates filter out pulses induced by radiation, allowing only those of a certain width to be correctly captured [15].

1.2 LHC Layout and Radiation Environment

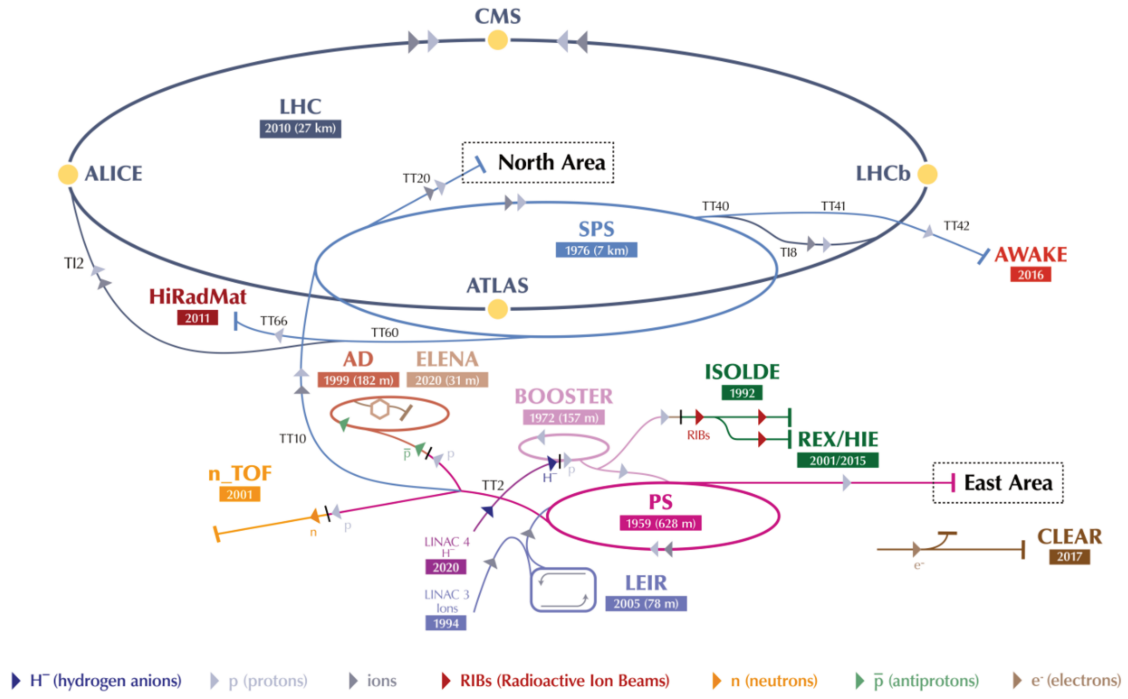

Figure 5 Overview of the CERN accelerator complex

Established in 1954, CERN now stands as the world's most important laboratory for particle physics, dedicated to research of new particles for better understanding of our universe. The laboratory operates a complex of several particle accelerators, which are used to accelerate and collide particle beams, allowing researchers to examine the particles resulting from the beam collisions.

Figure 5 presents an overview of the accelerator complex. The Proton Synchrotron Booster is the initial circular accelerator, whose task is to ramp up particles sourced from the Linear Accelerator from 50 MeV to 1.4 GeV. Subsequently, the beam progresses to the Proton Synchrotron, where its energy is further increased to 24 GeV before being directed to the third stage, the Super Proton Synchrotron (SPS), where they are accelerated up to 450 GeV and directed to the last stage, the Large Hadron Collider (LHC).

1.2.1 Layout

The LHC, currently the world's largest and most powerful particle accelerator, operates with a circumference of 27 kilometers. It aims to accelerate particles in two opposing directions up to 7 TeV each, resulting in collisions at 13.6 TeV within four designated points: ATLAS, CMS, LHCb, and ALICE. Guided by a magnetic field generated by superconducting magnets, the particles traverse through separate tubes, maintained under high vacuum conditions, enabling efficient electrical conduction without resistance or energy loss. To operate in a

superconductive state, the magnets must be cooled down to -271.3 °C, which is achieved through a liquid helium distribution system.

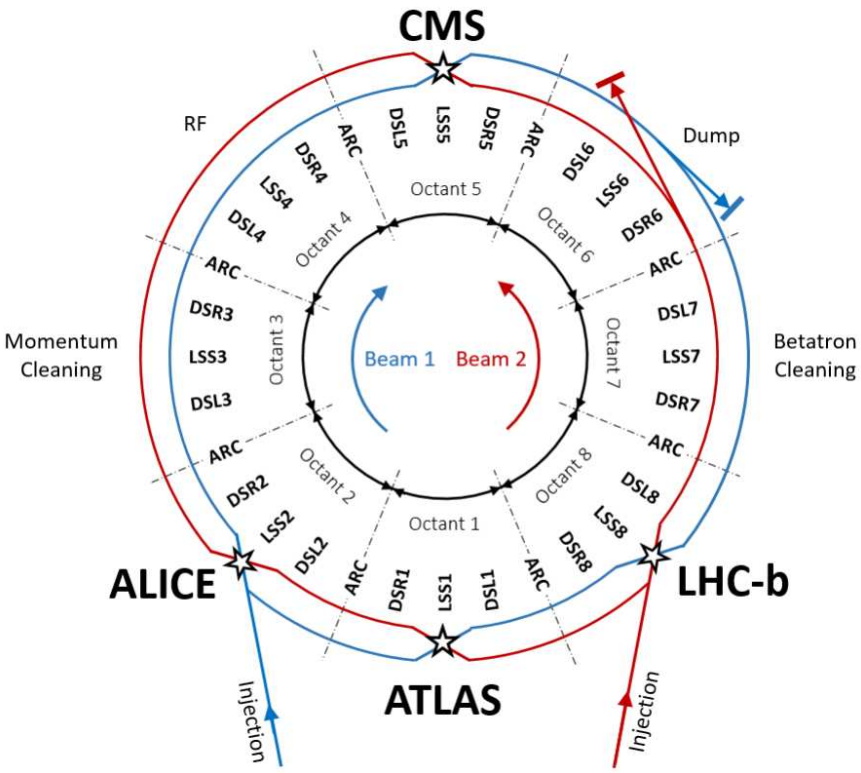

Figure 6 - LHC Layout

The LHC layout [16], shown in Figure 6, can be divided into eight octants with eight Insertion Regions (IR). Among these, four IRs are designated for experiments: IR1 for ATLAS, IR2 for ALICE, IR5 for CMS, and IR8 for LHCb. The remaining four IRs serve specific operational purposes: IR4 hosts Radio Frequency (RF) cavities, IR3 and IR7 are used for collimation, and IR6 for dump operations.

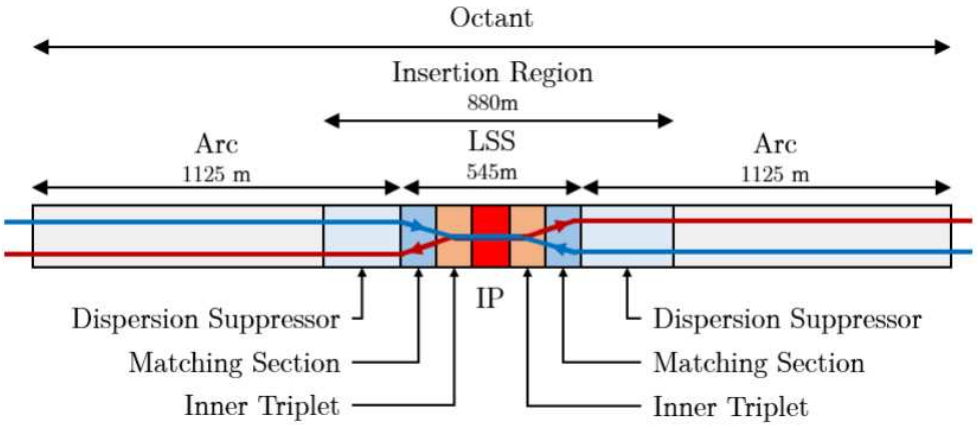

Figure 7 - LHC Octant

The structure of an octant is depicted in Figure 7. In an octant, an insertion region is situated between two ARCs region. Each ARC commences with a Dispersion Suppressor (DS) region aimed at minimizing dispersion in the vicinity of the insertion region. Following the DS is the Long Straight Section (LSS), comprising a Matching Section and the Inner Triplet. The Matching Section consists of a series of quadrupole and dipole magnets essential for

converging the two separate beams into a shared vacuum, while the Inner Triplet, found only in octants designated for interaction points, comprises a trio of quadrupole magnets necessary for beam focusing to facilitate collisions. Positioned at the midpoint of the octant is the interaction point.

1.2.2 Operating principle

The main objective of the LHC is to make the two beams collide, since the particles resulting from these collisions are of extreme interest for particle physicists. However, the interactions of interest are extremely rare, thus it is important to generate as many interactions as possible to generate them. The number of interactions per unit surface and time is quantified with the delivered luminosity, which in the case of the LHC is in the order of $10^{34} cm^{-2} s^{-1}$.

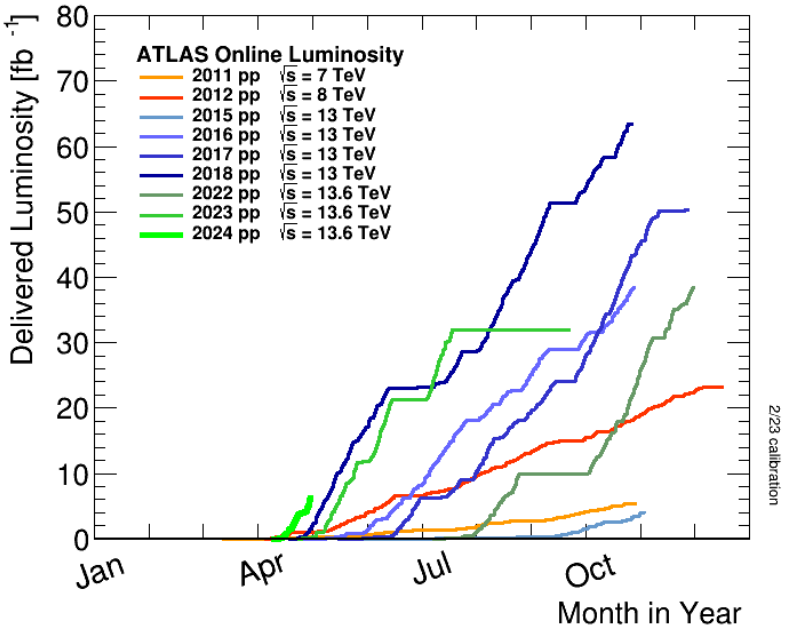

Figure 8 - Cumulative luminosity delivered to ATLAS over days of stable beam operation for high-energy proton-proton collision.

Figure 8 depicts the delivered luminosity [17] over recent years, showing a growing trend that can be attributed to the increased beam intensity. The number of useful collisions produced by the accelerator can be obtained by integrating the luminosity over a defined period of time, which is referred to as integral luminosity and it is expressed in inverse femtobarn fb^{-1}, which corresponds to $\sim 10^{12}$ proton-proton interactions. Thus, the integral luminosity can be seen as the final product of the collider. Consequently, for its improvement, either the luminosity or the operational time of the machine should be increased.

A luminosity increase is the objective of the next phase of the LHC, namely the High-Luminosity LHC (HL-LHC), which is foreseen to start in 2029, and includes technological upgrades that will allow to increase by five times the luminosity of the accelerator. The operational time instead, is linked to the operation cycle of the machine, which can be divided in three phases: fill, ramp, and stable beam. The fill phase involves filling the LHC pipes with protons sourced from the SPS. During the ramp phase, beams are accelerated to 7 TeV before entering the stable beam phase, where collisions occur, delineating the period suitable for scientific experimentation. However, luminosity diminishes gradually during this phase due to losses and collisions, prompting beam operators to “dump” the beam, i.e. interrupt the operation and restart. Ideally, a cycle lasts from 10 to 15 hours, however, very often the cycle is interrupted by uncontrolled beam dumps, due to system faults. Once the beam is interrupted,

some operations are necessary to reach again the stable beam phase, and the time required, known as turnaround time, depends on entity of the dump. In the case of uncontrolled dumps due to system faults, addressing such malfunctions requires additional time for diagnosis and repair, the duration of which varies based on the severity of the fault. Minor incidents may only require system restarts, whereas catastrophic events might require component replacements. In 2015, during LHC operations, beam dumps accounted for over 70% of operational time [18], significantly reducing overall availability. Within these dumps, a significant number were triggered by system failures induced by radiation effects.

1.2.3 Radiation effects impact on LHC operation

From what was mentioned above, a luminosity increase can be achieved with an high availability of the machine, thus its essential to keep the uncontrolled beam dumps at the minimum, including those related to radiation-induced failures.

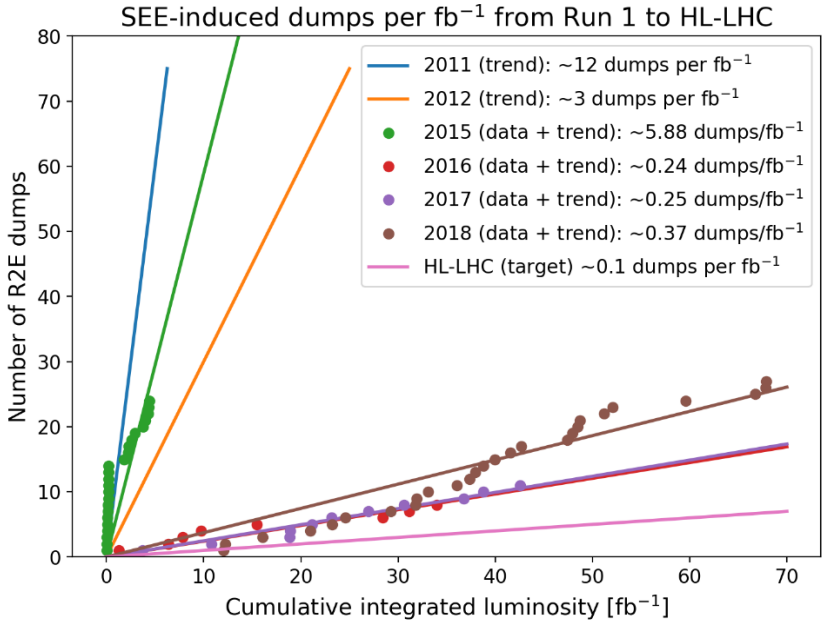

Figure 9 - LHC beam dumps induced by radiation effects as a function of the cumulative integrated luminosity [19]

Figure 9 illustrates the LHC beam dumps caused by radiation effects. Considering 2012, 300 hours of downtime were caused by system failures due to radiation effects, corresponding to ~3 dumps per fb^{-1}. The situation greatly improved in 2016, where it is evident how the number of dumps due to radiation effects notably decreased from 2015. During the 2016 run, LHC availability improved significantly, with only 53% of beam dumps attributed to failures, of which only 5% were induced by radiation effects [19]. This improvement was the result of implementing techniques to mitigate radiation effects on electronics. Initially, when the radiation environment was not fully understood, measures such as increasing shielding or relocating equipment to safer areas were taken. For this purpose, the Service Alcoves were utilized, which are designated areas along the LHC enclosed by concrete. However, not all equipment can be placed in these areas. Some must instead be positioned close to the LHC and exposed to radiation, such as the Quench Protection System (QPS) [20], which is crucial for safeguarding the machine in case of a magnet quench.

However, with advancements in radiation monitoring and simulation, knowledge of the radiation environment has expanded, enabling the development of strategies to qualify components. These efforts have led to the development of the CERN Radiation Hardness Assurance (RHA) procedure, which has evolved over the years to characterize electronic components effectively. The primary goal of RHA is to identify components capable of

withstanding the rigors of the LHC radiation environment, thus ensuring the construction of reliable systems for operation in the LHC tunnel.

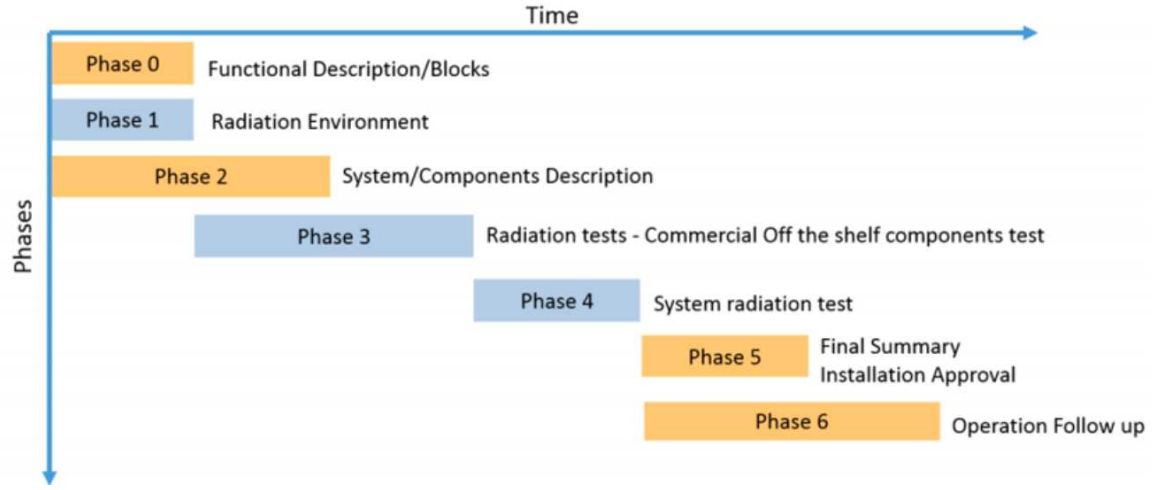

Figure 10 - Block diagram of the CERN RHA

The different steps of the current CERN RHA [10] are summarized in Figure 10. Specifically, the initial phase involves delineating system specifications and assessing radiation levels in the designated installation area, thereby establishing the radiation requirements for the system, for instance which is the maximum number of LHC dumps the system can cause. Once these requirements are established, they inform the system design, determining the components for development. Subsequently, the sensitivity of each component is assessed through radiation testing. Following the evaluation of all components, the performance of the entire system is evaluated under radiation to ensure compliance with the initial requirements. Upon successful verification, the system is ready for the installation. In case of any failure during phase 4, a new design solution should be proposed.

By following these steps, a high reliability for each system can be achieved, allowing to keep the number of radiation induced dumps at minimum. This is fundamental considering the HL-LHC upgrade, which aims at delivering an integrated luminosity of 3000 fb^{-1} in 12 years of operation [21]. Nonetheless, not only radiation levels will increase as it will be detailed in the following section, but it also imposes an higher availability of the machine. This arises mainly from the need to increase the peak luminosity to achieve the desired integrated luminosity. However, the higher peak luminosity would produce more interactions than what can be processed by the experiments. As a solution, it is planned to regulate the luminosity to a level below its peak performance. Nonetheless, reaching this levelled luminosity will only be possible after a certain fill duration. Consequently, the HL-LHC will necessitate very lengthy fills to operate efficiently, thus tolerating very few radiation-induced dumps. To maintain acceptable operation of the HL-LHC, the target failure rate due to radiation-induced effects is set at 0.1 dumps per fb^{-1} [21]. Achieving this target is only feasible by following the RHA procedure and considering radiation levels in different areas, as elaborated in the following section.

1.2.4 Radiation sources

In the LHC, there are three primary sources of radiation. Firstly, during proton-proton collisions in the interaction points (IPs), a cascade of secondary particles (protons, pions, kaons, electrons, photons and muons) ranging from thermal energies to GeV is generated. While detectors intercept most of these particles, a fraction still reaches surrounding elements. Near the IPs, photons and pions predominantly comprise the spectrum of emitted particles, with the distribution varying based on distance from the interaction point, thus constituting a

luminosity-driven radiation source. Secondly, direct losses occur between beam halo particles and collimators positioned around the experiments, initiating particle showers. These losses are correlated with beam intensity, making it an intensity-driven radiation source. Lastly, radiation is generated from the interaction between the beam and residual gas density within the vacuum pipes of the ARCs. Here, the radiation intensity is influenced by both beam intensity and pipe quality. The primary distinction between intensity-driven and luminosity-driven radiation sources lies in their emission period. Intensity-driven radiation is emitted continuously throughout each operation cycle, whereas luminosity-driven radiation is only emitted during the stable beam phase.

1.2.5 Current and future radiation levels

Given the complexity of the LHC structure, it is interesting to study the radiation levels depending on the LHC location. In fact, the radiation environment encompasses a range of particles spanning from meV to GeV, with varying distributions dependent on the area. Assessment of radiation levels involves the use of simulation tools like FLUKA [22], a Monte Carlo code designed for studying the interaction and transport of hadrons and photons within materials. Additionally, experimental data is obtained through various systems installed in different areas, such as Beam Loss Monitors (BLMs) [23] or Radiation Monitoring (RadMon) devices [24].
As outlined in the preceding section, electronic devices are susceptible to SEEs, TID, and DD effects. SEEs are quantified by the High Energy Hadron (HeH) fluence and Thermal Neutron (ThN) fluence. The HeH fluence (Φ_{HEH}) denotes the fluence of all hadrons with kinetic energy above 20 MeV, as they have equal probabilities of inducing SEEs. Conversely, hadrons with lower energy may induce SEEs but with reduced probability, rendering them negligible for SEE assessment. However, studies have shown that ThN, neutrons with an energy around 0.025 eV, significantly contribute to SEE, particularly in electronic devices containing boron-10 [25]. Consequently, ThN fluence (Φ_{ThN}) is monitored along the LHC, and a Risk factor ($R = \Phi_{HEH}/\Phi_{ThN}$) [26] is used to identify areas where ThN contributions are higher.

Thus, the main factors to consider are the Φ_{HEH}, R_{th}, TID, and DD. The HEH and ThN fluences are useful when determining the SEE risk, whereas the TID and DD quantifies the cumulative damage effects (lifetime perspective).
These factors can be analysed for the different positions thanks to simulation and experimental measurements. The areas can be grouped in two main categories: tunnel, and service alcoves. The tunnel is the space surrounding the LHC machine itself, where the number of electronic systems is minimized, thus including the ARC, the DS, the IR. The service alcoves instead, which are also referred to as shielded areas, are designated areas along the LHC enclosed by concrete, where most of the electronics is installed. Nonetheless, there are some instruments which cannot be placed there. For instance, the QPS protects the magnets in case one of the magnets suffer a quench, and thus it must be placed as close as possible to beam, therefore it is installed in the tunnel.

1.2.5.1 Tunnel

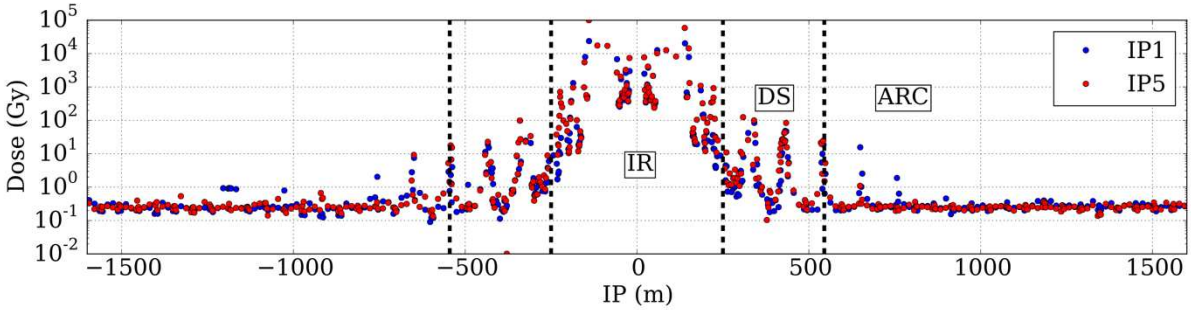

Figure 11 - TID measured by the BLM close to IP1 and IP5 during the LHC run in 2016 [27]

Concerning the tunnel, the TID measured by the BLM for the proton-proton run of 2016 is visible in Figure 11 for IP1 and IP5. The TID is reported as a function of the distance from the interaction point. As it is visible, in the IR, which is the closest to the IP, the electronics would be subject to a dose of 10 kGy/year. Thus, such radiation levels do not allow the use of COTS-based systems. In the DS instead, there are peaks of 100 Gy/year, however the use of electronics can be considered in those points below the dipoles magnet where it will be subject roughly to 10 Gy/y, which is instead a feasible target for qualified COTS-based designs. The same goes for the ARC, where there is around 1 Gy/year, which is acceptable from the lifetime point of view.

Concerning the SEE risk instead, the situation is different because the risk increases with the number of units deployed, therefore it must be carefully considered. The Φ_{HEH} is measured by the RadMon device [24], which is a radiation monitoring device developed and maintained at CERN, with more than 400 units currently deployed. For the proton-proton run of 2016, the measured Φ_{HEH} was of $5 \cdot 10^9 cm^2/y$ and $5 \cdot 10^8 cm^2/y$ for the DS and ARC respectively [27].

1.2.5.2 Service alcoves

As mentioned above, the radiation levels in the IR and in some parts of the DS are too high to host electronics, thus service alcoves (shielded areas by concrete) were built along the LSS. They can be divided in two main groups: the heavily-shielded and lightly shielded areas. The first, are placed close to the IPs. Two areas can be identified, the Underground Junction chambers (UJs) and the Underground Liason galleries (ULs). The lightly shielded areas instead, also known as RR, are placed at the intersection of the DS and LSS. For these areas, the measured radiation levels by the RadMon system during the proton-proton LHC run of 2016 are reported in Table 2.

Table 2 – Annual HEH fluence for the shielded areas measured during the LHC run of 2016 by the RadMON system [27]

2016 p-p run	**RRs**	**UJs**	**ULs**
Φ_{HEH}	3×10^8	5×10^8	1×10^8

While these numbers are lower compared to the HEH fluence of the tunnel, they are still relevant in terms of operation because they host a large amount of electronics which can suffer from SEE failures.

1.2.5.3 HL-LHC upgrade

The same parameters can also be examined for the HL-LHC phase, scheduled to commence in 2029 and span approximately 12 years. Assessments can be conducted for radiation levels in both the tunnel and shielded areas utilizing FLUKA simulation, with consideration given to scaling current radiation levels with luminosity. The anticipated yearly luminosity for the HL-LHC is approximately 250 fb^{-1}, translating to 3000 fb^{-1} over its 12-year lifespan. In contrast to the current LHC, areas near the BLM are expected to endure 10 kGy/y, making the use of COTS impractical. However, electronics are predominantly positioned beneath the dipoles and within the shielded zones.

Table 3 – Expected annual (250 fb^{-1}) radiation levels for the HL-LHC and average risk factor for the different areas of the tunnel and shielded areas.

Location	Φ_{HEH} $[p/cm^2]$	TID [Gy]	R_{Th}
Tunnel			
DS	5×10^9	100	4
ARC	1×10^9	2	3
Shielded areas			
UJ	5×10^9	10	50
UL	1×10^8	0.2	25
RR	3×10^9	6	6

The expected yearly radiation levels [27], derived from FLUKA simulations, are detailed in Table 3. Within the tunnel, especially in the DS region, electronics are expected to receive an annual dose of 100 Gy (1 kGy over the entire lifespan). This necessitates either relocating/substituting systems, which are currently qualified up to 200 Gy, or designing radiation-tolerant alternatives capable of withstanding such levels. Conversely, shielded areas will witness a significant rise in HEH fluences, a factor demanding careful consideration as they will accommodate most of the electronics.

The knowledge of these radiation levels is fundamental to assess the reliability of any instrument, including the wireless link designed as part of this doctoral work and described in Chapter 4 and 5. As elaborated in the following section, the control system comprises various endpoints situated in radiation-prone areas, necessitating the use of radiation-tolerant electronics. Given that certain systems are positioned in proximity to the accelerator, a portion of the wireless link will be placed in radiation zones. Therefore, a comprehensive examination of the impact of radiation effects on the devices and the consequences on the communication and the control system is imperative. This entails assessing how radiation exposure may affect device performance and reliability, its implications for the control system, and possible mitigation schemes to improve its reliability.

1.3 Radiation test facilities

To assess the impact of radiation on electronic components and determine their suitability for the LHC radiation environment, the components are tested at various facilities. Each facility offers different particle types and energy levels, with some presenting topologies that introduce additional challenges to the testing process, which will also be addressed in this doctoral work. Furthermore, since one of the main goals is to propose a radiation-tolerant wireless link for the

control network, all components have been characterized under irradiation to ensure a radiation-tolerant system. This section outlines the facilities used in this work to achieve that objective.

1.3.1 PSI Proton Irradiation Facility

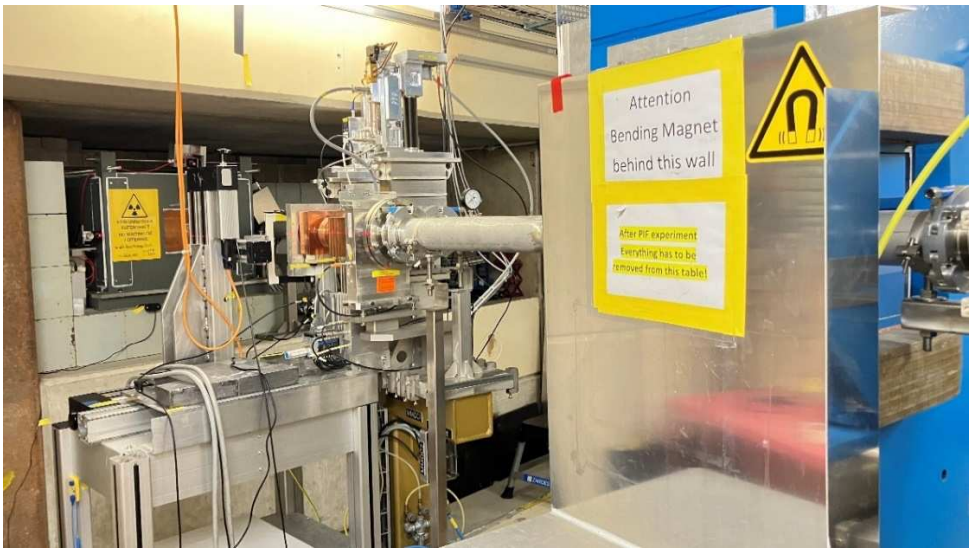

Figure 12 - Proton Irradiation Facility in Villigen, Switzerland

The Paul Scherrer Institute in Switzerland is a research centre dedicated to proton therapy. In collaboration with the European Space Agency (ESA), the institute developed the Proton Irradiation Facility (PIF), shown in Figure 12, by extracting the proton beam line from the existing cyclotron accelerator, COMET. The PIF provides a maximum beam energy of 230 MeV, which can be reduced to 10 MeV using degraders made of copper plates. The particle energy and flux can be further adjusted by varying the beam current between 0.1 nA and 10 nA; however, the beam becomes unstable outside this range. A current of 10 nA corresponds to an energy of 230 MeV and a proton flux of 2×10^9 p/cm^2/s, with the flux scaling linearly with the current. The beam size and shape can be modified using copper collimators, allowing for a maximum diameter of 9 cm.

To measure the beam flux, two ionization chambers are positioned upstream and downstream of the collimators. These chambers are calibrated by comparing their counts with the output of a scintillator detector, ensuring accurate flux measurement at the position of the Device Under Test (DUT).

This facility is used for TID, DD, and SEE testing as part of the CERN RHA procedure. One key advantage, as will be highlighted in the test setups described in this work, is that the radiation is focused as a beam with an adjustable size. This feature is particularly valuable because it allows for precise targeting of the component under test while avoiding surrounding circuitry. This capability enables the use of commercially available development kits for some DUTs without the need for custom boards to perform component characterization.

1.3.2 TRIUMF Neutron Facility

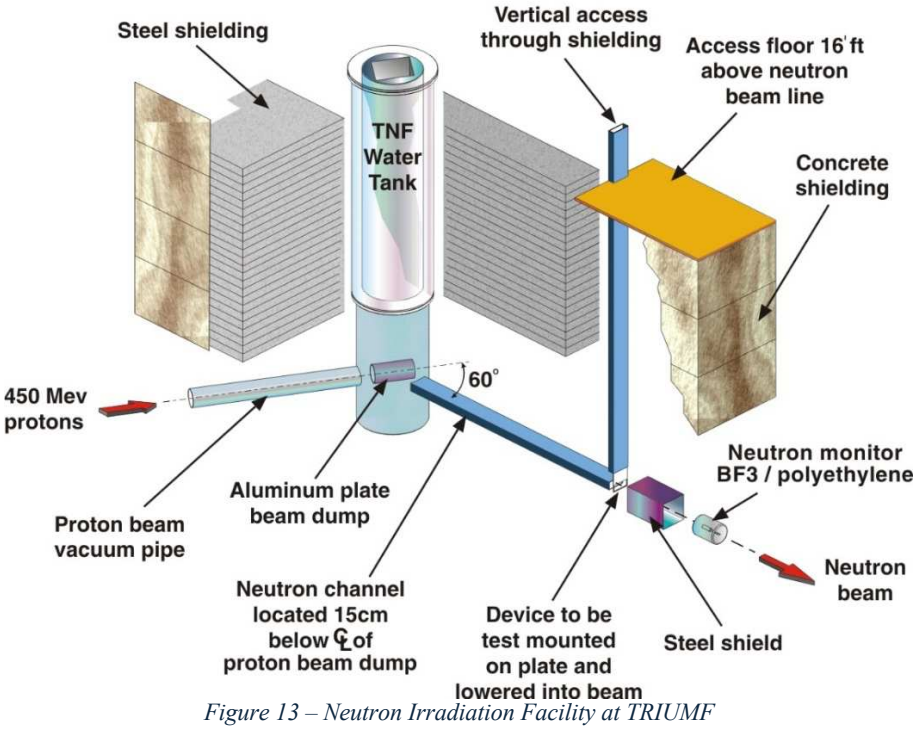

Figure 13 – Neutron Irradiation Facility at TRIUMF

TRIUMF is Canada's particle accelerator centre and offers three separate facilities for neutron irradiation. The TRIUMF Neutron Facility (TNF) specifically provides neutron energies ranging from thermal to 400 MeV at high intensity. The neutrons are generated through the spallation reaction of 100-150 μA of 400 to 450 MeV protons striking an aluminium plate absorber, which is surrounded by a water moderator. The neutrons are then channelled through the shielding surrounding TNF. An overview of the facility is shown in Figure 13.

The neutron spectrum at the TNF ranges from thermal energies up to 400 MeV. For energies above 10 MeV, the neutron flux ranges from 2×10^6 to 3×10^6 $\text{neutrons/cm}^2\text{/s}$, while for thermal energies, it is approximately 5×10^5 neutrons/cm²/s, with a homogeneity of ±10%. The flux is monitored using BF3 counters.

This facility provides a beam spot size of 5x12 cm. As shown in the figure, accessing the test area requires lowering the DUT into the access channel. This presents some challenges, as the entire area will be irradiated, potentially affecting surrounding components and compromising the radiation test.

1.3.3 CERN Co60 Facility

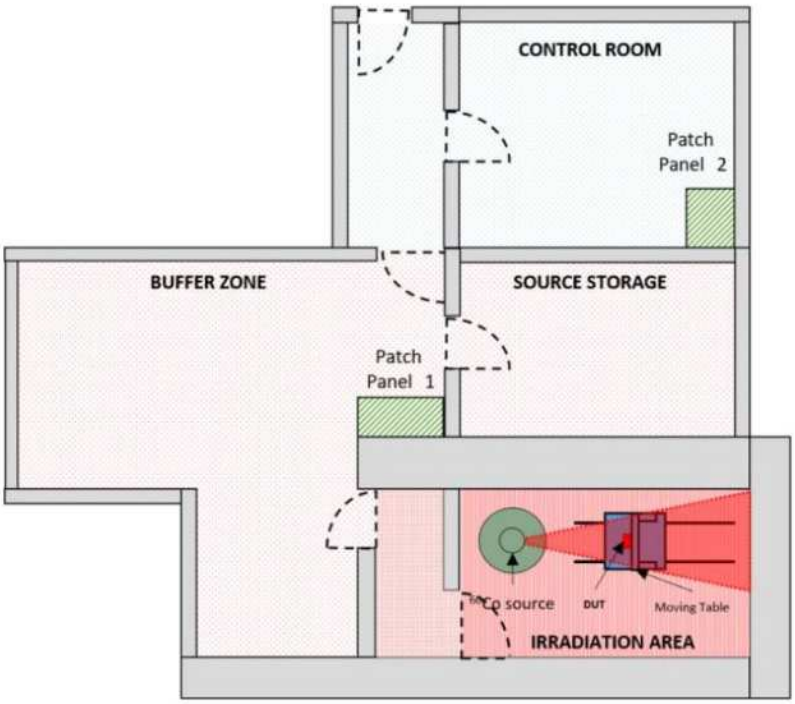

Figure 14 – Co60 irradiation facility at CERN

The Cobalt-60 facility at CERN, whose layout is shown in Figure 14, uses the decay of ^{60}Co isotopes which emits beta radiations into ^{60}Ni isotopes [28].

The highly activated nickel nucleus then emits two gamma rays with energies of 1.172 and 1.1332 MeV, which are powerful enough to penetrate the devices under test. The irradiated area forms a cone, with the dose rate decreasing according to the inverse square law as distance increases. The facility's minimum dose rate is 120 mGy/h over an area of 100 cm², and it can reach up to 20 Gy/h over an area of 20 cm². Additionally, the dose rate can be as high as 650 Gy/h over an area of 3 cm² when the device is mounted on a support at the aperture. An air-filled ionization chamber is used to measure the dose at the position of the DUT before starting the test.

1.3.4 CERN CHARM facility

CHARM is CERN's mixed-field irradiation facility, designed to qualify electronic systems for the LHC radiation environment. Its radiation field mimics the mixed-field conditions found in the accelerator.

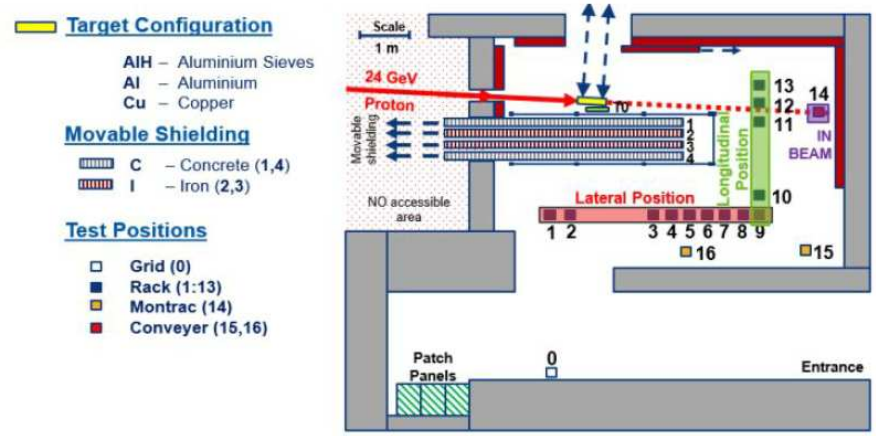

Figure 15 – CERN Mixed field facility layout

The facility layout is shown in Figure 15. The mixed field is produced by a mono-energetic 24 GeV proton beam interacting with a configurable target made of copper, aluminium, or aluminium sieves. The facility also features four movable shields, each 20 cm thick—two made of iron and two of concrete. It offers 14 different test positions within a 70 m² room, as highlighted in the figure.

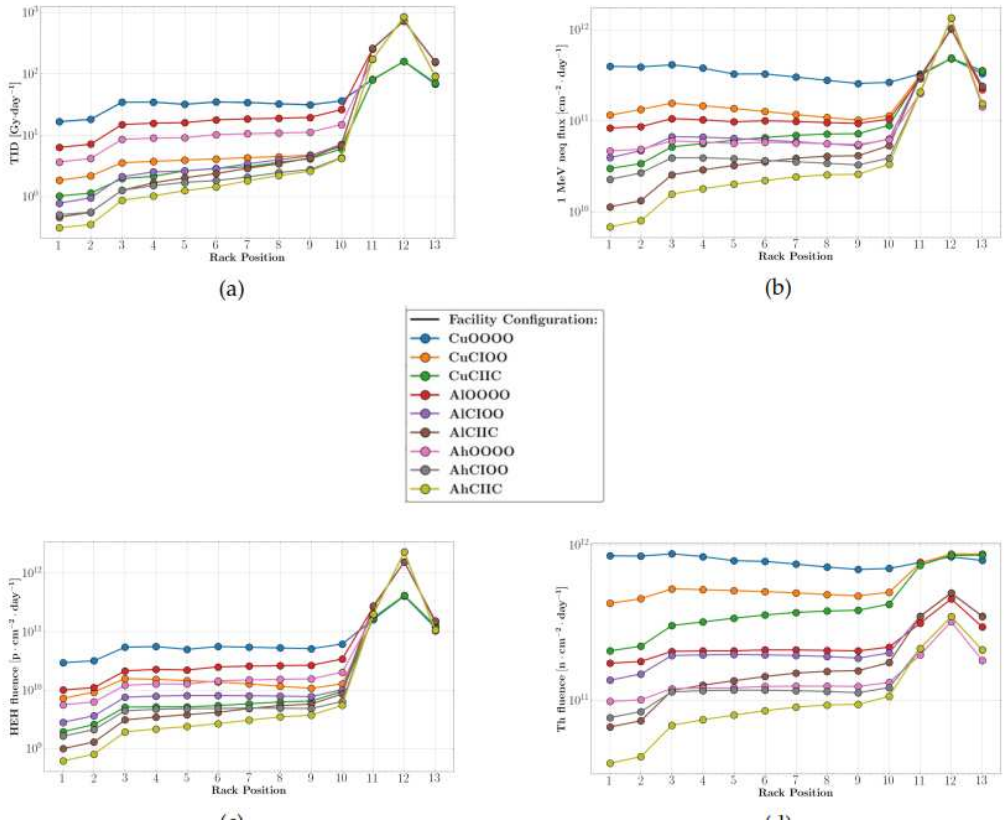

Figure 16 – Radiation levels in the different test position in the CHARM facility

Each position provides varying levels of TID, DDEF, HEH and thermal neutrons, which are presented in Figure 16, normalized per day [29]. Irradiation typically lasts for 5 days, and dosimetry is monitored using the RadMon system installed at the various test positions [28].

1.4 Chapter summary

The chapter provided an overview of the basic mechanisms of radiation effects on electronics and the different categories of events that can be generated in an electronic device. This is fundamental as the LHC particle accelerator necessitates many instruments to operate reliably, many of which cannot be placed in radiation-free zones but must be placed close to the accelerator, thus subject to radiation effects.

The first part of the chapter discussed how particles interact with matter and the different kinds of interactions. It was shown that when particles strike a material, they lose energy until stopped, and the energy losses can be classified as ionizing (IEL) or non-ionizing (NIEL). Hence, the consequences were classified into the three main effects on electronics, which are TID, SEE, and DD. Concerning TID, the mechanisms behind the degradation were discussed in the context of a MOS device, and formulas to compute the TID deposited by particles were presented.

Concerning SEE, the two categories of effects were presented: destructive and non-destructive events. For non-destructive events, an in-depth overview was presented along with different mitigation techniques that can be used in the design phase to counter SEE. The concept of cross-section was introduced, which is used to evaluate the reliability of a system. Then, DD was introduced, explaining the basic mechanisms of atom displacement, and providing a formula to evaluate the DDEF contribution.

The second part of the chapter described the LHC layout and its operating principles. The main metric for evaluating the performance of an accelerator, luminosity, was discussed. Luminosity, which measures the number of particle collisions the accelerator can generate, is crucial since the outcome of these collisions is of primary interest; therefore, higher luminosity is preferred. The different phases of an LHC run—ramp, fill, and stable beam—were outlined, with the stable beam phase being when useful data can be collected. It was explained how the beam can degrade, necessitating a beam dump, which means restarting the operation from the ramp phase. Additionally, it was shown that a beam dump can also be triggered by system failures due to radiation effects, as discussed in the first part of the chapter. Data from previous LHC runs highlighted how radiation-induced dumps can significantly reduce the accelerator's availability, underscoring the importance of building radiation-tolerant systems.

For this reason, electronic components are qualified under radiation in different facilities, which were presented in the last part of the chapter. In this context, the different the various radiation sources in the LHC were discussed together with the radiation levels for the current and future accelerators. This information is fundamental for the qualification process of systems, which is essential for building reliable systems in a radiation environment. This is a crucial aspect that will be covered throughout the rest of the manuscript, particularly in the design of a radiation-tolerant wireless link for a control system. As will be discussed in Chapter 2, many instruments that are part of the LHC control system are in radiation areas, necessitating a radiation-tolerant control interface to ensure reliable operation. The concepts presented in this chapter will be extensively used in next chapters to a) qualify all the electronics necessary for building the wireless link and b) evaluate its reliability in different areas of the LHC accelerator.

2 Control networks in the LHC accelerator

To operate the LHC successfully, various components are utilized: power converters for magnet manipulation, RF amplifiers for particle acceleration, cryogenic systems for superconductive magnet cooling, and instruments to measure the beam characteristics such as its profile and position, among others. Real-time monitoring and control are crucial for reliability, enabling operators to take prompt action, such as beam dumping, in case of failures. Therefore, the primary objective of the control system, linking high-level operators with the low-level instruments, is to establish values in active elements and retrieve data, status updates, and diagnose faults.

Currently, the LHC control system relies on a legacy wired infrastructure that was deployed in the 2000s, prior to the LHC's commissioning. This architecture was designed based on the commercial solutions available at the time, which met the requirements of all applications and provided necessary radiation tolerance. As discussed in the previous chapter, radiation tolerance is crucial not only to minimize radiation-induced dumps but also because parts of the control architecture are exposed to radiation. Ensuring radiation tolerance of the control interface is essential since its availability is critical for supervising the entire infrastructure effectively. The current CERN architecture controls more than 10’000 devices through a Fieldbus, a wired network for control system. Nonetheless, the solution does not offer flexibility, given that installation and updates are time and cost expensive. Moreover, cable deployment in the LHC tunnel is expensive, and cable trays are now full given the large quantity of devices. Eventually, it will be even more difficult for the FCC to rely on wired control system, given its length. Such conditions make wireless solutions appealing.

The objective of this doctoral research is to advance the state-of-the-art control system in the LHC accelerator by using wireless technologies, a topic that will be explored in the subsequent chapters. However, before delving into how this research addresses the radiation tolerance of a wireless link and its deployment in a control network, it is essential to understand the fundamentals of a wired control network, with a particular focus on the system currently in use at the LHC. These concepts are necessary to describe the challenges to face when extending such a network with wireless technologies.

Therefore, the first part of this chapter provides the basics of communication networks employed for control system, as well as the most widespread technologies used nowadays. The second part instead describes the architecture of LHC control system and its communication network. This overview is necessary to conceive a solution that can address the short-term cabling issue faced by the LHC by adding a wireless extension to the current architecture without any modification to the hardware already deployed, which will be the main challenge discussed in Chapter 4. Eventually, the various LHC control applications are presented, as well as their requirements. These will be used for the development of a custom wireless solution,

discussed in Chapter 5, that can be the basis for a wireless control network for future particle accelerators.

2.1 Communication networks for control system

A control system can be defined as an interconnection of components forming a system configuration that provides the desired response. This is achieved by reading the output of various processes through sensors, feeding these values to a controller (nowadays often a computer-based system) that executes an algorithm to process them, and then producing variables applied to the system through actuators.

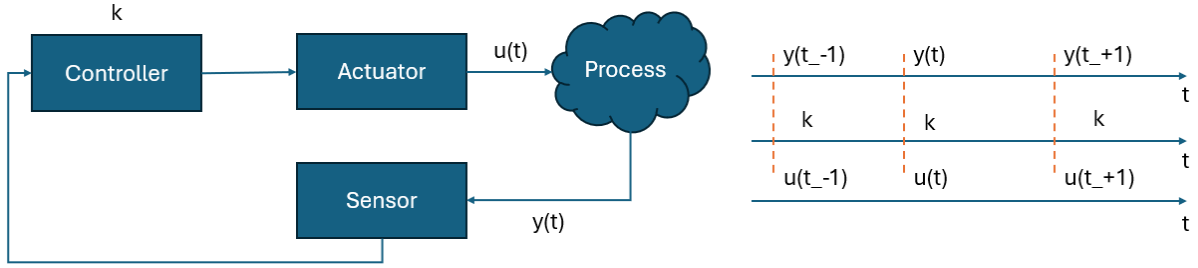

Figure 17 - Architecture of a discrete control system

Figure 17 illustrates the main actions of a control system, which involve three steps: a) sampling parameters, i.e. temperature or pressure, using a sensor, b) executing an algorithm to calculate the controlled variables, and c) applying these controlled variables to the process using an actuator, i.e. opening, or closing a valve. Typically, once configured, this process runs indefinitely in a periodic manner, which defines the sampling period of the control system.

Early control systems were based on point-to-point configurations, where sensors and actuators were directly wired to the controller through a 4-20 mA current loop [30]. In these systems, the three phases were performed simultaneously, as the delays introduced were negligible due to direct wiring. However, point-to-point networks lacked flexibility: adding new sensors or actuators was difficult, and remote monitoring was not possible. To address these limitations, networked control systems (NCS) were introduced [31] and have become the standard. An NCS connects the controller, sensors, and actuators through a common network, as shown in Figure 18.

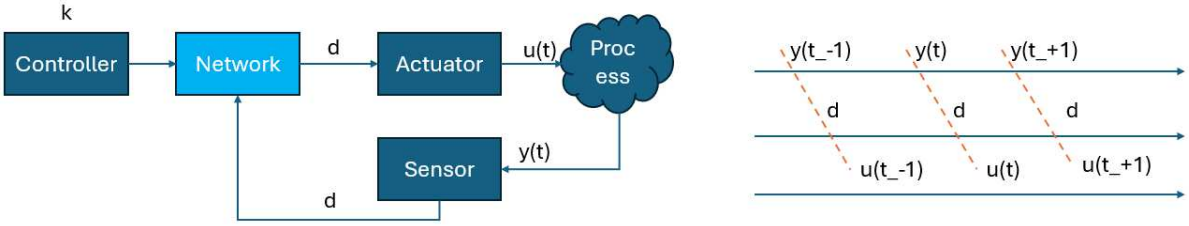

Figure 18 - Networked Control System (NCS)

In an NCS, data from sensors and to actuators is distributed via a shared network, forming a point-to-multipoint connection. This architecture allows new sensors and actuators to be added easily by connecting them to the network, and, if supported, enables remote monitoring. However, the network introduces non-negligible latency, meaning sampling, computation, and actuation cannot occur simultaneously as in point-to-point systems. As depicted in Figure 13, there are additional time delays for transferring data from the sensor to the controller and from the controller to the actuator. These delays are time-variant because data delivery depends on network nature, i.e. wireless or wired, and its characteristics such as data rate. Therefore, the network must be chosen according to the control application requirements.

Even though control systems are used in applications differently from each other, i.e. automation, chemical, etc., there are some common requirements that a network in control applications should follow [32]:

- Low latency, ideally close to zero, to resemble point-to-point architecture, as higher latency increases the difference between the actual and desired system response
- Deterministic delivery of data
- Successful delivery of the values

A network meeting these requirements is called an Industrial Control Network (ICN).

2.1.1 Industrial Control Network

An Industrial Control Network (ICN) differs significantly from traditional office or home networks. While the latter focuses primarily on transferring large quantities of data, ICNs are designed to meet the requirements of low latency, determinism, and reliability. These networks operate on the assumption that the nature of exchanged data, network topology, and traffic flow are consistent in control applications, unlike home or office networks where user and application patterns vary widely. Consequently, the protocol stack for an ICN differs from that of traditional network, which can be usually described through the Open System Interconnection (OSI) model.

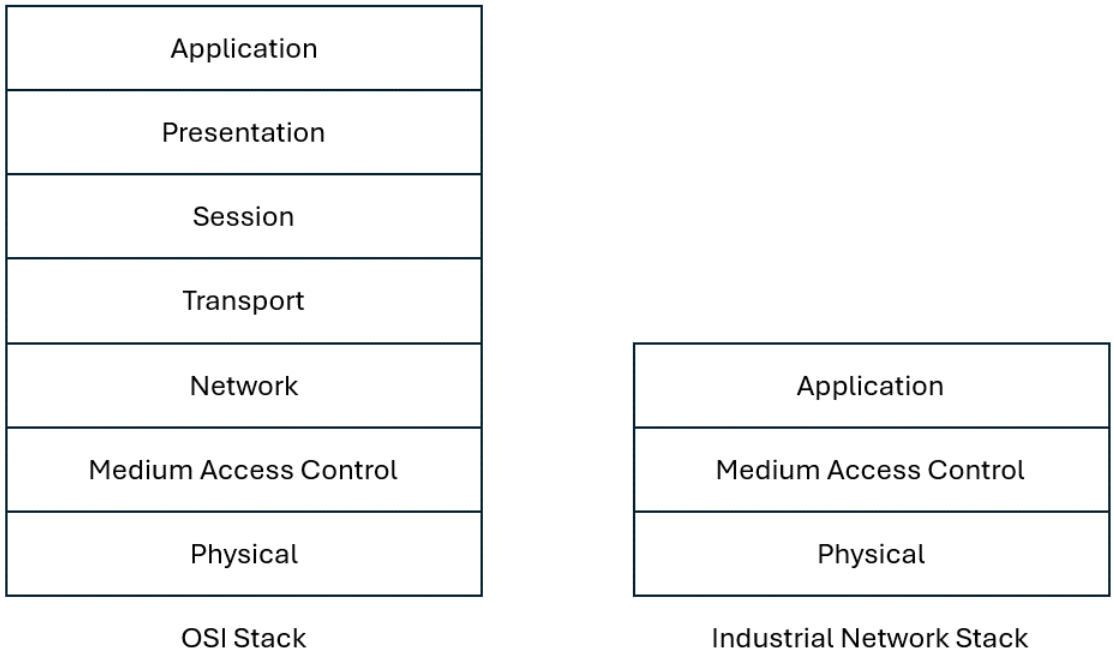

Figure 19 - OSI stack and ICN stack

The OSI model is divided into seven layers, as shown in Figure 14, and is typically implemented by traditional networks. These layers can be briefly described as follows:

- Application: Provides the Application Programming Interface (API) for users to perform tasks on the network.
- Presentation: Translates data between the application layer and the network services.
- Session: Manages communication sessions between two nodes.
- Transport: Manages data exchange between two nodes.

- Network: Handles multipoint networks, managing addressing and routing of traffic.
- Medium Access Control: Manages access to the physical medium, ensuring data exchange between endpoints.
- Physical: Responsible for transmitting and receiving raw bits over a physical medium.

In ICNs, only the Application, MAC, and Physical layers are implemented. The Physical Layer (PHY) typically remains unchanged from traditional networks, while significant differences are found in the MAC layer, which is specifically tailored to ensure deterministic and reliable data transfer in ICNs. Recently, there has been a growing interest in optimizing the physical layer as well [33] to meet the specific requirements of ICNs. In some cases, a custom physical layer may be necessary for certain control applications. This doctoral work follows this trend, since the LHC control system necessitates a custom physical layer to support wireless technologies. Hence, one of the challenges addressed by this work will be the design of a custom physical layer for a wireless control network.

2.1.1.1 Characteristics of Industrial control network

The need for Industrial Control Networks (ICNs) arises due to their distinct requirements compared to traditional home or office networks. While traditional networks prioritize high bandwidth and data rates, ICNs focus on reliability, low latency, and determinism. These specialized requirements often make it challenging to adopt standard technologies, but certain assumptions facilitate the design of communication networks for ICNs.

One key assumption stems from the architecture of a typical control network, as shown in Figure 17. The data exchanged consists of sampled values and control variables, which are typically small in size. This means that the data can be stored in just a few bytes, significantly reducing the amount of data that needs to be transferred compared to standard applications like video streaming.

Additionally, traffic in ICNs flows in a predictable manner. Once sampling times and element delays are known, network traffic can be anticipated even before deployment. This predictability contrasts with standard networks, where users may connect or disconnect at any moment.

Furthermore, ICN traffic is cyclic. In control applications, once configured, data exchange repeats indefinitely at a fixed sampling period. This cyclic nature simplifies the development of the Medium Access Control (MAC) layer. For instance, many ICNs use Time Division Multiple Access (TDMA) schemes, where time is divided into slots dedicated to individual transmissions. An example of such a MAC protocol for ICNs is illustrated in Figure 20, showing multiple sensors and actuators transmitting/receiving data only during their predefined timeslot. These time slots can also include idle periods to ensure packet retransmission in case of data loss. T_s defines the sample period of the control system. Such protocol makes it easier to implement a deterministic network, since data arrival is bounded to the duration of the time slot.

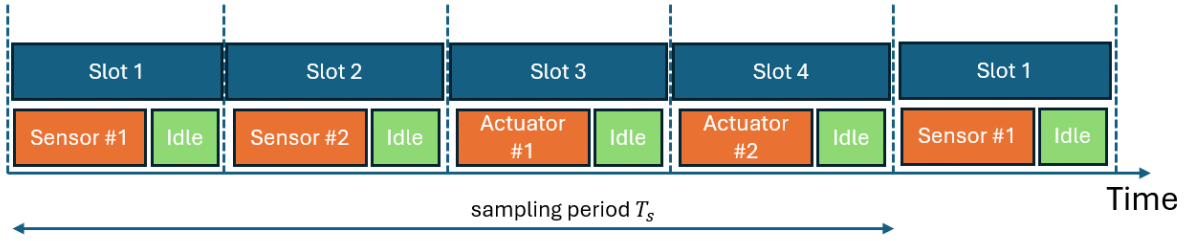

Figure 20 - TDMA example

Another important assumption involves network topology. Unlike traditional networks, ICNs typically remain stable once deployed, with changes mainly involving the addition of new sensors or actuators. Generally, ICNs are based on a single controller, such as an embedded computer, managing multiple sensors and actuators. This results in a consistent logical topology, usually a star topology (one controller managing many devices), even if the physical topology varies (e.g., devices arranged in a ring or along a bus).

The characteristics described above are crucial during the design phase of ICN communication networks. In addition to standard ICN requirements, the CERN control system impose additional demand, namely radiation-tolerance and flexibility. The first necessity arises because part of the instruments is placed in radiation-areas. As discussed, radiations can induce failures into electronics, which could result in the unavailability of some instruments, or wrong sampled/control values. Hence, radiation-tolerance is of paramount importance. The latter requirement instead is necessary to face technology aging. As will be discussed in the next section, the LHC control network was deployed before the LHC commissioning, and it counts thousands of devices. Nonetheless, upgrading it to more recent technology, such as wireless, is not practical because all the instruments would need to be upgraded, requiring too much effort in terms of time and cost. Hence, flexibility is important to avoid facing the same issue in the future accelerators, such as FCC.

Table 4 - Characteristics of a communication network for ICN

Characteristics
Short-packet size
Cyclic traffic
Fixed topology
Requirements
Low latency
Determinism
Reliability
Flexibility
Radiation-tolerance

Table 4 reports the characteristics of a ICN and the requirements of an ICN for the CERN control system, which will be considered throughout the manuscript for the development of a custom wireless solution.

2.1.2 Wired control network architecture

The concept of a control network is to link low-level equipment across all levels of a factory, ensuring integration with other layers. Figure 16 illustrates the pyramid scheme currently used to represent the different hierarchy levels within a factory [34].

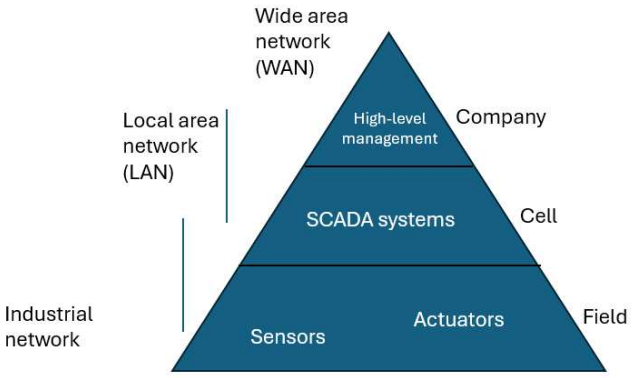

Figure 21 - Pyramidal scheme of a factory, showing the different levels of the hierarchy and how they are interconnected

At the top level, management applications run, providing an overview of the entire architecture through a Graphical User Interface (GUI) that conceals low-level details. This level typically connects to a Wide Area Network (WAN) to interconnect different factory sites or share data with suppliers. The cell level, dedicated to supervision systems (SCADA), connects via a Local Area Network (LAN) deployed across the entire site. The field level, responsible for controlling low-level equipment, is where the ICN is deployed. It is also notable that as one moves higher up the pyramid, the constraints on reliability and latency decrease, whereas these constraints increase towards the bottom of the pyramid.

2.1.3 Fieldbus

In the late 1970s, the first solutions for ICNs began to emerge, known as fieldbus systems. Fieldbus is a digital communication network designed for industrial control devices such as sensors, actuators, and controllers [35]. The main objective of fieldbus technology at the time was to replace point-to-point control networks, marking the advent of ICNs.

The 1980s saw the development and widespread adoption of several fieldbus systems, including Modbus, Controlled Area Network (CAN), P-NET, and ControlNet. These solutions were tailored to specific control scenarios, employing different methods for channel access and data transfer, such as Time Division Multiple Access (TDMA), polling, and random access [54].

As the use of fieldbus systems became more widespread, two common fieldbus protocols—Factory Instrumentation Protocol (FIP) and Profibus—gained prominence. This led to the establishment of a committee aimed at standardizing the fieldbus technology to address the lack of interoperability among various solutions, which created complexity for customers. Despite these efforts, the resulting IEC 61158 standard [35] was released as a multiprotocol standard, incorporating different fieldbus protocols without defining a single universal solution.

Each fieldbus has its own characteristics, mainly because it has been developed to fulfil requirements for a particular scenario.

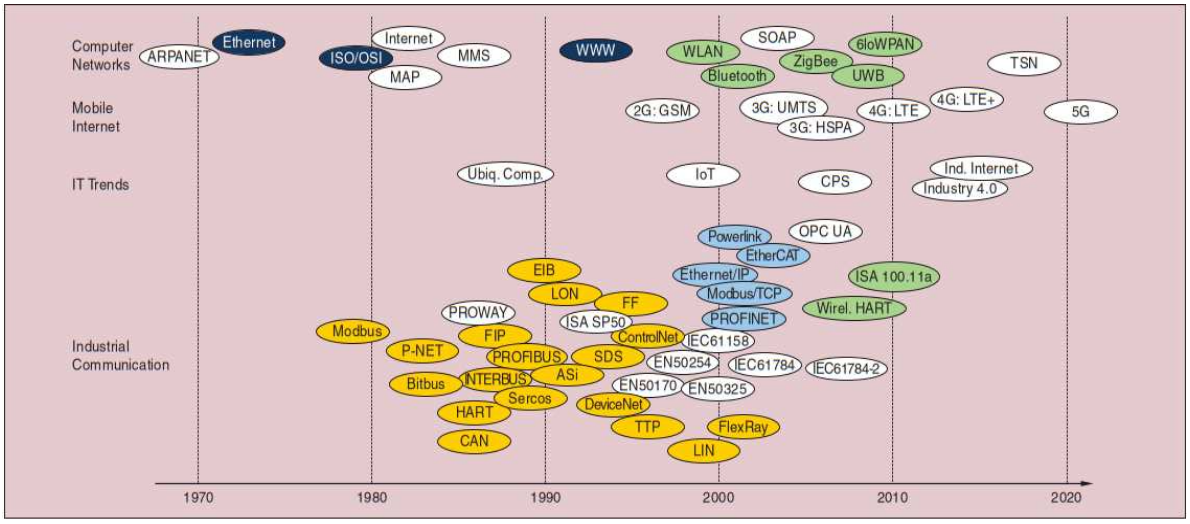

Figure 22 - Timeline of network evolution taken from [33]

Figure 22 illustrates a timeline depicting the evolution of digital communication networks over the past 50 years. The yellow highlights represent the various fieldbuses developed during this period. Notably, FIP, which was developed in the 80' continues to be used in the LHC today, as discussed in the following section that presents the LHC control architecture. Additionally, the timeline shows the recent emergence of wireless solutions, marked in green, which aim to address the lack of wireless options in industrial environments. A more detailed analysis of these wireless solutions will be provided in the subsequent chapters, when discussing the wireless extension for the LHC control architecture.

2.2 The LHC control system

Before the LHC commissioning, its "pre-injectors" were controlled with their own control systems. However, when the LHC project started, given its complexity, it was decided to develop a new control infrastructure by redesigning most of the technologies. Hence, in 2003 [36], a decision was made to implement a new architecture capable of overseeing the LHC and its injectors from a single control room through a unified control system.

2.2.1 Architecture

The current architecture of the CERN control system is mainly based standard components (computers, servers, networks), and uses industrial control for the supervision of the LHC equipment.

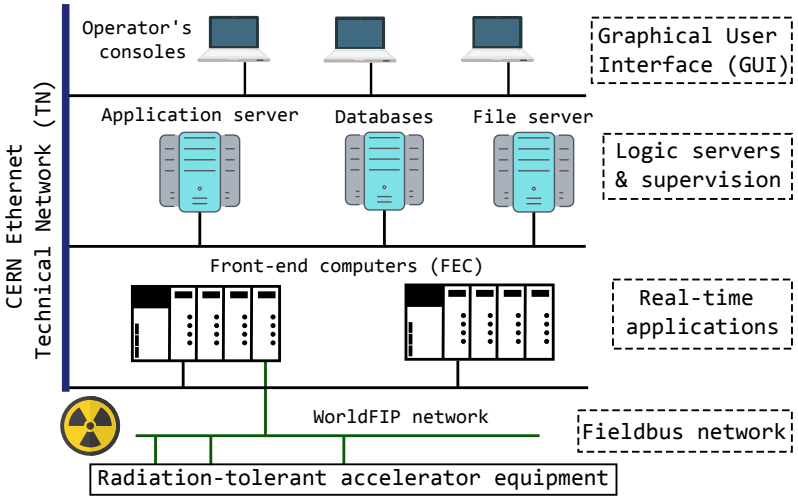

Figure 23 – Overview on the CERN control system architecture

Figure 23 provides an architectural overview of the control system [37], comprising three distinct layers. At the top are the operator consoles located in the CERN Control Center (CCC), running application graphical user interfaces for interacting with lower-level instruments. These consoles utilize standard Intel PCs operating on Linux.

The second layer consists of rack-mountable CPUs functioning as back-end servers, operational around the clock. Positioned in the CERN Control Room (CCR), adjacent to the CCC, these servers serve three primary functions: Network File System file servers, Front-End Computers boot servers, and high-level application servers.

The third layer encompasses embedded front-end computers (FECs), featuring CPUs, memory, and interfaces housed within mobile racks. FECs are primarily responsible for real-time control and data acquisition of accelerator equipment. All layers are interconnected with a database for data storage and retrieval. The entire architecture is interconnected via the CERN Technical Network (TN), dedicated to equipment control, which does not have access to the Internet. Utilizing a Gigabit Ethernet backbone with optical fiber distribution, the TN may also be connected to the CERN General Purpose Network (the standard network deployed in offices, giving internet access), via switches.

Hence, FECs play a pivotal role in low-level control and data acquisition for accelerator hardware. Designed for industrial environments, FECs are rack-mountable electronics enclosures comprising a chassis, backplane, power supply, and fan tray. In addition, they accommodate extension boards directly plugged into the backplane, serving as interfaces with accelerator equipment. The core component is the Single Board Computer (SBC), containing essential desktop computer elements (e.g., RAM, CPU) except for a hard disk (FECs boot from corresponding servers in the second layer via the network). The SBC, mounted within the backplane slot, controls other extension boards within the FEC enclosure. Figure 24 depicts a typical FEC with installed extension boards.

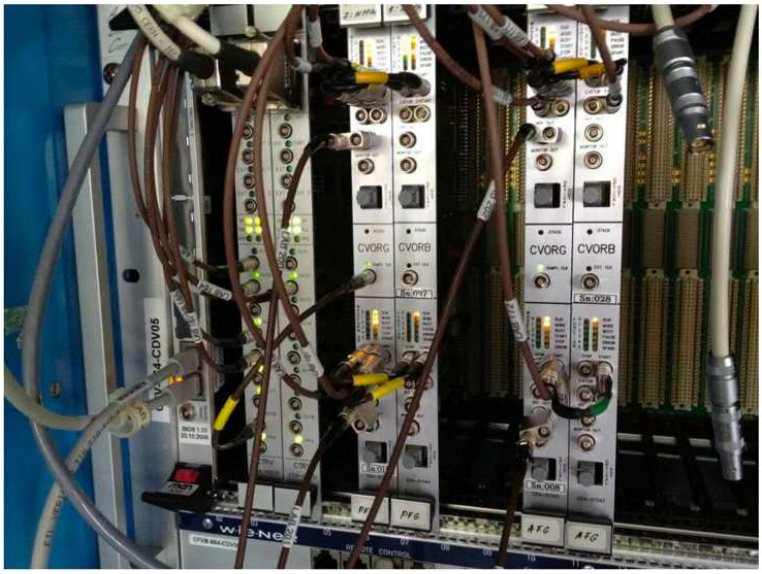

Figure 24 - Front-End Computer and some electronics module installed [37]

The FEC is versatile, accommodating various standards like PXI, PCIe, VXI, FMC, and others. Depending on specific application needs, different extension boards can be employed. In the context of this doctoral thesis, the primary focus is on the control application, particularly the slow control of LHC instrumentation, which is accomplished through an extension board that implements WorldFIP [38], an industrial Fieldbus for industrial control and monitoring.

2.2.2 WorldFIP fieldbus for slow control of LHC equipment

In various industrial environments, real-time and distributed control communication networks, namely fieldbus, play a vital role. These technologies are prevalent in settings such as nuclear power plants, chemical plants, automation systems, and manufacturing facilities. They are primarily utilized where precise control and monitoring of processes are necessary. This need for control and monitoring extends to the LHC as well, where numerous instruments must be effectively managed to ensure the accelerator's reliable operation.

Each fieldbus technology possesses distinct characteristics concerning data rate, latency, the number of devices supported, and network topology. These characteristics are tailored to specific applications. At CERN, in addition to meeting application requirements, fieldbus selection must consider the radiation environment, which can induce electronics failures leading to communication losses or errors. Such errors can be critical in control systems, where equipment must be continuously monitored to facilitate prompt action by beam operators. Consequently, an error may result in an instrument becoming invisible or data not being updated.

Taking these factors into account, a thorough study was undertaken prior to the commissioning of the LHC to identify the most appropriate fieldbus technology. The selected technology had to demonstrate robust radiation tolerance. In 2001, the Working Group on Fieldbus proposed three options for use in the LHC: Profibus, CAN, and WorldFIP. Consequently, investigations were initiated to assess the susceptibility of each fieldbus to radiation effects. Among them, WorldFIP showcased promising radiation tolerance [39], leading to its selection as the primary fieldbus for controlling LHC equipment.

2.2.2.1 CERN WorldFIP Technology

The WorldFIP protocol is part of the European fieldbus standard EN61158 [40], which serves as a network for industrial communications. Facilitating seamless interaction between sensors, actuators, and higher-level applications, it accommodates various network structures including centralized, decentralized, and master-slave configurations. The LHC at CERN employs a master-slave architecture, leveraging WorldFIP to establish deterministic and real-time links between low-level equipment and FECs within the tunnel. The master is installed as an extension board on the FEC, thus in radiation-safe areas. On the other hand, the slave is installed inside the tunnel, into radiation-exposed areas.

Initially, WorldFIP relied on a suite of components supplied by Alstom. These included the MicroFIP and FullFIP, implementing the slave and master functionalities respectively, along with the FielDrive chip and FieldTR. The FielDrive chip, employing a mixed design with a passive filter, regulates the network's data rate. The FieldTR, comprising an insulating transformer, ensures galvanic isolation between the line driver and the fieldbus medium, typically a copper twisted pair. The MicroFIP and FullFIP instead, are Application Specific Integrated Circuit (ASIC) which implements the WorldFIP protocol.

With over 10,000 instruments relying on this fieldbus, radiation tolerance of the slave device was a requirement from the beginning. Extensive testing of the MicroFIP ASIC revealed its robustness, demonstrating no functional failures, and lasting up to a total dose of 700 Gy [39], [41], [42], [43]. Consequently, it became the preferred fieldbus for numerous equipment groups and was therefore deployed in the tunnel.

However, Alstom discontinued the MicroFIP and announced the cessation of support for the WorldFIP protocol in 2009. Consequently, an internal task force was established at CERN in 2006 to address the situation and identify suitable replacements for WorldFIP technology that would meet requirements. However, it was found that no viable alternatives were available. Moreover, due to the extensive use of the protocol in the LHC tunnel, transitioning to a different solution would have been resource-intensive, requiring updates to all equipment. As a result, the decision was made to internalize WorldFIP technology at CERN to ensure the continued operation and maintenance of existing infrastructure [44]. Additionally, since the product was being discontinued, it was decided to replace both MicroFIP and FullFIP with an in-house solution for future development requiring the WorldFIP protocol. Thus, the initial step after internalization was the development of the nanoFIP agent [45], a radiation-tolerant solution based on FPGA, implementing a subset of MicroFIP functionalities, with consensus from equipment groups. This project was undertaken within the Open Hardware Repository (OHWR) collaboration [46], making it publicly accessible for modifications. Subsequently, following the development of the nanoFIP, it was decided to migrate the FULLFIP master devices provided by Alstom as well. Hence, the MasterFIP solution [47]was devised, based on an FMC solution mounted as an extension board on the FEC.

2.2.3 LHC equipment requirements

The operation of the LHC requires various systems, many of which utilize the WorldFIP fieldbus for communication with the FEC. These systems include Quench Protection, power converters, radio frequency cavities, radiation monitors, cryogenics, and beam instrumentation. Each of these applications has distinct demands regarding data rate, latency, and criticality.

There are thousands nodes installed in the LHC making use of WorldFIP, which extends over more than 300 km of cables spread around the LHC tunnel [48].

The main requirements influencing the design of a communication network, as explained in the previous chapter, include reliability, latency, determinism, and radiation-tolerance. Currently, the characteristics of the WorldFIP protocol meet the demands of all applications. However, as the following chapters discuss leveraging wireless technologies, a thorough understanding of the requirements for each application is essential for making the design choices for the wireless link. Therefore, the primary parameters to be analysed are:

1. The number of devices to be connected, which is crucial for assessing the reliability of the link in terms of radiation effects.
2. The sampling period of the control system, which helps determine the latency requirements that the wireless link must meet.

Understanding these parameters is vital for ensuring the wireless link meets the necessary performance and reliability standards.

2.2.3.1 The Quench Protection System

In the LHC, magnets operate in a superconductive state while cooled down to -271.3 °C. Nonetheless, if one of the magnet transits from superconductive to resistive state while it is powered, it induces a thermal runaway, also called quench. In this situation, the energy of the other magnets will be dissipated through the quenched magnet, destroying it. Therefore, the objective of the Quench Protection System is to detect the quenched magnet as soon as possible, bypass it, and de-excite the other magnets [20]. A QPS system involves a detection mechanism, which could be analog or digital, quench heaters power supplies to energize the quench heater strips mounted on the coils of the magnet, and a data acquisition and monitoring controllers which can transfer the data to the higher-level control system [49].
The data acquisition and monitoring system is composed by a series of local controllers which communicate with the corresponding systems via I2C or SPI. Data read from the devices is stored inside the controller SRAM and is then read back through WorldFIP and transferred to the higher-level applications. This allows monitoring of all the systems.
The QPS counts 2100 fieldbus-controlled data acquisition systems installed in the LHC, communicating with FEC. For each fieldbus segment there are up to 60 systems, and the sampling period is set to 200 milliseconds [49].

2.2.3.2 LHC Power Converters

The LHC power converters [50] instead are responsible of providing large currents and very low voltages to the superconducting magnets. Currently, there are 1700 power converters installed in the LHC. Each of them is equipped with an embedded computer control system named Function Control/Generator (FGC), which is responsible of controlling and monitoring of converter state (on, off, standby), generation of the current reference and monitoring of the regulation of the measured current [51]. Currently, for each WorldFIP segment there are 30 FCG systems connected to a single masterFIP, for a total of 70 networks to connect all the power converters to the LHC control system. Each segment works with a sampling period of 20 milliseconds.

2.2.3.3 Radiation monitoring

Radiation monitoring activities are important to understand the radiation levels in the different areas of the accelerator. As explained in the previous section, knowledge of the radiation levels is fundamental for the RHA because they allow to estimate the failure rate of a system in a specific area of the LHC. Radiation levels are currently monitored through the RADiation MONitoring (RadMon) [52] device, which is a device developed at CERN that provides measurements of TID, DD, and HEH fluence. Currently there are more than 400 devices deployed in the LHC tunnel divided into multiple segments. For each segment there are up to 32 devices, and the sampling period is set to 1 second [53]. The lower sampling rate is due to the nature of the application, which does not need high precision measurements since no control is involved.

2.2.3.4 Cryogenics system

The main goal of the cryogenic system is to maintain the pressurized superfluid helium bath of all magnets around the LHC at their nominal temperature of -271.3 °C, so that they can work in their superconductive state and steer and focus the particle beam [54]. Many front-end electronics for the cryogenic system which involves thermometry, pressure, have been custom developed at CERN to cope with the radiation-tolerant requirement, being they exposed to radiation. These components are accessed with WorldFIP. Each LHC sector has more than 100 electronic crates equipped with the WorldFIP interface to communicate with temperature, pressure, level gauges sensors, digital inputs (on-off for valves), and convey commands for electric heaters [55].
Per each sector, there are 100 WorldFIP crates, for a total of 800 crates covering the whole LHC. They are placed below the magnets in the ARC and DS area, or in shielded areas such as UJ and RR. In each sector there are 8 WorldFIP segments, where up to 32 devices are connected to a masterFIP on the FEC, with a sampling period of 500 milliseconds [55].

2.2.3.5 Beam instrumentation

Several parameters of the beam need to be monitored during the LHC operation. This is the objective of several instruments which are monitoring losses, orbit, and tuning [56].
However, most of the instruments is placed on the surface, keeping only front-end cards below the quadrupole. The beam parameters are monitored through 240 Beam Instrumentation crates, which include Beam Position and Beam loss monitors. The crates will process the beam information which are sent to the control room through WorldFIP, with a sampling period of 100 milliseconds per FIP segment [57].

2.2.4 Summary requirements

The requirements of the applications are summarized in Table 5.

Table 5 - Requirements summary for fieldbus instrumentation

System	Unit	Sampling period (per segment)	Max. Unit per segment
QPS	2100	20 ms	32
Power convert	1700	20 ms	30
Radiation Monitoring	400	1 s	30
Cryogenics	800	500 ms	32
Beam instrumentation	240	100 ms	32

2.3 Chapter summary

The chapter covers the basic concepts of control networks, focusing on the LHC control architecture.

The first part of the chapter explains in detail what ICNs are and how they differ from traditional communication networks. Control networks have unique requirements in terms of latency and reliability, as well as distinct characteristics related to data rate, topology, and network traffic. The OSI model, which describes communication networks, is presented, with a discussion of the relevant layers for ICNs. An overview of fieldbus networks, a family of networks for industrial control, is also provided.

The second part of the chapter presents the LHC control architecture, which connects high-level applications to the low-level instruments necessary for the reliable operation of the LHC. The structure managing the instruments is discussed, including FEC computers connected to equipment through WorldFIP, the primary fieldbus used in the LHC. An overview of the applications utilizing WorldFIP is given, along with their requirements.

The concepts introduced in this chapter are central to the rest of the manuscript, where wireless solutions are explored. The first issue addressed is the short-term cabling problem caused by the large number of devices already deployed and the difficulties in adding new ones due to high cable deployment costs and full cable trays. To overcome this problem, the current WorldFIP infrastructure at the field level (between FEC and instruments) should be extended with wireless capabilities. Chapter 4 discusses a feasibility study based on the architecture presented in this chapter, demonstrating how wireless capabilities could be added without modification. However, the study also highlights challenges in extending a wired network to wireless, which are then considered in Chapter 5, where the design of a custom wireless protocol is discussed. This development considers all the characteristics of industrial control networks and the requirements of the current applications described earlier.

Additionally, because the control interface will be exposed to radiation, its radiation tolerance is crucial. Thus, the first topic addressed by this doctoral research are the radiation effects on a radio frequency circuit, the core of any wireless system, and their impact on a wireless communication system.

3 Study of radiation effects on wireless communication

Wireless communications have become ubiquitous, especially in consumer electronics, due to their numerous advantages and versatile applications. Over time, communication systems have evolved significantly, with advancements in computing technologies, communication algorithms, and radio frequency (RF) hardware. These developments have resulted in modern wireless systems that are not only complex and high-performing but also capable of adapting to various scenarios. However, integrating these systems into the LHC control network at CERN presents challenges.

One of the primary concerns is the impact of radiation effects, which can compromise the reliability and availability of communication links within the network. As systems become more complex, the potential sources of failure multiply, making testing and validation increasingly challenging. Therefore, before deploying such technologies, it is essential to assess how radiation might affect the control network's performance.

Despite these challenges, wireless communications offer significant advantages for CERN's particle accelerators, particularly in terms of flexibility compared to wired solutions. Wireless systems can simplify installation and maintenance processes. Given the long operational lifetime of accelerators like the LHC, future adaptability is crucial, as new requirements may emerge over time. Switching to an entirely new solution would be resource-intensive due to the vast number of devices deployed, making it impractical. This has driven research into technologies that can be adapted as needs evolve. One promising technology that has gained traction in recent years, thanks to advancements in RF hardware, is Software-Defined Radio (SDR).

The concept of Software-Defined Radio (SDR) involves a wireless system whose functionalities—such as protocol, operating frequency, and bandwidth—can be adapted through software without altering the hardware. This flexibility makes SDR a viable solution for the CERN accelerator complex. With a wireless link based on SDR, the communication interface can be updated as new requirements emerge, without the need for hardware modifications to the systems deployed in the tunnel, which would otherwise be too resource-intensive. Consequently, this technology was further explored in this doctoral research to develop a wireless link, as discussed in Chapters 4 and 5.

Achieving an SDR architecture involves combining an RF front-end with programmable logic. The RF front-end manages RF transmission and reception, while the programmable logic handles signal processing to transmit and recover data. Therefore, it is crucial to assess the performance of both, evaluate how radiation effects impact subsequent stages of signal processing, and understand the broader effects on the communication link and, ultimately, the control network. This chapter addresses these issues on the RF front-end.

The first part of the chapter introduces the fundamentals of digital wireless communication. These concepts are central to the rest of the manuscript, as they are essential for understanding how radiation effects on RF circuits can impact wireless communication and for understanding the implementation details of the wireless link proposed in Chapters 4 and 5.

The second part of the chapter reviews the hardware used for implementation of wireless communications, and the related state-of-the-art testing methodologies, focusing on how their performance is evaluated outside and inside radiation environment. The limitations of current methodologies will be presented, followed by the introduction of a new methodology developed in the context of this doctoral thesis. This new approach overcomes existing limitations, allowing for better characterization of performance and a more accurate evaluation of radiation's impact on the communication link.

Finally, the last part of the chapter will discuss the test results of a state-of-the-art RF transceiver evaluated under radiation using the proposed methodology. Emphasis will be placed on how radiation effects influence the later stages of signal processing, the overall impact on the communication link, and potential mitigation strategies.

3.1 Background on wireless communication

Although wireless systems have evolved and become more complex over time, the fundamental principles of wireless communications have remained consistent.
The purpose of a wireless communication system is to transmit and receive data through electromagnetic waves over the air. In simple term, this involves a radio transmitter, which modifies the properties of a carrier waveform—such as amplitude, phase, and frequency—using a signal that contains the information to be transmitted. At the receiver, these modified properties are detected and processed, thereby recovering the original information. However, many other operations take place within a wireless communication system to effectively transmit and receive information.

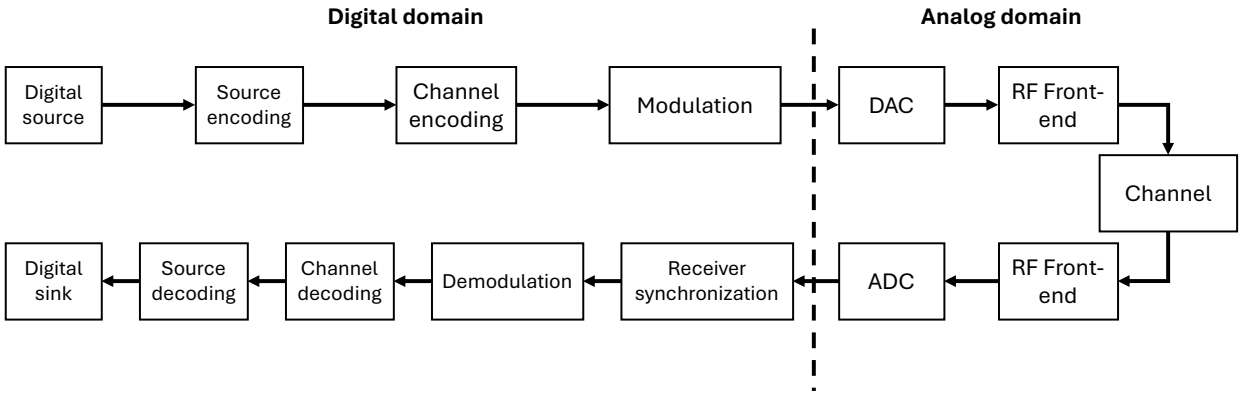

Figure 25 - Basic building blocks of a digital communication system

Figure 25 outlines the fundamental building blocks required to implement a digital communication system. The process begins with data in binary form, representing real-world information such as an image, sound, temperature reading, or radiation level from a sensor. This data is processed before transmission. The first block in the diagram, the source encoder, removes any redundant information to ensure a more efficient transmission by reducing the number of bits that need to be sent. Next, controlled redundancy is added by the channel encoder to protect the data from potential errors in a noisy environment. At this stage, the

binary information is converted into symbols, which define unique waveform properties, such as amplitude, carrier frequency, and phase. These discrete-time symbol samples are then resampled and converted into an analog waveform in the baseband domain using a digital-to-analog converter (DAC). In the analog domain, this waveform modulates a carrier wave at RF frequency, which serves as the vehicle to transport the information. At the receiver, the reverse operations take place: the analog signal at the RF carrier frequency is captured, down-converted to the original baseband frequency, and sampled by an analog-to-digital converter (ADC). The receiver then processes this signal to reconstruct the original information.

In the following, the different building blocks will be analysed with particular emphasis on digital modulation/demodulation, as these components are central to the development of the radiation-tolerant wireless PHY discussed throughout the manuscript.

3.1.1 Source encoding

The primary aim of a communication system is to efficiently transmit information across a medium. To achieve this, it's important to eliminate any unnecessary redundancy in the data, which can reduce the amount of information that needs to be sent. By doing so, the system can decrease transmission time, conserve resources, and lower power consumption. Source encoding plays a key role in this process by taking a sequence of symbols and transforming it into a new sequence that is nearly random and has minimal correlation between symbols. This transformation minimizes redundancy, ensuring that radio resources are used more effectively and reducing the waste that would occur with the original, more redundant data.

3.1.2 Channel encoding

While source encoding focuses on reducing data redundancy, channel encoding introduces intentional redundancy to protect information from corruption during transmission. Wireless channels are susceptible to various errors due to interference, fading, propagation loss, and other factors. To ensure reliable communication even in such conditions, a channel encoding scheme is employed. The redundancy introduced by the encoding process in the transmitter, is known at the receiver, so that it can be identified and removed to recover the original information.

A channel encoding scheme works by mapping the original message, which consists of a set of $v_l$words (where $l = 1,2,3, \dots, 2^K$ and K is the length of the source-encoded output vector), to a set of codewords c_N where N is the length of the code. During this process, the encoder adds $N - K = r$ additional bits, representing the redundancy used for error correction and recovery of the original message. In communication systems, the effectiveness of a coding scheme is often measured by its code rate, defined as the ratio of the number of information bits to the total length of the codeword, K/N.

3.1.3 Digital modulation

There are two primary types of modulation: analog and digital. In analog modulation, the carrier waveform is continuously varied in response to the analog information signal, such as an audio signal from a microphone. Common types of analog modulation include Amplitude Modulation (AM) and Frequency Modulation (FM), where the amplitude and frequency of the carrier signal are altered, respectively, to convey the information. These methods are among the earliest forms of modulation and are still used today, for example, in AM and FM radio broadcasts.

Over time, modulation techniques have evolved towards digital modulation, which allows for the transmission of more information compared to analog methods. Digital modulation requires data to be represented by binary digits, which are translated into specific properties of the waveform that uniquely represent each bit by a digital transmitter, and then recovered by a digital receiver.

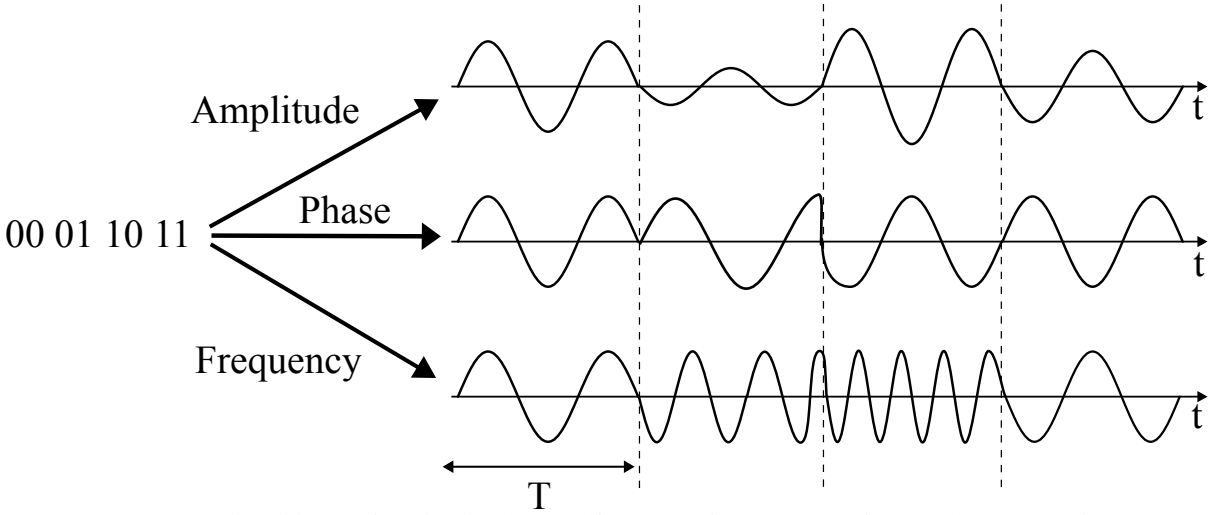

Figure 26 – Digital modulation of amplitude, phase and frequency of a carrier waveform used to convey information

For example, Figure 26 demonstrates how binary digits can be represented by altering different properties of the waveform within a time interval T. At the receiver, the reverse process, demodulation, involves detecting these waveform properties to recover the original binary information.

Many digital modulation schemes add an extra step: grouping a set number of binary digits, b, into a symbol. Each symbol then defines properties of the waveform. For every possible combination of 2^b values, there is a corresponding signal $s_i(t)$, where $1 \leq i < 2^b$, used to modulate the continuous waveform. Various methods exist for mapping binary data to symbols, which are then used to modulate the waveform. These methods differ in which property of the carrier waveform—such as phase, amplitude, or frequency—they modulate.

Nowadays, the most common way of generating the carrier waveform to represent a symbol is by combining cosine and sine waves, utilizing the principle of orthogonality. These waves are orthogonal because they are shifted by 90 degrees relative to each other, ensuring they do not interfere when added together. The cosine wave represents the in-phase (I) component, while the sine wave, shifted by 90 degrees, represents the quadrature (Q) component. When these two waves are combined, they produce a new sine wave with a different phase and amplitude. The phase of this resulting wave can be adjusted by varying the amplitudes of the I and Q components, according to the trigonometric identity $a\cos(x) + b\sin(x) = A\cos(x - \phi)$.

The primary advantage of this method is that both the phase and amplitude of the resulting sine wave can be controlled by simply adjusting the amplitudes of the I and Q components, without the need to directly modify the phase of the cosine or sine waves. This approach simplifies the design of both the transmitter and receiver as will be explained in the dedicated section.

Since the carrier waveform is defined by its magnitude and phase, the IQ representation allows it to be plotted on a complex plane, known as a constellation diagram, where I and Q represent

the real and imaginary part respectively. This diagram displays the signal as a two-dimensional scatter plot, where the angle of a point from the horizontal axis corresponds to the phase shift relative to a reference phase, and the distance from the origin indicates the signal's amplitude.

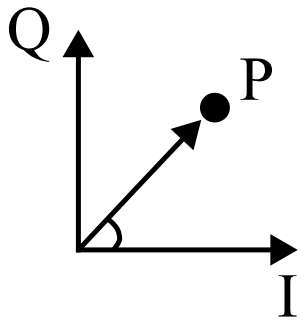

Figure 27 – Constellation diagram of one symbol corresponding to a signal whose amplitude is the distance from the origin, and whose phase is the angle from the horizontal axis

Since binary digits are grouped into symbols, and each symbol corresponds to a signal with a specific amplitude and phase, it means that the signals, hence the symbols, can be represented on the constellation diagram as points, known as constellation points. For instance, in the constellation diagram shown in Figure 27, the point *P* represents one possible symbol of a modulation, which corresponds to a signal with an amplitude given by $\mathrm{A} = \sqrt{\mathrm{I}^2 + \mathrm{Q}^2}$ and a phase given by $\phi = \arctan(\mathrm{Q}/\mathrm{I})$. The amplitude and phase of this signal can be adjusted by simply altering the I and Q components.
Thus, the constellation diagram provides a visual representation of all possible symbols that can be transmitted by a system using that particular modulation.

Depending on the properties of the signal and the number of bits grouped into a symbol, different modulation schemes are employed, each with distinct characteristics regarding power efficiency and transmission rate.

Regarding the transmission rate, it can be computed using the following formula:

$$\mathrm{R_b} = \mathrm{b_{sym}} \times \text{Symbol Rate}$$

where $\mathrm{b_{sym}}$ represents the number of bits per symbol, calculated as:

$$\mathrm{b_{sym}} = \log_2(\mathrm{M})$$

Here, M is the total number of symbols used in the modulation scheme. The Symbol Rate is the number of symbols transmitted per unit time. Thus, the transmission rate $\mathrm{R_b}$ reflects the total number of bits transmitted per unit time.

Power efficiency instead is a metric for evaluating the effectiveness of mapping bits to symbols (or signals) and understanding the transmit power required for each symbol.

To estimate power efficiency [58], the energy per symbol can be defined for a signal $\mathrm{s(t)}$ as

$$E_s = \int_0^T s^2(t)dt,$$

where T is the symbol period. For a modulation scheme with multiple symbols, the average symbol energy $\overline{\mathrm{E_s}}$ can be calculated as a weighted average:

$$\overline{\mathrm{E_s}} = \mathrm{P(s_1)} \cdot \int \mathrm{s_1^2(t)\ dt} + \ \ldots \ + \ P(s_\mathrm{M}) \cdot \int s_\mathrm{M}^2(t)\ dt$$

Where $\mathrm{P(s_i(t))}$ represents the probability of a given symbol occurring. The average energy per bit $\overline{E_b}$ can then be approximated by dividing the average symbol energy by the number of bits per symbol, $\mathrm{b} = \log_2(\mathrm{M})$, where M is the total number of symbols in the modulation scheme. Finally, the power efficiency can be expressed as:

$$\epsilon_\mathrm{p} = \frac{\mathrm{d_{min}^2}}{\overline{E_b}}$$

where $\mathrm{d_{min}}$ is the minimum Euclidean distance between two symbols on the constellation diagram, and $\overline{E_b}$ is the average energy per bit.

In the following, the different modulations and their characteristics will be presented and compared.

3.1.3.1 Pulse Amplitude Modulation

One of the most straightforward characteristics of a waveform that can be modified to convey information is its amplitude. Modulation techniques that encode information in the amplitude of the carrier waveform are known as amplitude modulation. In such schemes, demodulation involves detecting the varying amplitude levels of the carrier signal at each symbol period.

The simplest form of amplitude modulation is known as Pulse Amplitude Modulation (PAM), which uses a series of pulses to represent data. In its most basic form, a single binary digit is encoded into a symbol, associated to a waveform $\mathrm{s(t)}$ with a specific amplitude level. This leads to a simple modulation rule:

- $1 \rightarrow s_1(t)$
- $0 \rightarrow s_2(t)$

Here, $\mathrm{s_1(t)}$ and $\mathrm{s_2(t)}$ have distinct amplitude levels, with each waveform defined over a time period T and being zero otherwise. This form of modulation is called Binary PAM (B-PAM) because each bit corresponds to a unique waveform, resulting in only two symbols. Since the duration of each symbol is equivalent to the duration of a bit, the bit rate for B-PAM is given by $R_b = 1/T$ bits per second.
More complex forms of PAM can be implemented by grouping multiple bits (M) together into a single symbol, with each symbol uniquely identified by a waveform with a specific amplitude level. This approach is known as M-ary PAM, where the number of symbols increases as more bits are combined, allowing for more efficient data transmission. If M is number of symbols, each symbol groups $\mathrm{b} = \log_2(M)$ bits, hence giving a transmission rate of $R_b = b/T$.

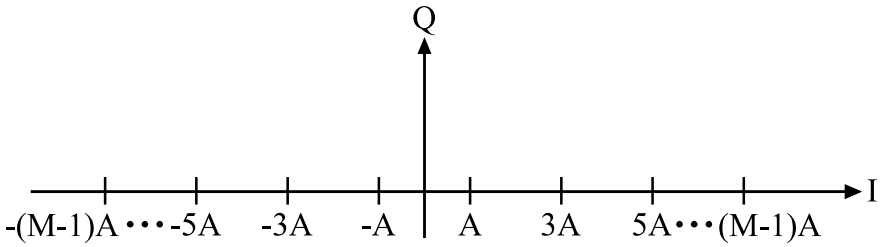

Figure 28 - M-ary PAM constellation diagram

Since each symbol in M-ary PAM modulation is associated with a signal that has a distinct amplitude level, which is a real number, the corresponding constellation point lies on the real axis of the constellation diagram. This occurs because there is no phase component to introduce an imaginary part. Therefore, in M-ary PAM modulation, where amplitude levels vary, the constellation diagram is represented as shown in Figure 28.

The power efficiency for this modulation scheme can be calculated as described previously, considering that the minimum Euclidean distance is between signals with amplitudes $-A$ and A. The average symbol energy can be computed by exploiting the symmetry of the constellation, resulting in a power efficiency given by:

$$\epsilon_p = \frac{12b}{2^{2b} - 1}.$$

3.1.3.2 Quadrature Amplitude Modulation

While Pulse Amplitude Modulation (PAM) is a one-dimensional modulation scheme where symbols are located only on the real axis, Quadrature Amplitude Modulation (QAM) is a two-dimensional modulation that utilizes both the in-phase and quadrature components. Since these components are orthogonal, QAM effectively doubles the transmission efficiency without additional cost.

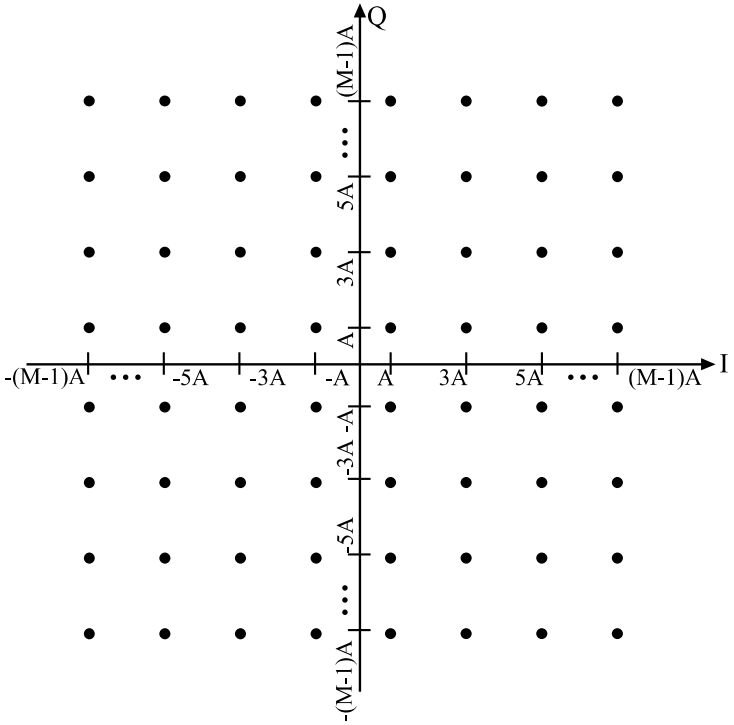

Figure 29 - QAM constellation diagram

Figure 29 illustrates the constellation diagram for M-QAM modulation. In this diagram, each constellation point is characterized by its I and Q amplitude levels, representing both the real and imaginary axes. The power efficiency for M-QAM can be expressed as:

$$\epsilon_p = \frac{3!\,b}{2^b - 1}$$

For the minimum number of symbols, $M = 4$, the power efficiency of QAM matches that of Binary PAM (B-PAM). However, QAM offers a doubled transmission rate due to the utilization of both the in-phase and quadrature components, highlighting the advantage of this technique.

3.1.3.3 Phase Shift Keying

Phase Shift Keying (PSK) is a modulation technique that conveys information by altering the phase of the carrier wave rather than its amplitude. In PSK, a finite number of phases are used, each corresponding to a specific pattern of bits. The receiver detects these phase changes by comparing the phase of the incoming signal with a reference signal, allowing the recovery of the original data. Unlike amplitude modulation, where the amplitude is crucial for information retrieval, PSK maintains a constant amplitude, making it particularly effective in environments with amplitude-related challenges, such as power-limited channels (e.g., deep space communications) or channels with amplitude distortion or fading (e.g., urban environments).

The simplest form of PSK is Binary PSK (BPSK), which uses two symbols to represent one bit of information. However, Quadrature PSK (QPSK), which employs both the in-phase and quadrature components, is more commonly used.

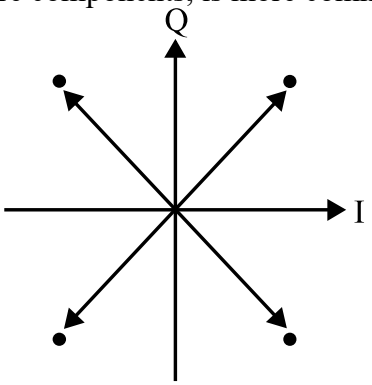

Figure 30(a) - QPSK constellation diagram

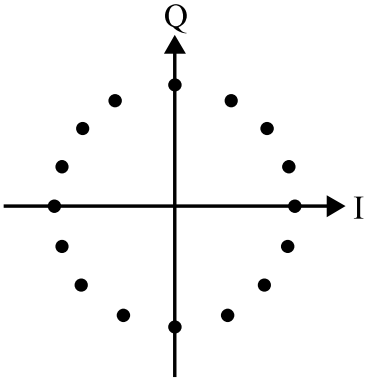

Figure 30(b) - 16-PSK constellation diagram

Its constellation diagram is shown in Figure 30a. QPSK achieves a power efficiency comparable to other modulation schemes but doubles the data rate by encoding two bits per symbol.

In general, M-PSK modulation uses M symbols, where the constellation diagram consists of M equally spaced points on a circle, as illustrated in Figure 30b for M = 16. While increasing M allows for higher data rates, it also decreases the spacing between points, reducing error robustness. However, since information is encoded in the phase, PSK is well-suited for power amplifiers and provides robustness against amplitude distortion in transmission channels.

3.1.3.4 Pulse Position Modulation

In Pulse Position Modulation (PPM), information is encoded by transmitting a single pulse at specific time shifts. Each pulse represents b bits, corresponding to one of 2^{M} possible time shifts. This process is repeated every T seconds, resulting in a transmitted bit rate of $R_b = M/T$. PPM offers distinct advantages and disadvantages compared to modulation schemes based on phase and amplitude.

One significant disadvantage of PPM is the complexity of receiver synchronization. The receiver must align its local clock with the start of each symbol, which can be challenging. To address this, Differential PPM (DPPM) is often used, where each pulse is encoded relative to the previous one, allowing the receiver to measure only the time difference between successive pulses. However, PPM is particularly sensitive to multipath interference, where echoes of the transmitted signal can complicate decoding since reception relies on precise timing.

Despite these challenges, PPM has notable advantages. It is relatively immune to noise affecting amplitude, making it power efficient as information is transmitted through timing, allowing power to be concentrated in short pulses and reducing average power consumption. Additionally, PPM simplifies receiver design, as the receiver only needs to detect the timing of the pulse, unlike other receivers that must also detect changes in amplitude or phase. This simplification is a key benefit, as will be further explored in the fourth chapter for the design of the wireless extension for the LHC control equipment.

3.1.4 Receiver synchronization

After data is modulated, it can be digitized, transmitted via RF, and captured by an RF receiver for processing. However, before demodulation can occur, several preparatory steps are necessary. Since the receiver operates on a different node from the transmitter, their clocks are not aligned. This misalignment introduces artifacts that must be compensated for in the receiver before demodulation occurs. Consequently, various techniques and algorithms are required.

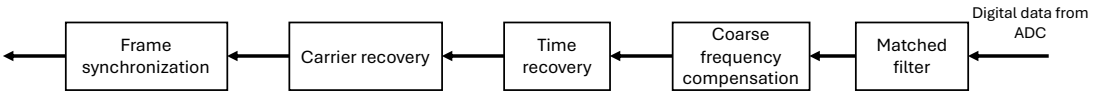

Figure 31 - Building blocks for the receiver synchronization

Figure 31 illustrates the basic building blocks to be implemented in a receiver to compensate for the artifacts introduced by the wireless channel. These data recovery algorithms operate on digitized data from the RF front end, meaning that if the front end is impacted by radiation, the digitized data can be distorted, affecting the recovery stages. The following sections provide an overview of the recovery techniques used in modern systems and adopted for developing the wireless links proposed in this work.

3.1.4.1 Matched filtering

When discussing matched filtering for a wireless system, it usually refers to the process of pulse shaping at the transmitter and matched filtering at the receiver. The processes serve three main purposes: making the signal suitable for transmission by limiting its effective bandwidth, increasing the SNR of the received waveform, and removing intersymbol interference.

Digital data is represented by 0s and 1s, but sharp transitions in the time domain result in a very wide frequency band in the frequency domain. To limit the used bandwidth, the signal is

filtered through a pulse-shaping filter, reducing sharp frequency transitions and resulting in a spectrally efficient signal. This process typically performs upsampling or interpolation of the signal, which however will reduce the effective bandwidth but make easier the recovery process at the receiver, since there will be more copies of the same symbol to make the decision. Hence, data will be produced at the same rate, but it will not use the whole bandwidth available. The filter commonly used for these operations is the square root-raised cosine (SRRC) filter, which is prevalent in communication systems [58]. This filter is also implemented in the receiver, so the combined response of the two forms a raised cosine filter, whose impulse response is described as

$$h(t) = \begin{cases} \dfrac{\pi}{4T} sinc\left(\dfrac{1}{2\beta}\right), & t = \pm\dfrac{1}{2\beta} \\ \dfrac{1}{T}\text{sinc}\left(\dfrac{t}{T}\right)\dfrac{cos\left(\dfrac{\pi\beta t}{T}\right)}{1-\left(\dfrac{2\beta t}{T}\right)^2}, & otherwise \end{cases}$$

where T is the symbol period and $\beta \in [0,1]$ is the roll-off factor, which determine the transition sharpness of the frequency response. If β is close to zero, the transition are very sharp, therefore reducing the effective bandwidth of the signal, but increasing ISI and signal distortion. With roll-off closer to 1, the transitions are smoother, however the bandwidth usage increases.

3.1.4.2 Timing synchronization

The timing error in wireless communication systems arises because the transmitter and receiver are implemented on different nodes, causing their clocks to be unaligned.

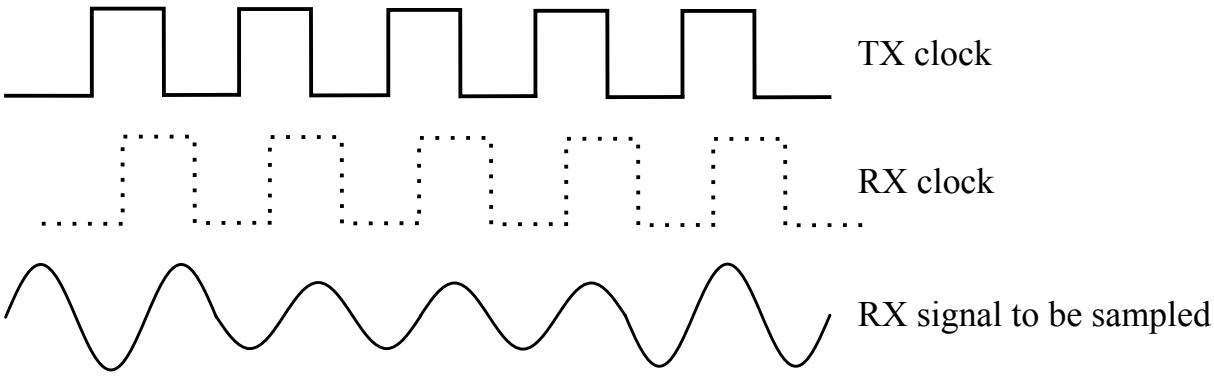

Figure 32 - Timing error in communication system

The impact of timing error can be visualized in Figure 32, where the transmitter and receiver clocks, along with the analog signal to be sampled at the receiver, are depicted. Assuming the receiver samples the signal on the rising edge of the clock, a small delay between the transmitter and receiver rising edges can lead to incorrect signal sampling, resulting in wrong binary digits

after demodulation. Therefore, a timing synchronization algorithm is necessary and must be implemented in the receiver system.

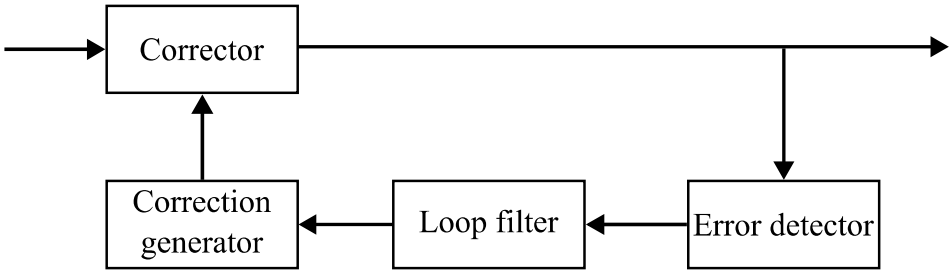

Figure 33 - PLL structure for a recovery algorithm

Synchronization can be performed using various methods that are well described for analog and digital implementations here [59]. Among them, the Phase-Locked Loop (PLL) based strategy was implemented in this doctoral work due to its robustness and simplicity. The PLL strategy involves a feedback loop, whose block diagram is shown in Figure 33, which aims to apply a correction to an incoming signal until the error is minimized. This strategy is the same used to correct not only timing errors, but frequency error as well, as it will be described later. The main concept involves first computing an offset between the incoming signal and a reference signal in the error detector block, which is designed based on the ideal signal. This offset is then passed through the loop filter, which governs the PLL loop's dynamics, including its range, lock time, and responsiveness. The correction generator block then produces a correction factor and applies it to the input signal for the next iteration. The error detector and correction generator blocks are designed depending on the factor to be corrected, such as timing or frequency offset, while the loop filter is shared and can be reused. The loop filter is commonly implemented as a proportional-plus-integrated filter with the transfer function $\mathrm{F(z)} = \mathrm{G1} + \frac{\mathrm{G2}}{1-\mathrm{z}^{-1}}$, where $\mathrm{G1}$ and $G2$ are gains calculated [59] based on the damping factor ζ and loop bandwidth B_{Loop}, expressed as

$$\theta = \frac{B_{Loop}}{M(\zeta + 0.25/\zeta)} \qquad \Delta = 1 + 2\zeta\theta + \theta^2$$

$$G1 = \frac{4\zeta\theta/\Delta}{M} \qquad G2 = \frac{4\theta^2/\Delta}{M}$$

Where B_{Loop} is a normalized frequency that can range $B_{Loop} \in [0,1]$ and ζ determines the responsiveness and the stability of the PLL.

Once the loop filter is implemented, the error detection algorithm is required. Key timing error detection algorithms used in digital communication systems include Zero-Crossing (ZC), Muller and Mueller (MM) [60], and Gardner [61]. The ZC method detects timing errors by monitoring when the received signal crosses zero, ideally at the midpoint between symbols. While simple to implement, its accuracy diminishes in noisy or distorted environments. The MM algorithm uses correlation between the received signal samples and expected symbol values to detect timing errors, making it more sophisticated. The Gardner algorithm, often used in systems based on BPSK or QAM, calculates timing errors by comparing two samples taken per symbol period. This blind method does not require knowledge of the transmitted symbols but may be less effective in channels with significant ISI.

Once the error is calculated with one of these algorithms, it is passed to the loop filter, whose output is then passed to the correction generator. The correction generator is responsible for providing the correction factor, representing the new sampling period to correct for the timing

offset. This block fractionally delays the signal by a factor μ. An example implementation for digital systems can be found in [59]. The correction process repeats indefinitely until the error approaches zero.

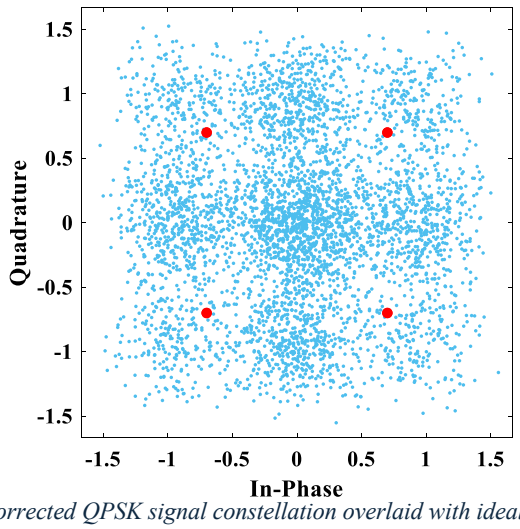

Figure 34 - Uncorrected QPSK signal constellation overlaid with ideal constellation in red

The effect of timing error can also be visualized on the constellation diagram. If the receiver samples the signal at the wrong instant, the timing error will result in clustering or scattering of the constellation points. For instance, Figure 34 shows the constellation diagram of a QPSK signal without timing correction, where the timing offset causes incorrect sampling, making the constellation at the receiver differ from that of a QPSK modulation.

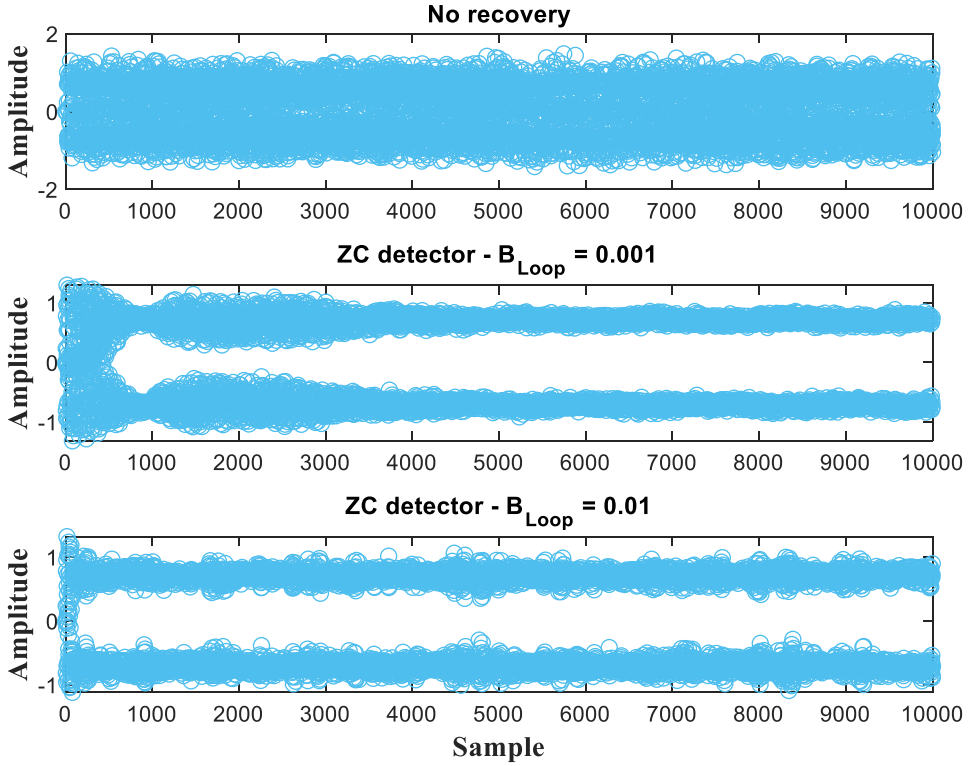

Figure 35 - example of different outcomes for error detection algorithms, ZC and different parameters

The outcome of error detection depends on the algorithm used and the signal type. Figure 35 shows the application of timing correction based on the ZC method for error detection, on a

QPSK signal (top). The bottom figures represent the QPSK signal after being processed with a ZC method with different bandwidth, illustrating how the system's response time changes and requires more samples to converge to a solution. As will be seen later, radiation effects also impact this process by disrupting the algorithm, requiring more time to reconverge to a solution.

3.1.4.3 Carrier synchronization

Since the transmitting and receiving nodes in a wireless communication system are distinct, their local oscillators (LOs) generally have an offset, even if set to the same carrier frequency. This offset can be dynamic, causing the LOs to drift relative to each other. The offset may include random phase noise, frequency offset, drift, and initial phase mismatches, all of which can be observed from the constellation diagram.

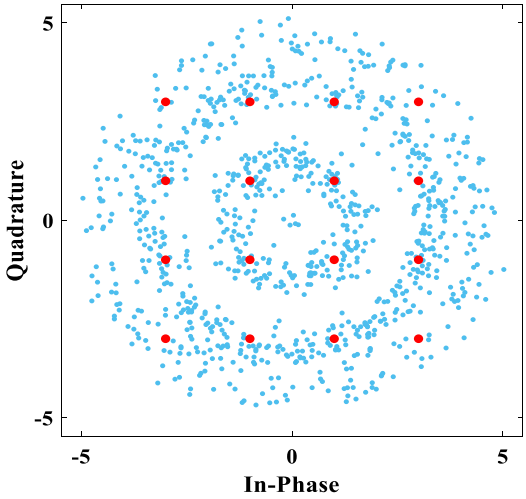

Figure 36 – 16-QAM signal constellation at the receiver with a 400 Hz offset compared to transmitter

For example, consider a 16-QAM signal transmitted between two devices with a 400 Hz offset between their LOs. The resulting constellation diagram at the receiver, shown in Figure 36, exhibits rotation over time compared to the expected QAM constellation, likely leading to demodulation errors. To address this, frequency offset compensation algorithms are implemented in the receiver.

These algorithms are divided into coarse and fine frequency corrections. Coarse frequency corrections handle large offsets, while fine frequency corrections address smaller offsets.

Coarse frequency offset correction can be data-aided or non-data aided (NDA) [58]. Data-aided methods rely on a known sequence to estimate the offset, which can limit system throughput but offer better performance than blind methods. Blind methods operate over the entire sequence without requiring a predetermined sequence. The most famous implementation is the FFT-based NDA coarse frequency estimation, which estimates the frequency offset by computing the FFT of the received signal. This method requires a large amount of data. Other techniques, such as the one by Luise [62], require less data and are more suitable for burst applications, although they provide biased estimates compared to FFT-based methods. After coarse frequency correction, a small frequency offset may still be present and must be corrected before demodulation.

Fine frequency correction schemes aim to produce a stable constellation for demodulation. If a frequency offset is present, the constellation points will rotate counterclockwise for a positive

offset and clockwise for a negative offset, with the rotation rate determined by the frequency offset.

The compensation algorithm for fine frequency correction follows the same principle as symbol correction, utilizing a PLL loop. The frequency offset error is calculated based on the known signal constellation, and a correction factor is applied to the input signal in the next iteration. This open-loop technique drives the frequency offset to zero after several iterations, with the number of required samples depending on the loop filter parameters, which determine the system's responsiveness.

3.1.4.4 Frame synchronization

The final stage of a digital receiver, necessary before retrieving the original transmitted data, involves detecting where the actual data resides within random noise. Once recovery algorithms have corrected for impairments between the transmitting and receiving nodes, the data is ready to be demodulated. At this stage, it can be assumed that each incoming sample represents a symbol that has been corrected for all impairments. However, in a realistic system, data is transmitted in frames or packets, so it is initially unknown when the frame starts. Therefore, it is essential for the receiver to distinguish useful data from noise.

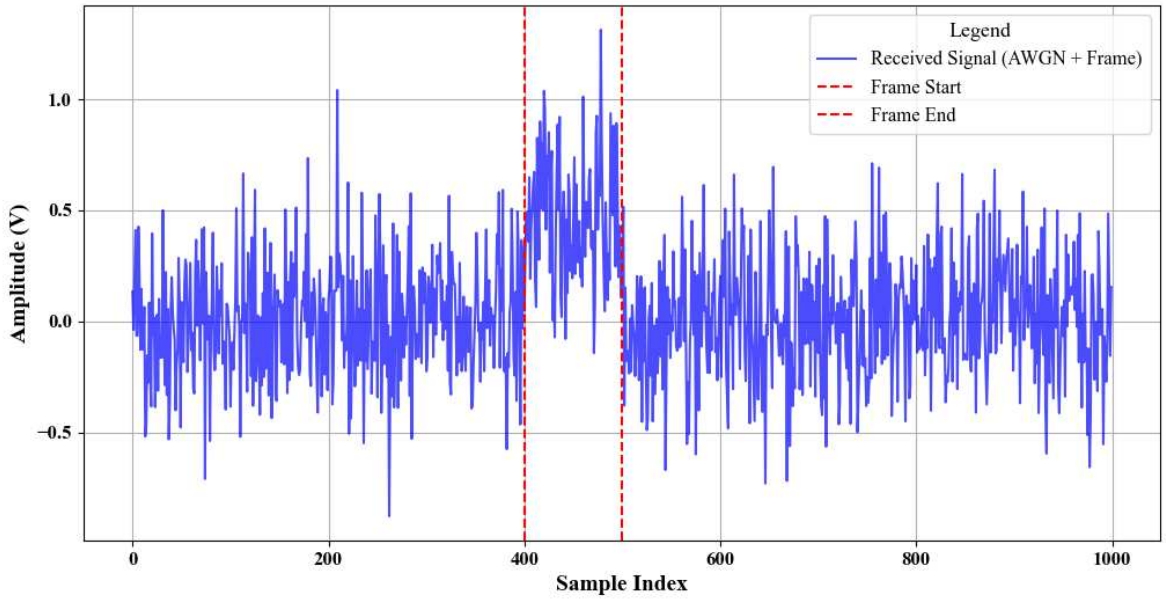

Figure 37 – Incoming signal at a wireless receiver. A frame containing valid data and marked within red lines is mixed with noise

Figure 37 illustrates this, showing an incoming signal at a wireless receiver where a frame containing valid data, marked within red lines, is mixed with noise. Depending on the receiver structure, frame synchronization can occur before or after demodulation. For instance, in a system that continuously transmits data, frame synchronization can occur after demodulation. A fixed sequence, known to the receiver, can be inserted into the data stream. In this case, the receiver can be designed so that symbols are demodulated continuously, and when the known sequence of bits is detected, the following data bits are treated as useful data.
On the other hand, in most modern wireless systems, packets are transmitted sporadically, making it more challenging to detect packets within the noise. Additionally, it may be the case that the start of a frame is modulated differently from the rest, with decoding information

contained at the frame's beginning. In this scenario, frame synchronization occurs before demodulation.

In both cases, the most widely adopted technique for frame synchronization is based on a marker sequence. This is a predefined sequence added by the transmitter to mark the start of useful data, and the same sequence, known to the receiver, is used to determine the start of the frame. Given a set of N binary sequences b_n of length L, and a known sequence d, the problem consists of finding the sequence dwithin the set of binary sequences. This can be done using cross-correlation. The cross-correlation can be expressed as

$$C_{(d,b)} = \sum_{m} \left(d^*(m) b_n(m+k) \right)$$

The value of $C(d, b)$ will be maximized for one $d = b_n$, distinguishing it from the other sequences b_{n-1}.

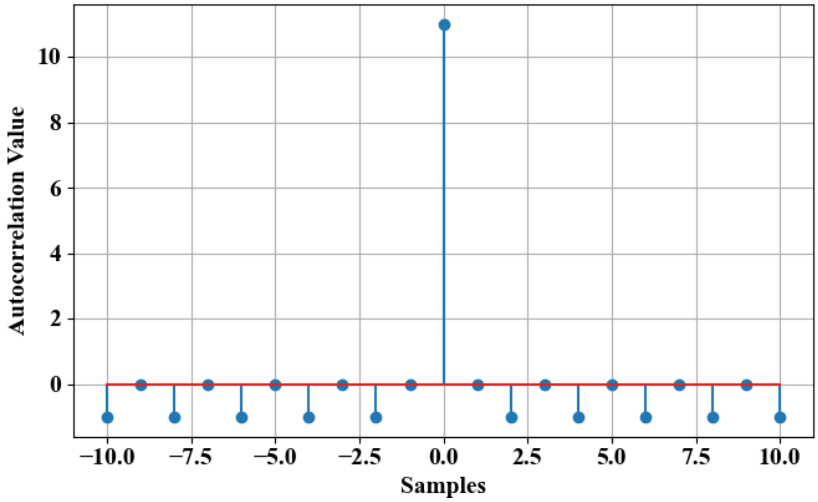

Figure 38 - Autocorrelation of a Berker sequence of length L = 11

The Barker sequence is commonly used for this purpose due to its autocorrelation properties, which minimize or idealize off-peak correlation. There are only nine known Barker sequences for L ∈ [1,2,3,4,5,7,11,13], and Figure 39 shows the autocorrelation for a sequence of length L = 11. In communication systems, one of these sequences is often replicated and included at the start of the transmitted packet to improve recognition at the receiver.

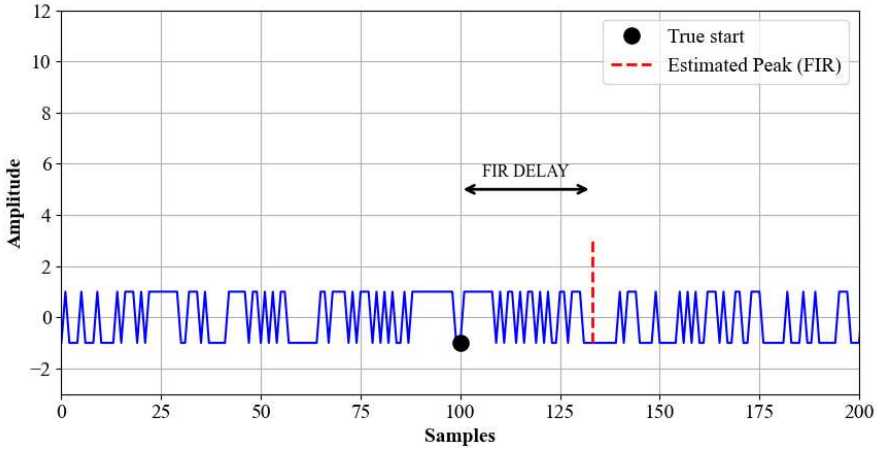

Figure 39 - Sequence of bits including the Berker sequence at the beginning. The dotted red line presents the output of a FIR filter that can be used to detect the position of the Berker sequence

This concept can be implemented in a receiver, where a correlator circuit is employed using a Finite Impulse Response (FIR) filter structure. The filter's coefficients are set to match the sequence of interest, typically a Barker sequence due to its ideal autocorrelation properties. The output of the FIR filter is continuously monitored, and when a peak is detected—corresponding to the Barker sequence—the start of the frame is identified, as shown in Figure 39. By accounting for the filter delay, this signal can then be used to determine when to begin demodulating the incoming symbols to get the actual data.

3.2 Radiation Testing methodologies for RF Integrated circuits

This section builds upon the theory of wireless communication by detailing the RF hardware used in the wireless system's implementation. While the fundamental components of an RF front-end have remained consistent, the architectures for implementing them have evolved. This section will provide an overview of these architectures, up to the most used design in modern systems. Following the architecture review, a survey of past radiation tests on RF devices will be presented. However, these past tests did not consider how radiation effects influence the actual transmission of data and impact downstream algorithms, such as those in radio receivers. To address this gap, a new test methodology has been developed as part of this doctoral research. This approach builds on traditional RF testing methods to assess the impact of radiation effects. The section will then present the application of this new methodology to a state-of-the-art RF transceiver, highlighting results from proton irradiation tests. Finally, the results will be discussed in the context of real-world applications, demonstrating how the methodology helps identify the impact of radiation effects on the RF front-end and its influence on the overall communication link.

3.2.1 Radio Frequency Front-End

Once digital information is modulated, it is converted into a waveform and transmitted over the air through an RF front-end device. This same device captures the waveform, allowing the retrieval of the original information.

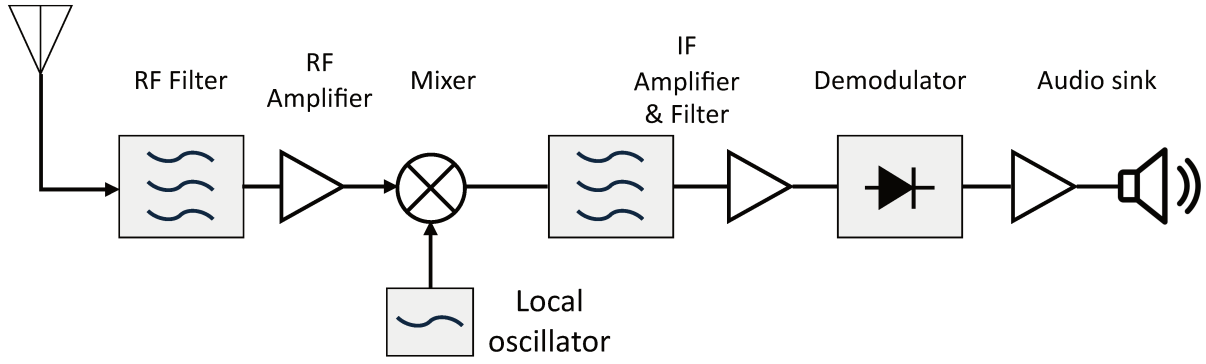

Figure 40 – Radio receiver based on the superheterodyne architecture

Traditionally, radio designs have relied on the superheterodyne architecture, as depicted in Figure 40. In this design, the antenna captures the radio signal, which is then tuned and amplified by an RF amplifier due to the weak signal output from the antenna. A band-pass filter follows, removing unwanted frequencies outside the desired reception range. The signal is then mixed with a sine wave generated by a local oscillator, producing an output that includes the RF signal at the local oscillator frequency and two additional frequencies at $f_{RF} + f_{LO}$. These

additional frequencies are filtered out by the next stage, leaving only the RF signal at the local oscillator frequency.

The signal then passes through an intermediate frequency (IF) stage, where it undergoes further processing via another mixer and band-pass filter, shifting it from the RF frequency to a fixed frequency, f_{IF}. This fixed IF stage allows for the optimal tuning of the IF amplifier stages, and since the center frequency is lower than the RF frequency, the design becomes simpler and more cost-effective. After this stage, the demodulator extracts the original information from the IF signal. However, a significant drawback of this architecture is the fixed IF stage, which complicates the creation of a common platform for multiple wireless systems.

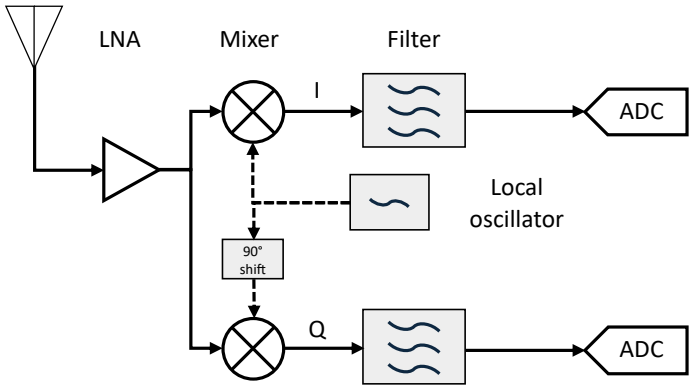

Figure 41 – Architecture of a Zero-IF radio receiver

To address this limitation, the zero-IF architecture emerged as an alternative. This architecture, illustrated in Figure 41, eliminates the IF stage by using a single mixer stage, where the local oscillator is set to the frequency of interest. This setup translates the signal directly to baseband into in-phase and quadrature components, where designing filters is simpler and less expensive compared to custom RF filters. Moreover, this architecture also allows to exploit the advantages offered by digital modulation exploiting quadrature components as previously described. However, maintaining a perfect 90° offset between the I and Q components is challenging due to factors such as process variation, leading to imbalances. Historically, this architecture was limited because its implementation with discrete components was difficult, leading to mismatches in monolithic devices and PCBs. Additionally, implementing the processor or DSP separately from the analog chain made it challenging to execute a quadrature algorithm effectively.

3.2.2 Modern wireless architecture: Software-defined Radio

Recent advancements have enabled the integration of both analog and digital chains onto a single piece of silicon, creating new opportunities for device development. Modern RF integrated circuits (RFICs) now incorporate the entire RF, analog, and DSP chains on a single CMOS device, with real-time correction algorithms embedded to account for process and temperature variations. These RFICs support a wide range of frequencies and bandwidths, making them versatile for various applications. This development has paved the way for Software-Defined Radios (SDRs), which are radio platforms where most functionalities are defined by software. By combining a flexible RFIC with programmable logic, it is theoretically possible to implement any wireless application without changing the hardware—only the code in the programmable logic needs to be updated.

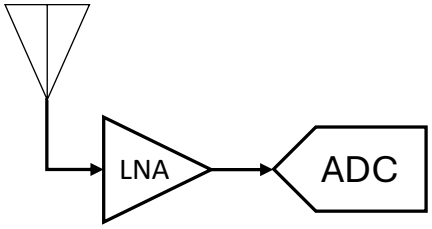

Figure 42 –Direct RF radio receiver

Ideally, a SDR would enable the digitalization of the entire RF bandwidth, allowing it to cover all radio channels supported by the device. A typical architecture of an ideal SDR architecture is depicted in Figure 42. However, the speed of ADC and DAC currently limits the front-end bandwidth.

As a result, most RF front ends incorporate a down conversion step from RF to baseband to ensure that ADCs and DACs operate at lower frequencies. Hence, RF front-end commonly used in SDR architectures are based on the Zero-IF receiver.

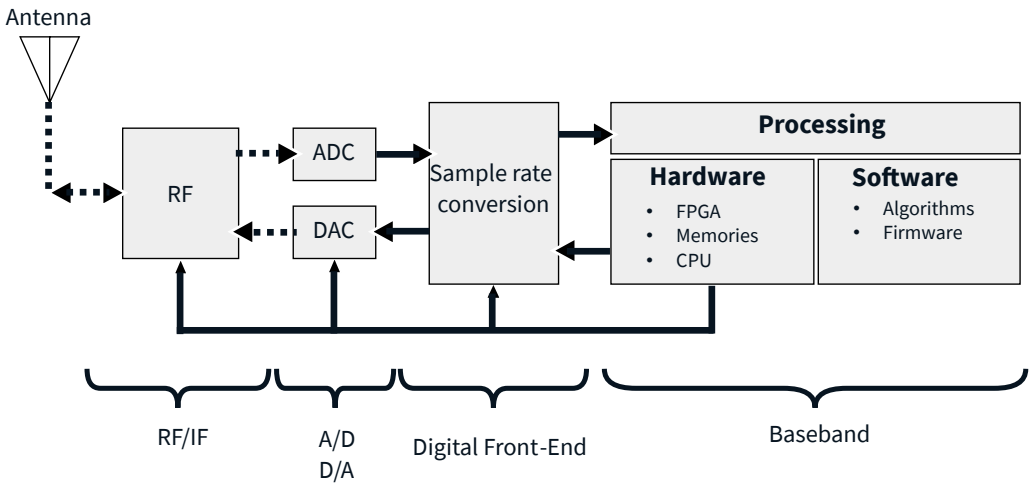

Figure 43 - SDR architecture

Figure 43 illustrates a practical SDR architecture. In this example, the ADC can be positioned after the IQ path, necessitating two ADCs—one for each component. Alternatively, the ADC can be placed immediately after the down conversion to baseband, with the IQ path implemented digitally. The latter approach is more resource-efficient, simplifies the architecture, and moves closer to the ideal SDR concept, where digitalization occurs as close to the antenna as possible.

Once the signal is digitized, it undergoes processing to reconstruct the original waveform and recover the data. This digital processing is typically performed using programmable logic, such as FPGAs. Consequently, this architecture allows communication parameters to be modified without changing the hardware configuration, offering flexibility and adaptability in various applications.

The flexibility of SDR is a major asset for long-term projects, such as the particle accelerators at CERN. As new systems, network topologies, and protocols may be introduced or modified, SDR enables these configurations to be updated without the need for hardware replacements. This is especially advantageous in environments like CERN, where thousands of devices are

in operation, leading to significant resource savings. As a result, the SDR architecture was selected for this project and will be a key focus in the following sections of the manuscript.

Currently, two architectures are available for SDR: the first combines an RFIC with an FPGA, while the second uses a Radio Frequency System-on-Chip (RFSoC). The RFSoC solution is based on a direct-RF architecture, where high-speed ADCs and DACs are positioned immediately after the antenna. Examples of such systems include recent products from AMD [63]. However, RFSoCs are more complex, as they integrate a processor for additional flexibility and functionalities, which can increase the potential for failures in radiation environments. Additionally, RFSoCs are quite expensive, with costs around 10,000 euros per device, making them unsuitable for this work.

In contrast, the combination of RFIC and FPGA is a more cost-effective solution, priced at less than 1,000 euros. This technology is already mature and offers strong performance. Furthermore, radiation test data is already available for some RFICs, confirming their robustness, making them a viable and practical choice for the development of this thesis.

3.2.3 Review on radiation effects on RF circuits

To fully understand the impact of radiation effects on an RF front-end, it is essential to first examine its structure and components. Independently of the structure, the components building the RFIC are the same. The main components usually are filters, mixers, amplifiers, oscillators, ADC, DAC, which can all be affected by radiation effects. Even though impairments are addressed during the design stage [64], radiation effects can still impact the device characteristics.

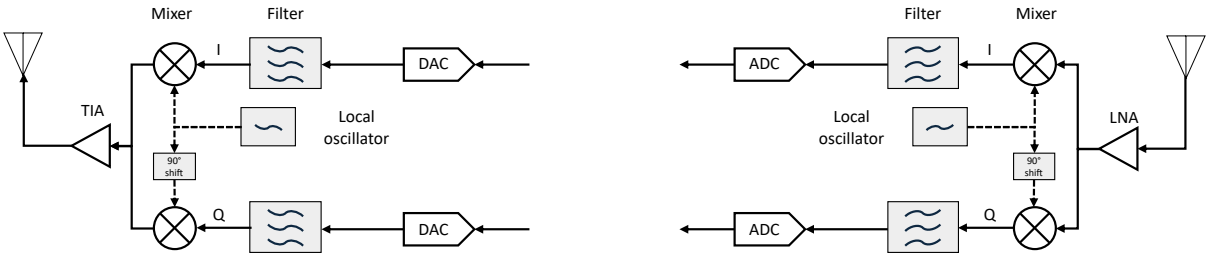

Figure 44 - Transmitter and Receiver structure for a RF front-end based on the Zero-IF architecture

Figure 44 shows the simplified architecture of a transmitter and receiver based on the Zero-IF interface, which was adopted in this work as previously discussed.

One of the key components in this architecture is the local oscillator, which provides the frequency required for down conversion. A critical element of the local oscillator circuit is the Voltage Controlled Oscillator (VCO), as it plays a fundamental role in ensuring a stable LO for up/down conversion, depending on whether it's a receiver or transmitter. Any noise in the VCO can degrade the transmitted or received signal, making it a sensitive component in radiation-prone environments. In [65], the increase in noise on the VCO due to TID effects is documented. Additionally, VCOs have been shown to be vulnerable to SET, as reported in [66] and [67]. Similarly, certain Phase-Locked Loop (PLL) circuits have experienced degradation, as highlighted in [66] and [67], with [68] discussing the effects on a SEE-hardened PLL.

Following the local oscillator, the Low Noise Amplifier (LNA) plays a critical role in boosting the signal gain while minimizing noise on the received signal. This is especially important because the signal strength at the antenna is very low and must be amplified for subsequent stages of signal processing. However, the LNA's gain and noise characteristics can degrade due to TID exposure, as shown in [69].

In addition, ADCs and DACs are integral components of the RFIC, responsible for converting signals between the analog and digital domains. However, both can be susceptible to SEE, as discussed in [70]. In such cases, corruption of an ADC sample can result in incorrect signal modulation of the carrier waveform, leading to a distorted symbol in the received waveform and, ultimately, corrupted data.

All these impairments result in changes to the amplitude and phase of the transmitted signal, thereby affecting the performance of the communication link in various ways. Consequently, understanding the impact of radiation on radio frequency transceivers is crucial for developing a radiation-tolerant physical layer, which will be the fundamental component of the SDR platform.

In the past, multiple RF front ends have been tested under radiation for various purposes, ranging from radiation monitoring to more complex space applications. Among RFICs, two categories can be distinguished: integrated transceivers and generic transceivers. The first category includes RFICs that integrate digital processing functionalities directly into the IC, without allowing custom processing in programmable logic. These transceivers are specifically designed to comply with a specific protocol, such as WiFi. As a result, they are not suitable for the development of an SDR, since communication algorithms cannot be modified through software. Nonetheless, past radiation tests on these transceivers provide valuable insights into potential failure mechanisms of a wireless link.

RF transceivers tested include those related to specific protocols, which were assessed for SEE and TID, showing performance degradation and link failures. For example, radiation tests on RF transceivers dedicated to LoRa, WirelessHART, and ZigBee revealed significant link reliability degradation when exposed to radiation [71]. ZigBee transceivers failed at 600 Gy, while WirelessHART and ISA100 transceivers failed at 167 Gy and 200 Gy, respectively. However, these tests did not include SEE analysis, which is also crucial. In fact, other studies demonstrated that different LoRa transceivers experienced communication losses, degraded performance, and functionality loss at varying dose levels [72]. At CERN, a wireless IoT platform is used to monitor radiation levels around the accelerator sector, and previous work showed that the system suffered from SEE-induced failures in the RF transceivers [73].

These tests highlight two main factors. First, when using integrated RFICs designed for a specific protocol, identifying the source of failure is challenging. Failures often manifest as a loss of link, requiring a recovery mechanism, or as missed packets. However, pinpointing the exact cause—such as whether the failure is due to a SEE affecting the device configuration or RF modulation—becomes difficult. This is because such integrated RFICs offer limited observability, making it hard to detect the source of failure. As a result, radiation tests for these devices typically focus on system-level performance, viewing the system from a higher perspective.

The second category includes transceivers that provide an interface to programmable logic, making them suitable for an SDR platform, which is the platform used in this work. For this

category, relatively few devices have been tested under radiation, and therefore, radiation test results are limited.

Majority of the tests were conducted on a specific RF transceiver from Analog Devices, the AD9361, which was integrated into the development of an SDR platform for space radiation environments. In previous work [74], the transmitter and receiver performance of the AD9361 was evaluated at different radiation dose levels by transmitting and receiving a tone signal at an RF frequency of 2.4 GHz. For the receiver, the signal amplitude was measured at various gain settings, while for the transmitter, intermodulation distortion and output power were evaluated at different dose levels. These metrics were then assessed at three different frequencies: 800 MHz, 2.4 GHz, and 5.5 GHz [75]. Additionally, the full platform, the Generic SDR (GSDR), was tested at the system level at CHARM [76], highlighting the effects of radiation at the system level, where COTS components experienced SEE-induced resets or reconfigurations of the platform.

The AD9361 provides a wide range of functionalities accessible through its registers, which are SRAM-based and, therefore, susceptible to SEUs, potentially leading to configuration changes or SEFIs that interrupt the device's operation. This type of event was also tested previously in [77] and [78] under exposure to protons and heavy ions. Furthermore, the authors used a sine wave to detect SETs in the integrated ADC, whose occurrence were reported in [78].

Another device from the same family is the AD9364. In a prior experiment, this device was tested by transmitting an image while being irradiated with heavy ions, while another device received the image. The test showed that radiation effects could induce corruption in the transmitted signal, causing some frames of the received images to be distorted [79].

However, these tests are not comprehensive for RFICs. Depending on their complexity, RFICs may offer various features through analog and digital circuits, such as LNAs and analog/digital filters, which are critical for the quality of transmitted and received signals. In fact, these radiation tests do not fully assess the combined impact of radiation effects, and the consequence for the DSP algorithms. Therefore, a different test methodology was put in place in the context of this doctoral work, which aims at studying radiation effects using standard metrics for the evaluation of RF circuits.

3.2.4 Using standard metrics to evaluate radiation effects on RFIC

When designing an RFIC, one key metric used during the design stage to evaluate the device's characteristics is the Error Vector Magnitude (EVM). EVM is a system-level metric that provides an overview of various system impairments, such as local oscillator phase noise and amplifier compression, and their impact on overall system performance. This metric is typically used during the design phase to identify and correct issues [80].

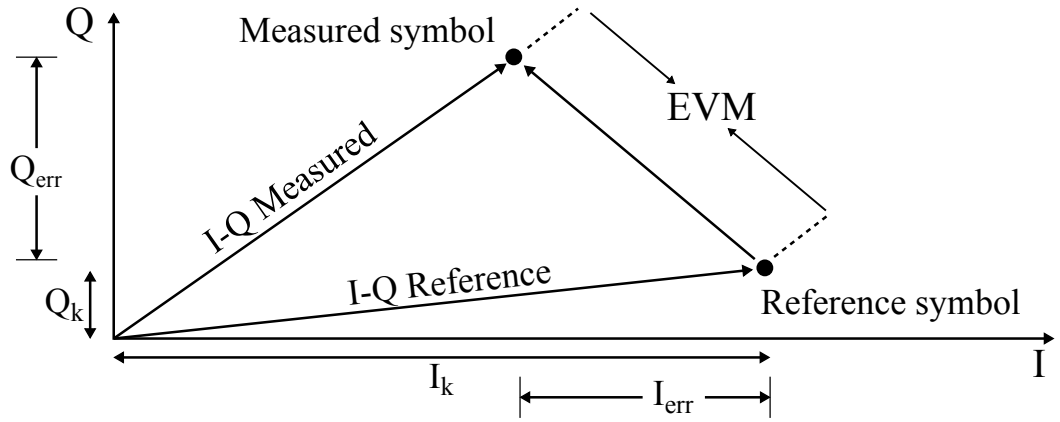

Figure 45 – Visual representation of EVM using constellation diagram. The reference symbol represents the ideal location when there are no impairments

EVM is defined as the deviation of the received symbol from its ideal location, as shown in Figure 45. It is measured as

$$EVM_{rms}(\%) = \sqrt{\frac{\frac{1}{N}\sum_{k=1}^{N} I_{err}^2 + Q_{err}^2}{\frac{1}{N}\sum_{k=1}^{N} I_k^2 + Q_k^2}} \times 100$$

In this case, 'N' represents the number of symbols considered for measurement. Clearly, the acceptable EVM range varies depending on the modulation scheme. For example, BPSK modulation is highly robust to errors since the decision algorithm decodes a binary digit based on the quadrant where the point resides. In contrast, QPSK has tighter boundaries, making it less error-tolerant and leaving less margin for the EVM parameter. The EVM requirements also depend on the application in which the RF front end is used. For instance, in standardized communication protocols like 5G and LTE, the minimum EVM requirements for compliance are specified by the European Telecommunications Standards Institute (ETSI) [81]. Notably, the higher the modulation order, the stricter the EVM requirements, as the symbol boundaries become closer.

The EVM measures modulation accuracy, offering insight into the system's actual performance when used in a wireless link. This is particularly important for systems based on digital modulation, where EVM plays a critical role. Regarding radiation effects, EVM has already been used as a metric to evaluate performance under radiation, validating its use in such contexts. For instance, in [82], the authors used EVM to assess the impact of SETs on a silicon-germanium bipolar CMOS receiver. In [83] and [84], radiation hardening techniques for LNA and VCO circuits were proposed based on EVM evaluations of the system. However, these studies focused on the receiver side and did not consider the transmitter side. Nonetheless, SET on the transmitter will impact the data at the receiver, hence affecting the DSP algorithms and therefore the demodulation process. Therefore, its evaluation is fundamental on the transmitter as well.

The modulation accuracy, measured through EVM, can provide an overview of the system's performance and help determine if the overall modulation accuracy of the device has degraded after radiation exposure, such as from TID. As previously discussed, TID can affect the performance of components like local oscillators and amplifiers, which are integral to the RFIC architecture, and in turn, impact the device's modulation accuracy and the quality of the

communication link. The quality of the transmitted and received signal is crucial for correct demodulation, and therefore for accurate data decoding.

In addition to EVM, other metrics can be analyzed to evaluate the impact of radiation effects. As mentioned earlier, the transmitter and receiver are implemented on different nodes, resulting in clock signals that are not perfectly synchronized in phase and frequency. This requires various recovery stages at the receiver before data demodulation can occur [58]. These stages adjust the timing of the incoming signal for accurate sampling or correct offset differences between the nodes. Since these recovery methods rely on PLL-based structures, they can be disrupted by SEE, or their performance can degrade due to TID, causing the recovery algorithm to fail.

These effects primarily concern the RF section of the device. However, modern RF front ends have achieved a high level of configurability, allowing adjustments to RF frequency, gain, bandwidth, and more. This high level of configurability is what enables the implementation of SDR, as these settings can be controlled via software. Nonetheless, such configurations must be stored in the device itself, often using SRAM-based register maps. These areas of the device are susceptible to SEE, which can impact the device configuration and influence the communication link. Depending on which configuration parameters are affected by an SEU, the impact on the link will vary. For instance, registers controlling critical settings like frequency may lead to a loss of link, as the devices will no longer be tuned to the same carrier frequency. On the other hand, registers that store less critical settings, such as status information, may not affect the link, and their impact could be masked.

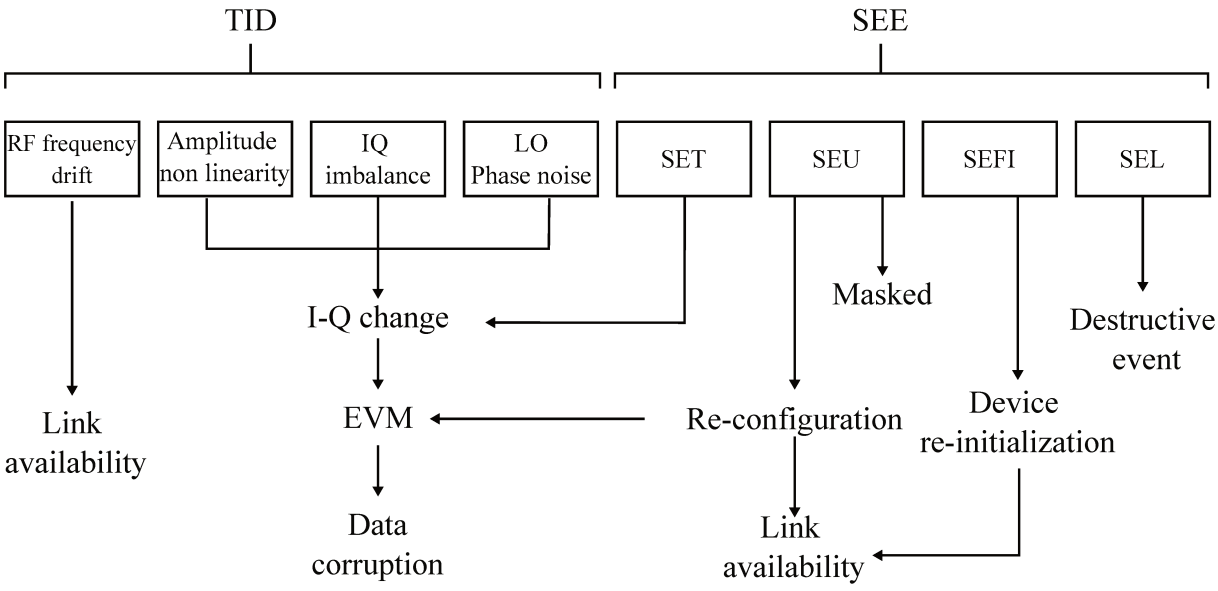

Figure 46 - Source of errors in a RFIC

Hence all these effects must be taken into consideration when evaluating its performance before the usage in a wireless system. Figure 46 summarizes the effects discussed, and their impact on the system's performance.

In conclusion, understanding the impact of radiation effects, particularly SEE and TID, on modulation accuracy and system performance is critical for evaluating RF front-end devices in harsh environments. Metrics like EVM, along with other system-level considerations such as clock recovery and configuration resilience, offer valuable insights into how these effects

degrade communication links. The configurability of modern RFICs, while enabling the flexibility of SDR platforms, also introduces vulnerability to radiation-induced errors in critical device settings.

In the next section, these metrics will be used to study radiation effects on a state-of-the-art RF transceiver, the AD9361.

3.3 Exploring radiation effects on a RFIC

The selected candidate for applying the radiation testing methodology is the AD9361 RF transceiver. This choice is motivated by its proven robustness in past radiation tests, particularly its resistance to SEL, making it a strong candidate for developing a radiation-tolerant wireless link, which will be detailed in the following chapters. Additionally, the availability of previous test results offers a valuable benchmark for comparison. For these reasons, the AD9361 has been tested using the methodology outlined earlier.

3.3.1 Device under test

The AD9361 RF Transceiver is based on a direct conversion architecture. It combines a RF section, with a mixed-signal baseband section.

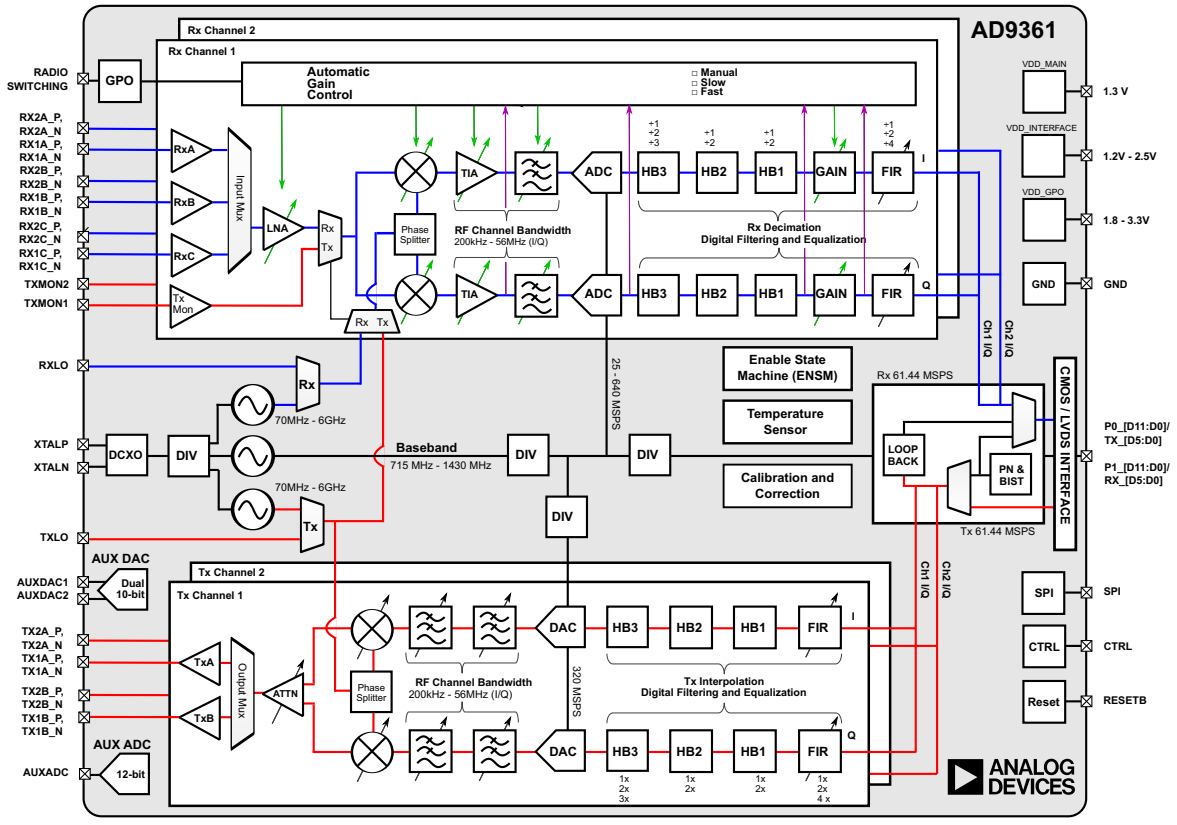

Figure 47 - AD9361 architecture [85]

The architecture of the AD9361 is shown in Figure 47. It features two independent transmitter and receiver chains, each equipped with its own local oscillator. The local oscillators for the transmitter and receiver chains can be tuned to different frequencies, covering a range between 70 MHz and 6 GHz, making the device versatile for various applications.

The receiver includes all the necessary components to receive an RF signal, downconvert it to baseband, digitize it, and transfer it to a baseband processor through a standard CMOS/LVDS interface for further processing. The RX chain has two independent channels, allowing reception from two different antennas, making it suitable for multiple input and output (MIMO) systems while sharing a common frequency synthesizer. Each RX channel can multiplex between three inputs, ideal for diversity systems with multiple antenna inputs. The RX chain is a direct conversion, or zero-IF receiver, which downconverts RF signals to baseband in a single step. It includes a low-noise amplifier, followed by 90-degree IQ mixers that separate the in-phase and quadrature components, downconverting the RF signal into I and Q signals for baseband processing. The signal paths for I and Q are identical and consist of a transimpedance amplifier and a programmable analog low-pass filter, configurable with bandwidths between 200 kHz and 56 MHz, which helps reduce adjacent channel interference. A 12-bit sigma-delta ADC follows, with a sample rate adjustable via a local oscillator, shared with the transmitter sampling rate. After digitization, the signal passes through four stages of digital decimation filters (which can be bypassed) to reduce the signal rate from up to 640 MSPS to 64.11 MSPS for the digital interface. A 128-tap FIR filter is also available to compensate for magnitude and phase distortions introduced by the analog and digital filters in the passband. The gain throughout the chain is configured according to system requirements, with gain distributed across blocks based on a preprogrammed map for optimal performance at each stage. Gain adjustment can be handled by an internal Automatic Gain Control (AGC), which operates in fast or slow mode depending on the signal environment, or manually configured by the user. Each channel also includes an independent Receive Signal Strength Indicator (RSSI) measurement circuit. For IQ imbalance correction, the AD9361 includes an RX quadrature calibration mechanism that continuously tracks the signal and compensates for phase and gain errors in real time.

Similarly, the transmitter section contains two identical channels based on a direct conversion architecture that handles all digital, mixed, and RF processing. The TX chain receives digital data from the baseband processor in IQ format (12-bit, 2’s complement data) via a CMOS/LVDS interface. This data passes through a 128-tap user-configurable FIR filter before undergoing interpolation to increase the data rate for the following 12-bit DAC, which operates at rates as low as 320 MSPS. The DAC output passes through a configurable low-pass filter (200 kHz to 56 MHz), after which the I and Q signals are combined and modulated onto the carrier waveform at the RF frequency. The combined waveform is further shaped by an analog filter before passing through the output amplifier, which provides a range of attenuation options to optimize the device for various RF systems. Both I and Q signal paths in the transmitter are identical, and the section includes self-calibration circuitry for automatic real-time adjustments. Additionally, the transmitter features a TX monitor for each channel, which routes the transmitted signal back to an unused receiver channel for monitoring.

All device settings are stored in an SRAM-based register array, accessible by the baseband processor through a Serial Peripheral Interface (SPI). The device's functionalities are configured using Application Programming Interfaces (APIs) provided by the manufacturer.

3.3.2 Application under test

The AD9361 must be integrated with a baseband processor to handle data generation before transmission and data processing after reception for decoding information. Therefore, to conduct radiation testing of the AD9361 and to analyse the metrics previously, a full SDR platform had to be set up, including transmission/reception DSP algorithms.

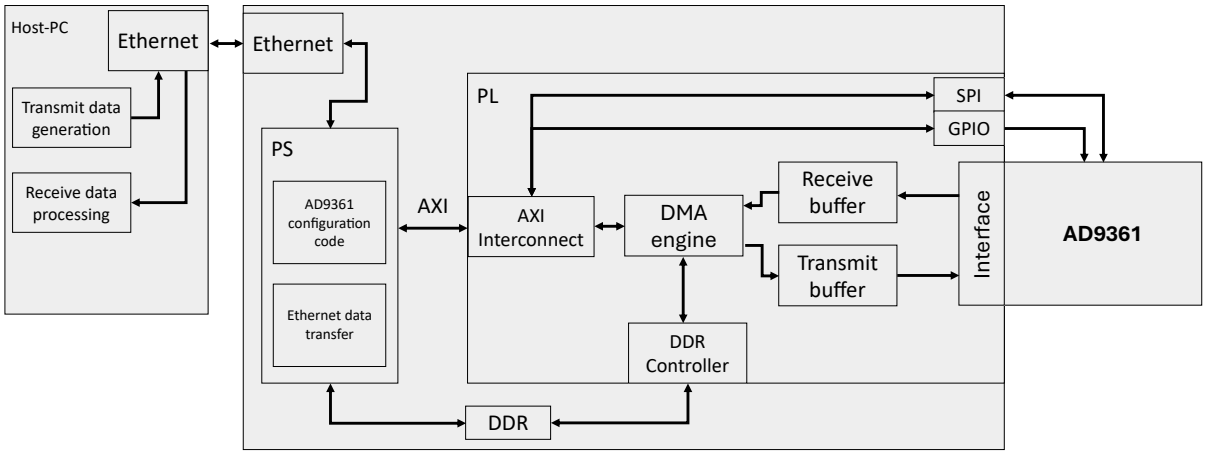

Figure 48 - Architecture of the SDR implementation

Figure 48 illustrates the architecture of the SDR implementation. The platform is built around the Zynq-7000 SoC from AMD-Xilinx, which integrates both programmable logic (PL) and an ARM Cortex-A9 CPU in the processing system (PS). The AD9361 RF transceiver is connected to the PL via an LVDS interface for high-speed data exchange. The PL also implements SPI communication, which allows for accessing and modifying the AD9361's configuration.

The high-level configuration of the AD9361 is managed by the ARM Cortex-A9 processor through APIs running on the PS. Data to be transmitted is generated on the host PC and transferred to the SoC over Ethernet. The PS manages the Ethernet data transfer and AD9361 configuration code, while the PL handles real-time data movement between the AD9361 and memory via a DMA engine. The DMA engine, connected to a DDR controller, manages both receive and transmit data buffers for the AD9361. Data received from the AD9361 is stored into the memory, and streamed to the host PC through Ethernet where it is post-processed with the DSP receiver algorithms.

In this approach, the DSP algorithms required for data transmission and reception—such as modulation, demodulation, and data recovery—are not implemented on the FPGA but rather in software on the host PC. This method was chosen over a full FPGA implementation because the focus is not on developing a complete wireless link but on understanding the impact of radiation effects on the RFIC and its influence on the communication link. Therefore, the FPGA is only used to stream raw samples to and from the AD9361, while all data processing is handled on the host PC. In this way, additional processing could be performed to better analyse the experimental data.

Regarding the wireless algorithms under test, two wireless applications were implemented using different modulation schemes. One application uses QPSK modulation, commonly employed in Time Division Multiple Access (TDMA) systems, such as satellite communications. The second application is based on the Long-Term Evolution (LTE) protocol, a standard for wireless communication in mobile devices, utilizing QAM modulation.

3.3.2.1 QPSK application under test

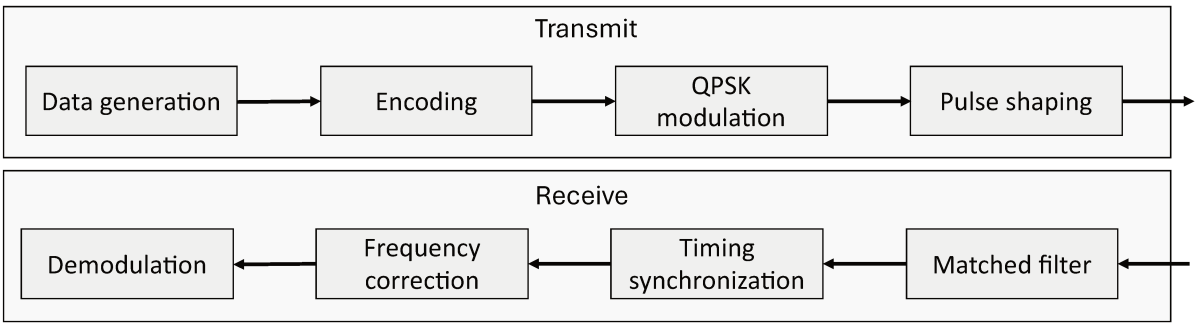

Figure 49 - QPSK application

The QPSK application is representative of standard systems using this type of modulation. Its block diagram is shown in Figure 49. The process begins with data generation, where any data the user wants to transmit can be used. In this case, a fixed payload of randomly generated data is transmitted. A 32-bit Cyclic-Redundancy-Check (CRC) signature is also included. The total packet size is 6400 bits. The data is encoded using a convolutional coding scheme [86] with a code rate of 2/3. Afterward, the data is modulated using QPSK and passed through a pulse shaping filter. The packet is then transferred to the FPGA via Ethernet and stored in a buffer. This buffer continuously feeds the AD9361 during transmission. Therefore, once transmission begins, the same buffer is transmitted cyclically to the AD9361.

On the receiver side, the process implements all the previously discussed steps to correct timing synchronization and frequency offset. The first stage is matched filtering, where the data passes through an SRRC filter that removes inter-symbol interference and improves the signal-to-noise ratio for optimal sampling. Then, symbol recovery is performed using a PLL-based algorithm, specifically the Gardner implementation [61]. Following symbol recovery, carrier recovery is executed, which is divided into two stages: coarse and fine frequency correction. The coarse correction uses an FFT-based method to eliminate large frequency offsets between the two nodes, while fine correction is performed using a PLL-based loop that stabilizes the constellation before demodulation.

Afterward, the data is demodulated. The demodulation process involves detecting the start of a packet using a predefined sequence, the preamble, as described earlier, and then decoding the rest of the packet, which contains the fixed payload.

3.3.2.2 QAM application

A second application was implemented based on QAM modulation, specifically following the LTE protocol, which is a wireless communication standard for mobile devices. This application employs Orthogonal Frequency-Division Multiplexing (OFDM) [87], where multiple orthogonal carriers are transmitted to send data in parallel. These carriers are separated by guard frequencies to prevent interference. By transmitting multiple carriers, the system can increase data throughput by sending more data in parallel. The LTE standard is defined in the TS36.104 document [81], which outlines the waveform characteristics required for compatibility with LTE networks. It also specifies EVM requirements to ensure a device meets the standard's performance criteria. In this implementation, the waveform is based on 64-QAM modulation, with a subcarrier spacing of 15 kHz.

3.3.3 Radiation test setup

The SDR platform described earlier was implemented using commercially available development kits for both the AD9361, the Device Under Test (DUT), and the Zynq-7000 SoC. The DUT was mounted on the FMCOMMS-3, an evaluation kit designed for the AD9361 on an FMC mezzanine card, which includes a power scheme and wide-range RF balun to cover the full frequency range of the AD9361. The FMC form factor allows the board to connect seamlessly with the Zynq-7000 development kit, specifically the ZedBoard.

The radiation experiments were conducted at the Proton Irradiation Facility (PIF) at PSI in Villigen, Switzerland, using a 200 MeV proton beam. One advantage of this facility is the beam size, which allowed precise targeting of only the DUT, while the rest of the SDR platform was kept outside the irradiation zone. This setup prevented any potential data corruption in the system and focused the radiation effects solely on the DUT. The DUT was connected to the ZedBoard with an FMC-to-FMC cable. The DUT output was connected to a second SDR, implemented in the same way and located outside the beam, which acted as the receiver. The two SDR units were connected via SMA-to-SMA cables to isolate and evaluate radiation-induced errors in the communication link.

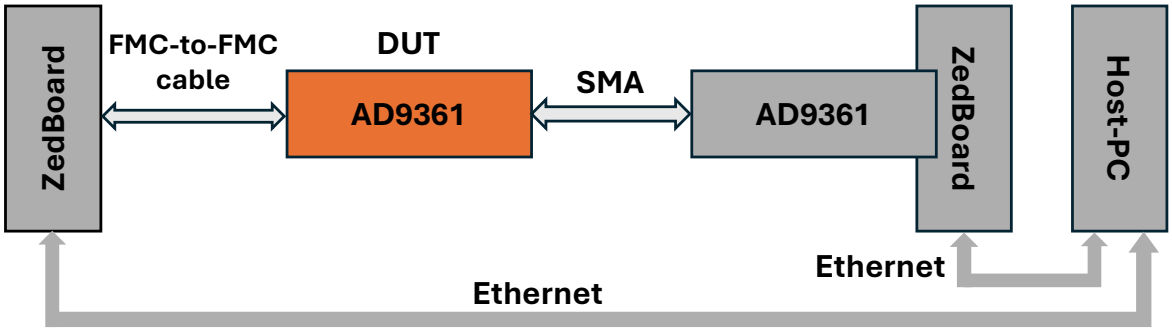

Figure 50 - Test setup for the AD9361 irradiation test

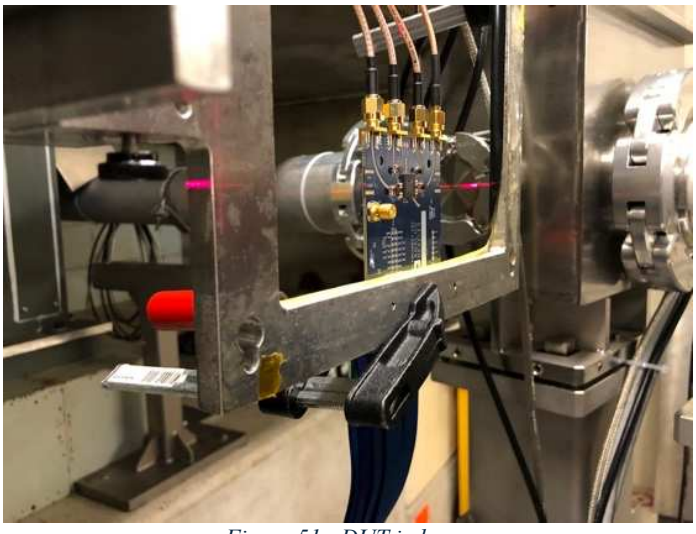

Figure 51 - DUT in beam

Figure 50 shows the test setup used during irradiation, while Figure 51 displays the DUT positioned within the beam. Data for the test was generated on the host PC, based on the application under test (QPSK or QAM), and then transferred to the transmitter SDR. During irradiation, the receiving SDR was set up to acquire data continuously, store it on the computer, and then transfer it again without performing real-time data recovery or packet demodulation. This approach was chosen to maintain a higher acquisition rate, as performing data processing

between receptions would have significantly slowed down the process. All data processing and analysis were conducted post-irradiation.

For the QPSK application, as discussed earlier, the main parameters for the PLL-based algorithms are the damping factor and the loop bandwidth, which control the loop filter's responsiveness and convergence time. In this case, a damping factor of 0.7 and a loop bandwidth of 0.01 were used. With this configuration, the communication performance achieved an EVM of less than 5% prior to irradiation. The QPSK test involved continuously transmitting the same data until the irradiation ended or until a device misconfiguration caused the communication to stop.

For the LTE application, waveforms were generated in compliance with the LTE standard specified in the TS36.104 document using MATLAB's LTE Toolbox. This toolbox was employed for both data generation and analysis. The generated waveform adhered to LTE standards, and the toolbox analyzed the incoming waveform, computed the EVM, and estimated the frequency offset. In this case, the frequency offset was calculated using the cyclic-prefix technique, as defined in [88], instead of the FFT-based method. The sampling error, however, was computed using an FFT-based approach.

For both the QPSK and QAM applications, a background script was continuously running to read all configuration registers during the irradiation experiment. This allowed for correlating SEUs in the device's configuration registers with any communication link issues, helping to better understand the impact of SEUs on the link. Additionally, the device was periodically reconfigured using the high-level APIs on the PS. Specifically, for the QPSK application, the device was reconfigured every 200 packets, while for the QAM-based LTE application, the device was reconfigured before transmitting each LTE waveform.

3.3.4 Radiation test results

This section discusses the results obtained from the 200 MeV proton irradiation of the device. The findings are divided between the QPSK and QAM applications, showing how SEE and TID impacts the RF transceiver and how this, in turn, affects each downstream algorithm differently.

3.3.4.1 QPSK-based application

The QPSK application was tested at two different sampling rates, 1 MSPS and 5 MSPS, while the device accumulated a total dose of 500 Gy. For each decoded packet, the root-mean-square EVM was calculated.

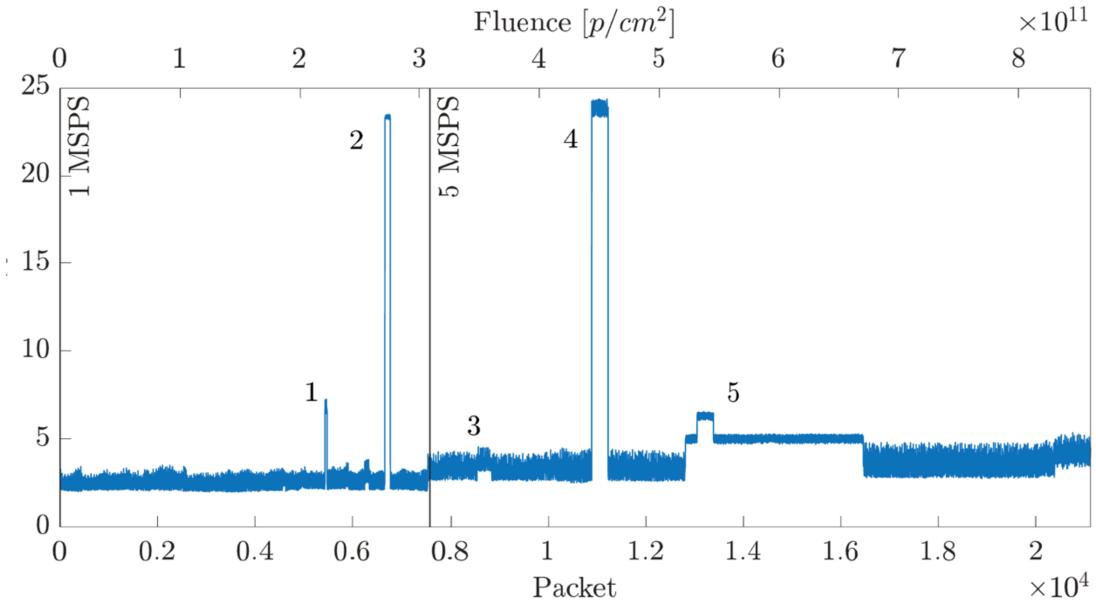

Figure 52 - EVM of QPSK system during irradiation

Figure 52 shows the EVM for all packets transmitted under irradiation as a function of the accumulated fluence for both sampling rates. The EVM for each packet was calculated across all the symbols in the packet (6400 symbols, corresponding to 12,800 data bits, since QPSK maps two bits per symbol) and is expressed as a percentage. Initially, the EVM stayed below 5%, consistent with the pre-irradiation performance. However, during the rest of the irradiation, shifts and abrupt changes in EVM were observed, which can be attributed to various factors. One noticeable shift occurred when the sampling rate was changed, indicating that this shift was related to the rate change rather than TID degradation. Toward the end of the irradiation, after the device had accumulated 500 Gy, a slight EVM shift was observed without any changes to the device parameters. This might suggest a degree of degradation due to the accumulated dose.

As for the abrupt changes in EVM, these are marked with numbers in the figure and are likely related to SEUs affecting the device's register space, which altered its configuration. Each time a reconfiguration was triggered, the EVM returned to its nominal performance.

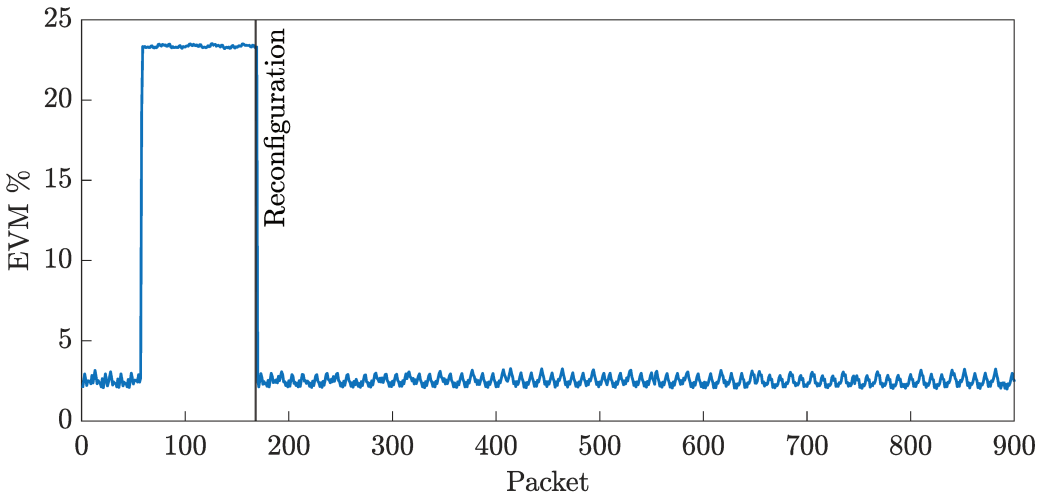

Figure 53 - EVM before and after reconfiguration

Figure 53 highlights one of these events in detail. When an SEU altered the device configuration, the EVM increased; however, after 200 packets, the device was reconfigured, and the EVM returned to normal. Despite the temporary increase in EVM, the worst-case scenario did not prevent packet decoding, and the packet was retrieved correctly without errors. Upon analysing the post-irradiation register configuration, it was found that the registers controlling the device gain were no longer configured as they were before irradiation, suggesting that their change affected performance. This was the case for events 1 through 4 in the figure. However, during the last event, the issue persisted despite reconfigurations after 200 packets. Further analysis of the registers revealed that many had changed their values, making it difficult to pinpoint the exact cause. Nevertheless, this event can still be attributed to SEUs, as the nominal operation was eventually restored after an additional reconfiguration, likely due to further operations performed by the low-level drivers.

Another observation can be made by analyzing the EVM within each packet, rather than considering the overall EVM for the entire packet. As previously discussed, QPSK demodulation occurs after symbol recovery and frequency recovery, both of which are PLL-based algorithms. These algorithms react differently depending on the input signal, and it was observed that abrupt changes could disrupt their performance.

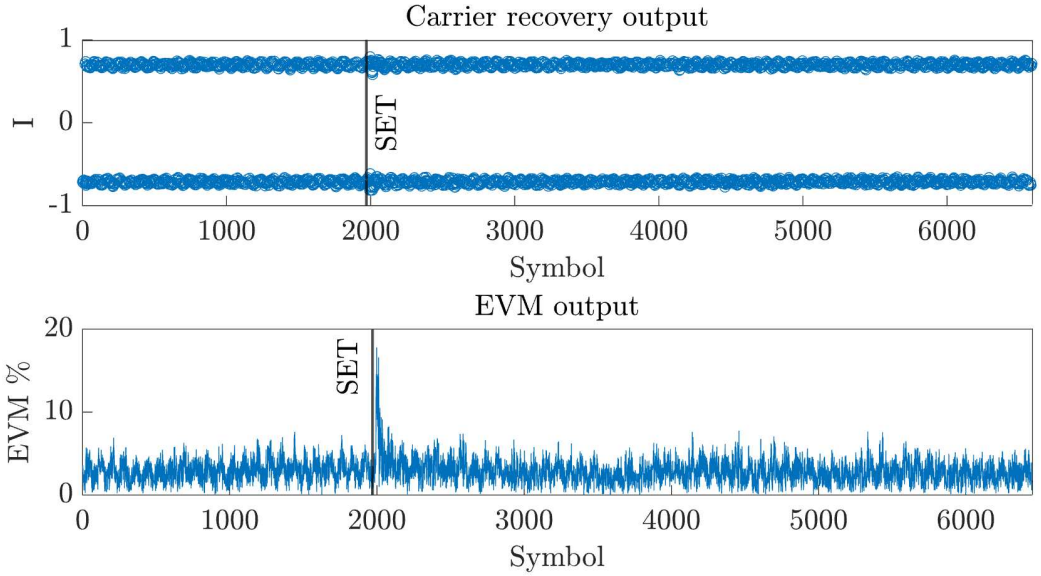

Figure 54 – EVM output

Figure 54 (bottom) shows the EVM measured for each individual symbol, compared to the expected value. In certain instances, the EVM spiked to nearly 20%. This spike is the result of the disruption of the carrier recovery algorithm, as shown in Figure 54 (top), where the I and Q components of the signal are plotted. When the algorithm operates normally, the I and Q components are stable around 0.7, which corresponds to the nominal QPSK constellation values. However, when an abrupt change occurs, the algorithm becomes disrupted, and the loop needs time to recover before stabilizing the constellation again. Once the algorithm reconverges, the EVM drops back below 5%. Despite these disruptions, no errors were introduced into the packet, which was still retrieved correctly.

Table 6 – Cross-sections of the SEE observed on the QPSK application

Type	**Fluence** $[p/cm^2]$	**Events**	$\sigma\ (cm^2/device)$		
			Lower limit	σ	Upper limit
SET	8.6×10^{11}	1	2.34×10^{-14}	1.16×10^{-12}	6.47×10^{-12}

SEU		5	1.85×10^{-12}	5.81×10^{-12}	1.36×10^{-11}

The occurrence of these events is summarized in Table 6, which details the type of events, the fluence, and the cross-section with the 95% confidence interval [89].

3.3.4.2 QAM applications

A different DUT was used to test the QAM application, focusing on evaluating EVM and Carrier Frequency Offset. The test was paused at various dose levels to assess these two parameters. Between each dose step, the device was reconfigured to run the QPSK application. After reaching 400 Gy, the device was reconfigured again for the QAM-based application, and its response was assessed for SEE.

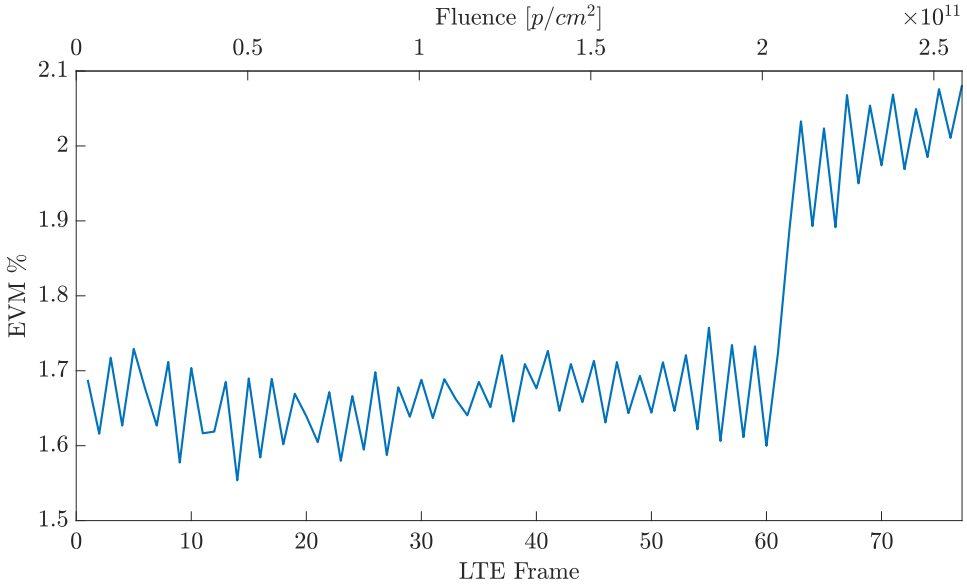

Figure 55 - EVM evolution for LTE frames

Regarding SEE effects, Figure 55 shows the EVM for each LTE frame received as a function of the fluence. The device did not exhibit significant degradation, and no major SEE events were observed. The device was exposed to a lower fluence compared to the QPSK application due to an issue at the testing facility, which could explain the absence of noticeable events during the QAM tests.

As for TID effects, the device showed no significant degradation, with EVM measured between 1.6% and 2.1% from 0 Gy to 520 Gy. However, a slight increase in EVM was observed towards the end of the test, around 500 Gy, similar to the QPSK application. In this case, a reconfiguration was triggered before each packet was sent, but since the EVM did not return to nominal values, the increase is likely due to TID-induced degradation.

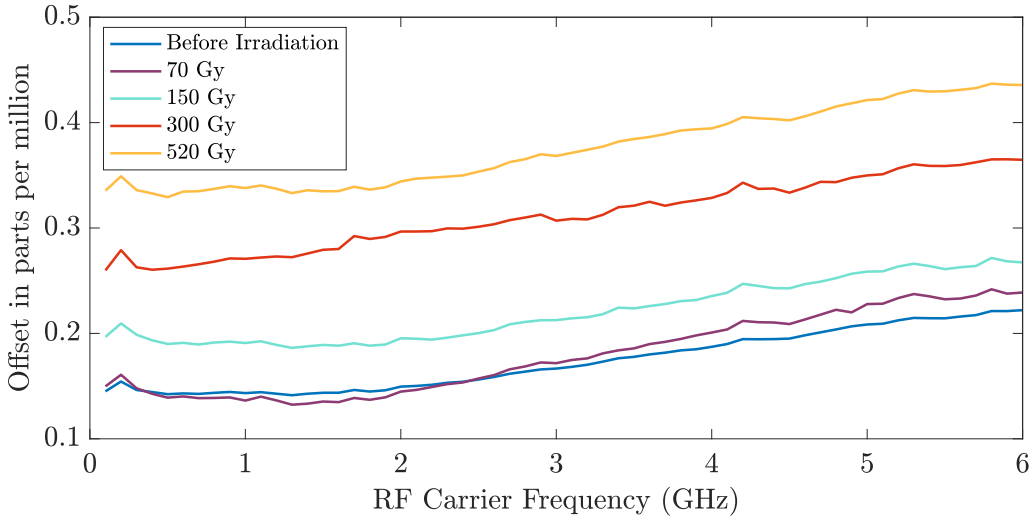

Figure 56 - CFO over frequency

Regarding the frequency offset, it was measured at various dose steps and is presented in Figure 56, expressed in parts per million (ppm), which is the standard unit for representing oscillator uncertainty as a function of carrier frequency. The results show that the frequency offset increases with accumulated dose, indicating degradation in the local oscillator's performance. Additionally, the variation is more pronounced at higher RF frequencies, further demonstrating the impact of dose on the oscillator's stability.

3.3.4.3 SEE Analysis

The AD9361 configuration is managed through a set of registers accessible via SPI. Understanding the content of these registers and the impact of SEUs on them allowed to assess how SEUs could affect the communication link and to explore possible mitigation techniques to restore proper functionality. To this end, a background script was run during the entire test to continuously monitor the register values via SPI. A golden copy of the register values was taken before the start of irradiation, and throughout the test, periodic reads of all registers were performed and compared to this reference.

The AD9361 contains 1024 registers, which can be classified into configuration and status registers. Status registers refer to parameters that change continuously and are read-only, as they are intended to be used for monitoring purposes. Since they do not influence the functioning of the device or the communication link, any SEU affecting these registers would not lead to a communication failure.

On the other hand, configuration registers directly manage the device's functionalities and can impact the communication link if affected by SEUs. These registers were further categorized based on their impact. The first category includes registers that, if altered by an SEU, cause performance degradation but not total communication failure. The degree of degradation also depends on the modulation scheme in use, as seen with the EVM increase in the QPSK application. The second category includes registers that, if affected, lead to a complete loss of communication by altering critical parameters, such as those managing the LVDS interface or line swapping. Lastly, the third category includes registers whose corruption causes the device to become unrecognizable by the host PC, effectively halting communication entirely.

Table 7 –SEE cross-sections for the three categories of events observed

Type	Fluence $[p/cm^2]$	Events	σ $(cm^2/device)$		
			Lower limit	σ	Upper limit
Reconfiguration - QPSK	8.60×10^{11}	5	1.85×10^{-12}	5.81×10^{-12}	1.36×10^{-11}
Reconfiguration – QAM	2.58×10^{11}	0	-	-	1.43×10^{-11}
Re-Initialization	1.12×10^{12}	13	6.06×10^{-12}	1.16×10^{-11}	1.99×10^{-11}
Permanent failure		3	5.35×10^{-13}	2.68×10^{-12}	7.83×10^{-12}

These three categories of events, along with their occurrence during the test, are summarized in Table 7, including the 95% confidence interval.

For each category, a different recovery mechanism was identified. In the case of the first category, where configuration changes only degrade performance, simply restoring the correct configuration is sufficient to return the device to proper working conditions. For the second category, where critical parameters are modified leading to a loss of communication, the device can be reinitialized using the initialization pin on the IC, followed by reconfiguration. However, in the third category, where the device is no longer recognized by the controlling FPGA, reconfiguration is not possible. In this scenario, a power cycle is required to fully restore the device's functionality.

3.3.5 Discussion

The experimental results showed that the EVM did not exhibit significant variations with increasing radiation dose in both the QAM and QPSK-based communication links. This indicates that the device's modulation accuracy remains unaffected by radiation up to 500 Gy, where a small degradation was observed in both cases, which could be attributed to TID effects.

In the QPSK application, sharp changes in EVM were observed due to modifications in the device configuration, likely caused by SEUs affecting the RFIC's configuration registers. For instance, manual analysis revealed that the transmitter gain had increased, leading to receiver saturation since the devices were connected with cables. As a result, the EVM measured after the recovery algorithm was higher. Despite this, the data was still retrieved correctly. In a typical communication system, EVM is related to the BER [90], with higher EVM generally leading to higher BER. In this experiment, however, BER could not be evaluated due to the intermittent nature of data reception, as data was stored on the host PC upon reception, limiting real-time processing. Nonetheless, EVM degradation can be used to estimate BER in communication systems.

Regarding SET, similar events have been previously categorized in earlier work where the AD9361 was tested using a different method. In a previous experiment, a sine wave was used to test the device under proton irradiation [77]. The results showed that the received waveform was distorted due to SETs affecting the ADC of the device. However, with IQ-modulated data, as used in practical applications, it is challenging to directly observe the impact of SETs because the transmitted waveform oscillates between different amplitude and phase levels. Despite this, SETs can be detected by analyzing EVM spikes, allowing for correlation between SETs and their impact on the communication link.

This EVM-based approach is also useful in understanding how SETs affect the downstream algorithms in the receiver, such as PLL-based algorithms, which adjust the signal based on

error estimations computed at each sample. The response of these algorithms depends on the characteristics of the input signal. For example, in a packet where an SET occurs, the behaviour of the Gardner-based algorithm used in the receiver can be analyzed as a function of loop parameters, such as bandwidth and damping factor, which determine the convergence speed and response time of the algorithm.

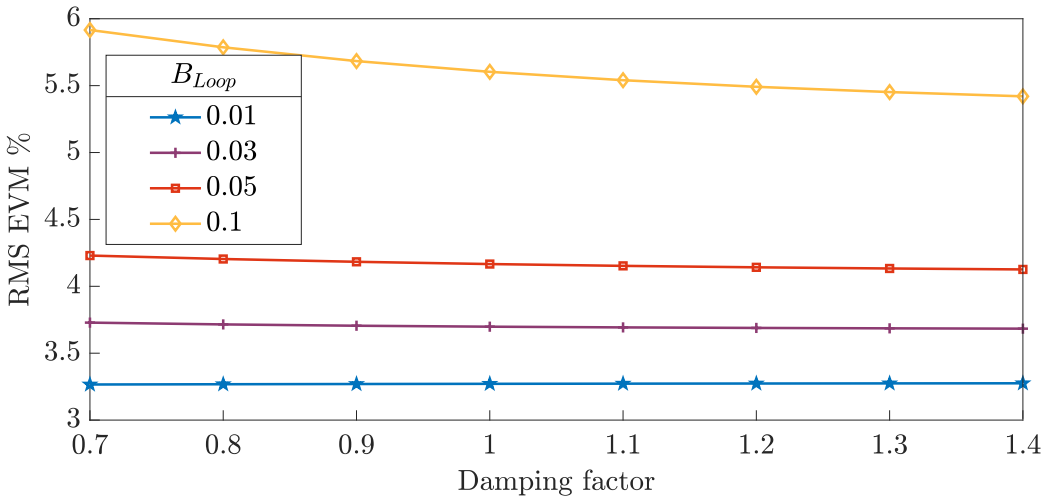

Figure 57 - EVM of the packet affected

The results of this analysis are shown in Figure 57. These results were obtained by processing the same input packet while varying the loop filter parameters. As seen in the figure, the RMS EVM increases as the loop bandwidth is increased. This occurs because, after the occurrence of the SET, the loop struggles to converge to a stable and clean solution with a higher bandwidth. In contrast, with a lower loop bandwidth, the signal correction is more effective, resulting in a less noisy output and a better EVM. This observation can be used to study the system's response to such events and guide the design of the receiver to optimize performance under different conditions.

Another important consideration involves the frequency offset correction algorithm. In the QPSK application, the algorithm employed a two-stage process: coarse frequency correction followed by fine frequency correction. Initially, the large frequency offset is detected and corrected using an FFT-based method, and any remaining offset is then removed by a PLL-based mechanism. However, this approach may not always be suitable, as the choice of recovery stages depends on the specific application.

Additionally, part of the data acquired during the QPSK tests showed that the frequency offset could become too large for the algorithm to compensate. An SEU in the device configuration caused a significant frequency shift, which exceeded the correction capabilities of the algorithm. In a system without coarse frequency correction, this could result in a loss of the communication link until the device is reconfigured.

Furthermore, considering the TID results and the variation of frequency offset with radiation dose, these factors must be taken into account when designing recovery stages. The trade-off between coarse and fine correction methods should also be carefully evaluated, particularly in terms of available system resources. The advantage of the PLL-based method is that it is an iterative approach, meaning it does not reduce system throughput and requires fewer resources. Since these algorithms will likely be implemented on an FPGA, one must weigh the trade-off

between available resources and the required lifetime of the communication link before deciding on the appropriate implementation.

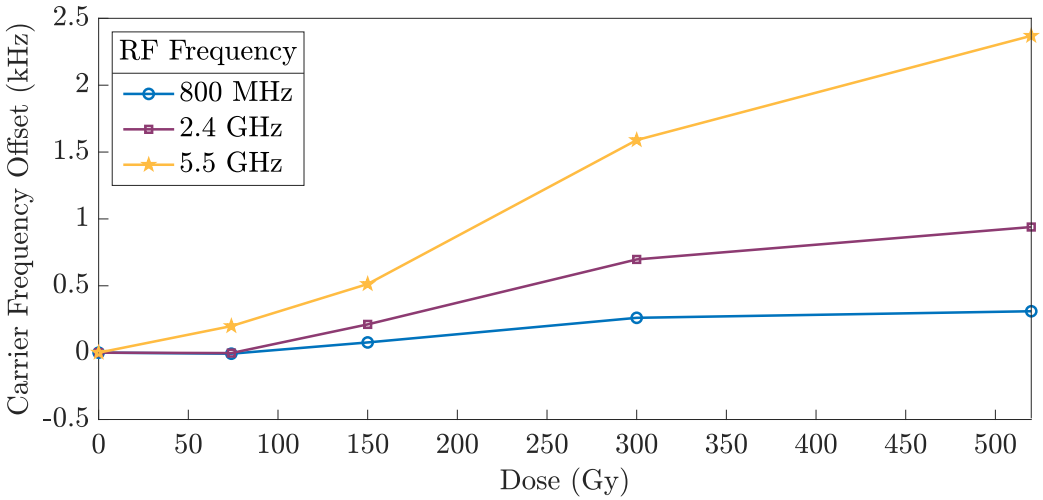

Figure 58 - CFO estimation

The same applies to the carrier frequency offset measured during the QAM-based application. This offset, shown in Figure 58 for three different carrier frequencies at various dose levels, is compensated before data demodulation. The compensation algorithm is based on the Cyclic Prefix technique, where each symbol is extended with a prefix, essentially a repetition of the last few symbols. This technique is commonly used in OFDM-based systems and helps estimate and correct the frequency offset between two devices. However, it only works if the frequency offset is less than half of the subcarrier spacing. In this test, the subcarrier spacing was 15 kHz, meaning any larger offset could not be corrected and would have caused a loss of data, as the DSP algorithms would have failed.

Additionally, as discussed in [91], CFO can impair the orthogonality of the subcarriers, leading to intercarrier interference and performance degradation. However, this was not observed with the device under test, as only 25% of the correction threshold was reached at 5.5 GHz after 520 Gy. Had the offset been higher, a possible solution would have been to employ alternative recovery algorithms, such as the one described in [92], which first uses a pilot carrier to estimate the larger offset, followed by the Cyclic Prefix algorithm to correct the smaller offset.

3.3.5.1 Remarks

Considering the experimental results, the AD9361 has proven to be an ideal candidate for wireless applications in the LHC environment due to its robustness against both TID and SEE. Although this experiment tested the device up to 500 Gy, previous radiation tests demonstrated the AD9361's resilience up to 450 kGy [93], after which signs of degradation and then loss of functionalities were observed. Given these results, the AD9361 is capable of surviving beyond the expected 12-year operational lifetime of systems in the LHC tunnel.
Furthermore, the device does not suffer from destructive events, and all SEE effects observed were recoverable, as previously discussed. Its robustness against SEE, combined with its wide range of functionalities and frequency coverage, makes it an ideal candidate for applications in the LHC tunnel and future particle accelerators. As a result, it has been selected as one of the key building blocks for the radiation-tolerant wireless links developed in the framework of this doctoral work.

3.4 Chapter summary

This chapter covers the fundamentals of wireless communication systems, starting with the theory behind wireless communication, the data processing required for sending and receiving data over the air, and the hardware components essential for implementing these systems.

It also examines the current use of wireless systems in radiation environments, focusing on their failures and limitations under radiation. After highlighting the advantages of SDRs as the most suitable solution for long-lifespan environments like CERN's particle accelerators, the chapter reviews current SDR platforms used in radiation environments and discusses how they have been tested under radiation, identifying the limitations of traditional testing methods.

To address these challenges and gain deeper insights into the impact of radiation on wireless links, a new methodology using real applications was presented, showing how traditional metrics, typically used to assess communication link performance, can be applied to investigate radiation-induced effects on radio frequency front-ends.

This methodology was applied to the AD9361, a promising candidate for the implementation of a radiation-tolerant wireless link, given the availability of existing radiation test results that already proved its robustness. Two applications, based on QPSK and QAM modulation, were implemented using the AD9361 and tested. The methodology enabled the identification of SEU, SET, and TID effects on the device and allowed for an analysis of their impact on the communication link. For the QPSK-based application, SETs in the analog circuitry were detected via EVM analysis, which disrupted downstream algorithms in the receiver. However, the EVM did not degrade with radiation doses up to 500 Gy for either application. Additionally, while the device exhibited slight frequency drift with increasing dose, it did not lead to a failure of the communication link. Concerning SEE effects, various categories could be identified depending on the effect on the communication link. Nonetheless, all the failures were never permanent.

Given the positive results, the wireless solutions explored in the remainder of this thesis will be based on this front-end for the SDR platform. In the next chapter, the AD9361 will be used to build a wireless link to extend CERN's current control network. Then, in Chapter 5, a complete radiation-tolerant wireless PHY for control applications incorporating the AD9361 will be presented.

4 A radiation-tolerant wireless extension for the WorldFIP LHC control network

The WorldFIP LHC control network, as discussed in Chapter 2, is responsible for ensuring the reliable operation of the LHC by controlling critical systems. This network consists of thousands of nodes connected through 400 km of cables using the WorldFIP protocol. However, expanding the network through additional wired connections is no longer feasible due to fully occupied cable trays, and maintaining these cables is labour-intensive and costly. While wireless solutions offer a potential solution to these challenges, their implementation is complex. One major hurdle is that the electronics, including the control interface, are exposed to the LHC's radiation environment, meaning any wireless interface must also be radiation tolerant.

There are two possible approaches for integrating wireless technologies into the LHC control network. The first, explored in this chapter, is to extend the current infrastructure by adding wireless capabilities without modifying the existing hardware or protocol. This approach is more suitable for addressing immediate cabling issues since no hardware modifications are required. The second approach, which will be covered in Chapter 5, involves developing a wireless link specifically tailored to the LHC control network's unique requirements.

Extending the WorldFIP infrastructure to include wireless capabilities presents significant challenges. WorldFIP is a fieldbus designed to facilitate deterministic data exchange, but incorporating it over a wireless channel is complicated, as wireless channels are inherently more error-prone than wired ones. This issue extends beyond WorldFIP, as fieldbus systems are widely used in industrial settings to control instrumentation, and the integration of wireless technologies into such systems has been explored in the past. The first part of this chapter reviews past attempts to integrate wireless technology with fieldbus systems, identifying challenges and solutions from those efforts.

Building on these insights, the second part of the chapter addresses how to extend CERN’s WorldFIP control infrastructure wirelessly. A major contribution of this doctoral work is the development of a wireless extension for the WorldFIP network, moving from a high-level overview to a proof-of-concept demonstration using real CERN hardware. The proposed solution leverages the SDR based on the AD9361, highlighting the benefits discussed in the previous chapter.

Finally, the third part of the chapter explores the radiation tolerance characteristics of the proposed wireless solution. It presents experimental results from proton and neutron irradiation tests, demonstrating the effects of radiation on the wireless extension and its impact on the LHC control network. The chapter concludes with an analysis of the pros and cons of the proposed solution in the context of the LHC control system.

4.1 Integration of wireless technologies into fieldbus systems

Fieldbus systems have been widely used in industrial settings for many years to interconnect instruments within plants to higher-level applications, enabling remote monitoring and control. Initially, custom solutions were used to connect instruments to sensors and actuators, but as the number of devices grew, costs and complexity increased, prompting the need for standardization. In the 1980s, the fieldbus standardization process began to provide a common solution for interconnecting field instruments with controllers [94]. By the 1980s and 1990s, several fieldbus protocols were standardized and made available, including WorldFIP, CAN, PROFIBUS, HART, and Foundation Fieldbus. Each protocol offers distinct characteristics in terms of network topology, distance, data rate, and the number of supported devices, serving various industrial applications, such as manufacturing, automation, and nuclear power plants.

As wireless technologies began emerging, offering significant advantages to consumer electronics, interest in integrating wireless solutions into industrial environments also grew [95]. The motivation for these studies remains strong today, as wireless adoption in industrial settings could provide numerous benefits. Wireless technology would enable new applications that are not feasible with wired systems, such as mobile instrumentation or the installation of devices in areas prone to cable damage, like harsh environments with chemicals, vibrations, or moving parts. Additionally, the most apparent advantages include the elimination of cables, faster prototyping, quicker instrument installation, and easier plant reconfiguration, leading to reduced costs and time.

Given the widespread use of fieldbus systems and the growing potential of wireless technologies, integrating the two has been a topic of great interest. Various research groups and institutions have already pursued this goal in the past, facing different challenges and constraints depending on the specific fieldbus system, with varying degrees of success.

When integrating wireless technology into a fieldbus system, the main goal is often to maintain compatibility with the existing devices. Replacing all the equipment in an industrial plant, already set up and configured for a fieldbus, would be too costly and time-consuming. This is relevant for CERN's control network as well, where thousands of nodes are already deployed, which is why the first approach in this doctoral work focuses on extending the current infrastructure rather than replacing it.

4.1.1 Challenges

As discussed in Chapter 2, a fieldbus implements three layers of the OSI model: the physical, MAC, and application layers. Maintaining compatibility with existing devices primarily means replacing only the physical layer while keeping the MAC and application layers unchanged. This allows wireless technology to be integrated seamlessly, making the transition transparent to the overall system.

However, when running an existing fieldbus over a wireless channel, certain issues can arise due to the different nature of the wireless medium, which is more prone to errors compared to wired connections. These errors can lead to bit errors and packet loss. Several factors must be considered when dealing with a wireless medium.

First is the signal strength, which decreases as the signal travels from the transmitter to the receiver. This reduction in strength is called path loss, and its magnitude depends on several parameters, such as the antenna, frequency, and environment. Path loss can be modelled using various mathematical approaches, one of the most widely used being the log-distance model. In this model, the path loss is expressed as $P_r = P_t \times (d/d_0)^y$, where $d_0$is a reference distance (which depends on the antenna), d is the distance to the receiver, and y is the path loss exponent. The exponent varies based on the environment, typically ranging from 2 to 6, where 2 represents free-space path loss [96]. In industrial settings, path loss exponents between 2 and 3 have been reported [96].

Additionally, when a signal propagates over the air, it spreads in multiple directions and can be subject to diffraction, reflection, and scattering. As a result, multiple copies of the same signal may reach the receiver after traveling different paths and arriving at different times. This time variation is measured by the RMS delay spread, and depending on its magnitude, the signals can interact in various ways. If the delay spread is small, also known as multipath fading, the signals can either interfere constructively or destructively. When destructive interference occurs, the channel is said to be in a deep fade state, which leads to a high rate of errors during decoding [96]. On the other hand, if the delay spread is large, signals corresponding to different symbols may overlap at the receiver, causing ISI, making signal recovery much more challenging [96].

Finally, other issues in wireless communication can arise from co-channel interference from nearby wireless devices, Doppler shifts, and environmental noise [96] —common in industrial settings due to strong motors, electrical discharges, and other machinery. For instance, measurements taken in industrial environments have shown that the delay spread can be significant, leading to strong ISI in wireless links operating at Mbit/s rate [95].

The main consequence of these characteristics is a higher error rate compared to wired channels, resulting in more bit errors and packet losses. For example, packet loss can occur if the receiver fails to achieve bit synchronization, potentially causing delays if the packets need to be retransmitted.

To address these challenges, wireless communication protocols often use training sequences to help the receiver more easily achieve bit synchronization. These sequences vary between protocols, but their presence can pose challenges when implementing a fieldbus over a wireless medium, as will be shown later. For instance, the IEEE 802.11 protocol uses a training sequence lasting 128 microseconds [97], which, as will be shown, is too long to maintain compatibility with the upper layers of some fieldbus protocols.

Moreover, fieldbus protocols were not designed to handle the challenges of a wireless channel, meaning their MAC layers may not account for bit errors or packet losses caused by the channel. This can lead to various issues depending on the specific fieldbus protocol. For example, in the case of PROFIBUS, which relies on a token-passing mechanism where a token circulates between nodes to initiate communication, a deep-fade state in the wireless channel could prevent the token from reaching the next node, leading to severe network instability [98].

4.1.2 Past attempts

Despite these challenges, the potential benefits of wireless technology have motivated multiple attempts to replace the PHY of existing fieldbus systems with a wireless one. These efforts aimed to understand the consequences for the fieldbus protocol, including whether the

network's properties could be preserved with the new wireless PHY. The following section reviews some of the most significant attempts, highlighting the challenges encountered during feasibility studies and the outcomes of these efforts.

4.1.2.1 Simulation of PROFIBUS over the 802.11 DSSS PHY

A widely used fieldbus for which wireless integration experiments have been conducted is PROFIBUS. Developed in Germany and later standardized in Europe [99], PROFIBUS targets a range of manufacturing and control industries. It comes in two versions to accommodate different protocols and services depending on the application. The first is PROFIBUS Decentralized Periphery (PROFIBUS-DP), which connects a single controller to multiple sensors and actuators, where only one node acts as the master, controlling the medium. The second version is PROFIBUS Fieldbus Message Specification (PROFIBUS-FMS), used in distributed control applications, where the role of the master station changes dynamically through a token-passing protocol. In this system, the node holding the token controls the bus, becoming the master and initiating communication with the slave nodes.

Research has been conducted to determine whether the token-passing properties of PROFIBUS can be maintained when the protocol is used over a wireless physical layer. One notable study by Willig explored the challenges of implementing the PROFIBUS MAC layer over the 802.11 PHY particularly the version using Direct Sequence Spread Spectrum (DSSS) [100]. The study identified significant performance issues due to the token-passing nature of the PROFIBUS protocol. Willig measured packet loss and bit error rates (BER) in an industrial environment using the 802.11 DSSS PHY [101], finding that errors occurred not only due to bit-level issues but also from packet losses, such as when the receiver failed to acquire bit synchronization. Moreover, packet losses tended to occur in bursts, alternating with error-free periods.

Using these measurements, a simulation of a PROFIBUS network with 10 stations was conducted. The results showed that even with just a 6% packet loss, the token ring was incomplete 50% of the time. When a BER of 10^{-3} was considered, the ring was incomplete 78% of the time. These results underscore a critical issue when replacing the physical layer of a fieldbus: the protocol is not designed to handle the unreliability of wireless channels, and as a result, the real-time properties of the fieldbus are no longer satisfied [98].

One proposed solution to this problem was an alternative MAC layer for PROFIBUS that replaced the token-passing protocol with a polling-based mechanism [102], which improved the performance of the fieldbus when a wireless PHY was used.

4.1.2.2 The R-Fieldbus project

The R-Fieldbus project [103], funded by the European Union in the early 2000s, aimed to demonstrate the feasibility of using wireless technologies in industrial environments. One of the primary objectives was to explore the development of a wireless fieldbus that could meet several key requirements: support for real-time applications, multimedia applications, and interoperability with existing fieldbus systems, particularly PROFIBUS.

After analyzing these requirements, the project assessed the performance of various wireless technologies, including Bluetooth, UMTS, and 802.11. However, the analysis revealed that while these technologies had sufficient performance to meet the project's requirements, their MAC layers introduced excessive delays (greater than 10 ms), which were unacceptable for fieldbus applications [104]. As a result, none of these protocols were adopted.

Instead, the solution involved developing a custom radio device based on Direct Sequence Spread Spectrum techniques, which utilized the PROFIBUS MAC layer. To support multimedia applications, the device incorporated Internet Protocol (IP) functionalities, and an additional layer was introduced to tunnel PROFIBUS packets through the IP network [105]. To demonstrate the project's feasibility, a custom repeater was also built based on this architecture [106] [107].

4.1.2.3 CAN

The Controlled Area Network (CAN) was introduced in the 1980s as a communication standard for in-vehicle systems and has since been widely adopted in various industries, including aerospace and industrial automation. One of CAN's main advantages is that it enables communication between microcontrollers, devices, and sensors without the need for a host computer. Its key feature is that it is a message-based protocol, meaning messages are not directed from one node to another based on an address. Instead, all nodes receive the same message and decide whether to process or discard it.

However, CAN operates on a Carrier Sense Multiple Access (CSMA) mechanism. In this type of network, a node wishing to transmit first senses the medium. If the medium is idle, the node transmits its message. If the medium is busy, the node waits for the transmission to end before attempting its own. This setup can lead to collisions, as multiple nodes might start transmitting simultaneously when the medium becomes idle. To prevent this, CAN uses a deterministic method to resolve contention, ensuring that the highest-priority message wins and gets delivered.

Past research has focused on how this mechanism could be applied in wireless networks, leading to efforts to transition from traditional CAN to Wireless CAN (WCAN). Various ideas have been proposed, primarily focused around the development of alternative MAC protocols, such as Wireless MAC (WMAC) and Radio Frequency MAC (RFMAC) [108] [109], or token-frame scheme [110].

4.1.2.4 Past trials on WorldFIP extension

A feasibility study exploring the integration of wireless technologies into the WorldFIP fieldbus was conducted in the 1990s. The goal of the study was to determine whether it was possible to extend access to distant sensors and actuators in a FIP network using commercially available wireless solutions. Two theoretical approaches were considered: a repeater-based solution and a gateway-based solution. In both cases, the objective was to ensure compatibility with existing devices, allowing the wireless integration to be added transparently, without modifying the WorldFIP protocol.

The study proposed two potential architectures. In the repeater approach, the focus was on designing a repeater that would bridge the WorldFIP PHY and a wireless PHY. The main idea was to complete the wireless transmission and reception of WorldFIP frames within the strict timing constraints of the WorldFIP protocol. The authors presented the necessary performance metrics that a wireless PHY, particularly one based on Direct Sequence Spread Spectrum (DSSS), would need to meet in order to comply with WorldFIP’s timing requirements [111].

The second scenario, the gateway approach, introduced a new device to the network—a gateway—that would handle communication with distant sensors independently of the

WorldFIP protocol [112]. In this setup, the gateway acts as the sole device on the FIP bus, responding to requests from the WorldFIP master. Simultaneously, it manages communication with distant sensors over the wireless channel, separate from the WorldFIP bus. This approach, unlike the repeater method, relaxes the constraints on the wireless PHY, as it no longer must meet WorldFIP transmission and reception timing requirements.

4.1.3 Summary

Given the past attempts, it is evident that fieldbus performance degrades when the same protocol is implemented over a wireless PHY. The wireless channel, being more prone to errors, makes it difficult to maintain the reliability and real-time requirements of a wired fieldbus [113], as the original protocols were not designed to handle its unreliability.

For PROFIBUS, alternative MAC layers were proposed to address performance issues due to the wireless channel. Concerning CAN protocol, alternative MAC are necessary to use the same protocol features. The studies on WorldFIP instead, did not consider the impact of the wireless channel on fieldbus performance, but provided behavioural descriptions for implementing the transition to wireless.

Due to these challenges, transitioning an existing fieldbus over a wireless channel may not be the best option for most scenarios. This has led to the development of fieldbus systems specifically designed for wireless communication, such as WirelessHART [114] and ISA100.11a [115], which were built from the ground up to handle the challenges of wireless environments.

However, for CERN, transitioning to such wireless fieldbuses in the short term is not feasible. The number of deployed devices is too high, and the equipment is custom-designed by various equipment groups, making a complete upgrade to a new fieldbus resource-intensive. As a result, extending the current network remains a practical solution for addressing the immediate cabling issues in the LHC control network.

Building on previous work with WorldFIP, this doctoral research conducted a feasibility study to explore this challenge, which is discussed in the following sections.

4.2 Extending the WorldFIP LHC control network using SDR

Building on previous work conducted for the WorldFIP protocol, this section extends the repeater and gateway approaches. Specifically, for the repeater approach, it moves beyond the behavioural description to provide a practical implementation of a repeater extension that can be transparently integrated into the WorldFIP network at CERN.

Three key requirements were considered: compatibility with CERN's existing hardware to ensure a plug-and-play solution that can seamlessly connect with the current WorldFIP infrastructure, radiation tolerance to withstand the harsh radiation environment, and full transparency to the WorldFIP protocol.

The section begins with a brief overview of the WorldFIP protocol, followed by a discussion of the design choices that led to the proposed solution. It then describes the implementation of the repeater extension using a SDR platform. This solution was validated with real CERN hardware and software, and the results of the evaluation are discussed in the context of the LHC control network, weighing the advantages and disadvantages.

Finally, the section explores the gateway approach in greater detail, presenting a gateway device that could address the performance limitations identified in the repeater approach. A comparison of the two solutions, highlighting their pros and cons, is provided, along with a discussion of their potential application within the LHC control network.

4.2.1 The WorldFIP protocol

WorldFIP is an industrial fieldbus communication protocol designed for real-time, deterministic data exchange in automation and control systems. It facilitates communication between sensors, actuators, and PLCs/controllers at the lowest level of a factory.
Like other fieldbus systems, it operates using a simplified OSI stack, focusing on just three layers: Application, Data Link, and PHY.

As discussed in Chapter II, the WorldFIP project was insourced at CERN to maintain long term support to the LHC infrastructure, since the company producing the hardware dropped the support. Therefore, after the insourcing, only the subset of functionalities used by CERN equipment’s groups were implemented and are described thereafter.

In the following, only the PHY and the Data Link Layer will be described, as they are the ones imposing the constraints on the wireless extension.

4.2.1.1 Physical Layer

The PHY layer is responsible for all aspects related to transmitting and receiving bits of information between devices. In the case of WorldFIP, the specification dictates that the physical medium used for connections is either twisted copper pair or optical fiber.
WorldFIP uses Manchester encoding to encode digital bits, which relies on transitions between logic levels to represent data.

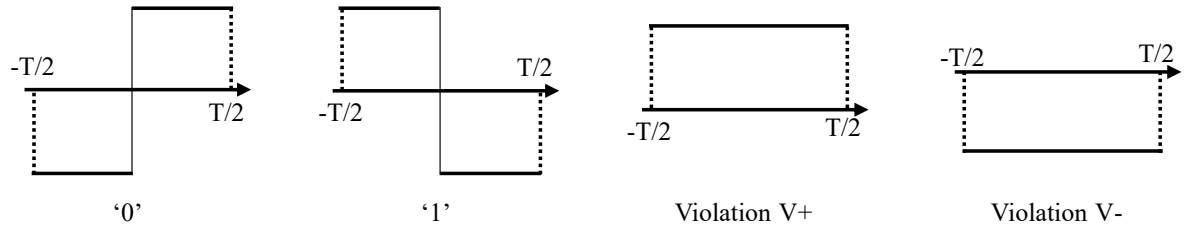

Figure 59 - Manchester encoding used in the WorldFIP PHY to encode binary digits

This encoding, represented in Figure 59, enables the simultaneous transmission of synchronization information and data. Each time interval used to encode a bit is split into two equal parts. A bit is represented by a transition between two voltage levels, while a constant voltage during the entire interval indicates a violation, signaling an error.

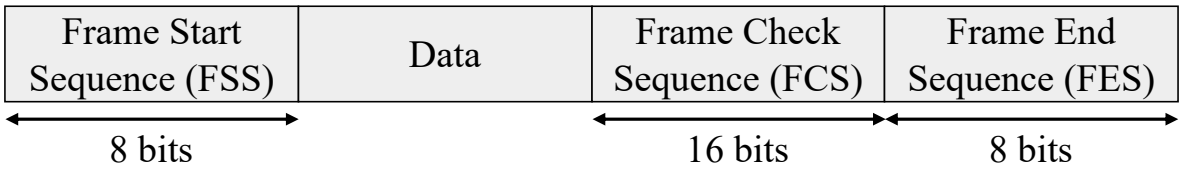

Figure 60 - Structure of a WorldFIP frame

When data is transmitted over WorldFIP, the physical layer adds 24 extra bits to the data and constructs a WorldFIP frame. The structure of this frame, shown in Figure 60, includes three fields in addition to the data field. The Frame Start Sequence (FSS) is an 8-bit field that helps the receiver detect the start of a valid WorldFIP frame. The Data field contains the actual binary data being transmitted. The Frame Check Sequence (FCS) consists of 16 bits calculated based on the data field and allows the receiver to verify whether any errors occurred during transmission. Lastly, the Frame End Sequence (FES) indicates the end of a valid WorldFIP frame.

The protocol supports three transmission speeds: 31.25 kb/s, 1 Mbit/s (the standard speed), and 2.5 Mbit/s.

4.2.1.2 Data Link Layer

The Data Link Layer in WorldFIP facilitates the cyclical transmission of variables. The system is pre-configured with a set of variables and their transmission periodicity, ensuring that data exchange happens automatically, without user intervention. The Data Link Layer is divided into two sub-layers: the Logical Link Control (LLC) and the Medium Access Control (MAC). The LLC provides services for the application layer and utilizes the PHY layer for communication, while the MAC manages access to the communication medium.

Logic Link Layer

Each variable in WorldFIP is identified by a unique identifier, and the participating nodes do not address each other directly. Instead, they reference the variable's identifier, which they recognize either as a producer or a consumer. WorldFIP operates on a producer-consumer model, where each variable has one producer and one or multiple consumers. Variables are exchanged through broadcasting, meaning all nodes receive the broadcasted variable, but only the producer sends it and consumers store the data.

The Data Link Layer consists of a set of produced and consumed buffers, which hold the most recent variable values. These buffers are updated either by the application layer (input from the user) or by the physical layer (when the station receives new data via the WorldFIP bus).

Stations can be designated as producers (the sole producer of a specific variable) or consumers, and the resources for these buffers are allocated during the system's initial configuration. Access to the buffers is managed through the unique producer or consumer identifiers.

Medium Access Control

A WorldFIP network consists of multiple stations with two possible roles:

- Bus management (handling access to the communication bus).
- Production/consumption of variables (data exchange).

While a station can perform both functions, only one station—the Bus Arbiter—manages access to the bus, and all other stations, known as agents, focus on producing and consuming variables.

The Bus Arbiter is responsible for scanning the variables according to the configured periodicity. This information is stored in the Bus Arbiter's table, which lists variables, their types, and their scan intervals. For each variable in the table, the Bus Arbiter sends a request on the bus using a question frame, ID_DAT, which contains the variable's identifier. The frame is broadcasted to all stations. Each station checks the variable identifier, but only the producer responds using the RP_DAT frame. Stations identified as consumers of the variable will then store the content of the RP_DAT frame.

The Bus Arbiter continues this question-answer cycle for each variable in the table, ensuring that data is transmitted in a systematic and periodic manner. This process guarantees the real-time, deterministic exchange of information in the network.

Table 8 – Example of Bus Arbitrator table containing variables identifier with their corresponding periodicity

Variable	Periodicity (ms)
A	5
B	10
C	15

Table 8 presents an example of a Bus Arbiter table containing variables and their respective periodicities, which are scanned continuously, following the same cycle.

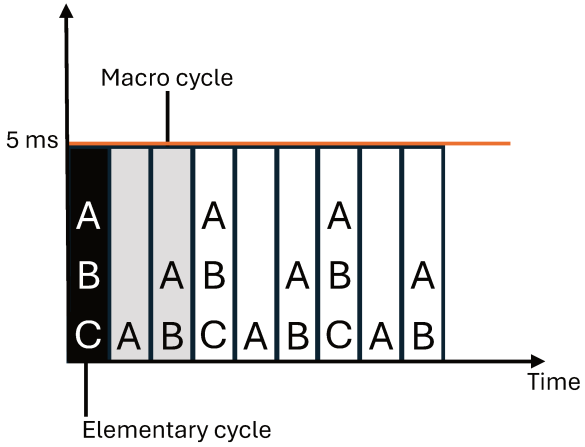

Figure 61 - Macrocycle performed by a bus arbiter

Figure 61 illustrates an example of a cycle executed by the Bus Arbiter. In this example, the Bus Arbiter first scans variables A, B, and C, forming the first elementary cycle with a total duration of 5 ms. For each variable, the Bus Arbiter sends the ID_DAT question frame and receives a corresponding RP_DAT response frame. Any remaining time within the 5 ms cycle is left unused. After this elementary cycle, a second elementary cycle begins, where only variable A is scanned. Finally, a third cycle scans variables A and B. These three elementary cycles form a macrocycle, which is repeated indefinitely by the Bus Arbiter, until stopped by the user.

When the Bus Arbiter requests a variable on the bus, the designated agent responds after a pre-configured turnaround time t_r. If no response is received, the Bus Arbiter waits for a maximum pre-configured silence time t_s before moving on to the next variable. Thanks to this structure, WorldFIP ensures deterministic network communication, as variables are scanned according to their configured periodicity.
Both turnaround and silence times are fixed and depend on the bus speed. With the standard bus speed of 1 Mbit/s used at CERN, the turnaround time is 14 µs, and the silence time is 150 µs.

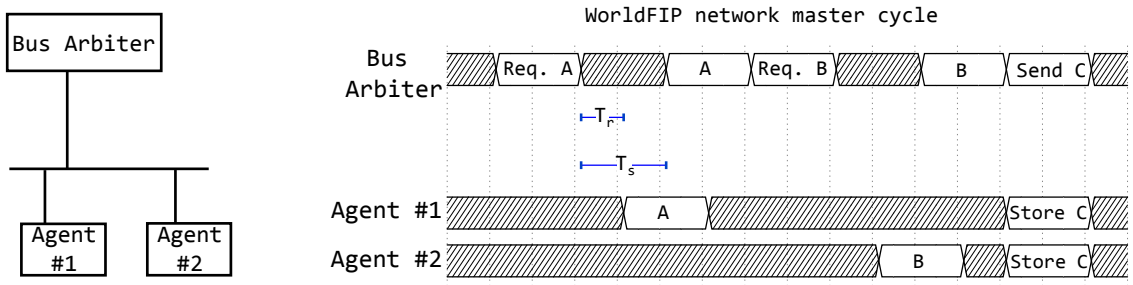

Figure 62 - Timing diagram for a WorldFIP network with a Bus Arbiter and two agents

Figure 62 shows the execution of a macrocycle in a WorldFIP network with two agents and a Bus Arbiter. In this setup, agents #1 and #2 are configured as the producers of variables A and B, respectively, and as consumers of variable C. The Bus Arbiter is also configured as a consumer for both variables B and C. The pre-configured macrocycle scans variables A, B, and C. First, a request for variable A is sent, and agent #1 replies with the RP_DAT frame after the turnaround time. If no response were received within the silence time, the Bus Arbiter would proceed to the next variable. The process continues as the Bus Arbiter scans variable B, followed by variable C, before repeating the macrocycle indefinitely.

4.2.2 Design choices

The development of the wireless repeater extension for the WorldFIP network was guided by three main requirements:

1. **Compatibility with the existing hardware** in the LHC: no modifications to the current hardware interfaces are allowed;
2. **Radiation tolerance**: the electronics used in the system would be exposed to the radiation present in the LHC tunnel.
3. **No modifications to the WorldFIP protocol:** the solution needed to be completely transparent to the devices already deployed in the tunnel, as well as to the control software used in CERN's control room;

The first constraint—requiring compatibility with existing hardware—limits the wireless extension to be implemented using an external device that can easily connect to the current

WorldFIP interfaces at CERN as a plug-and-play solution. The second constraint dictates that the hardware used to implement this extension must be radiation-tolerant. The third constraint, involving protocol transparency, imposes limitations on the wireless PHY being used.

4.2.2.1 CERN WorldFIP implementation

At CERN, each instrument that communicates over the WorldFIP bus implements the WorldFIP PHY and MAC layers using an FPGA. Following the insourcing of the project, CERN goal was to develop a radiaton tolerant agent to replace Alstom ASIC, the MicroFIP chip. The result was nanoFIP [45], an FPGA-based WorldFIP node, that is responsible for implementing the functionalities of WorldFIP MAC and PHY layers. It is the standard interface for any WorldFIP-compatible device at CERN.

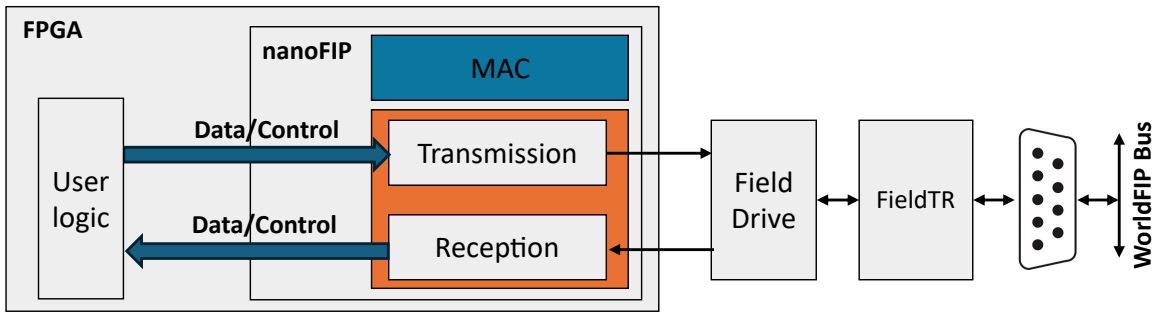

Figure 63 - WorldFIP PHY implemented on CERN WorldFIP devices

Figure 63 shows its block diagram. It consists of two main parts: the physical interface and the WorldFIP logic. The physical interface includes the FielDrive and the FieldTR insulation transformer. The FielDrive is a line driver chip to interface a protocol component to a copper twisted pair using a standard Sub-D 9 pins, through an insulating transformer, the FieldTR. Both the FielDrive and FieldTR components are produced by Alstom.

The WorldFIP logic, implemented on the ProASIC3 FPGA, contains the nanoFIP HDL core, which was developed internally at CERN. This core handles both the PHY and MAC layers of the WorldFIP protocol, enabling communication between devices over the network. The core interfaces on one side to the FielDrive chip through transmit and receive lines, and on the other side to the user logic in the FPGA.

For transmission, the nanoFIP core connects to the transmit side of the FielDrive chip. On the user logic instead, it is interfaced through one data signal and three control signals. Data to be transmitted over WorldFIP is sent to the core byte by byte. The control signals are used to specify the beginning, validity, and end of the byte stream. Once the core receives the signal indicating the beginning byte, it starts by sending the Frame Start Sequence, as specified by the WorldFIP protocol. After this, the core serializes the data bytes received as input, processing each byte until it detects the last-byte signal, which indicates that the next byte will be the last one. At this point, the core embeds the Frame Check Sequence (FCS), followed by the Frame End Sequence (FES), completing the WorldFIP frame.

For reception, the FieldDrive receive line is connected to the nanoFIP core. The core deserializes the incoming bits and provides an interface to the user logic, consisting of data lines and three control signals. These control signals indicate the detection of a WorldFIP frame, the validity of the data, and the end of frame signal, which notifies when the complete

frame has been received. Additionally, an error flag is provided, which, if raised alongside the end of frame signal, indicates that the received frame contains an error.

In addition to the transmission and reception functions, the core implements the functionalities defined by the MAC layer. Specifically, when the Bus Arbiter sends a WorldFIP frame during an elementary cycle, it transmits an ID_DAT frame, which is followed by the response RP_DAT from the WorldFIP agent acting as the variable producer. The WorldFIP protocol sets a fixed silence time (ts), which is the maximum time the Bus Arbiter waits for the RP_DAT frame after sending the ID_DAT. This time depends on the network speed, and in this case, with a speed of 1 Mbit/s, the corresponding silence time is 150 µs.

4.2.2.2 Solutions explored

Considering the requirements, the first point impose that the WorldFIP PHY previously described must remain unchanged. The third point instead, impose that the WorldFIP MAC layer must be fully respected, meaning that the wireless data exchange must occur within the silence time window, otherwise the nanoFIP core will automatically ignore it. With these considerations, two approaches were considered. The first, based on traditional wireless standards, i.e. Bluetooth, LoRa, through the use of COTS devices. The second, based on a custom wireless PHY.

The first solution, the most straightforward, consisted of developing the custom logic to interface the wireless standard with the existing WorldFIP devices at CERN. In this scenario, the WorldFIP frame would be embedded within the wireless standard's frame and transmitted over the air. This option seemed attractive because existing wireless standards are already designed to handle the challenges of wireless channels.

However, two key limitations made this solution unsuitable. First, various COTS wireless transceivers have shown to be unreliable in radiation environments [71], [72], as discussed in the previous chapter. Second, and more critically, the timing constraints imposed by wireless standards introduce delays that are too long to comply with WorldFIP's 150 µs limit. Previous studies on extending fieldbus systems over wireless channels have highlighted that existing wireless standards introduce unacceptable delays [113].

As a result, the second solution was the only viable, and a custom PHY that adheres to the timing constraints of the WorldFIP MAC layer was developed. Previous work already discussed the repeater approach from the behavioural point of view [111], analysing the timing necessary for an hypotetical wireless PHY to respect the timing of the WorldFIP protocol. Nonetheless, the effects of the wireless channel, the radiation effects and the limitation imposed on the hardware compatibility, were not considered in that work. Therefore, they have been addressed in this doctoral research to study the feasibility of such a solution for the LHC control network.

4.2.3 Repeater extension

The repeater extension is described here using a top-down approach. This extension has been designed as a plug-and-play solution that can be connected to any device within the network. Because the extension is fully transparent to the MAC layer, it can be seamlessly integrated with any WorldFIP device, regardless of whether the device is a bus arbiter or an agent.

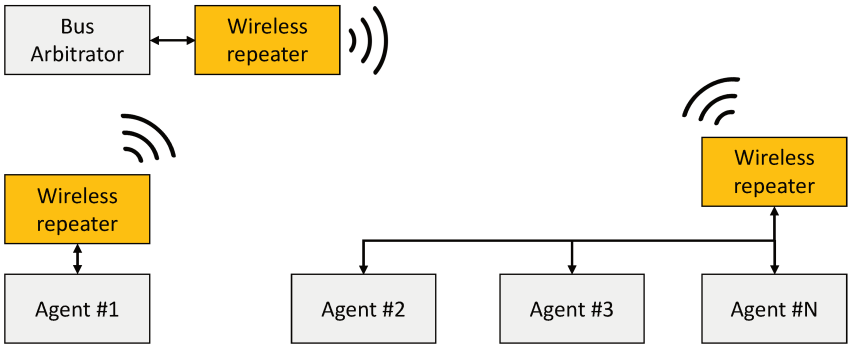

Figure 64 - WorldIFIP network using wireless repeaters. The same device can be connected either to the bus arbitrator, to any agents, or to the bus interconnecting multiple agents, since its totally transparent to the WorldFIP protocol MAC layer

Figure 64 illustrates a WorldFIP network incorporating the repeater extension. It is based on a SDR approach, selected for its long-term flexibility, customizable implementation, and radiation tolerance, as discussed in Chapter 3.

The repeater consists of two primary components: an FPGA and an RF front-end. The FPGA implements the logic required to interface with WorldFIP devices, implementing the CERN WorldFIP PHY discussed earlier. It implements the wireless PHY, developed for transmitting and receiving WorldFIP frames over the air within the maximum timings imposed by the WorldFIP MAC. Finally, the FPGA handles the logic necessary to interconnect the wired WorldFIP domain with the wireless domain, ensuring seamless communication between the two.

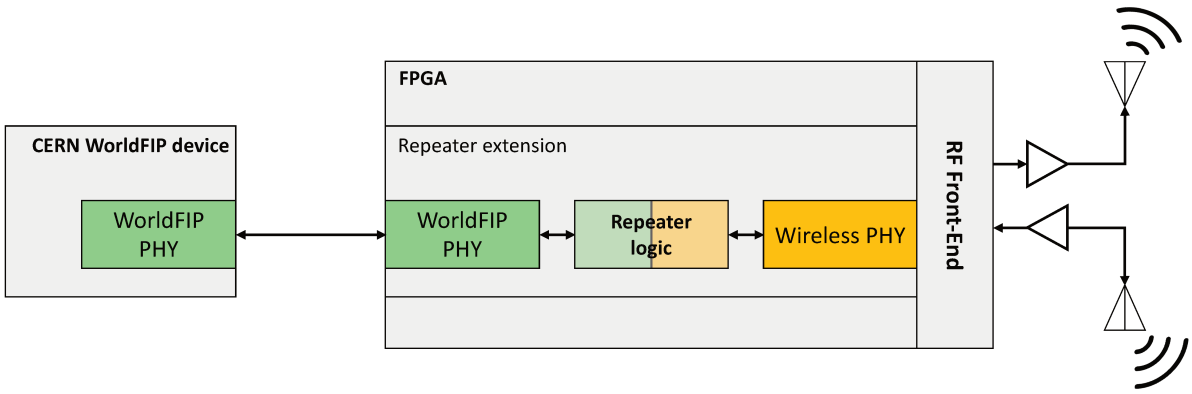

Figure 65 - Top level architecture of a CERN WorldFIP device and its corresponding SDR-based wireless repeater

Figure 65 shows the top-level architecture of the SDR-based repeater. The WorldFIP PHY in this design is identical to the one used in CERN devices, consisting of a physical interface (FieldDrive and FieldTR insulator) that connects to CERN equipment. The CERN nanoFIP core is included but modified to retain only PHY functionalities, such as serialization and deserialization of WorldFIP frames, omitting MAC layer functionalities, i.e. the silence or turnaround time. This adjustment ensures that the repeater extension remains protocol-agnostic, allowing it to connect seamlessly to either a bus arbiter or an agent.

With the WorldFIP PHY, the extension is capable of transmitting and receiving WorldFIP frames to and from CERN equipment. Then, the wireless PHY is necessary to transfer the frames over the air, and the repeater logic is necessary to transfer the frames from one domain to the other.

4.2.3.1 Wireless PHY implementation

To meet the timing requirements of the WorldFIP protocol, the Wireless PHY implementation was kept as simple as possible, avoiding any unnecessary processing in the receiver logic. It uses PPM modulation, which was introduced in Chapter 3. PPM modulates data by varying the position of a pulse within a predefined time slot. This modulation scheme is advantageous because it simplifies receiver design, making detection easier and reducing processing overhead.

The frame structure was also kept as simple as possible. It contains a preamble of 8 bits, followed by the payload (the data field of the WorldFIP frame), and ending with a frame delimiter of 8 bits. No additional error detection mechanisms are included in the wireless frame. This choice was made to minimize computation, relying instead on the Frame Check Sequence already included in the WorldFIP protocol. As a result, any erroneous frame received by the WorldFIP PHY will be discarded, ensuring data integrity without adding complexity to the wireless PHY.

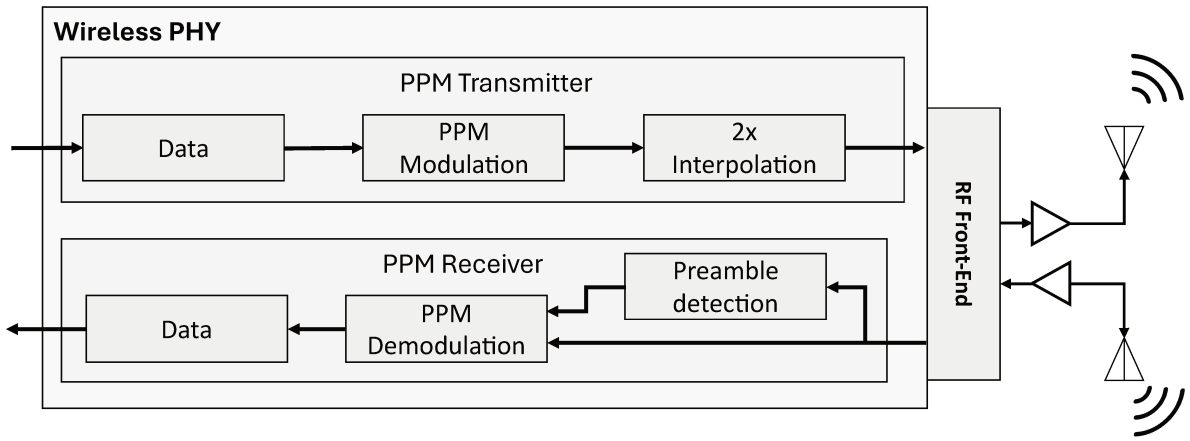

Figure 66 - Top level architecture of the Wireless PHY

Figure 66 shows the block diagram of its implementation. The wireless PHY consists of transmitter and receiver blocks, which handle the DSP functions necessary for wireless transmission and reception. These blocks connect to the RF front end on one side and interface with the repeater logic through control and data signals on the other.

The transmitter includes two Look-up Tables (LUTs) containing sequences for the preamble and frame delimiter, a PPM modulator, and a filter that performs 2x interpolation. The transmitter's control interface consists of a start signal, an end signal, and a data input. Upon receiving a start signal from the repeater, the transmitter begins by sending the 8-bit preamble to avoid excessive overhead. Once the preamble is transmitted, the transmitter modulates incoming data bits using PPM: a bit of value 0 is encoded as a transition from 0 to 1, while a bit of value 1 is encoded as a transition from 1 to 0.
After modulation, the output is upsampled by a factor of 2 to facilitate reception. Incoming bits are modulated in real time, with no idle periods. When the end signal is raised, the frame delimiter sequence is transmitted, finalizing the packet. This approach ensures that, after the initial delay for preamble transmission, bits can be sent as they are received, minimizing idle time. The data rate is set to 1 Mbit/s, matching the WorldFIP PHY's rate. At this rate, the preamble requires 8 μs to transmit, introducing a minimal delay in the wireless transmission.

The receiver includes a preamble detection block, a PPM demodulator, and a frame delimiter block. It receives data from the RF interface and outputs demodulated bits for the repeater logic, along with start and end signals. Data from the RF front-end is routed to both the preamble detection and demodulation blocks. The preamble detection block estimates the signal's average power and compares it against a threshold determined by a FIR filter. When the filter's output surpasses the threshold, indicating the start of a packet, the signal is sent to both the demodulator and the repeater logic.
The demodulator then begins processing incoming data, with the frame delimiter block simultaneously monitoring for the end sequence. Upon detecting this sequence, the end frame signal is triggered, stopping demodulation and signaling the end of the frame for the repeater logic. Given the 1 Mbit/s data rate, the only delay in this process is the 8 µs preamble detection time, after which bits are continuously demodulated and passed to the repeater.

4.2.3.2 Repeater logic

The repeater logic serves to control frame transmission on both the WorldFIP and Wireless PHYs. Essentially, the repeater listens to the WorldFIP PHY for incoming frames, and upon detection, triggers transmission on the Wireless PHY. Conversely, it simultaneously monitors the Wireless PHY for any detected preamble, prompting the transmission of a corresponding WorldFIP frame. This operation relies on using the control signals provided by both interfaces to synchronize and manage data flow between the two PHYs effectively.

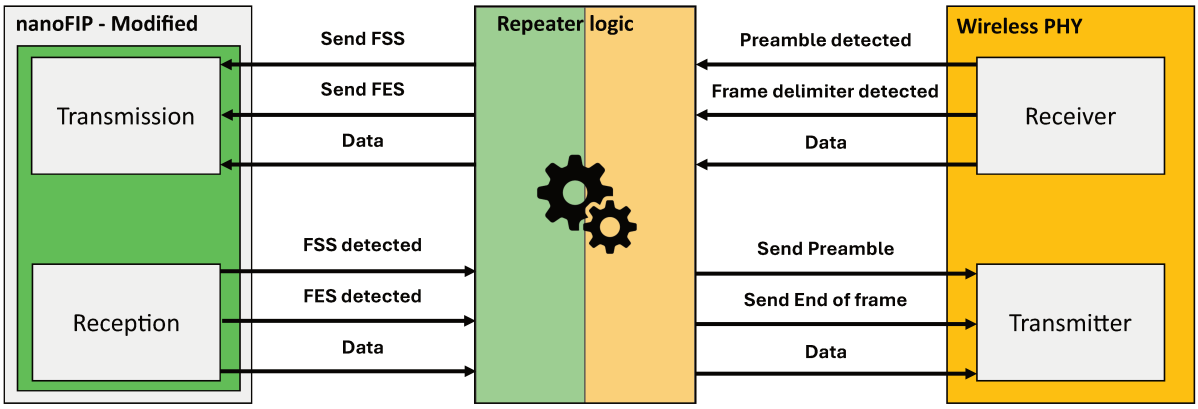

Figure 67 – Illustration of how the repeater logic interconnects the Wireless and WorldFIP domain

Figure 68 highlights the main connections with the WorldFIP PHY and the Wireless PHY. Particularly, its behaviour can be summarised as follows. Concerning the forwarding of packets coming from WorldFIP:

1. When the WorldFIP PHY detects the start of a WorldFIP frame it raises the FSS detect signal. The repeater triggers a wireless transmission using the Send Preamble. Simultaneously, the incoming WorldFIP data is buffered.
2. After the preamble, the buffered data from the WorldFIP PHY is transmitted through the repeater logic and the Wireless PHY without any modifications.
3. When the WorldFIP PHY detects the end of the WorldFIP frame and raises the FES detected signal, the repeater informs the Wireless PHY through the 'Send End of Frame' signal. Then, the Wireless PHY appends an end-of-frame sequence to complete the wireless transmission.

When forwarding packets received from the Wireless side to the WorldFIP instead:

1. When the Wireless PHY detects an incoming packet, the repeater starts a new WorldFIP transmission by raising the ‘Send FSS’ signal. The WorldFIP PHY therefore starts sending a new WorldFIP frame with the FSS on the WorldFIP Bus
2. The data received from the Wireless PHY is then forwarded to the WorldFIP PHY.
3. When the Wireless PHY detects the end of the frame, the repeater logic alerts the WorldFIP PHY through the ‘Send FES’ signal. The WorldFIP PHY then appends the FES, completing the transfer of the frame.

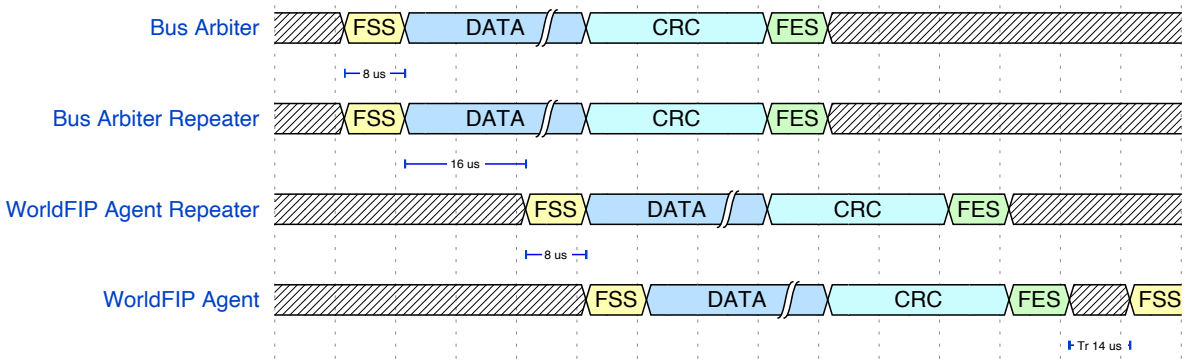

Figure 68 - Timing diagram for a WorldFIP transaction using repeater extension

Figure 68 shows a timing diagram for the exchange of an ID_DAT frame between a Bus Arbiter and a WorldFIP agent, both with their own repeater. The process begins as the Bus Arbiter initiates the transmission of a WorldFIP frame with the Frame Start Sequence (FSS). The FSS is sent over the bus to the repeater, where the WorldFIP PHY detects the FSS after 8 µs, triggering the start signal for radio transmission of the wireless preamble. While this preamble is transmitted, the incoming data bits are buffered. Once the preamble is complete, data bits are sent over the air in real-time. The Bus Arbiter continues with the Frame Check Sequence (FCS) and Frame End Sequence (FES), which the repeater forwards to the Wireless PHY until the FES is detected. At this point, the repeater raises the end signal for the Wireless PHY, which appends the frame delimiter to finalize transmission.

Meanwhile, the repeater on the agent side detects the wireless preamble, triggering the WorldFIP PHY to begin transmission by sending the FSS on the WorldFIP bus towards the WorldFIP agent. The Wireless PHY adds a 16 µs delay before incoming bits are relayed over the bus through the WorldFIP PHY to the agent. When the Wireless PHY detects the end sequence, it raises the end signal, prompting the WorldFIP PHY to append the FES. At this point, the agent has received the full ID_DAT frame and initiates its turnaround time of 14 µs before replying with the RP_DAT frame.

Including the turnaround time of 14 µs and the Wireless and WorldFIP PHY delays of 16 µs each, the repeater starts sending the RP_DAT frame 46 µs after the initial FSS. Clearly, on the opposite direction, the same delays apply, without the turnaround time. Therefore, with an additional 32 µs delay due to the WorldFIP and Wireless PHYs, the RP_DAT frame arrives at the Bus Arbiter after a total of 72 µs. This timing is well within the 150 µs constraint imposed by the MAC layer, confirming it as a feasible solution.

4.2.3.3 Solution validation

Before proceeding with the radiation characterization, it is important to understand the potential impact of wireless communication on the fieldbus. It is essential to keep in mind that this solution does not incorporate any mechanisms for correcting impairments introduced by the wireless channel, such as error correction codes. This means that the fieldbus's performance

may be affected by issues like signal degradation, noise, and interference inherent to wireless communication, potentially leading to bit errors or packet loss.

The described solution has been implemented on an SDR platform and tested using actual CERN hardware. The SDR platform is based on the Zynq-7000 SoC. Specifically, the SoC includes a programmable logic and a processing system. The programmable logic implements the WorldFIP PHY, the Wireless PHY, and the repeater logic described before. The processing system instead runs the code for the configuration of the RF front-end, the AD9361.

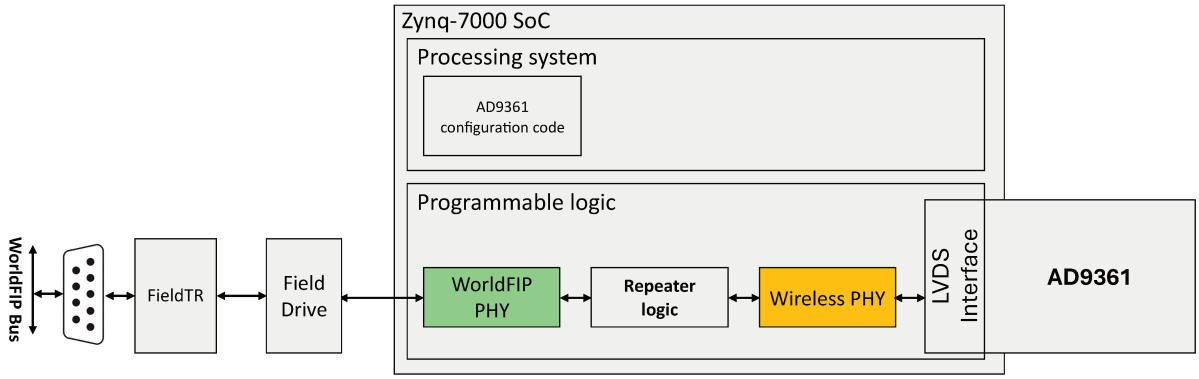

Figure 69 - Architecture of the SDR-based repeater extension used for the system validation

Figure 69 shows a top-level diagram of the full architecture, including the physical interface with the WorldFIP bus. Such architecture has been implemented in hardware using the commercially available development kits: the ZedBoard, which mounts the Zynq-7000 SoC and the FMCOMMS3, which mounts the AD9361. Concerning the WorldFIP physical interface instead, a custom board has been developed. This board adapts the WorldFIP physical interface (DB9 connector) to a PMOD input available on the Zed Board evaluation kit.

For what concern WorldFIP hardware, the standard CERN hardware was utilized for validating the solution. Specifically, the WorldFIP Bus Arbiter was implemented using the FIP USB devkit. This devkit, developed internally at CERN, is based on the legacy Alstom components (the FullFIP2) and runs a standard macrocycle configuration. The macrocycle used for testing was set to a 2-second duration, cycling through 32 WorldFIP agents. For each agent, the bus arbiter requests three variables for which they act as producers and sends one variable for which they are the consumer. The WorldFIP agent employed in this setup was the RadMON device, which is CERN's radiation monitoring device used to track radiation levels in CERN accelerators.

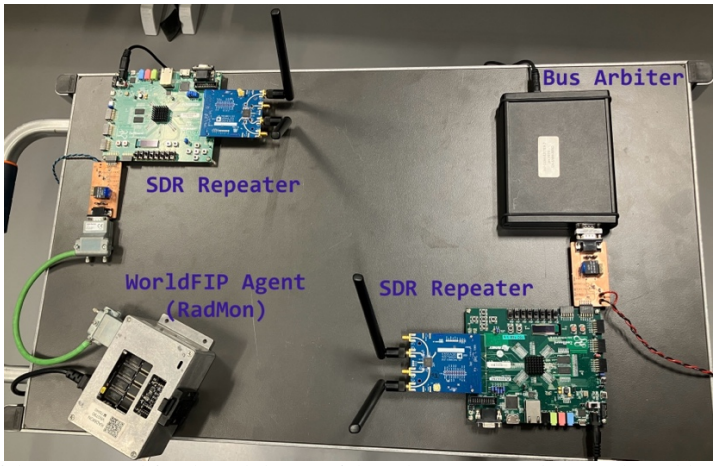

Figure 70 – Overview of the setup used for the validation of the solution. Bus Arbiter and WorldFIP agent are connected to their corresponding repeater, which is SDR-based, made of ZedBoard and FMCOMMS3

Figure 70 shows the complete setup used for the evaluation. The Bus Arbiter and the WorldFIP agent were each connected to their respective SDR-based repeaters using a custom interface board.

As previously mentioned, the repeater design is protocol-agnostic, meaning that the same implementation was used for both the agent and bus arbiter repeaters. At power-up, the processing system runs the initialization code to configure the AD9361 RF front-end. Once this configuration is complete, the system is ready for operation, and the macrocycle can start. From this point, the system follows the operation flow detailed in the previous section. With the system running, various measurements were conducted in a laboratory environment to evaluate performance.

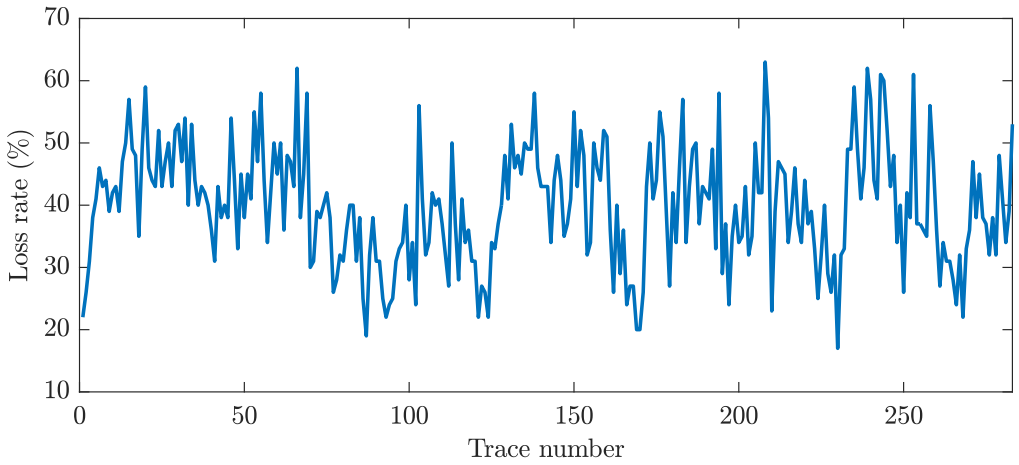

Figure 71 – Loss rate of the variable updates taken with a measurement of 16.5 hours

Figure 71 presents the results from a 16.5-hour measurement run. Although only a single agent was used for this test, the system supports multiple agents as the repeater implementation does not differentiate between them. The figure displays the updates correctly registered by the Bus Arbiter, indicating that for each ID_DAT frame sent by the Bus Arbiter, an RP_DAT frame was successfully received from the agent. The results are organized into traces, each containing 100 updates from the Bus Arbiter, and the figure shows the percentage of missing updates within each trace.

The results reveal that the percentage of missing updates varies over time. In some traces, fewer than 20% of updates were missing, while in others, losses exceeded 50%. These losses stem from missed packets or bit errors during wireless transmission. Packet loss can occur if the receiver fails to acquire bit synchronization, which may be due to various factors such as the short preamble, lack of a training sequence, or other channel-related issues. This can happen on both repeaters: if the repeater connected to the agent misses or incorrectly decodes a packet, it will not detect the request from the Bus Arbiter and thus will not send a response. Conversely, the repeater connected to the Bus Arbiter might miss the agent's reply even if it was correctly transmitted.

Bit errors, on the other hand, result from data corruption during wireless transmission. These errors cannot be corrected because the communication setup only incorporates error detection, not correction. When a radio packet is detected, regardless of its accuracy, the transmission to the WorldFIP starts, directed either toward the Bus Arbiter or the agent. If the packet is corrupted, it will be automatically discarded due to an incorrect FCS in the WorldFIP frame.

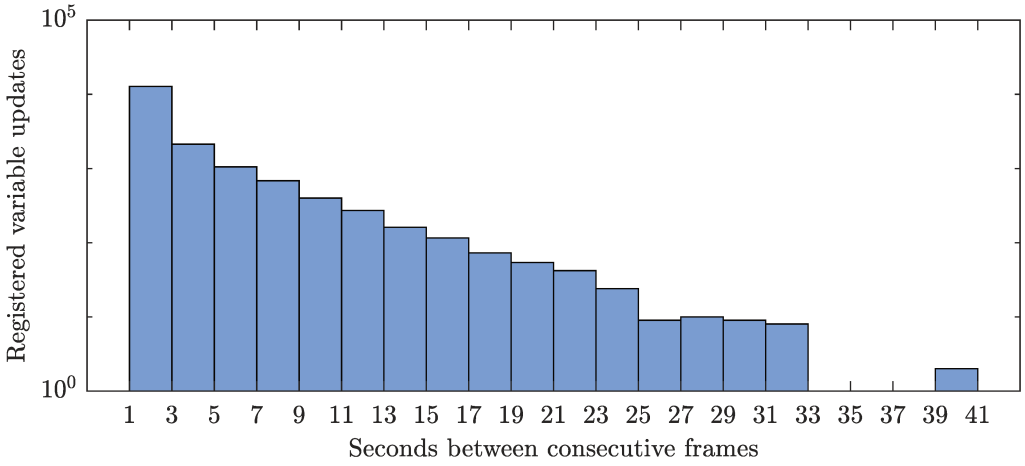

Figure 72 - Distribution of the variable updates period on logarithm scale.

Figure 72 also shows the distribution of time intervals between registered variable updates throughout the entire measurement period. The graph indicates that most of the time during the test, updates occurred without packet loss. However, there were instances where updates were missing for extended periods, exceeding 40 seconds. This performance variability can be attributed to several factors inherent to wireless communication.

First, the modulation scheme and DSP algorithms used in this work are susceptible to multipath interference. This results in multiple echoes of the transmitted signal arriving at the receiver, which can lead to decoding failures and packet losses. The situation is exacerbated by the high transmission rate of the bus arbiter, which scans variables with a latency of 150 µs. The macrocycle configuration scans 32 agents, each with three variables, meaning the bus arbiter repeater sends a new frame over the air every 150 µs, aligning with the WorldFIP protocol's silence time. This rapid burst communication can disrupt the repeaters, affecting wireless transmission quality.

Additionally, synchronization issues contribute to packet loss. The repeaters operate on separate devices with independent clocks, making decoding more complex. The decoding relies on consecutive pulses, making the receiver algorithm sensitive to signal variations. The 8 µs preamble used in this implementation may not be sufficient to ensure reliable packet decoding.

These observations align with findings from previous studies: using wireless communication in a fieldbus system can compromise fieldbus properties due to the inherent challenges of wireless transmission, particularly when no mechanisms are implemented to address these challenges. In this case, the design was intentionally kept simple to avoid adding overhead and to comply with WorldFIP timing constraints.

To improve performance, additional solutions could be considered, such as more advanced signal processing, alternative modulation schemes, error correction codes, and packet retransmission. However, these approaches introduce delays that could conflict with the strict timing requirements of the WorldFIP protocol, necessitating modifications to the protocol itself, which was outside the scope of this work given the requirements.

4.2.4 Gateway extension

Given the challenges posed by the wireless channel and the performance limitations of the repeater extension, this solution might not be suitable for applications demanding higher reliability. In such cases, an alternative approach could be the gateway extension. Although this approach was not fully explored with a hardware implementation and validation in this work, a preliminary analysis was conducted to understand its potential advantages and drawbacks compared to the repeater extension.

The SDR-gateway functions by independently handling agents on behalf of the Bus Arbiter. Its architecture can be split into two main entities: the wireless entity and the WorldFIP entity, both running independently. The gateway is pre-configured with the list of agents and the corresponding variables managed by the Bus Arbiter. It maintains a local copy of the latest available value for each variable in a gateway table, which both entities can access. Once set up, the gateway acts as an agent from the perspective of the Bus Arbiter. The working principle of the WorldFIP entity is as follows:

1. For each variable requested by the Bus Arbiter, the gateway replies with the stored local copy.
2. For variables directed to agents, the gateway stores them in the gateway table.

To achieve this, both the WorldFIP PHY and MAC layers need to be implemented within the gateway.
The wireless entity's role is to update the values from the remote agents and send updates back to them. This entity follows a cyclic behavior:

1. Requesting updated values for each variable from the corresponding agent.
2. Sending variables to the agents as needed.

On the agent side, the same SDR-repeater described earlier can be used, provided it supports the same wireless protocol as the SDR-gateway. Its function remains to forward requests between the gateway and the agent.

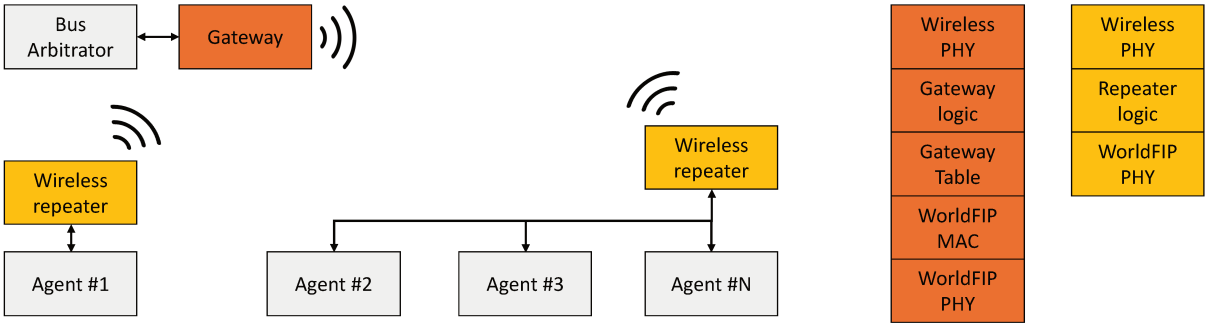

Figure 73 - WorldFIP network with the gateway approach

Figure 73 illustrates an example of a WorldFIP network utilizing the gateway approach. Unlike the repeater implementation, this approach does not impose strict timing constraints on the wireless packet exchange between the wireless entity and the agents. This flexibility allows for the use of either custom or standard wireless protocols for communication. However, the time taken for these exchanges will determine the minimum interval needed for the gateway to update its internal table. Consequently, the overall time required to scan all the agents must be such that the deterministic properties of the protocol are preserved. The selection of the wireless standard will directly influence the minimum cycle time of the fieldbus, which is the shortest duration necessary to cycle through and update all the agents.

4.2.5 Discussion

Two different approaches were examined to assess the feasibility of extending CERN WorldFIP architecture with wireless capabilities. Each approach has its own advantages and limitations that must be evaluated in the context of CERN control system.

The repeater approach showed that the strict timing constraints of the protocol, especially the silence times, impose significant constraint on the implementation of the wireless physical layer. A practical demonstration using PPM modulation showed that while it is feasible to meet these timing constraints, the impact on link reliability limits the usability of the approach without any protocol modification. For example, extending the silence time would allow for techniques such as longer preambles and acknowledgment packets, which could enhance link reliability to acceptable levels. However, adjusting the timing constraints could slow down the variable consumption rate. While systems like the RadMon, which require a 1 second variable consumption rate, could benefit from this adjustment, more critical systems demanding faster rates might not. Additionally, the impact on link reliability should be assessed for each specific application. For instance, the RadMon is not critical for real-time equipment control, could be well-suited to the repeater approach.

The gateway approach, on the other hand, eases the constraints on the Wireless PHY by decoupling the wireless communication from the WorldFIP protocol. This decoupling allows for the use of error correction codes, packet retransmission, and longer training sequences etc. However, the primary challenge in this approach is to ensure that variable updates remain synchronized with the WorldFIP macro-cycle to keep the deterministic nature of the protocol, which is essential for LHC operations. In this approach, the wireless communication must still complete the variable exchange between the gateway and remote agents within one WorldFIP macro-cycle. However, being decoupled from the WorldFIP protocol allows for more flexibility in meeting this requirement. For example, a MAC layer based on Frequency-Division Multiple Access could be used, enabling simultaneous transmissions from agents and significantly reducing the cycle time compared to Time-Division Multiple Access, where agents transmit one at a time. This approach may be better suited for more critical applications where reliability is necessary. However, the complexity increases as a new device—the gateway—must be deployed and configured for each Bus Arbiter.

4.3 Implementation and characterization of a radiation tolerant repeater extension

One of the main objectives of this work was the implementation of a radiation-tolerant solution that could withstand the LHC radiation environment. Once the solution was validated, the next step was to conduct its radiation characterization, which is discussed in this section.

However, conducting radiation tests of the SDR platform in some facilities required the development of a custom test platform. In fact, as described in Chapter 1, in some facilities the irradiation area is wider than the DUT size, which complicates the test when using commercially available development kits. Usually, on a devkit, the DUT is surrounded by active components that could be subject to radiation effects and interfere with the test. Therefore, a custom radiation tolerant tester board was used to carry out the experiments, and it is described hereafter.

Following the introduction of the tester platform, the second part discusses the radiation characterization of the repeater extension. First, the hardware choices are presented. In fact, the SDR platform differs from the one used during the system validation because of the FPGA. In particular, the FPGA part that was investigated to implement the radiation-tolerant repeater is the flash-based PolarFire FPGA, from Microchip [116]. The PolarFire had already been qualified by previous experiments demonstrating great robustness to radiations [117]. Therefore, it represents a good candidate for systems operating in the LHC radiation environment and its use was further explored in this doctoral work, both for the repeater extension, and for the custom wireless physical layer described in Chapter 5.

4.3.1 CRaTeBo: A radiation tolerant testing platform for FPGA-based systems

The test platform was designed to ease radiation testing of FPGA-based systems in facilities where the beam size cannot be set to target the DUT only. Specifically, the main facility for which this platform was developed is the CHARM facility at CERN, which replicates the LHC radiation environment. Its main characteristic is that it comprises an entire room where various positions correspond to different LHC locations in terms of radiation environments. However, the entire room is irradiation, which poses a challenge for radiation testing, because all the components on a board are exposed to radiation, limiting the use of commercial development kits due to the presence of active components that could fail and interfere with test results. This adds a lot of complexity to test setup preparation, because users must design custom radiation-tolerant boards, increasing both time and cost.

To overcome these challenges, the CHARM Radiation Tolerant Tester Board (CRaTeBo) was developed. The primary goal is to provide a complete platform with essential features for radiation testing, including a communication link, power supply, and interfaces for FPGA and additional front-end electronics. CRaTeBo is designed with a high degree of flexibility, allowing it to support various testing configurations. Moreover, all components used in CRaTeBo are either radiation-qualified COTS or radiation-hardened by design. This combination of radiation tolerance and versatility offers a ready-to-use platform that significantly reduces the preparation effort and complexity involved in setting up radiation tests.

4.3.1.1 Board overview

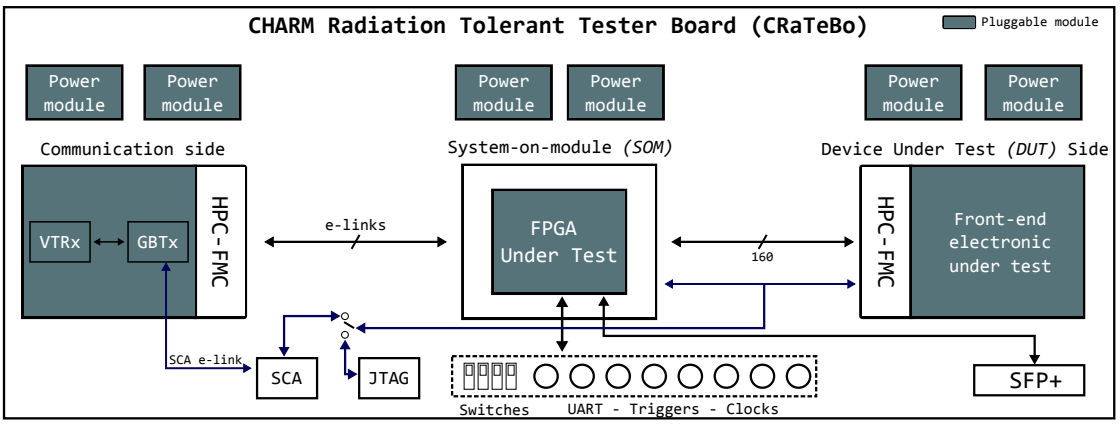

Figure 74 - CHARM Radiation Tolerant Tester Board (CRaTeBo) block diagram

The board, whose block diagram is shown in Figure 74, is meant to be as versatile as possible, allowing to test different kind of applications, with different requirements in terms of interface and power supply. The board features a custom System on Module (SOM) socket to host FPGAs. It mounts two High-Pin Count FPGA Mezzanine Card (HPC-FMC) connectors, compliant with the ANSI/VITA 57.1 standard [118]. On one of the two connectors, the board features a custom radiation tolerant link for the communication between the DUT and the host computer. The second FMC can host additional front-end electronics necessary for the system. Finally, it features a radiation-tolerant and flexible power scheme that can be adapted to provide different power supply according to the system requirements.

4.3.1.2 Communication interface

On one of the FMC connector, CraTeBo features the Versatile Link [119], which is a bidirectional optical link operating at rates up to 4.8 Gb/s. The link can be used to communicate with the host computer, to get test results, or to send stimulus to the FPGA under test. The link uses the Giga-Bit Transceiver ASIC (GBTx) [120], a radiation-hardened serial transceiver designed to implement the GBT protocol [121], developed at CERN. The GBTx ASIC can withstand up to 1 MGy of TID. The ASIC is mounted on the L-GEFE FMC mezzanine card, developed at CERN in the framework of the the GBT-based Expandable Front End (GEFE) project [122].

On one side, the GBTx communicates with the FPGA under test via low-speed electrical links called e-links, which operate at 40, 80, 160, or 320 MHz. On the other side, it connects to an FPGA-based back-end through the Versatile Link through the VTRx [119], a radiation-tolerant optical transceiver that can whitstand up to 500 kGy of TID. This communication link was adopted on CraTeBo because it is extensively used in the LHC systems for high-energy physics experiments and beamline control, hence its presence will simplify the development of test setups for systems that already use this communication standard.

The link offers significant advantages. First, it provides a reliable communication with the DUT, ensuring data integrity during radiation testing, and supports high-data rate processing for more complex designs, such as those involving high-speed ADC/DAC components in detector acquisition systems. Furthermore, the FMC connector on CRaTeBo allows for the potential development of alternative communication interfaces if needed.

4.3.1.3 Power supply

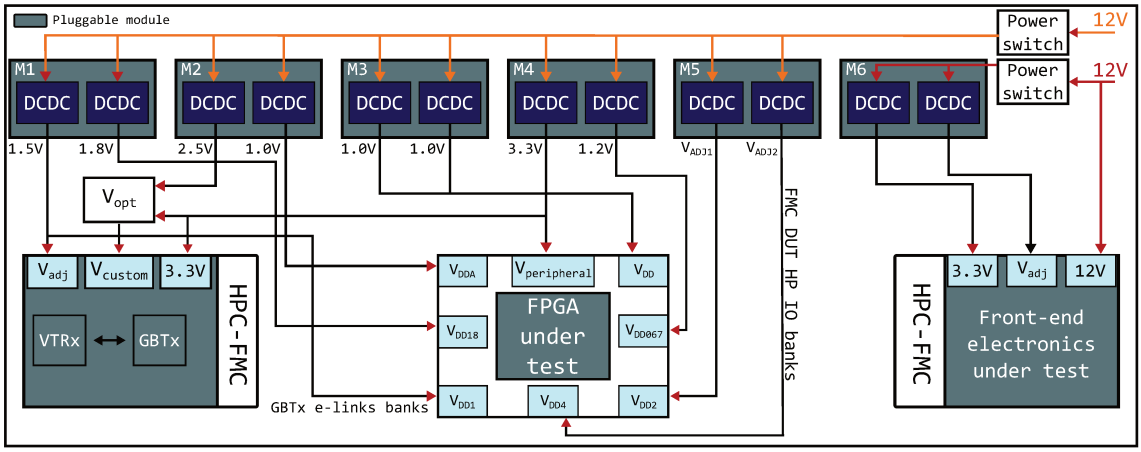

Figure 75 - CRaTeBo power supply scheme

The power supply scheme, shown in Figure 75, can be configured according to the requirements of the system under test. It uses radiation-tolerant power modules, made of two possible solutions: the BPOL12, a radiation-hardened DCDC converter developed at CERN [123], capable of withstanding up to 500 kGy of TID, or radiation-qualified COTS components, tested up to 1 kGy, that can be replaced before reaching the TID threshold. Moreover, the power modules come with a parallel or single output, depending on the amount of current they need to drive. For instance, the FPGA core is powered by a two-output regulator considering that some FPGA applications may be more demanding in terms of current consumption.

Concerning the input voltage, the board features two separate 12V inputs, to supply the carrier and the HPC-FMC for the front-end electronics respectively. The two supplies were decoupled to prevent propagation of radiation-induced destructive event on the mezzanine card to the carrier board.

4.3.1.4 FPGA interface

To host FPGAs, CRaTeBo features a SOM socket. Since no standard SOM solution was available on the marked, a custom solution was adopted. It uses three SAMTEC ADF6-60-03.5-L-4-2-A high-density connectors, which provides 260 I/Os for the FPGA signals. On one side, the FPGA connects to the serial e-links of the Versatile link, whereas on the other side it connects to the HPC-FMC connector for interfacing any front-end electronics. Moreover, it is also connected to LEDs, buttons, switches, and six external signals through SMA connectors.

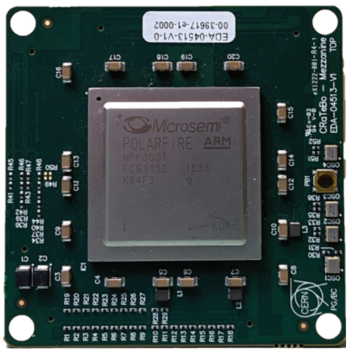

Figure 76 - PolarFire System-on-Module for CRaTeBo

Figure 76 shows the SOM developed for the PolarFire FPGA, which has been used to carry out the radiation tests of the wireless PHYs, as will be described later.

4.3.1.5 Front-end interface

CRaTeBo features a second HPC-FMC connector compatible with the FMC standard, which can be used to connect any front-end electronics mounted as an FMC mezzanine. This choice was done for two main reasons. First, an FPGA-based application is usually coupled with a digital or mixed-signal circuit, i.e. an ADC that samples the analogue signals from a beam position monitor system. Therefore, the slot allows to test the full system under radiation. Secondly, the availability of the FMC enables testing of any other kind of commercial module that comes with this form factor.

4.3.1.6 Slow Control Adapter

The board mounts the Slow Control Adapter (SCA) integrated circuit [124], developed at CERN as part of the GBT chipset. This device distributes control and monitoring signals to front-end electronics. It is radiation-hardened by design, providing robustness against SEE and TID up to 1 MGy. It is hardwired to a dedicated port on the GBTx ASIC through dedicated e-links. It embeds different peripherals: a JTAG master connected to the SOM, allowing to program the FPGA through the GBT/Versatile link, an I2C interface to the DUT mounted on the HPC-FMC, and a 12-bit ADC monitoring the on-board voltages.

4.3.1.7 JTAG Configuration

Different configurations are available for the JTAG interface. The desired programming chain can selected set through a set of jumpers. The source can be chosen between the on-board 10-pin JTAG connector and the JTAG master on the JTAG-SCA. The first can be used outside of the radiation environment with standard JTAG programmers, for debugging/development purposes. The latter instead, allows programming the FPGA through the optical link, therefore it is essential when the board is exposed to radiation. Both the sources can program either the FPGA on the SOM or the DUT on the FMC. Moreover, the FPGA can be used as a master to configure the DUT on the FMC.

4.3.1.8 Other features

The board features an SFP+ connector, 4x user LEDs, 4x switches, 6x SMA connectors, all connected to the FPGA. Two SMA connectors are compatible with the trigger signal of the CERN accelerator, which is based on Transistor/Transistor Logic (TTL). In fact, the board is thought to target CHARM, which is a pulsed beam facility where the particle spill lasts around 300 ms, with an average of two spills per minute. The trigger is synchronised with the spill, hence it can be used to trigger data transfer from the DUT to the host-PC only when there is radiation. This reduces the necessary bandwidth and storage, important for high data-rate applications where the amount of data would be very high in case of continuous acquisition.

4.3.2 Repeater extension: component selection and platform description

The platform implementing the repeater extension targeting the LHC radiation environment differs from the implementation used for its in-lab validation. While the AD9361 was selected as the RF front end because of its robustness to radiation effects, the AMD Zynq-7000 SoC was not. The PolarFire FPGA was selected instead. There are different reasons that led to this choice. When choosing an FPGA for an accelerator environment, the choice usually falls on flash-based FPGA because their configuration memory is immune to SEU. At CERN, the flash-based PROASIC3 was used for all the past developments, whereas the IGLOO2 is currently being used for new systems. The PolarFire FPGA, is a flash FPGA based on SONOS 28-nm technology node, and it had already been characterized under radiation for use in CERN applications [117]. The PolarFire showed resilience up to 5 kGy, which represents a very good

results considering the radiation levels expected during the HL-LHC. During that phase, the electronics will be subject to roughly 100 Gy/year in some locations and considering a lifetime of 12 years for a system, the PolarFire represents an attractive solution to be embedded in future systems. However, the experiments highlighted the susceptibility to SEFI, which cause a temporary failure of the application. Therefore, particular attention was devoted to this category of effects during the radiation tests carried out in the framework of this doctoral work.

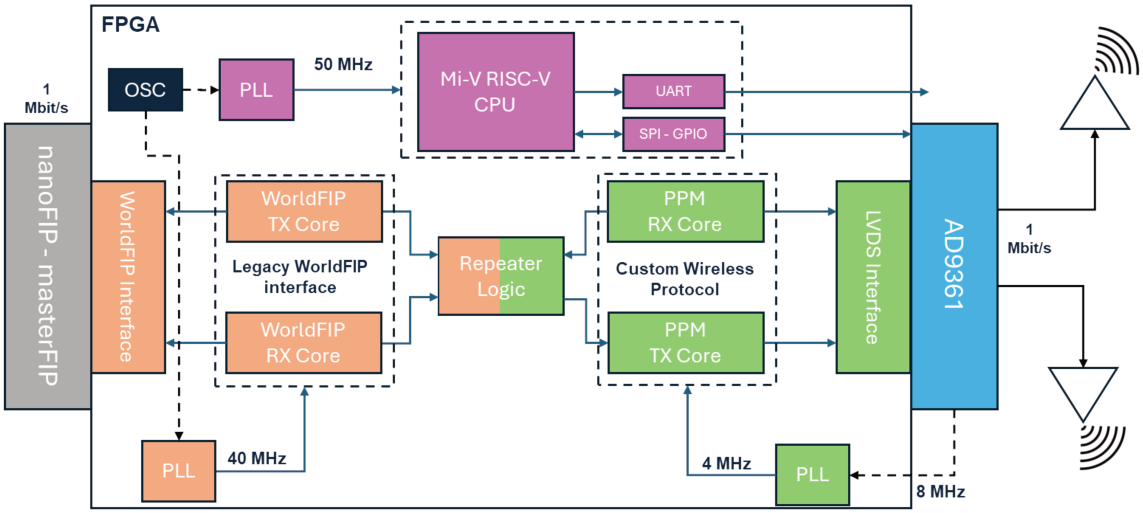

Figure 77 - Repeater extension block diagram implemented on the PolarFire FPGA and AD9361

Figure 77 shows the block diagram of the SDR platform implementing the repeater extension. It includes the AD9361 as RF front-end and the PolarFire FPGA implementing the WorldFIP/Wireless PHY, and the repeater logic described before. Moreover, since the PolarFire FPGA does not include a processor as in the Zynq-7000 SoC, a configuration unit was embedded in the design. The unit includes the Mi-V RISC-V soft-CPU, a General-Purpose Input/Output (GPIO) peripheral, a Serial Peripheral Interface (SPI) and a UART interface. The soft-CPU executes the application code to configure the AD9361. However, this requires storing the application code in a radiation-tolerant memory, which will be discussed in the next chapter. For this platform, the application code was stored on the host-PC and downloaded through JTAG at every boot. The application was stored inside the FPGA memory blocks. The SDR communicates to a host-computer through UART.

The radiation tests have been conducted on a reduced version of the platform that involved only the wireless protocol, the configuration unit, and the RF front-end. The WorldFIP domain was excluded from the test because a) it had already been tested under radiation during the characterization of the nanoFIP core, and b) to evaluate the impact of radiations on the wireless PHY only. The tests were carried out in a neutron and proton environment. The following sections describe the two experiments and the corresponding results.

4.3.3 Neutron radiation testing

The characterization of the device in a neutron environment was performed at the neutron facility in TRIUMF, which was described in Section 1.3.2. The facility covers energies ranging from thermal neutrons up to 400 MeV. The tests aimed at evaluating the sensitivity of the SDR platform to neutrons. Previous experiments showed that the PolarFire FPGA is sensitive to thermal neutrons, whereas the AD9361 had not been tested prior to this study. Therefore, the evaluation was helpful to determine the robustness of the platform considering its application in the LHC tunnel.

4.3.3.1 Radiation test setup

The SDR platform under irradiation consisted of CRaTeBo, mounting the PolarFire SOM, and the FMCOMMS3. In particular, the SOM mounts the PolarFire MPF300T FPGA, whereas the latter mounts the AD9361 RF. The SDR acted as a transmitter only, whereas a second SDR based on the ZedBoard platform acted as a receiver. Only the transmitter SDR was irradiated, while the receiver was kept outside of the beam.

At TNF, the beam area is accessible through a channel, where the components can be lowered down in the beam. The channel is accessible from the test area, which is 6 meters above the neutron channel. Therefore, the transmitter SDR was mounted on a plate, and then lowered six meters down in the beam area.

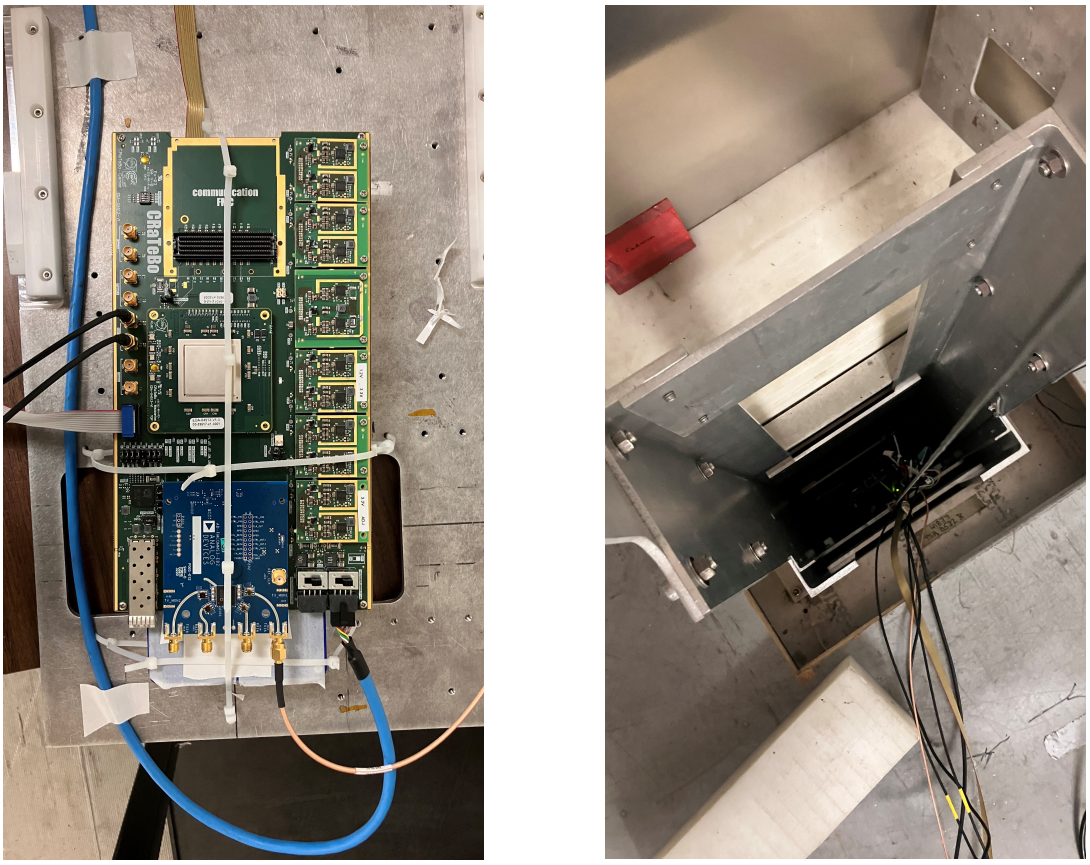

a) b)

Figure 78 – a) SDR under test, with CRaTeBo mounting the PolarFire MPF300T SOM and the FMCOMMS3 with the AD9361. b) The setup, mounted on a plate is, lowered in the beam area

Figure 78 a) shows the CRaTeBo-based SDR mounted on the plate, while b) the mechanism through which the plate with the DUT is lowered down to the beam area. The receiver SDR instead, stays inside the test area.

The transmitter SDR is connected through an SMA cable to the receiver SDR, and to the host PC via UART and via JTAG, to download the application code inside the FPGA memory. The receiver SDR instead, is connected to the host-PC via Ethernet. Both the SDRs are powered by an Agilent power supply.

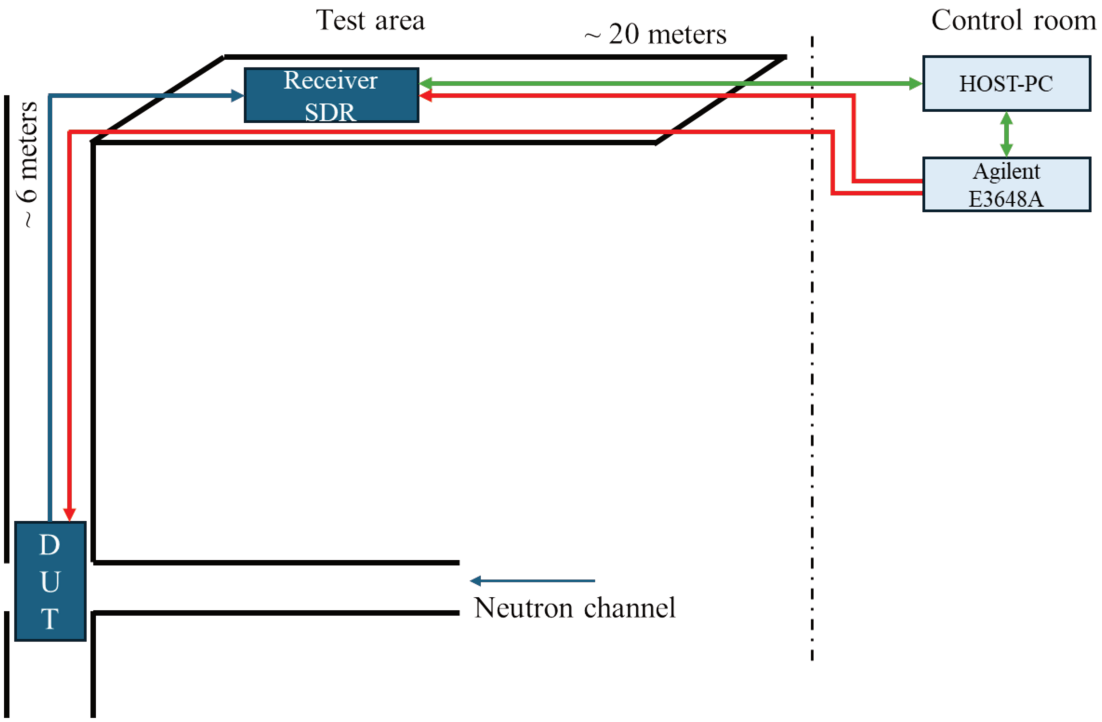

Figure 79 - Irradiation test setup of the SDR platform at TNF facility

Figure 79 shows the schematic overview of the test setup with the different connections and corresponding distances.

The use of CRaTeBo in this test was essential for two reasons. First, at TNF the irradiation area of 5x12cm is wider than the DUT size. Therefore, using commercial kits for the PolarFire FPGA could have compromised the test. Secondly, even though the AD9361 would have suffered from any destructive event, this would have not propagated to the carrier board since the power is decoupled.

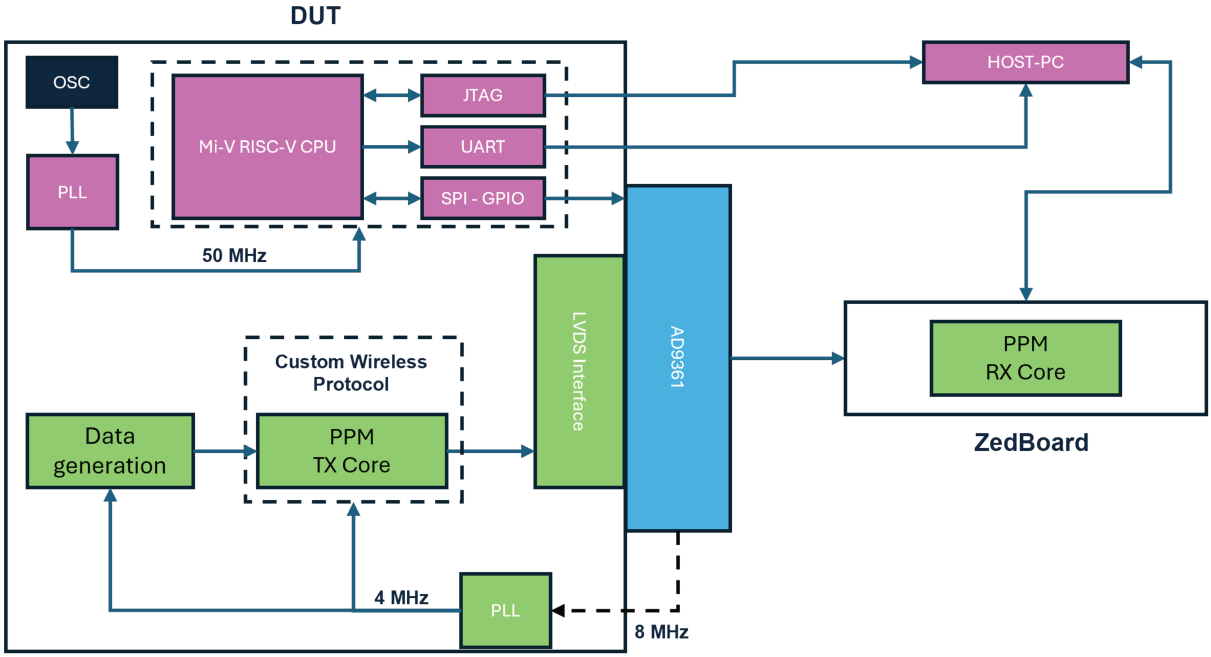

Figure 80 - FPGA implementation during the radiation test

Concerning the application under test, the transmit unit of the Wireless PHY and the configuration unit were implemented in the FPGA. Given that no WorldFIP logic was implemented, a fixed payload was used to create a packet that was transmitted continuously by the Wireless PHY, with a transmission rate of 1 packet/second. Moreover, after the configuration of the AD9361, the RISC-V within the configuration unit was set to send a fixed message to the host-pc through UART. The message was used to understand whether the configuration unit or the Wireless PHY were impacted by any radiation effects. The RF transmission occurred at 2.4 GHz. The SDR receiver, outside of the beam, implemented the receiving part of the Wireless PHY. Therefore, after decoding the packets, the payload was sent to the host-pc through the Ethernet interface via UDP, where they were logged using MATLAB software, which was used to receive, in parallel to packets logging, a dump of the RF data before the demodulation. Figure 80 shows a schematic of the application under test.

The main objective of the test was to identify a) the susceptibility of the Wireless PHY to SEE, meaning errors in the packets or transmission interruptions and b) the communication losses due to the AD9361.

4.3.3.2 Experimental results

The two components of the SDR were irradiated separately to discriminate the source of errors, nonetheless the test setup and the application under test stayed the same for all the irradiations.

Concerning the AD9361, in the previous proton experiments described in Chapter 3, the configuration registers were categorized based on the impact of SEU on the communication. Registers, if affected by SEU, can interrupt the communication, degrade performance, or remain masked, in the case of status register. However, differently from what has been observed during the proton irradiation, no loss of communication was observed in this experiment, for a cumulated fluence of $3.55 \times 10^{11}\ neutrons/cm^2$. For what concern the other category of SEUs, in this case the test setup did not allow to monitor them.

During the irradiation of the PolarFire instead, an issue in the host-pc prevented logging of the received packets, hence statistics like packet error rate or bit errors could not be analysed. Only the loss of communication is discussed in the following.
Three communication losses were observed, where the receiver SDR did not receive any new packets, indicating that an SEE occurred in the transmitter SDR. However, it was not possible to exactly identify which part of the design failed or was still functional through the UART interface. Still, a hypothesis on the possible source of failure could be done considering the current consumption of the SDR platform.

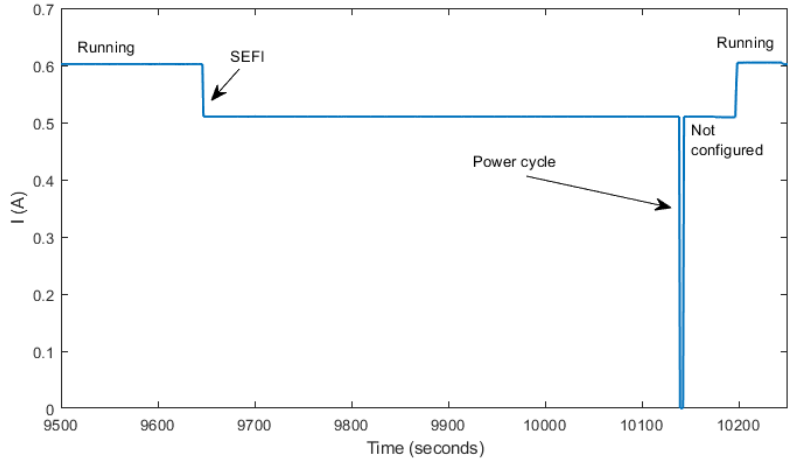

Figure 81 - Current consumption of the Transmitter SDR during irradiation

Figure 81 shows the current consumption of the transmitter SDR during irradiation. In each of its state, the transmitter SDR has a different current consumption. Immediately after power on, hence while the AD9361 is not configured yet (indicated on the picture as Not configured) the current consumption is around 0.5 A. When the SDR becomes operational, which means after the RISC-V CPU completed the AD9361 configuration, the platform consumes in average 0.6 A. When the loss of communication occurred (indicated as SEFI on the figure), the current consumption dropped again to 0.5 A, at the same level as before the configuration. One of the possible reasons for this behaviour is the radiation induced reset, categorized as SEFI, already shown by past radiation tests on the PolarFire FPGA. When it occurs, the FPGA logic goes in a reset state, including the GPIOs connected to the AD9361. Among the GPIOs, one of them is connected to the reset pin of the AD9361, necessary for its configuration. Therefore, if a radiation-induced reset occurs, this will initialize the AD9361, and when the PolarFire recovers from the SEFI, since the AD9361 is not configured anymore, the transmission does not recover automatically.

The device cumulated a fluence of 3.22×10^{11} $neutrons/cm^2$, and three losses of communications were recorded.

Table 9 – Cross-sections for the communication loss due to SEFI of the AD9361 and the PolarFire FPGA.

DUT	Event	Occurrences	Fluence $[n/cm^2]$	$\sigma\ (cm^2/device)$		
				Lower limit	σ	Upper limit
AD9361	SEFI	-	3.55×10^{11}	-	-	
PolarFire		3	3.22×10^{11}	1.86×10^{-12}	9.32×10^{-12}	2.72×10^{-11}

Table 9 reports the fluences cumulated, and the cross-sections of the link interruptions due to the two components, with the corresponding lower and upper limit cross-sections with the 95% confidence interval.

4.3.4 Proton radiation testing

The device was characterized using the 200 MeV proton beam at the PSI facility. Since the AD9361 had already been characterized in this facility, the experiment focused on the study of the SEE on the Wireless PHY, also taking as input the results retrieved from the neutron irradiation to improve the setup. Different configuration of the Wireless PHY have been implemented to investigate possible mitigations to SEFI, and additional observability points have been added to discriminate the source of failure.

4.3.4.1 Radiation test setup

The radiation test setup consisted of a single SDR platform implementing both the PPM transmitter and the receiver, the RISC-V subsystem and the RF front-end. A fixed payload was generated inside the FPGA, and it was sent at a fixed data rate of 1 Mbit/s. Data was transmitted to the AD9361, which was configured to implement an internal loopback to send the data back to the FPGA, where it was decoded by the receiver. A custom UART module was used to report diagnostic information to the host-PC.

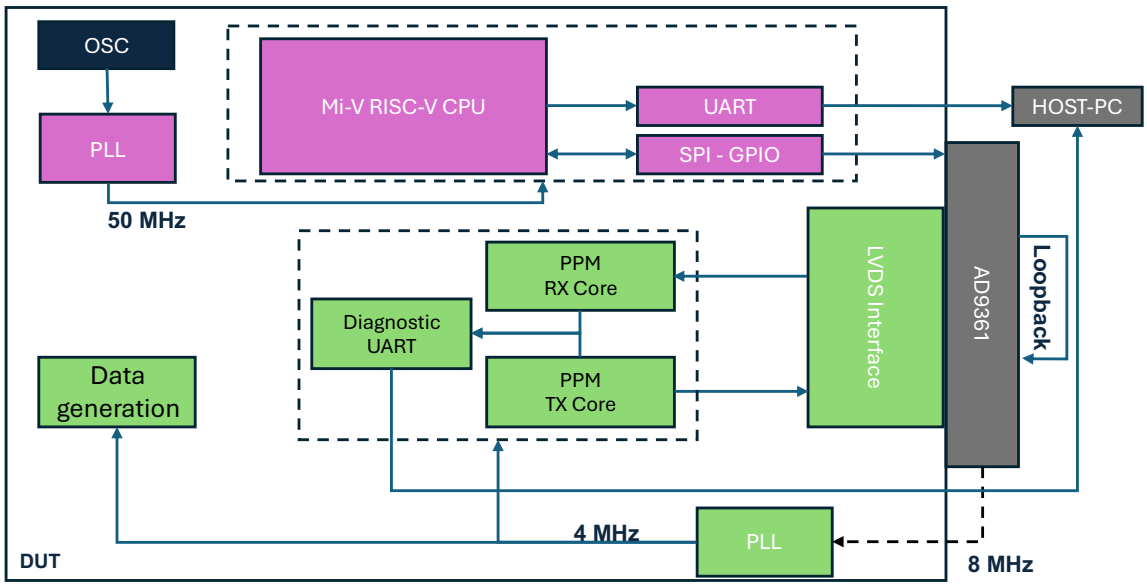

Figure 82 - Radiation test setup at PSI

Figure 82 shows an overview of the test setup used during the irradiation. The test setup was used to evaluate different metrics of the Wireless PHY: number of incorrect packets, SEFIs, wireless core crashes, and CPU crashes. The metrics could be collected using the two UART interfaces available in the two separate clock domains, indicated in green and purple in the figure.

For what concern the CPU crashes, these could be identified thanks to the UART peripheral inside the configuration unit the monitors the execution of the application code on the RISC-V CPU during the operational phase, right after the configuration of the AD9361 is completed. If the CPU crashes, the related UART interface is lost as well, while the UART interface in the wireless domain stays active. In the event of a CPU crash, however, the test can continue because the configuration unit is not necessary anymore after the AD9361 configuration.

Communication errors instead are reported by the diagnostic UART integrated inside the wireless unit. Each core, transmitter and receiver, include statistics about the number of transmitted and correctly received packets respectively. Since the AD9361 implements a digital loopback, and it is not irradiated, communication errors can be due only to radiation induced effects in the FPGA. Hence the counters reported through UART can be used to understand how many packets have not been received correctly. However, there is no distinction among lost or corrupted packets. In fact, a SEU can prevent the receiver logic to detect a packet, but it can also induce an error in the payload during the transmission, which will not be counted as a received packet by the receiver. This is the expected behavior considering the repeater approach, because in both cases the protocol does not foresee any retransmission.

A Wireless core crash is identified when the diagnostic UART connected to the transmitter/receiver core starts reporting a constant number of transmitted/received packets. This indicates an SEU inside the wireless core which caused the transmission/reception to halt.

For what concerned communications losses categorized as SEFIs, according to the results of the neutron campaign, these were likely due to radiation induced reset, causing the initialization of the Ad9361. Therefore, an initialization message was sent by the two UART peripherals as soon as they are out of reset. This, together with the current signature, could confirm whether the loss of link was due to a radiation induced reset of the FPGA.

In addition of the evaluation of these metrics, one of the objectives was to mitigate the loss of the link, i.e. the SEFIs, and to eventually reduce the number of corrupted communications or CPU/wireless crashes. Two mitigation techniques were adopted for this purpose. Concerning

the SEFI, one possible source of radiation induced reset was identified in the PolarFire System controller. According to the user manual, the controller monitors the power rail state of the FPGA and releases the FPGA from reset when everything is ready. After that, it provides additional features such as security or online updates. However, as suggested by the user manual, it should be suspended when operating in radiation environments. Therefore, one approach consisted in suspending the system controller to see if it impacted the SEFIs occurrence. In addition to that, the TMR mitigation technique was implemented to understand if packet losses or CPU crash could be reduced. The triplication was implemented by replicating all the registers in the design by using the functionalities available in the Libero tool (the version used for this experiment is Libero 2021.3).

Table 10 - Naming convention used for the four implementations of the design under test

Design name	TMR	System controller
RT-PPM1	No	Yes
RT-PPM2	No	No
RT-PPM3	Yes	Yes
RT-PPM4	Yes	No

Table 10 summarizes the four different designs under test with the corresponding feature enabled and disabled.

4.3.4.2 Experimental results

Table 11 - Events and fluence reached for the four designs under test using a proton flux of 3.80×10^8 p/cm²/s

Error signature	RT-PPM1		RT-PPM2		RT-PPM3		RT-PPM4	
	Events	Fluence	Events	Fluence	Events	Fluence	Events	Fluence
SEFI	16	4.28×10^{11}	8	2.61×10^{11}	27	5.43×10^{11}	18	3.11×10^{11}
CPU Crash	16	7.91×10^{10}	2	1.23×10^{10}	23	1.40×10^{11}	18	1.07×10^{11}
Corrupted communication	18	4.13×10^{11}	13	2.11×10^{11}	38	4.95×10^{11}	5	2.56×10^{11}
Wireless core crash	6	3.60×10^{11}	4	3.62×10^{10}	12	1.64×10^{11}	13	9.32×10^{10}

Table 11 reports the error signatures that were observed during the experiment for each design under test, together with the number of events, and the cumulated fluence. The experiment was carried out with a proton flux of 3.80×10^8 p/cm²/s.

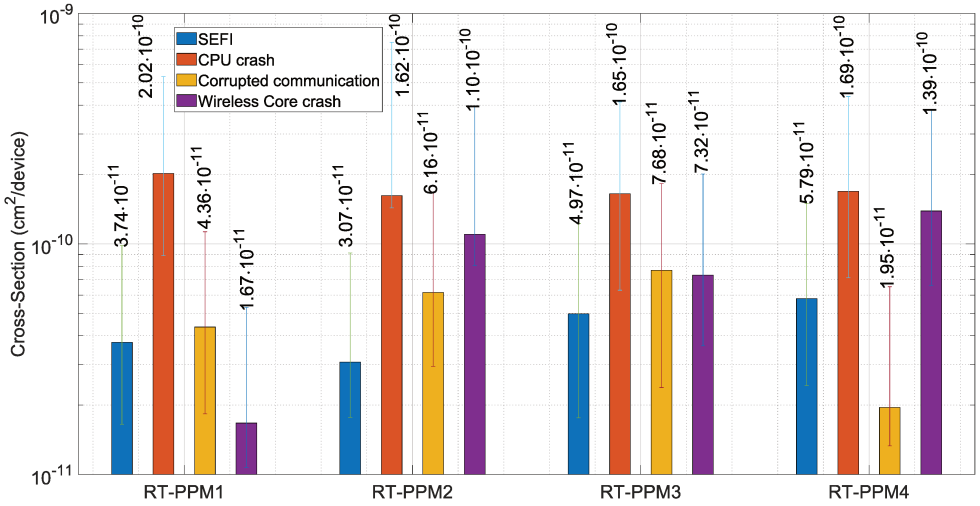

Figure 83 - Proton cross-sections for the different error signatures observed during the irradiation

The cross-section of the errors observed is reported in Figure 83 with the 95% confidence interval. Different observations can be made with these results.
The cross-sections for CPU crashes were similar across the four designs, regardless of the use of TMR. This behaviour is reasonable because the application code executed by the CPU is stored in the internal FPGA memory, which was unprotected in all tested designs. If the memory gets corrupted by a SEU, the CPU execution halts, leading to crashes. This suggests that CPU reliability is the main factor influencing the CPU operation, and clearly TMR does not mitigate this issue. To improve the situation, error correction schemes could be implemented.
Concerning the crashes of the wireless core, the designs with TMR (RT-PPM1 and RT-PPM2) showed higher cross-sections compared to the designs without. This could indicate that TMR introduces additional logic that increases potential points of failure in the FPGA design.
The use of TMR also increases the cross-section for the corrupted communications compared to the standard versions, which could be due to the additional voting logic. However, surprisingly, using TMR in combination with a disabled system controller shows a 4x reduction in the cross-section compared to the version with the system controller.
Finally, all the designs exhibited a SEFI cross-section of the same order of magnitude, with the TMR designs showing a slightly higher cross-section. Despite the suggestion in the user manual that system controller should be disabled to avoid unintended reset when operating in radiation environments, these radiation tests indicate that this is not the case.

During the irradiation of the four designs, only the FPGA was exposed to radiation, therefore no degradation in the link was expected given the robustness of the PolarFire FPGA, as demonstrated by previous experiments, where it was tested up to 5 kGy. Nonetheless, the device cumulated a total dose of 1 kGy and was still operational and programmable.

4.4 Discussion

The impact of radiation effects on the wireless logic could be divided into four main categories: CPU crashes, wireless core halt, corrupted communications, and SEFIs. Their effect can then be analyzed on the main application of this wireless link, i.e. the repeater inside the WorldFIP infrastructure in the LHC.

The crash of the CPU after the configuration phase does not have any impact on the wireless link, since the communication is not affected at all. Therefore, the reliability of the WorldFIP fieldbus is not affected neither. Nonetheless, the lack of an operational CPU could represent an

issue in case a reconfiguration of the AD9361 is necessary. In Chapter 3, the experimental results from the proton irradiation of the RF transceiver showed that SEU on the configuration registers could decrease the modulation accuracy or interrupt the communication, requiring a re-configuration of the device. Depending on the source of the crash, there are different solutions that could be adopted to increase the CPU availability. For instance, one hypothesis is that the crash is due to the corruption of the application code stored inside the FPGA memory, which was unprotected. In this case protecting the memory with ECC could improve the situation. However, the ECC may not prevent double errors, which could still cause the crash of the application. Therefore, one could think about other mitigations technique such as scrubbing, which continuously refresh the content of the application memory. Another possible reason could be the corruption of program counter, which stores the address of the instruction to execute. If the register is corrupted and the processor goes into an unforeseen state, it will stay there until reset. Hence, a watchdog can be implemented to monitor the processor status, and in case of anomaly, issues a reset which restarts the application.

The watchdog would be a good solution even in the case of a halt of the wireless core. The halt could be due to different reasons. For instance, an SEE occurring inside the preamble detection logic could prevent packet detection, whereas an SEE inside the transmission could stop the transmission of the packets. During the test, it was not possible to discriminate the source of error. Nonetheless, a possible point of failure could be the FSM inside the transmitter/receiver wireless cores, which go into an unknown state until the next reset. For this reason, a watchdog monitoring the status of this cores could restore their functionalities.

For what concern the communication errors, they are due to packet losses, or payload errors. The first category contains all those packets that were not decoded at all because the preamble was not detected. The second includes those packets for which the payload is wrong. However, because of the internal loopback inside the RF front-end, errors can be due only to SEE in the logic. The source of errors can be multiple. For instance, data is generated by using a fixed payload, which is streamed through a FIR filter for interpolation before going to the RF interface, where is then looped back to the receiver side. An SEE on any of the logic element could flip the content of any bit, resulting in a wrong payload being sent, resulting in a wrong packet at the receiver. The same could happen in the receiver section. Among the mitigations that could be applied to reduce the cross-section for these events, only TMR was evaluated. However, it did not bring many advantages considering that the cross-section for these events was consistent among the different designs, except for the version with system controller in suspended mode. However, it is very unlikely that the system controller can have an impact on the SEE rate inside the FPGA logic elements. The higher cross-section reported for the other designs, could be the indication that local TMR, applied at register level, is not efficient in these occasions. In addition to TMR, the payload storage could be protected with ECC, since it is stored in a memory blocks (even though this will not be the case for operation, where the payload will be provided by WorldFIP device), or TMR could be applied at different hierarchy levels to evaluate possible improvements. The rate of communication errors will impact the control application. In fact, no error correction is available neither in the wireless link, nor in the WorldFIP protocol (which only provides a checksum for integrity check). Therefore, in case of any error, the packet will be discarded, resulting in a missing updated for the variable concerned.

For what concern SEFI instead, they cause a permanent link loss, even though the FPGA itself recovers from the SEFI. This is due to the loss of configuration of the RF front-end, first observed during the neutron testing campaign with the current consumption, and then

confirmed by the diagnostic interfaces during the proton testing. After the loss of configuration, the AD9361 does not provide any more the clock to the wireless core domain, which cannot resume operation. The diagnostic interface instead, which is on another clock domain, works correctly. To cope with this effect, the RF must be reconfigured by the soft-CPU. However, after the SEFI, the FPGA memory blocks are erased, hence the soft-CPU has no code to execute until it is manually uploaded through JTAG. Nonetheless, the link loss can be considered just temporary. In fact, by storing the code necessary for the RF configuration in an external memory, a bootloader could be implemented that automatically copies the code from the external memory to the FPGA and releasing the reset of the processor, which will configure the RF device. In fact, this last aspect will be covered in the next chapter, which will discuss the bootloader implementation together with the radiation characterization of the memory storing the application code.

4.4.1 Link availability and lifetime across the LHC

In section 4.2.5, the pros and cons of the repeater approach were discussed. Even though the PPM wireless PHY may not reach adequate requirements for all the applications, the platform used for its implementation will eventually be the same. Since it will be integrated in the LHC control network, it is interesting to estimate what could be the impact of the SEFI on the platform for the LHC control network. Among the failures, two of them will cause the loss of the link: the SEFI on the PolarFire, and the SEFI on the AD9361, reported in Section 3. Nonetheless, the higher-cross section is represented by the SEFI on the PolarFire. Considering the upper-limit cross-section retrieved from the RT-PPM2 design, the number of link losses can be evaluated for the different areas of the accelerator.

Table 12 - Link availability across the LHC considering the radiation levels of the HL-LHC. The calculation assumes that all the nodes make use of the repeater extension.

LHC Area	Annual fluence [p/cm^2]	Communication nodes	Upper Limit Annual failure rate	Upper limit Annual link losses
DS	5×10^{10}	~600	1.53	~918
ARC	1×10^{9}	~9000	0.03	~270
UJ	5×10^{9}	~400	0.15	~60
RR	3×10^{9}	~2000	0.09	~180

Table 12 reports the approximative number of devices installed in the LHC areas, the expected annual failure rate based on the radiation levels expected during the HL-LHC. The computation considers that all the device in the network will make use of the repeater extension. In the DS, the area with the highest failure rate, roughly 900 link losses per year can be expected. On the other hand, the ARC contains most of the instrumentation, and only <300 failures per year are expected. It must be noted that, once a recovery mechanism is implemented, the link losses will be just temporary, causing the instrument unavailability for a short amount of time. This amount of time includes the application reloading and the RF configuration, which will be better discussed in the next chapter when presenting the solution proposed to cope with this issue.

The platform is very robust against TID effects as well. The PolarFire FPGA was tested up to 5 kGy in a previous experiment [117], whereas the AD9361 was tested up to 500 Gy in this work, but up to kGy [93] without degradation or loss of functionalities. Considering that during the HL-LHC run the yearly dose expected for electronics is roughly 100 Gy/year in the DS, the platform will be operational for the whole lifetime of the system (usually 12 years of operation).

4.5 Chapter summary

A solution for the short-term cabling issue for the WorldFIP network in the LHC has been discussed in this chapter. Previous research on the integration of wireless technologies in fieldbus systems was analysed, and a deeper study was carried out for the specific case of the WorldFIP network in the LHC.

In the first part, two implementations were discussed, a repeater and a gateway approach. For the repeater approach, a full implementation was given considering the specific requirements of the LHC, i.e. no change to the current infrastructure, and radiation tolerance. For its seamless integration, a custom wireless transmission algorithm was implemented, tailored to the timing requirements of the WorldFIP protocol. The solution, conceived as a plug-and-play device, can be connected to the WorldFIP interface of CERN instrumentation, and add wireless capabilities. The wireless repeater extension was implemented on a SDR made of the flash-based PolarFire FPGA and the AD9361 RF front-end.

The second part discussed the radiation characterization of the repeater extension. The CRaTeBo platform was introduced. It consists of a radiation tolerant carrier card, aiming at reducing the development time for a radiation test in those facilities where the beam size is wider than the DUT size, which complicates the use of development kits. The platform was used to carry out experiments on the repeater extension in two facilities.

Experimental results show that the repeater extension is subject to SEFI, i.e. link losses, which will cause the instruments unavailability in the control network. Nonetheless, the results retrieved considering the HL-LHC fluence are promising. For what concern TID effects, the extension can survive for the whole lifetime of a system during the HL-LHC run.

There are two main outcomes in this chapter. First, a seamless implementation of a wireless link within a fieldbus protocol can be challenging, with a very high chance of not keeping the original characteristics of the fieldbus itself. This is mainly due to the strict timing requirements of the protocol, which does not foresee a wireless channel, and thus the challenges associated to it. Nonetheless, an implementation is still feasible, if performance losses are allowed by the application. Secondly, the radiation test shows the feasibility of implementing a reliable wireless link in the LHC control network. Clearly, the link losses pose a challenge for the control network.

Starting from these results, the next chapter will cover the implementation of a different wireless algorithm, not constrained anymore by the WorldFIP protocol, and will propose mitigation technique for the link losses.

5 Design and development of a wireless physical layer for control applications

Integrating wireless technologies in a wired fieldbus system has proven to be challenging for two main reasons: a) the impairments of the wireless channels, and the strict requirement of the fieldbus protocol, which is not made to operate on a wireless medium. The repeater extension discussed in Chapter 4 could be useful to cope with the short-term cabling issues, but a different solution is necessary for future particle accelerators, such as the FCC.

Multiple efforts have been made to fill the gap of wireless link in industrial settings, especially because of the transition to Industry 4.0, which envisions entire plants connected and manged on the cloud. The transformation required new developments to cope with the wireless channel, since the needs are not addressed by wireless technology used in consumer electronics. Therefore, wireless fieldbus protocols have emerged. Differently from a fieldbus examined in the previous chapter, they embed the necessary mechanisms to operate reliably on a wireless channel. Nonetheless, they are mainly targeting applications with medium-to-high requirements in terms of latency, such as monitoring applications. Control applications, on the other hand, have much more strict requirements, but a commercial solution is not available yet.

However, recent research has shown that by developing a custom wireless physical layer, and building the protocol on top of it, one can reach performance comparable to that of a wired fieldbus. However, radiation effects were not considered, but they are a main concern for CERN control network. The last objective pursed during this doctoral work aims to bridge the gap between the two domains, building a custom radiation-tolerant wireless PHY for control applications. The main idea is to set the foundation for a radiation-tolerant wireless link to be used within the tunnel for control applications, which is currently non-existent.

The first part of the chapter presents an overview of the current wireless solutions available for industrial plants, their limitations under radiations and for the control systems. An overview of the trend in developing custom wireless PHYs is presented, emphasizing the benefits of this approach.

The second part of the chapter discusses the custom wireless PHY developed during this doctoral work. The hardware choices are presented, together with the implementation of the DSPs algorithm for wireless transmission/reception. The proposed implementation reaches a latency in the order of microseconds, which fits in the requirements of the LHC control network.

Afterwards, the radiation characterization of the proposed wireless-PHY is presented, and its reliability under radiation is discussed in the context of the LHC control network. Moreover, a preliminary performance test of the link has been performed in the LHC tunnel, showing the capability of the link in terms of achievable range and bit error rate.

5.1 State-of-the-art wireless control networks in radiation environments

As mentioned in the previous chapter, the challenges posed by wireless channels made it difficult to directly adapt a wired fieldbus protocol for wireless communication without modifications. However, the adoption of wireless technologies in industrial environments had gained significant interest due to the numerous benefits they offered, and therefore more and more industries were attracted by their potential benefits. However, employing standards used for wireless personal area network (WPAN) was not an option because they lacked support for industrial requirements such as determinism and low latency. To address these challenges and enable the use of wireless networks in control industries, many efforts have been done by different organizations, which led to the realization of different solutions still used nowadays.

A significant milestone in this effort was the ratification of WirelessHART [125] in 2007, the first open standard specifically designed for control and measurement applications. WirelessHART was developed to address challenges, such as strict timing requirements and determinism, that wireless standards like ZigBee and Bluetooth could not meet in industrial control environments at the time. WirelessHART builds on the IEEE 802.15.4-2006 physical layer, while defining its own MAC layer. The physical layer specifies the modulation technique, which is Direct Sequence Spread Spectrum (DSSS), the operating frequency of 2.4 GHz in the ISM band, and a data rate of 250 kbps. On top of this, the MAC layer is based on Time Division Multiple Access (TDMA). WirelessHART presents two main features. First is the use of a 10 ms time slot, during which a node can transmit its data. This ensures collision-free communication and enables deterministic behaviour, which is essential for a fieldbus. Second is its support for mesh communication, allowing each WirelessHART device to forward packets on behalf of others. This enables packets to reach their destination through multiple paths, enhancing network reliability in the event of outages caused by issues on the wireless channels.

Another example of a wireless fieldbus is ISA100.11a, released by the International Society of Automation (ISA). This standard was designed to support various applications, including noncritical monitoring, alerting, supervisory control, open-loop control, and closed-loop control. Its physical layer is based on the IEEE 802.15.4 standard. It operates in the 2.4 GHz ISM band and employs two modulation schemes: DSSS and QPSK, achieving a maximum data rate of 250 kbps. Even in this case, the protocol uses only the PHY layer of the IEEE 802.15.4 standard, defining a custom MAC layer on top of it to account for the unreliability of the wireless channel. As for WirelessHART, it supports TDMA for avoiding collision, with time slots with a constant duration, configurable between 10 and 12 ms during network setup. The collection of time slots repeats cyclically. The standard supports various network topologies, including star, mesh, and hybrid configurations, offering flexibility for different industrial use cases.

WirelessHART and ISA100.11a are currently the leading standards in the wireless control industry. They can meet the millisecond-level requirements of control applications, since a key feature of both standards is their use of the TDMA access mode, with a time slot of 10 ms allocated for each node to transmit its data. While this approach is comparable to the mechanisms used by traditional fieldbuses, it also impacts the sampling time for each device.
As discussed in Chapter 2, a control network typically consists of a master node that communicates with multiple slave nodes, either sending commands or requesting data. A

critical performance metric is the network's cycle time—the duration required for the master node to interact with all devices in the network, i.e. to set or get a parameter.
For traditional fieldbuses like WorldFIP, the cycle time is relatively short, typically in the microsecond range. However, as it scales linearly with the number of devices, the use of a fixed 10 ms time slot for wireless fieldbuses such as WirelessHART and ISA100.11a results in significantly longer cycle times.

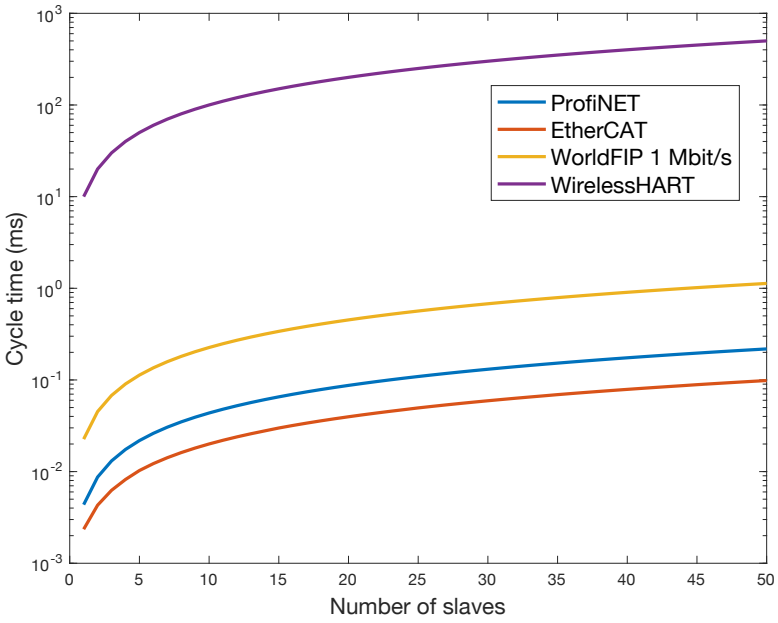

Figure 84 - Cycle time for a multi-slave network of standard fieldbus compared to WirelessHART

The cycle time for networks with varying numbers of slave devices can be compared across different fieldbuses, as illustrated in Figure 84. The comparison highlights that the cycle time of WirelessHART, is two orders of magnitude higher compared to that of standard wired fieldbus like WorldFIP, because of the time slot duration. Even though, this performance could be acceptable for some applications, they are still far from those of standard fieldbus therefore remaining the preferred choice for applications with stricter requirement.

WirelessHART and ISA100.11a rely on the wireless PHY layer defined by the IEEE 802.15.4 standard, which specifies the modulation techniques and data rates for communication. However, the PHY layer plays a critical role in the design and performance of the upper MAC layer. For example, the time slot duration in TDMA must account for the PHY capabilities, i.e. ensuring enough time for transmitting a payload, handling potential retransmissions, and so on. For instance, if data rates higher than 250 kbit/s, hence customizing the PHY layer, the slot duration could be reduced. As such, there has been growing interest in developing a custom wireless PHY to meet the stringent requirements of industrial control systems. Recent research has explored this bottom-up approach, aiming to create a wireless PHY that aligns with the performance of traditional wired fieldbuses, such as Profibus or EtherCAT. The goal is to offer a robust wireless alternative for critical applications—like motor control and power converters—that demand microsecond-level latencies.

To achieve this, researchers proposed a custom wireless physical layer called Wireless High-Performance (WirelessHP) [126]. This design minimizes unnecessary overheads to keep the

latency at the minimum. Unlike the single-carrier modulation techniques used in standards like IEEE 802.15.4, WirelessHP employs OFDM, a multi-carrier technique already utilized in IEEE 802.11 (WiFi), which enables very high data-rate and robust communication. The key difference lies in the optimization possible by targeting only control applications. For example, the IEEE 802.11 standard foresees a preamble with information, like modulation type, encoding length, and packet size and more, to adapt to varying and dynamic urban environments. In contrast, WirelessHP targets control networks, where communication is cyclical, payload sizes are predefined, and the environment remains relatively static once configured. These characteristics eliminate the need for headers, control information, allowing decoding to begin immediately after packet detection, since the receiver already knows all the information.

After theoretical analysis, researchers implemented WirelessHP on an SDR testbed for preliminary evaluation. The results demonstrated WirelessHP can achieve transmission latencies in the microsecond range, making it a promising candidate for high-performance industrial control applications.

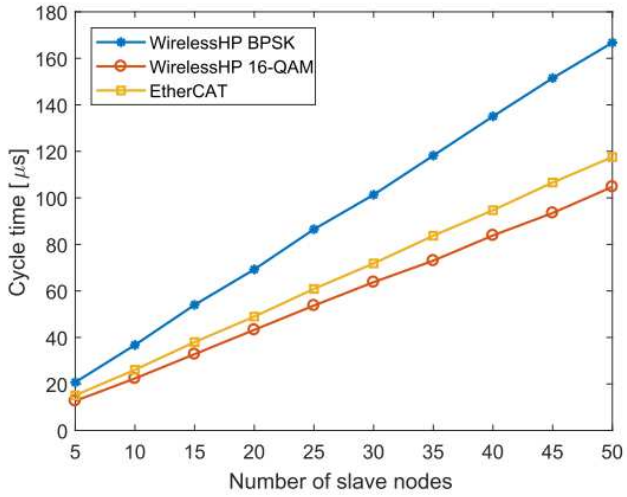

Figure 85 - WirelessHP performance compared to EtherCAT [126]

Figure 85 compares the cycle time for a multimode network using EtherCAT, a wired fieldbus, with that of WirelessHP employing BPSK and 16-QAM modulation. As shown, the cycle time of WirelessHP remains in the microsecond range, even for a network with fifty nodes, demonstrating performance comparable to that of EtherCAT.

This comparison highlights the feasibility and the benefits of developing a custom physical layer tailored to the requirements of control applications. Unlike reusing the PHY layer of an existing wireless standard, which often includes overhead for such scenarios, a custom PHY can achieve performance comparable with wired fieldbuses. Once such a physical layer is established, it can serve as the foundation for the upper layers of the protocol stack.

The remainder of this doctoral work builds on this principle, focusing on the development of a custom physical layer that aligns with the requirements of control applications. Additionally, a critical objective is the radiation tolerance to enable reliable operation within the environment of CERN accelerators.

5.2 Design and evaluation of a custom wireless PHY for LHC control applications

The design of the wireless PHY started with selecting the hardware, followed by implementing the wireless DSP algorithms to achieve microsecond latency for packet transmission and reception, meeting the requirements of the LHC control application described in Chapter 2.

5.2.1 Hardware platform

The chosen architecture for the implementation is the SDR, for the benefits that have already been explained throughout the manuscripts: flexibility, possibility to achieve more radiation tolerance, and custom implementation. The last characteristics is fundamental, given the custom design necessary for this work. The two main components of the base platform are the AD9361 and the PolarFire FPGA. The AD9361 was selected for its robustness against radiation effects, as demonstrated in its characterization discussed in Chapter 3. The PolarFire FPGA, on the other hand, was chosen for the reasons outlined in the previous chapter, particularly its robustness to TID, which supports the development of highly radiation-tolerant systems for the FCC, where TID effects will be even more pronounced.

As described earlier, configuring the AD9361 requires executing C code, necessitating the use of a processor. For the wireless physical layer design, a soft-CPU was chosen despite the availability of the PolarFire SoC, which includes a hard CPU. This decision was driven by the flexibility offered by a soft-CPU, which can be customized for specific applications, unlike a fixed-architecture hard CPU. Additionally, while the hard CPU can rely only on software-based radiation mitigation techniques, a soft-CPU enables more mitigation strategies like TMR. The Mi-V was selected as the starting point for the wireless PHY, with plans to explore other options in the future.

In addition to the core components that are the same as the repeater extension, additional peripherals were included in the hardware platform of the custom wireless PHY. As noted during the development of the wireless repeater, non-volatile memory is required to store executable code. For this purpose, different storage options were available, but a magneto-resistive random-access memory (MRAM) was investigated for this purpose due to its high resilience to TID and SEE compared to silicon-based memories. The memory chosen is the EM008LXQ from Everspin Technologies, based on their industrial STT MRAM technology. This memory is SPI-compatible, with an interface running up to 200 MHz, ensuring very fast loading times. This feature helps minimize downtime in case of a link crash, the main source of failure observed during the characterization of the PolarFire FPGA. Furthermore, MRAM is known for its strong resistance to TID and immunity to SEE, making it a suitable choice for this application. While MRAM may be sensitive to magnetic fields, this is not an issue in the tunnel where the link will be deployed.

Finally, an RF platform also requires transmitter and receiver amplifiers to extend the communication range. To meet this need, two solutions were investigated: the ADL5521 and the TRF37C73EVM. The ADL5521 is a low-noise amplifier from Analog Devices, based on GaAs pHEMT technology. It operates within a frequency range of 400 MHz to 4 GHz, which does not cover the entire range of the AD9361 (4 GHz to 6 GHz). However, it was chosen as a starting point since the current evaluation focuses on the 2.4 GHz ISM band.

The TRF37C73EVM, on the other hand, is a power amplifier from Texas Instruments, utilizing silicon-germanium (SiGe) technology. This device covers instead the entire frequency range of the RF Front-end, operating from 1 MHz to 6 GHz.

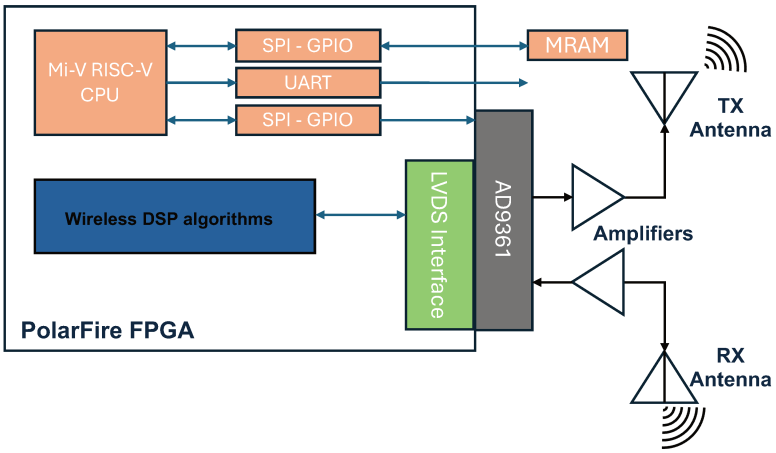

Figure 86 – HW platform used for the implementation of the custom physical layer

Figure 86 shows the block diagram of the hardware platform.

5.2.2 Gateware implementation

The gateware interfaces with the AD9361, implements the complete DSP chain required for wireless communication, which is developed to fit CERN's control system requirements. It also provides an interface to a host computer for accessing diagnostic information and connects to the external MRAM.

The gateware is divided into two independent units: the configuration unit and the communication unit. The configuration unit handles platform setup. It operates on a 50 MHz clock generated by a PLL available in the PolarFire FPGA, with the input provided by the PolarFire internal RC oscillator. This unit includes the Mi-V soft CPU, 256 KB of storage implemented using the PolarFire internal memory blocks, a UART peripheral, a JTAG interface for debugging, and SPI and GPIO peripherals for interfacing with the AD9361 and MRAM.

The soft CPU executes the commands needed to configure the AD9361, which include a sequence of SPI instructions for calibration and setting parameters like frequency, bandwidth, and gain. The MRAM interface is necessary for loading the application into the internal memory, both at startup and in case of a SEFI, as observed during the characterization of the repeater extension.

Independently of the configuration unit, the communication unit handles all communication-related tasks. This unit uses a separate clock provided by the AD9361 after its configuration. It includes an LVDS interface with the AD9361 for sending and receiving IQ samples, and the transmission and reception DSP algorithms.

The transmission blocks handle all functions required for data transmission, including data generation. In a real application, data would come from an external source, such as a sensor, and be processed by the transmission block before being sent to the AD9361. For this design,

a data generation unit was included to prepare the data for transmission. The packet structure is kept as simple as possible: a preamble, followed by a fixed payload, and ending with a CRC sequence. Since the modulation, packet structure, and data rate are kept fixed, the receiver will know how to decode the packet after detection, and there is no need for additional information to be included in the packet. This ensures less time in the reception.
The preamble consists of a Barker sequence repeated twice, essential for packet detection at the receiver. This is followed by a payload with a fixed size of 80 bytes (640 bits). This size was chosen considering that the current payload used in the WorldFIP frame of the RadMON agent is in the same order. Clearly, the final size will depend on the requirement of the applications which have to be defined. A 32-bit CRC sequence, calculated from the payload, is appended to the packet. Once generated, the packet samples are processed one at a time by a scrambler unit to balance consecutive zeros and ones, which will simplify the clock recovery process at the receiver.
Next, the data is modulated using QPSK modulation, chosen as a starting point for its balance between spectral efficiency and noise robustness. Previous characterization of PPM modulation for the repeater extension highlighted its drawbacks, including high sensitivity to signal amplitude. QPSK, in contrast, encodes information in the phase of the carrier, making it more robust. However, other modulation schemes should be explored. For instance, BPSK might provide better tolerance to radiation effects and a more reliable link, but with a reduced data rate. Conversely, QAM could lower latency but at the cost of reduced link reliability.
After modulation, the data is filtered through a root-raised cosine filter to minimize intersymbol interference and perform pulse shaping. The filter has a roll-off factor of 0.5, resulting in a bandwidth 1.5 times the symbol rate, and operates with four samples per symbol, achieving a sample rate four times the symbol rate. Finally, each sample is sent to the RF front-end through the LVDS interface.

On the receiver side, the inverse processing is implemented, along with additional DSP algorithms necessary to compensate for timing and frequency offsets, as the transmitter and receiver operate on separate devices. Once the signal is digitized at the receiver by the AD9361, it is sent to the FPGA for processing.
The first step involves a gain control block to stabilize the signal amplitude, which is crucial for the subsequent algorithms. The samples are then passed through an RRC filter, which, combined with the transmitter's RRC filter, performs matched filtering. Following filtering, the samples are processed by a series of recovery algorithms to address timing and frequency offsets. Specifically, the models described in Chapter 3 were implemented.
The recovery chain starts with a loop-based filter to compensate for the timing offset between the transmitter and receiver, ensuring proper sampling even if their clocks are not synchronized. Right after, frequency offset compensation is performed to address differences between the local oscillators of the transmitter and receiver front ends, which may not operate at exactly the same frequency. This process is divided into coarse and fine steps, where the coarse step corrects most of the offset, and the fine step eliminates any residual offset.
After offset compensation, the signal proceeds to the packet detection algorithm. This algorithm exploits the autocorrelation properties of the Barker sequence used in the preamble, as described in Chapter 3. It employs a filter with coefficients that are the inverse of the Barker sequence. The filter's output peaks when the sequence passes through, generating a control signal that triggers the downstream processing.
The demodulator then performs QPSK demodulation to recover the raw bits, which are sent through a descrambler to restore the original data stream. The payload length and the packet structure is encoded in the receiver, therefore it can decode the expected number of bits for the

payload and CRC. These bits are passed to a CRC block to check for errors. The final outputs of this block include the recovered data stream, a control signal indicating the start and stop of the packet, and an error signal that asserts if any errors are detected from the CRC check
This output interface serves as the wireless PHY's connection to the rest of the application. However, in this case the work was focused on the physical layer only and therefore a full protocol stack and the full application were not implemented. For evaluation purposes, the current platform includes additional logic to count bit errors in the received packets and track the number of correctly and incorrectly received packets.

Regarding the performance of the proposed algorithm chain, the default symbol rate is set to 7.68 MHz, resulting in an occupied bandwidth of 11.52 MHz. The sample rate is four times the symbol rate, because the filter operates on four samples per symbol, giving a 30.72 Msps sample rate. This is the sample rate of the communication unit.
With QPSK modulation, which encodes 2 bits per symbol, the final bit rate is 15.36 Mbit/s.
The execution latency of the chain was measured in simulation as 280 µs, which accounts for the time taken from the data from the transmitter FPGA to the point where the receiver FPGA processes the data and detects the packet. This performance meets all the timing requirements outlined in Chapter 2, ensuring suitability for the various applications under consideration.

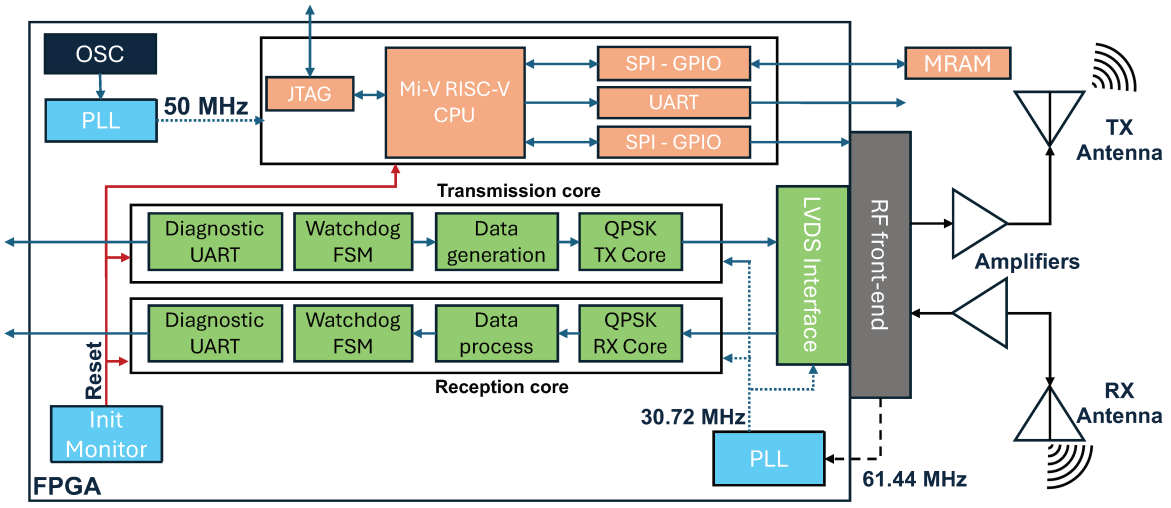

Figure 87 - Hardware and gateware implementation of the wireless physical layer

The full architecture of the wireless PHY, including the gateware described above are represented in Figure 87.

5.2.3 Software description

The application code for the proposed wireless PHY is structured into two primary components: a first-stage bootloader application and a main application.

First stage bootloader

It handles the loading of the main application from the external MRAM into the PolarFire internal memory. Its code is stored in the PolarFire non-volatile memory. The working principle is the following:

1. Upon boot, the bootloader executes a series of read operations to retrieve the application code from the MRAM.
2. The retrieved code is copied into the PolarFire memory blocks, mapped to the Mi-V soft processor.

3. After the transfer is complete, the final instruction of the bootloader application modifies the processor's instruction pointer to the start of the newly copied main application.

At this stage, the Mi-V starts the execution of the main application.

Main application

It performs the configuration of the AD9361 and controls the SPI communication needed for its initialization. Once it starts, the application computes the configuration parameters for the AD9361 and sends a series of SPI commands to apply these settings.
After the configuration is complete, the software enters an infinite idle loop. In fact, in its current state, the application is very simple; after the AD9361 configuration, the soft-CPU remains unused. However, it could be used for a more significant role when a complete stack is implemented. For example, it could implement higher-level functionalities, as those required by the MAC layer, or manage additional tasks related to network coordination or diagnostic.

5.2.4 Platform working principle

The working principle of the wireless PHY can be categorized into three phases: initialization, configuration, and operational.

During the initialization phase, the PolarFire performs its boot process, where essential components such as I/Os, voltage banks, and memory blocks are brought to an operational state. Once this is complete, the system controller copies the first-stage bootloader application code from the non-volatile memory into the FPGA's internal memory associated with the soft-CPU. After this, the system controller releases the device reset signal, ensuring all FPGA fabric components are properly initialized. The soft-CPU then executes the first-stage bootloader, transferring the main application code from the external MRAM into the internal memory. This phase concludes once the memory transfer is completed.

In the configuration phase, the soft-CPU runs the main application code, which calculates the necessary settings for the AD9361 RF transceiver and issues the required SPI commands to configure it. Once the AD9361 is fully configured, it generates a clock signal for the PLL driving the communication unit, activating this component and bringing the system to the operational phase, allowing the system to transmit and receive data.

5.2.5 Performance characterization of the wireless PHY

The proposed wireless PHY has been implemented on hardware, with the transmitter and receiver deployed on two separate SDR platforms comprising a PolarFire FPGA and an AD9361 RF transceiver. An initial evaluation was performed by streaming over 1×10^7 bits of continuous data to compute the system BER. In this setup, the transmitter continuously generated packets, while the receiver reported the total received bits and bit errors to a host computer via a UART interface. The tests were conducted at an RF frequency of 2.4 GHz, with the RF front-end operating in automatic gain mode.

Measurements were carried out in different laboratory conditions. In one configuration, the transmitter and receiver were connected via SMA-to-SMA cables, with the transmitter power set to -10 dBm to avoid saturating the receiver. In another configuration, the transmitter and receiver were connected to omnidirectional antennas. Two tests were performed with the devices placed at distances of 1 meter and 5 meters in the two different tests. Across all these measurements, a consistent BER of 10^{-3} was observed, demonstrating that external interference

was negligible. Tests with the AD9361 configured for fixed gain control instead of automatic gain control also showed similar performance.

Additional things could be done in this matter to improve the BER performance. First, evaluate the possibility to add to the system channel equalization mechanism, which aims to reduce amplitude, frequency and phase distortion in a radio channel. Secondly, introduce FEC mechanism, which aims to correct bit errors at the receiver by using additional bits added at the transmitter. Even though the current implementation shows a BER that could be too high for some applications, it represents a good starting point to be used for the evaluation of its reliability under radiation, and to understand its performance inside the tunnel.

In fact, the second step of the evaluation was performed in the LHC tunnel, specifically in the LHC Point 1, to assess the system performance in a real-world deployment scenario. The transmitter was fixed at a designated location, while the receiver was positioned at various distances for approximately two seconds of data acquisition per position. The data was processed offline to determine BER, PER and EVM.

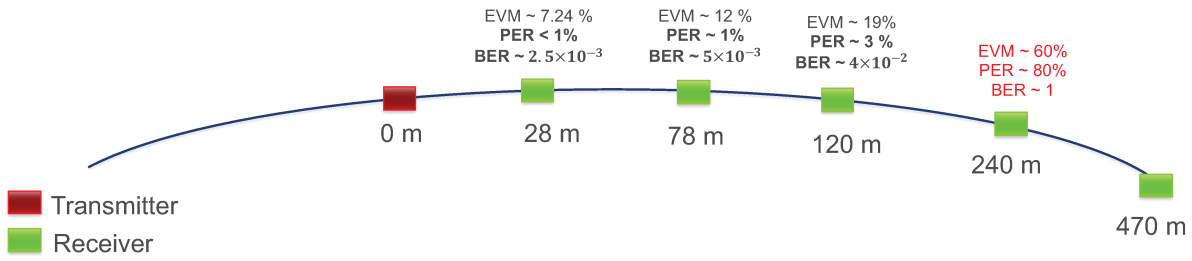

Figure 88 - Link performance measured in the LHC point 1

Figure 88 illustrates the measurements at varying distances within the tunnel. At 30 meters, the BER was 2.5×10^{-3} and doubled at around 80 meters. At 120 meters, the BER increased to approximately 10^{-2}, which correlated with a rise in EVM from 7% at shorter distances to 20% at 120 meters. The PER remained below 1% up to 80 meters but rose to 3% at 120 meters. At 240 meters, the signal's EVM was too high for effective measurement, though some packets could still be decoded. Beyond this, at 470 meters, no packets were captured.

These results reveal that the system's PER remains low (<1%) at shorter distances despite a higher BER. Analysis of the raw data indicates that most bit errors stemmed from receiver synchronization loss. This desynchronization caused error spikes within a few consecutive packets, causing increase of the BER, while packets received post-recovery were decoded correctly. Consequently, enhancing the receiver's synchronization algorithm could significantly reduce BER before implementing more complex FEC mechanisms.

5.3 Radiation characterization of the wireless physical layer

Following the implementation and evaluation of the wireless PHY, the next critical step was its characterization under radiation. The primary objective of this phase was to assess the reliability of the communication link in radiation environments. Specifically, the aim was to determine whether radiation effects have a measurable impact on the BER of the system and, by extension, to evaluate how such effects might influence the performance of the control network when deployed in its intended environment.

5.3.1 Test methodology

The proposed wireless PHY was characterized at PSI using the 200 MeV proton beam. Each component of the system was irradiated individually prior to their integration. This approach allowed evaluating the components in different operating modes, including those not foreseen in operation. With this approach, distinct failure mechanisms could be identified and analyzed, facilitating the implementation of mitigation techniques to enhance the system's reliability.

5.3.1.1 MRAM

The DUT was the EM008LXQ from Everspin Technologies. The test consisted in irradiating nine units mounted on a single DUT board. All memory units were irradiated simultaneously. For testing, an FPGA-based tester implemented on the Zynq-7000 SoC was employed. This SoC features both programmable logic and an ARM Cortex-A9 processor. The tester connected to the memory units via a uCoax cable, linking the FPGA to all nine memories.

An SPI master controller was implemented on the FPGA to perform read/write operations and issue commands to the DUT. The test sequence was executed on the ARM processor using Python within the PYNQ framework developed by AMD.
The memories were tested in various operating modes to investigate potential failure mechanisms. Initially, a fixed pattern was written into each memory before irradiation. For one memory unit, continuous read operations were performed during irradiation, with the original pattern rewritten in case of detected errors, without halting irradiation. For the other eight memory units, data was read only after irradiation, and errors were corrected by rewriting the pattern before resuming irradiation. Additionally, three of these eight memories were placed in deep-sleep mode before irradiation to assess potential failures in this state. After irradiation, these memories were reactivated, read, and checked for errors. Before each new irradiation session, a different pattern was written to all memories to verify the functionality of the write operation.

5.3.1.2 RF Amplifiers

For the RF amplifiers, being purely analog devices, the focus of the investigation was on TID and DD effects. The primary objective was to evaluate the impact of cumulative effects on the gain of the devices over the entire frequency range of operation. Measurements were conducted using the E5071C vector network analyzer, with the devices connected via RF cables. Due to the long distance between the instrument and the DUT, it was challenging to measure the device response in real-time during irradiation. As a result, the characteristics of the devices were measured prior to irradiation, and the devices were kept powered throughout the irradiation process. After irradiation, their responses were measured again to assess the impact of radiation on their performance.

5.3.1.3 PolarFire gateware

The gateware of the wireless PHY was characterized using the Everest evaluation kit, an evaluation board for the PolarFire MPF300 FPGA. The RF front-end used in the setup was the FMCOMMS2, featuring the AD9361 RF transceiver. During testing, the transmitter was configured for continuous data transmission, while the receiver was implemented on a second SDR setup. This second SDR also used an FMCOMMS2, with its gateware implemented on the Xilinx-7000 SoC mounted on the ZedBoard. The receiver was not subjected to irradiation and reported the number of correctly received packets and the system's BER through an Ethernet connection. Transmitter and receiver were interconnected with cables to isolate radiation-induced effects.

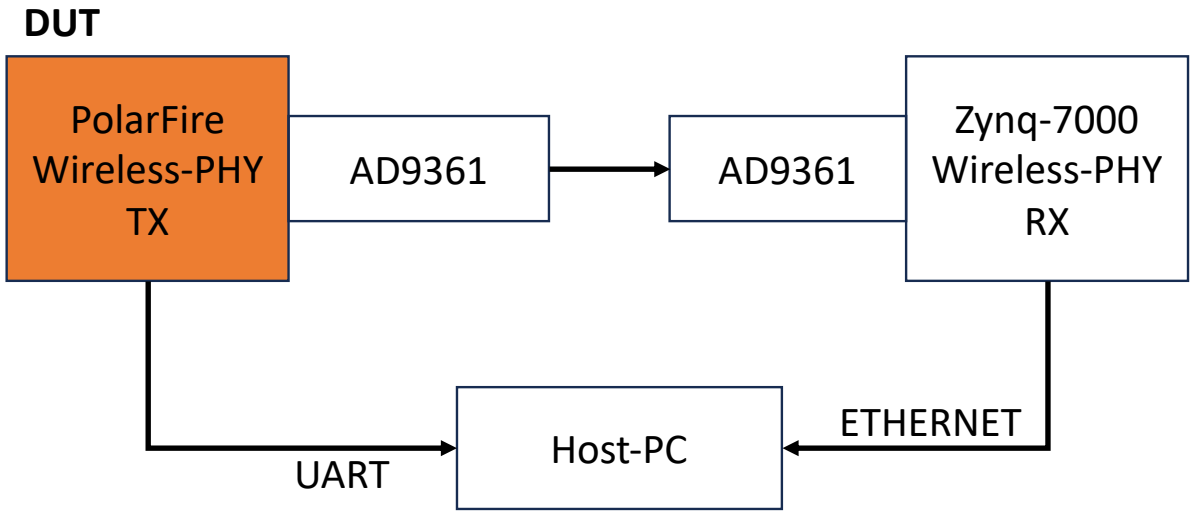

Figure 89 - Test setup used for the characterization of the Wireless PHY gateware

The complete test setup is shown in Figure 89.

To evaluate the impact on link reliability, different versions of the gateware with different features were tested. These features are:

- TMR mitigation technique, implemented at register level using the functionality of Microchip libero tool.
- Suspended system controller
- A modified power-on-reset (POR) scheme.

Regarding the system controller, it powers up the FPGA fabric and I/Os during initialization. It remains in reset until required voltage levels are confirmed by dedicated sensors. Once these thresholds are met, the system controller releases the reset, enabling the initialization of PLLs, memory, fabric, and I/Os. However, as stated in the PolarFire datasheet [127], the system controller should be disabled during operation to prevent unintended behaviour. This feature was already valuated during the characterization of the repeater gateware, but its impact on the link reliability was not noticeable. An additional change instead was done in this characterization, that consists in implementing a different POR scheme. The standard POR uses the PolarFire initialization monitor, which provides flags when particular event like the threshold voltages of the banks are correct, or when the device is fully operation. In particular, this last flag is used to bring the fabric out of reset and is used as reset signal for the configuration and communication unit of the FPGA gateware. In addition to the standard scheme, a different POR was implemented where the initialization monitor is not used. The PLL lock signal of the configuration and communication unit are used instead as reset for their corresponding unit.

Table 13 - Characteristics of the designs used for the characterization of the Wireless PHY

Design name	TMR	System Controller	Init Monitor
PHY-1	NO	Enabled	YES
PHY-2	YES	Enabled	YES
PHY-3	YES	Suspended	NO
PHY-4	NO	Suspended	NO
PHY-5	YES	Enabled	NO

Table 13 summarizes the characteristics of the different gateware versions that were tested.

5.3.2 Test results

This section presents the experimental results from the radiation testing of the memories, amplifiers, and the gateware.

5.3.2.1 MRAM

Concerning the memories, they exhibited distinct failure modes categorized as follows:

- E1: When configured in deep-sleep mode, the memory failed to exit this mode using the corresponding command after irradiation. A power cycle was necessary to restore functionality.
- E2: Errors detected during the reading process, but the reading return the correct values after a power cycle, without rewriting the memory.
- E3: Errors were detected during the reading process, but the correct value was retrieved after issuing a software reset.
- E4: Errors detected after a first reading operation, but performing a second read returned the correct pattern.
- E5: Errors detected even after three consecutive read operations. The memory was rewritten, after which the read operation returned no errors.

In particular, failure category E4 suggests that memory locations were not corrupted; instead, the errors may have originated in the reading circuitry, as the correct content was retrieved without rewriting. For category E5, the correct pattern was retrieved only after rewriting the memory. Nonetheless, a power cycle was not performed, therefore either the memory content was really corrupted, or it represents the same error type as E2. Despite these events, no bitflip was observed in any of the nine samples under test.

Table 14 – MRAM failure modes and cross-section for a cumulated fluence of 1.29×10^{-12} *p/cm²*

Failure type	No. of irradiated samples	Error occurrences	$\sigma\ (cm^2/device)$		
			Lower limit	σ	Upper limit
E1	3	1	5.22×10^{-15}	2.58×10^{-13}	4.80×10^{-13}
E2	9	3	5.16×10^{-14}	2.58×10^{-13}	7.56×10^{-13}
E3		1	1.74×10^{-15}	8.61×10^{-14}	4.80×10^{-13}
E4		1	1.74×10^{-15}	8.61×10^{-14}	4.80×10^{-13}
E5		6	1.86×10^{-13}	5.17×10^{-13}	1.13×10^{-12}

The cross section of all the events is reported in Table 14.

Following this phase, the memories were powered off and subjected to an additional 2 kGy dose in steps of 500 Gy. After each step, the memories were powered on, read, and rewritten, and then powered off. None of this operation reported any errors at the different steps.

5.3.2.2 RF amplifiers

Concerning the RF amplifiers, they exhibited high resilience to radiation, with no noticeable degradation in their performance before and after irradiation. The devices accumulated a total dose of 1 kGy during the tests.

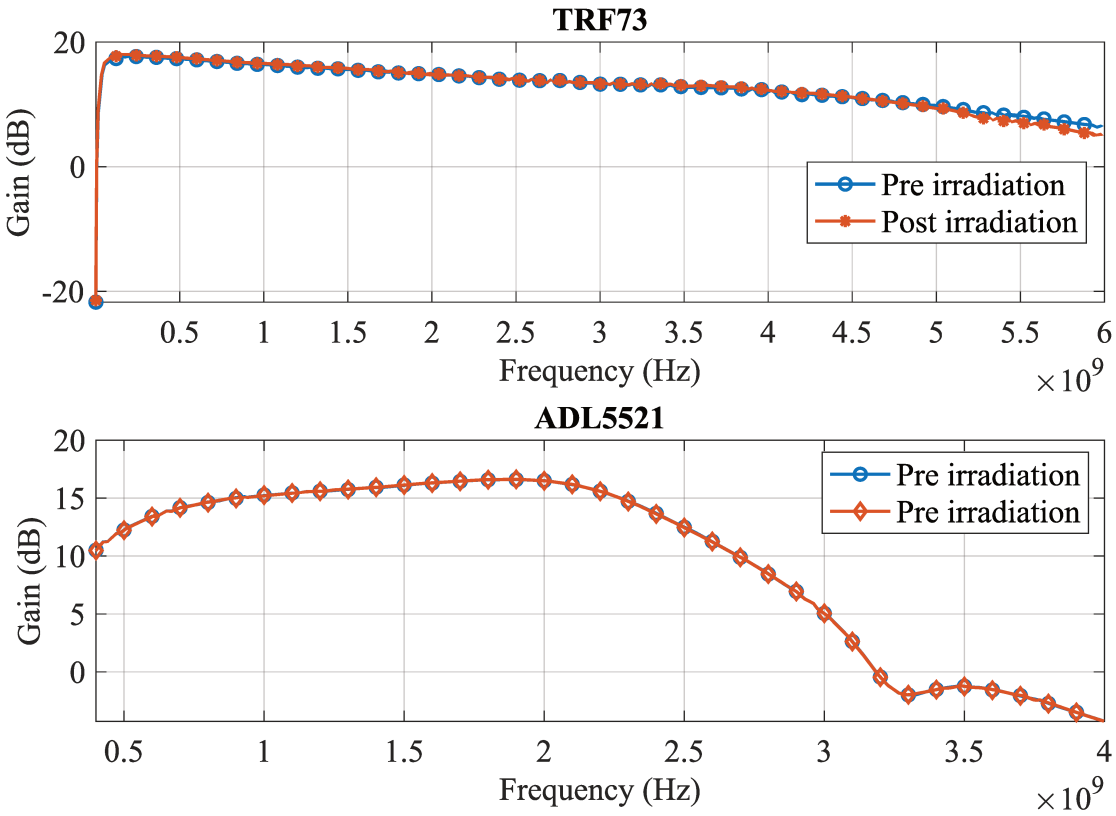

Figure 90 - Gain measurements before and after irardiation for the TRF37C73EVM (top) and ADL5521 (bottom).

Figure 90 illustrates the measurements of the gain across the entire operating frequency range of the devices. The ADL5521 did not report any degradation over the entire frequency range. The TRF73 instead reported a slightly degradation between 5 GHz and 6 GHz.

5.3.2.3 Gateware

The impact of radiation effects on the gateware can be summarized in terms of performance and reliability of the link.

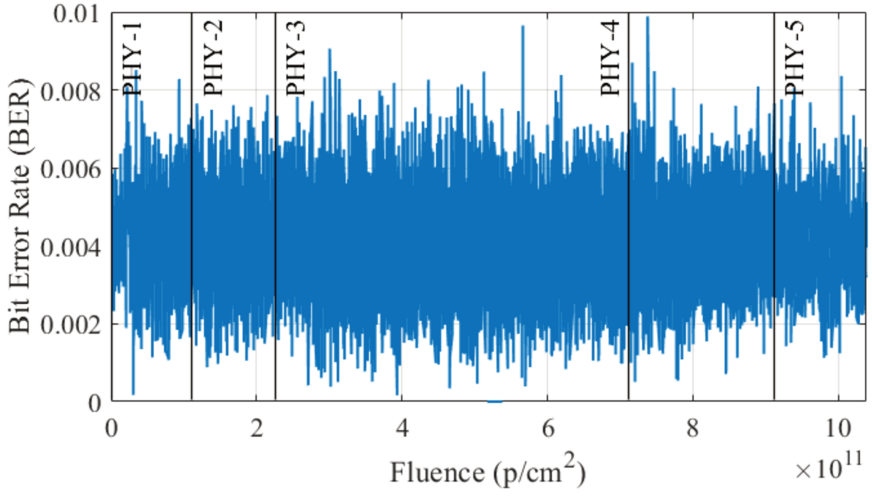

Figure 91 - BER evolution with the fluence of all the PHY-x designs.

For what concern the performance, Figure 91 illustrates the BER as a function of the fluence accumulated across all the implementations listed in Table 13. The measurement considers only the link in the operational phase. The BER is displayed in this way because the bit corruption can occur only in the baseband unit that is identical to all implementations. As it is

visible from the figure, the BER was stable during the whole irradiation, around the nominal performance of 10^{-3}. No spikes or significant change have been observed, up to a total fluence of 1.3×10^{11} p/cm^2. An increase of BER would indicate that either a SET or a SEU occurred in the baseband unit of the wireless PHY, resulting in the transmission of wrong data. However, none of such events were detected. A possible explanation is that a SET occurred, but it did not lead to a bit error, as it was already shown in Chapter 3 during the EVM evaluation. Another possibility is that even though it occurred, the bit error was not detectable because of the high baseline BER. In this case, if the baseline BER of the system is reduced, it is more likely to observe the impact of SETs.

Table 15 – SEFI cross-sections of the Wireless-PHY implementations tested with a proton flux of 3.80×10^8 p/cm^2/s.

Design name	Fluence [p/cm^2]	SEFI	Lower limit	σ [cm^2/device]	Upper limit
PHY-1	8.74×10^{10}	4	1.21×10^{-11}	4.58×10^{-11}	1.17×10^{-10}
PHY-2	1.29×10^{11}	2	1.79×10^{-12}	1.56×10^{-11}	5.60×10^{-11}
PHY-3	5×10^{11}	-	-	-	7.38×10^{-12}
PHY-4	2.39×10^{11}	5	6.34×10^{-12}	2.09×10^{-11}	4.89×10^{-11}
PHY-5	1.32×10^{11}	3	4.54×10^{-12}	2.28×10^{-11}	6.65×10^{-11}

For what concern the reliability aspect, the DUT were affected SEFI, as already expected considered the outcomes of the experiment reported in Chapter 4. The SEFI causes the loss of the link, but it has no impact on the BER. Table 15 reports the cross-section for this event for the different implementations of the DUTs. As it can be seen, the designs exhibit different sensitivities, showing the impact of the initialization monitor and system controller on the link reliability. Starting from PHY-1, it exhibited the highest cross-section, primarily due to failures associated with the initialization monitor and system controller. PHY-2 instead showed a cross-section approximately four times lower than PHY-1, indicating that the reset source might originate from the initialization monitor's reset logic. PHY-3, which employs TMR, disables the initialization monitor, and suspends the system controller, demonstrated the lowest cross-section with no errors detected. PHY-4, without the initialization monitor and system controller but with TMR disabled, exhibited a higher cross-section than PHY-3 but lower than PHY-1, likely due to TMR. Finally, PHY-5 reported a higher cross-section, likely due to the system controller not being in suspended mode. These results highlight that the system controller and initialization monitor have a great impact on the system reliability, thus their usage should be considered carefully.

For what concern TID effects, no degradation was observed as expected, since only the FPGA was exposed to radiation, and it was already known from previous test the robustness of the FPGA. Nonetheless, the different PHYs were tested on two separate development kits cumulating 330 Gy and 525 Gy respectively, without any performance degradation.

5.3.3 Results analysis and discussion

This section discusses the experimental results, focusing on the link reliability in the radiation environment of the accelerator. Mitigation techniques are proposed to improve it, along with an analysis of the impact on the control network and its lifetime within the accelerator.

5.3.3.1 Link reliability

The reliability of the link is crucial, as any failure would render instruments within the LHC control network unusable. The experimental results showed that even with an alternative POR scheme, the FPGA remains susceptible to SEFIs, which consist in a radiation-induced reset of

the fabric. Although the user manual suggests that disabling the system controller should prevent unintended resets, the experimental results indicate that this is not always the case. Moreover, the system controller is not the only source of reset, as a different power-on reset scheme, implemented without the initialization monitor, also affects the cross-section. Because of this, it may be beneficial to evaluate additional power-on reset schemes that do not use PolarFire resources, instead employing external circuitry, to determine if this can reduce the cross-section. In such a case, the reset scheme would need to be designed to ensure that all PolarFire resources exit the reset in the correct order, guaranteeing proper functionality.

However, despite the SEFI causing a link loss, the bootloader serves as a mitigation technique for this event. When a SEFI occurs, the application is reloaded from the external MRAM, and the processor restarts from the configuration phase. During this phase, the AD9361 is reconfigured, and the system returns to the operational phase. Therefore, with this implementation, the SEFI results in only a temporary downtime of the link. Based on the results from the MRAM characterization, it was observed that data can be corrupted during the reading process. However, the data can be successfully retrieved after a reset, power cycle, or additional read attempt. These methods can be employed to handle such failures and proceed to the configuration phase. No other failures on the MRAM that could affect the reliability of the link were observed. The same goes for the RF amplifiers, since the TRF37C73EVM exhibited a only small gain degradation after 5 GHz, which could affect the performance but not the link reliability.

Another important aspect to consider is the link loss caused by the RF front end. As shown in Chapter 3, SEUs on the AD9361 can lead to misconfiguration, which requires re-initialization. If this occurs, the result will be a link loss. To mitigate this type of failure, a watchdog circuitry that periodically checks functionality is necessary. A simple implementation would involve monitoring the loss of lock of the PLL in the communication unit. If a SEU occurs on the AD9361, causing a loss of configuration, the AD9361 will no longer provide a clock signal to the PLL, which will trigger the loss of lock signal. At this point, the processor could initiate a device reconfiguration, bringing the link back to the operational phase.

The main failure mode observed was link loss, which is mitigated either by the first-stage bootloader or by the watchdog circuitry. Despite various techniques, such as disabling the system controller or removing the initialization monitor, link loss is inevitable, but its impact has been reduced to a temporary downtime of the link. In the current implementation, this downtime lasts around 4 seconds, from the occurrence of the SEFI to when the link returns to the operational phase. Most of this time is spent during the first-stage bootloader process, particularly when loading the main application from the external memory. Currently, SPI transactions are performed at 10 MHz; however, this can be sped up by implementing the interface at the maximum speed allowed by the MRAM, which is 200 MHz.

5.3.3.2 System performance

Regarding the performance of the link, the BER showed no degradation with the FPGA gateware exposed to radiation. In fact, the performance measured during irradiation aligns with the performance observed during the system's lab evaluation. While the BER may seem high for some applications, it is important to note that it represents only a baseline, which can be further improved. Moreover, as discussed in the previous section, the link's performance could be enhanced by adapting the recovery algorithm or implementing FEC codes.

5.3.3.3 System lifetime

Regarding the lifetime of the system, all components demonstrated good robustness to TID. Starting with the gateware test, the PolarFire FPGA was exposed to a total of 525 Gy without any degradation. However, in previous experiments, the PolarFire had already proven to operate up to 5 kGy. Concerning the AD9361, the device did not show any degradation up to 500 Gy, as reported in Chapter 3. Additionally, previous experiments have also confirmed its robustness up to several kGy of dose. As for the MRAM, after accumulating 3.250 kGy, it was still functioning correctly. Regarding the RF amplifiers, the ADL5521 did not show any degradation after 1 kGy, while the TRF37C73 exhibited a slight gain degradation above 5 GHz. Therefore, all components demonstrated good resilience against TID. Considering the radiation levels for the HL-LHC, as reported in Chapter 1, the harshest location in the LHC where instruments will be deployed is the DS area. In this location, the electronics will be subject to approximately 100 Gy per year. Given the robustness of the wireless PHY components, this suggests that the solution could operate for more than 10 years without the need for replacement, which is a very promising result.

5.3.3.4 Link availability

The link availability can be computed considering the link loss cross-section. Within the system, a link loss can be caused either by a SEFI on the PolarFire, or by the AD9361. Nonetheless, the SEFI of the PolarFire has the higher cross-section among the two and therefore should be used for the estimation. Considering the fluence levels expected during the HL-LHC phase and the number of instruments in each LHC location, the number of link losses can be estimated by using the upper limit cross-section observed for PHY-3.

Table 16 – Annual failure rate and link availability in the different LHC areas considering the radiation levels expected during the HL-LHC phase and if all the nodes use the wireless link.

LHC area	Annual fluence [p/cm^2]	Communication nodes	Annual failure rate [Upper limit]	Annual link losses [Upper limit]	Annual Link availability
DS	5×10^{10}	~ 600	0.37	222	> 99%
ARC	1×10^{9}	~ 9000	0.0074	67	> 99%
UJ	5×10^{9}	~ 400	0.04	16	> 99%
RR	3×10^{9}	~ 2000	0.02	40	> 99%

Table 16 reports the number of failures expected, per year, if all the nodes make use of the wireless link. As visible, the DS is the area with the highest probability of observing a failure, and therefore up to 222 link losses could occur per year. On the other hand, the ARC is where most of the instruments are installed, but only a maximum of 67 loss of links are expected. Moreover, each link loss, in the current implementation, corresponds to 4 seconds of downtime, which can be used to compute the availability of the link during the LHC run. In fact, considering a complete LHC run going from March to October, the availability of the link would still be greater than 99% in all the areas of the LHC. These numbers demonstrate that the solution is a promising candidate for radiation environment of the LHC.

5.4 Chapter summary

This chapter discussed the development and the characterization of a custom wireless physical layer whose characteristics could fit the requirements of the current LHC control network.
A review of existing solutions for wireless industrial control environment was presented, focusing on two wireless fieldbuses designed to address the challenges of the wireless channel, which were highlighted in the previous chapter. On the other hand, existing wireless fieldbuses, do not fit those control applications whose requirements are very strict, in the order of microseconds. In addition to available standards, recent research has showed the benefits of developing a physical layer tailored to the application requirements, reducing unnecessary communication overhead and simplifying the design to meet the timing constraints.

Following the same principle, the second part of the chapter presented the development of a custom physical layer tailored to the LHC timing requirements. The PHY is based on the SDR approach, which as discussed throughout this manuscript, provides flexibility and the possibility to achieve more radiation tolerance, both critical for systems operating in the LHC. The hardware and gateware implementations necessary for the wireless communication were presented. The PHY employs single carrier QPSK modulation, achieving a latency of roughly 280 us, fitting the requirements of the LHC, a data rate of 15 Mbit/s, and an occupied bandwidth of 11 MHz. Performance tests conducted in the LHC tunnel demonstrated the link capability of communicating over distances of up to 80 meters, achieving a packet error rate of less than 1%, and a BER of 10^{-3}, serving as the baseline for the radiation characterization.

The third section of the chapter focused on the radiation characterization of the wireless PHY. All hardware components were tested under radiation. The gateware tests showed no performance degradation during irradiation, with the system's BER remaining consistent with nominal performance. Regarding the hardware, the primary radiation-induced effect was the SEFI observed in the PolarFire FPGA, also discussed in Chapter 4. As a mitigation technique, a bootloader mechanism was implemented, making link losses into temporary downtimes. Considering the HEH fluences expected during the HL-LHC and assuming all instruments rely on wireless PHY in all the different areas of the tunnel, the proposed link reaches an availability greater than 99% with respect to the total run time of the LHC for one year. Furthermore, all the components demonstrated excellent resilience to TID, qualifying the system for more than 10 years of operation in the HL-LHC environment.

This chapter demonstrates that the proposed wireless PHY represents a promising solution for introducing wireless network into the LHC control system and the LHC radiation environment. It represented the first exploratory work, which sets the foundations for future research paths, as will be outlined in the next and final chapter.

Conclusions

This doctoral research aimed to advance the state of the art in wireless technologies for the control networks of particle accelerators at CERN.

The need for wireless communications arises from the complexity of the accelerator system and the high number of instruments required for its reliable operation. Currently, these instruments are operated through a control network requiring thousands of kilometres of cables installed in the tunnels. However, the infrastructure has reached its capacity, creating a bottleneck for new installations. It also imposes significant costs on the infrastructure and limits the flexibility of upgrades.

The challenge is even greater when considering the new accelerator proposed by CERN, the Future Circular Collider, which will be around 100 kilometres long. A cabled infrastructure for its control network would be impractical, given the huge number of systems and thus cable required. In contrast, wireless technologies offer greater flexibility, reduce infrastructure costs, and simplify the installation and relocation of instrumentation.

Despite the advancements in commercial wireless solutions, their adoption in the accelerator presents two primary challenges. First, there is a lack of commercial solutions that meet the stringent requirements of control networks, included CERN control network. Second, radiation effects pose a significant concern compared to other sector, as they can degrade system performance or lead to unintended behaviour. For a control network, this is critical, as the unavailability of instruments could impact the reliable operation of the accelerator.

This research evaluated the feasibility of implementing wireless technologies in the accelerator control network, focusing on three key aspects: the flexibility of the solution, compliance with CERN's control application requirements, and radiation tolerance.

The need for flexibility led to the adoption of the SDR approach to develop a wireless platform capable of hosting a wide range of wireless applications just by reconfiguring the programmable logic, i.e., the FPGA. In this design, most of the wireless link functionalities are implemented within the FPGA, while the RF front-end handles the RF operations and analog/digital conversion.

This approach is particularly advantageous in an environment like the LHC, where thousands of instruments are deployed. Future modifications to the system would otherwise be impossible, being too much resource intensive. With SDR, the same hardware can be easily reconfigured to support different protocols, offering a high degree of adaptability.

Implementing the SDR approach requires a highly configurable RF front-end capable of adjusting operating frequency, bandwidth, and sampling rate according to the specific application. Furthermore, the front-end must be robust against radiation effects to ensure reliable operation. Previous studies in the literature identified the AD9361 as a promising candidate for use in radiation environments. Based on these findings, it was selected for further investigation in this doctoral work.

The experimental results available at the start of this research provided limited insights into the impact of radiation effects when the device is used in practical applications. Therefore, the first step was to develop a methodology to characterize the radiation effects on RF transceivers, using the AD9361 as a candidate example. The device was characterized under proton irradiation following the proposed methodology. The results demonstrated the effectiveness of this approach in investigating radiation effects on an RF front-end. Furthermore, the findings revealed the AD9361's robustness against both TID and SEE and therefore was used throughout this research to study the feasibility of adopting wireless solutions in CERN control network.

The first step in this context consisted in evaluating whether the existing infrastructure, the WorldFIP network, could be extended using wireless technologies. The reasoning behind this approach was to explore the benefits of wireless technologies while minimizing changes to the existing system. Replacing the communication interfaces of the thousands of already deployed instruments with wireless links would be resource intensive. Instead, the proposed solution involved a plug-and-play repeater extension that could connect to existing instruments without requiring hardware/software modifications.

The proposed repeater, designed with an SDR approach, could interface with existing CERN hardware transparently to the protocol. To reach this goal, a custom wireless physical layer was developed to manage data transfer within the timing constraints imposed by the WorldFIP protocol. The repeater extension was characterized under radiation at various facilities, including tests using CRaTeBo, a radiation-tolerant tester board developed specifically for FPGA-based systems.

The radiation characterization of the repeater extension revealed different errors on the communication due to SEE, and it highlighted the susceptibility of the PolarFire to SEFI causing the loss of the link. Outside radiation environments, the repeater extension was also validated using real CERN hardware, specifically the RadMON. The evaluation showed the feasibility of the solution, but the repeater exhibited poor performance. In fact, the wireless communication algorithms were intentionally kept simple to meet the timing constraints imposed by the WorldFIP protocol and maintain transparency. However, this simplicity compromised performance, highlighting the challenges of extending a wired fieldbus protocol over a wireless channel.

While this approach offers short-term benefits for addressing cabling issues, its adoption is unlikely for long-term projects like the FCC. For such scenarios, wireless solutions should be integrated into the control network from the beginning, as the high volume of cabling required would be impractical. However, despite the significant interest in integrating wireless technologies into industrial control scenarios, efforts to standardize wireless fieldbuses have not yielded solutions capable of meeting the requirements of critical applications, mainly because wireless fieldbuses adopt the physical layers of general-purpose wireless standards. Recent research has demonstrated the advantages of developing custom physical layers tailored

to the specific requirements of control networks, because the overhead can be minimized compared to traditional physical layers built for general-purpose applications. The results in the literature indicate that performance comparable to traditional wired fieldbuses can be achieved.

Building on these findings, the final phase of this doctoral work focused on implementing a custom physical layer designed to meet the requirements of the LHC control network while being radiation-tolerant for the accelerator environment. The proposed physical layer is based on an SDR approach, ensuring flexibility. It satisfies the control application requirements of the LHC, achieving a latency of approximately 300 µs for packet transmission and reception, and it was fully characterized under radiation. Considering the results of the repeater extension characterization, where SEFI emerged as the most critical failure mode, a mitigation technique was implemented, consisting of a bootloader that automatically restarts the link in case of failure. The other hardware peripherals of the physical layer include an MRAM for application code storage and RF amplifiers to enhance communication range. All components were qualified under radiation, and the implemented mitigation techniques ensure that link loss results in only temporary downtime. The proposed solution achieves an availability of over 99%, even under the radiation levels of the HL-LHC, the next accelerator upgrade. All the hardware components demonstrated high robustness against TID, with the PolarFire up to 5 kGy, the AD9361 tested up to 500 Gy, the MRAM tested up to 3 kGy, and the RF amplifiers up to 1 kGy. This will be very important for the FCC, where TID will be a significant concern. The performance of the physical layer was also evaluated inside the LHC tunnel. Communication with a packet error rate below 1% was achieved over distances of up to 80 meters, despite a bit error rate of 10^{-3}. This performance could be further improved by enhancing the receiver algorithms or FEC codes.

The work carried out lead to a first solution that addresses all the challenges stated at the beginning for implementing a wireless link for control network in CERN accelerators: latency, flexibility, and radiation tolerance. However, it represents the first exploratory study, which opened multiple research pathways that will be further investigated.

One of the primary areas for future exploration involves evaluating different modulation techniques to understand their impact on communication performance. While this work focused on QPSK, other digital modulation schemes, such as BPSK and QAM, should be implemented and characterized under radiation conditions. BPSK, for instance, may offer greater reliability at the cost of data rate and latency, while QAM could provide higher speeds but might be more susceptible to radiation-induced errors.

Additionally, this doctoral work primarily explored single-carrier communication techniques. However, multi-carrier transmission methods, such as OFDM, have become prevalent in modern wireless links. OFDM employs Fast Fourier Transform and Inverse Fast Fourier Transform algorithms which may have varying sensitivity to radiation effects. While multi-carrier systems could provide lower latencies compared to single-carrier techniques, as demonstrated in the WirelessHP project, their reliability under radiation must be carefully evaluated.

Another critical step is the clear definition of requirements for future control networks. The current work was based on the timing needs of existing instrumentation in the LHC control network. However, future accelerators may incorporate new instruments with different requirements. Defining these requirements will guide the development of wireless algorithms

tailored to meet specific needs. Although the proposed physical layer may not fully satisfy the requirements of future particle accelerators, it serves as a solid starting point for improvement. Thanks to its SDR-based hardware platform, the same hardware can be reused to implement and evaluate different algorithms without significant modifications. Furthermore, the proposed components, with their robustness to TID, make this hardware a promising candidate for the FCC, where dose will be the main concern.

Moreover, the work carried out in this thesis is related to the physical layer only. Its implementation is just the first step toward establishing a wireless link for the control network. The next objective is the definition of the MAC layer, which will define access to the wireless medium and introduce mechanisms to enhance reliability and determinism. Different aspects should be considered in its definition:

- Reliability: Critical messages could be retransmitted until successfully delivered, while non-critical messages might be dropped if lost. These strategies will impact the overall reliability of the network;
- Determinism: Deterministic communication is a fundamental characteristic of fieldbuses. Mechanisms like TDMA can be implemented to allocate time slots for devices, ensuring an upper bound on message delivery time;
- Frequency management: Multiple frequencies could be assigned to different devices to prevent interference and allow simultaneous data transmission, thus reducing cycle time of the network.

Clearly, all these possibilities should be also evaluated considering the impact of radiation effects and will also play an important role in the definition of the network topology. A robust network topology will be essential to support the highly distributed systems of the LHC and future accelerators. A simple point-to-point topology is impractical given the large scale and spatial distribution of instruments. Different approaches could be explored:

- Mesh networks: they allow each node to act as a relay, forwarding messages to their destinations. This approach enables coverage of larger distances and increases network reliability, as messages can take alternate paths in case of node failures due to radiation effects
- Subnetwork Architecture: The control network could consist of multiple subnetworks distributed throughout the tunnel. Each subnetwork, managed by a local master device, would employ the wireless PHY for low-latency communication. The subnetworks could then be interconnected via a high-speed backbone network, such as Ethernet or optical fiber.

This doctoral research marks an important milestone in evaluating wireless technologies for CERN's control networks, providing a foundation for future advancements. The exploration of a radiation-tolerant wireless physical layer tailored to the stringent requirements of the LHC environment highlights the potential for replacing traditional cabling with innovative, high-performance solutions.

List of publications

The work carried out at CERN during these years of doctoral studies led to the publication of different articles presented at international conferences and published in international journals:

1. A. Scialdone et al., "Exploring Radiation-Induced Vulnerabilities in RFICs Through Traditional RF Metrics," in IEEE Transactions on Nuclear Science, vol. 70, no. 8, pp. 2068-2075, Aug. 2023, doi: 10.1109/TNS.2023.3270206.
2. A. Scialdone et al., "CRaTeBo: a high-speed, radiation-tolerant and versatile testing platform for FPGA radiation qualification for high-energy particle accelerator applications." Journal of Instrumentation 19(01), 2024: P01012.
3. A. Scialdone et al., “Feasibility Study of a Radiation-Tolerant Wireless Extension for the CERN WorldFIP infrastructure”, *2024 24rd European Conference on Radiation and Its Effects on Components and Systems (RADECS)*, Maspalomas, Spain, 2024, to appear
4. A. Scialdone et al., "Design and Characterization of a Radiation Tolerant Wireless Physical Layer for Control Applications in CERN Particle Accelerators," in IEEE Transactions on Nuclear Science, doi: 10.1109/TNS.2025.3537402.

List of Figures

Figure 1 - Particle interactions [1] 21
Figure 2 Radiation-induced charge generation in a MOS [3] 23
Figure 3 - Charge generation and collection in a reverse-biased junction [10] 27
Figure 4 - Ion strike on a memory cell 27
Figure 5 Overview of the CERN accelerator complex 29
Figure 6 - LHC Layout 30
Figure 7 - LHC Octant 30
Figure 8 - Cumulative luminosity delivered to ATLAS over days of stable beam operation for high-energy proton-proton collision. 31
Figure 9 - LHC beam dumps induced by radiation effects as a function of the cumulative integrated luminosity [19] 32
Figure 10 - Block diagram of the CERN RHA 33
Figure 11 - TID measured by the BLM close to IP1 and IP5 during the LHC run in 2016 [27] 35
Figure 12 - Proton Irradiation Facility in Villigen, Switzerland 37
Figure 13 – Neutron Irradiation Facility at TRIUMF 38
Figure 14 – Co60 irradiation facility at CERN 39
Figure 15 – CERN Mixed field facility layout 40
Figure 16 – Radiation levels in the different test position in the CHARM facility 40
Figure 17 - Architecture of a discrete control system 43
Figure 18 - Networked Control System (NCS) 43
Figure 19 - OSI stack and ICN stack 44
Figure 20 - TDMA example 46
Figure 21 - Pyramidal scheme of a factory, showing the different levels of the hierarchy and how they are interconnected 47
Figure 22 - Timeline of network evolution taken from [33] 48
Figure 23 – Overview on the CERN control system architecture 49

Figure 24 - Front-End Computer and some electronics module installed [37].....50
Figure 25 - Basic building blocks of a digital communication system.....56
Figure 26 – Digital modulation of amplitude, phase and frequency of a carrier waveform used to convey information.....58
Figure 27 – Constellation diagram of one symbol corresponding to a signal whose amplitude is the distance from the origin, and whose phase is the angle from the horizontal axis.....59
Figure 28 - M-ary PAM constellation diagram.....61
Figure 29 - QAM constellation diagram.....61
Figure 30(a) - QPSK constellation diagram.....62
Figure 31 - Building blocks for the receiver synchronization.....63
Figure 32 - Timing error in communication system.....64
Figure 33 - PLL structure for a recovery algorithm.....65
Figure 34 - Uncorrected QPSK signal constellation overlaid with ideal constellation in red.66
Figure 35 - example of different outcomes for error detection algorithms, ZC and different parameters.....66
Figure 36 – 16-QAM signal constellation at the receiver with a 400 Hz offset compared to transmitter.....67
Figure 37 – Incoming signal at a wireless receiver. A frame containing valid data and marked within red lines is mixed with noise.....68
Figure 38 - Autocorrelation of a Berker sequence of length L = 11.....69
Figure 39 - Sequence of bits including the Berker sequence at the beginning. The dotted red line presents the output of a FIR filter that can be used to detect the position of the Berker sequence.....69
Figure 40 – Radio receiver based on the superheterodyne architecture.....70
Figure 41 – Architecture of a Zero-IF radio receiver.....71
Figure 42 –Direct RF radio receiver.....72
Figure 43 - SDR architecture.....72
Figure 44 - Transmitter and Receiver structure for a RF front-end based on the Zero-IF architecture.....73
Figure 45 – Visual representation of EVM using constellation diagram. The reference symbol represents the ideal location when there are no impairments.....76
Figure 46 - Source of errors in a RFIC.....77
Figure 47 - AD9361 architecture [85].....78
Figure 48 - Architecture of the SDR implementation.....80

Figure 49 - QPSK application 81
Figure 50 - Test setup for the AD9361 irradiation test 82
Figure 51 - DUT in beam 82
Figure 52 - EVM of QPSK system during irradiation 84
Figure 53 - EVM before and after reconfiguration 84
Figure 54 – EVM output 85
Figure 55 - EVM evolution for LTE frames 86
Figure 56 - CFO over frequency 87
Figure 57 - EVM of the packet affected 89
Figure 58 - CFO estimation 90
Figure 59 - Manchester encoding used in the WorldFIP PHY to encode binary digits 99
Figure 60 - Structure of a WorldFIP frame 99
Figure 61 - Macrocycle performed by a bus arbiter 100
Figure 62 - Timing diagram for a WorldFIP network with a Bus Arbiter and two agents 101
Figure 63 - WorldFIP PHY implemented on CERN WorldFIP devices 102
Figure 64 - WorldIFIP network using wireless repeaters. The same device can be connected either to the bus arbitrator, to any agents, or to the bus interconnecting multiple agents, since its totally transparent to the WorldFIP protocol MAC layer 104
Figure 65 - Top level architecture of a CERN WorldFIP device and its corresponding SDR-based wireless repeater 104
Figure 66 - Top level architecture of the Wireless PHY 105
Figure 67 – Illustration of how the repeater logic interconnects the Wireless and WorldFIP domain 106
Figure 68 - Timing diagram for a WorldFIP transaction using repeater extension 107
Figure 69 - Architecture of the SDR-based repeater extension used for the system validation 108
Figure 70 – Overview of the setup used for the validation of the solution. Bus Arbiter and WorldFIP agent are connected to their corresponding repeater, which is SDR-based, made of ZedBoard and FMCOMMS3 109
Figure 71 – Loss rate of the variable updates taken with a measurement of 16.5 hours 109
Figure 72 - Distribution of the variable updates period on logarithm scale. 110
Figure 73 - WorldFIP network with the gateway approach 111
Figure 74 - CHARM Radiation Tolerant Tester Board (CRaTeBo) block diagram 114
Figure 75 - CRaTeBo power supply scheme 115

Figure 76 - PolarFire System-on-Module for CRaTeBo 115
Figure 77 - Repeater extension block diagram implemented on the PolarFire FPGA and AD9361 117
Figure 78 – a) SDR under test, with CRaTeBo mounting the PolarFire MPF300T SOM and the FMCOMMS3 with the AD9361. b) The setup, mounted on a plate is, lowered in the beam area 118
Figure 79 - Irradiation test setup of the SDR platform at TNF facility 119
Figure 80 - FPGA implementation during the radiation test 119
Figure 81 - Current consumption of the Transmitter SDR during irradiation 120
Figure 82 - Radiation test setup at PSI 122
Figure 83 - Proton cross-sections for the different error signatures observed during the irradiation 124
Figure 84 - Cycle time for a multi-slave network of standard fieldbus compared to WirelessHART 130
Figure 85 - WirelessHP performance compared to EtherCAT [126] 131
Figure 86 – HW platform used for the implementation of the custom physical layer 133
Figure 87 - Hardware and gateware implementation of the wireless physical layer 135
Figure 88 - Link performance measured in the LHC point 1 137
Figure 89 - Test setup used for the characterization of the Wireless PHY gateware 139
Figure 90 - Gain measurements before and after irardiation for the TRF37C73EVM (top) and ADL5521 (bottom). 141
Figure 91 - BER evolution with the fluence of all the PHY-x designs 141

List of Tables

Table 1 - e-h pair generation rate 23
Table 2 – Annual HEH fluence for the shielded areas measured during the LHC run of 2016 by the RadMON system [27] 35
Table 3 – Expected annual (250 fb-1) radiation levels for the HL-LHC and average risk factor for the different areas of the tunnel and shielded areas. 36
Table 4 - Characteristics of a communication network for ICN 46
Table 5 - Requirements summary for fieldbus instrumentation 53
Table 6 – Cross-sections of the SEE observed on the QPSK application 85
Table 7 –SEE cross-sections for the three categories of events observed 88
Table 8 – Example of Bus Arbitrator table containing variables identifier with their corresponding periodicity 100
Table 9 – Cross-sections for the communication loss due to SEFI of the AD9361 and the PolarFire FPGA. 121
Table 10 - Naming convention used for the four implementations of the design under test. 123
Table 11 - Events and fluence reached for the four designs under test using a proton flux of 3.80 × 108 p/cm2/s 123
Table 12 - Link availability across the LHC considering the radiation levels of the HL-LHC. The calculation assumes that all the nodes make use of the repeater extension. 126
Table 13 - Characteristics of the designs used for the characterization of the Wireless PHY 139
Table 14 – MRAM failure modes and cross-section for a cumulated fluence of 1.29 × 10-12 p/cm^2 140
Table 15 – SEFI cross-sections of the Wireless-PHY implementations tested with a proton flux of 3.80 × 10^8 p/cm^2/s. 142
Table 16 – Annual failure rate and link availability in the different LHC areas considering the radiation levels expected during the HL-LHC phase and if all the nodes use the wireless link. 144

Bibliography

[1] R. Ferraro, "Development of Test Methods for the Qualification of Electronic Components and Systems Adapted to High Energy Accelerator Radiation Environments," École doctorale Information, Structures et Système (I2S), Montpellier, 2019.

[2] J. R. Schwank, M. R. Shaneyfelt, D. M. Fleetwood, J. A. Felix, P. E. Dodd, P. Paillet and V. Ferlet-Cavrois, "Radiation Effects in MOS Oxides," *IEEE Transactions on Nuclear Science,* vol. 55, no. 4, pp. 1833-1853, 2008.

[3] T. R. Oldham, Ionizing radiation effects in MOS oxides, vol. 3, World Scientific, 1999.

[4] C. Leroy and P.-G. Rancoita, Principles of Radiation Interaction in Matter and Detection, World Scientific, 2016.

[5] S. Wood, N. J. Doyle, J. A. Spitznagel, W. J. Choyke, R. M. More, J. N. McGruer and R. B. Irwin, "Simulation of Radiation Damage in Solids," *IEEE Transactions on Nuclear Science,* vol. 28, no. 6, pp. 4107-4112, 1981.

[6] F. Sexton, "Destructive single-event effects in semiconductor devices and ICs," *IEEE Transactions on Nuclear Science,* vol. 50, no. 3, pp. 603-621, 2003.

[7] P. Dodd and L. Massengill, "Basic mechanisms and modeling of single-event upset in digital microelectronics," *IEEE Transactions on Nuclear Science,* vol. 50, no. 3, pp. 583-602, 2003.

[8] J. Brews, M. Allenspach, R. Schrimpf, K. Galloway, J. Titus and C. Wheatley, "A conceptual model of a single-event gate-rupture in power MOSFETs," *IEEE Transactions on Nuclear Science,* vol. 40, no. 6, pp. 1959-1966, 1993.

[9] R. Baumann, "Radiation-induced soft errors in advanced semiconductor technologies," *IEEE Transactions on Device and Materials Reliability,* vol. 5, no. 3, pp. 305-316, 2005.

[10] A. Zimmaro, "Radiation Effects and System-Level Testing: Application to Wireless Communications," University of Montpellier, 2024.

[11] L. Sterpone, M. Sonza Reorda, M. Violante, F. L. Kastensmidt and L. Carro, "Evaluating Different Solutions to Design Fault Tolerant Systems with SRAM-based FPGAs," *Journal of Electronic Testing,* vol. 23, p. 47–54, 2007.

[12] H. Quinn, Z. Baker, T. Fairbanks, J. L. Tripp and G. Duran, "Software Resilience and the Effectiveness of Software Mitigation in Microcontrollers," *IEEE Transactions on Nuclear Science,* vol. 62, no. 6, pp. 2532-2538, 2015.

[13] M. Berg, C. Poivey, D. Petrick, D. Espinosa, A. Lesea, K. A. LaBel, M. Friendlich, H. Kim and A. Phan, "Effectiveness of Internal Versus External SEU Scrubbing

Mitigation Strategies in a Xilinx FPGA: Design, Test, and Analysis," *IEEE Transactions on Nuclear Science,* vol. 55, no. 4, pp. 2259-2266, 2008.

[14] R. L. Shuler, A. Balasubramanian, B. Narasimham, B. L. Bhuva, P. M. O' Neill and C. Kouba, "The Effectiveness of TAG or Guard-Gates in SET Suppression Using Delay and Dual-Rail Configurations at 0.35 um," *IEEE Transactions on Nuclear Science,* vol. 53, no. 6, pp. 3428-3431, 2006.

[15] S. Rezgui, R. Won, J. Tien and J. George, "SET characterization & mitigation in 65-nm test structures,," in *12th European Conference on Radiation and Its Effects on Components and Systems*, Seville, Spain, 2011.

[16] L. Evans and P. Bryant, "LHC Machine," *JINST,* vol. 3, p. S08001, 2008.

[17] CERN, "ATLAS EXPERIMENT — PUBLIC RESULTS," 2024. [Online]. Available: https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun3. [Accessed 23 May 2024].

[18] A. Apollonio, O. Rey Orozko, R. Schmidt, M. Valette, D. Wollmann and M. Zerlauth, "Lessons Learnt from the 2016 LHC Run and Prospects for HL-LHC Availability," in *8th International Particle Accelerator Conference*, Copenhagen, Denmark, 2017.

[19] Y. Q. Aguiar, G. Lerner, R. G. Alía, D. Prelipcean, A. Apollonio, F. Cerutti, S. Danzeca and M. S. Gilarte, "Radiation to Electronics Impact on CERN LHC Operation: Run 2 Overview and HL-LHC Outlook," in *12th Int. Particle Acc. Conf. IPAC2021*, Campinas, SP, Brazil, 2021.

[20] L. Coull, D. Hagedorn, V. Remondino and F. Rodriguez-Mateos, "LHC magnet quench protection system," *IEEE Transactions on Magnetics,* vol. 30, no. 4, pp. 1742-1745, 1994.

[21] CERN, "High-Luminosity Large Hadron Collider (HL-LHC): Technical design report," CERN, Geneva, 2020.

[22] T. Böhlen, F. Cerutti, M. Chin, A. Fassò, A. Ferrari, P. Ortega, A. Mairani, P. Sala, G. Smirnov and V. Vlachoudis, "The FLUKA Code: Developments and Challenges for High Energy and Medical Applications," *Nuclear Data Sheets,* vol. 120, pp. 211-214, 2014.

[23] B. Dehning, E. Effinger, J. Emery, G. Ferioli, G. Guaglio, E. B. Holzer, D. Kramer, L. Ponce, V. Prieto, M. Stockner and C. Zamantzas, "The LHC Beam Loss Measurement System," in *22nd Particle Accelerator Conference* , Albuquerque, NM, USA, 2007.

[24] G. Spiezia, P. Peronnard, A. Masi, M. Brugger, M. Brucoli, S. Danzeca, R. G. Alia, R. Losito, J. Mekki, P. Oser, R. Gaillard and L. Dusseau, "A New RadMon Version for the LHC and its Injection Lines," *IEEE Transactions on Nuclear Science,* vol. 61, no. 6, pp. 3424-3431, 2014.

[25] M. Cecchetto, R. García Alía, F. Wrobel, M. Tali, O. Stein, G. Lerner, K. Bilko, L. Esposito, C. Bahamonde Castro, Y. Kadi, S. Danzeca, M. Brucoli, C. Cazzaniga, M. Bagatin, S. Gerardin and A. Paccagnella, "Thermal Neutron-Induced SEUs in the LHC Accelerator Environment," *IEEE Transactions on Nuclear Science,* vol. 67, no. 7, pp. 1412-1420, 2020.

[26] D. Kramer, M. Brugger, V. Klupak, C. Pignard, K. Roeed, G. Spiezia, L. Viererbl and T. Wijnands, "LHC RadMon SRAM Detectors Used at Different Voltages to

Determine the Thermal Neutron to High Energy Hadron Fluence Ratio," *IEEE Transactions on Nuclear Science,* vol. 58, no. 3, pp. 1117-1122, 2011.

[27] R. García Alía, M. Brugger, F. Cerutti, S. Danzeca, A. Ferrari, S. Gilardoni, Y. Kadi, M. Kastriotou, A. Lechner, C. Martinella, O. Stein, Y. Thurel, A. Tsinganis and S. Uznanski, "LHC and HL-LHC: Present and Future Radiation Environment in the High-Luminosity Collision Points and RHA Implications," *IEEE Transactions on Nuclear Science,* vol. 65, no. 1, pp. 448-456, 2018.

[28] F. Pozzi, R. G. Alia, M. Brugger, P. Carbonez and S. Danzeca, "CERN IRRADIATION FACILITIES," *Radiation protection dosimetry,* vol. 180, pp. 120-124, 2018.

[29] A. Thornton, "Charm facility test area radiation field description," CERN, Geneva, 2016.

[30] P. Martí, R. Villà, J. Fuertes and G. Fohler, "Networked Control Systems Overview," in *The Industrial Information Technology Handbook*, CRC Press LLC, 2005, pp. 1-16.

[31] A. Ray, "Introduction to networking for integrated control systems," *IEEE Control Systems Magazine,* vol. 9, no. 1, pp. 76-79, 1989.

[32] J. Decotignie and P. Pleinevaux, "A survey on industrial communication networks," *Annals of Telecommunications,* vol. 48, pp. 435-448, 1993.

[33] M. Wollschlaeger, T. Sauter and J. Jasperneite, "The Future of Industrial Communication: Automation Networks in the Era of the Internet of Things and Industry 4.0," *IEEE Industrial Electronics Magazine,* vol. 11, no. 1, pp. 17-27, 2017.

[34] T. Sauter, "The Three Generations of Field-Level Networks—Evolution and Compatibility Issues," *IEEE Transactions on Industrial Electronics,* vol. 57, no. 11, pp. 3585-3595, 2010.

[35] IEC, "Industrial communication networks - Fieldbus specifications - Part 1: Overview and guidance for the IEC 61158 and IEC 61784 series," 2023.

[36] B. Frammery, "The LHC Control System," in *10th ICALEPCS Int. Conf. on Accelerator & Large Expt. Physics Control Systems*, Geneva, 2005.

[37] ATS Department, "Introduction to the BE-CO Control System," CERN, Geneva, 2020.

[38] AFNOR, "FIP bus for exchange of information," 1990.

[39] T. Wijnands and C. Pignard, "Radiation tolerance of the WorldFIP fieldbus interface," CERN, Geneva, 2006.

[40] Afnor EDITIONS, "Réseaux de communications industriels - Spécifications des bus de terrain," 2015.

[41] J. Casas-Cubillos, P. G. F Faccio, M. Martin and M. Rodriguez-Ruiz, "SEU tests performed on the digital communication system for LHC cryogenic instrumentation," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment,,* vol. 485, no. 3, pp. 439-443, 2002.

[42] J. A. Agapito, J. Casas-Cubillos, F. J. Franco, B. Palan and M. A. Rodríguez-Ruiz, "7th European Conference On Radiation And Its Effects On Components And Systems, , 15 - 20 Sep ,," in *RADECS 2003*, Noordwijk, The Netherlands, 2003.

[43] R. Denz and F. Rodríguez-Mateos, "Radiation Tolerance of Components Used in the Protection System of LHC Superconducting Elements," in *RADECS2002*, Padua, Italy, 2002.

[44] BE DEPARTMENT, "WORLDFIP INSOURCING PROJECT," CERN, Geneva, 2009.

[45] G. Penacoba, P. Alvarez, E. Gousiou, J. Palluel, J. Serrano and E. V. d. Bij, "Design of an FPGA-based radiation tolerant agent for WorldFIP fieldbus," *Journal of Instrumentation,* vol. 7, no. 03, p. C03027, 2012.

[46] CERN, "Open Hardware Repository," [Online]. Available: https://ohwr.org/project/cern-fip/-/wikis/home. [Accessed 24 May 2024].

[47] CERN, "MasterFIP OHWR," [Online]. Available: https://ohwr.org/project/masterfip/-/wikis/home. [Accessed 24 May 2024].

[48] R. Brun, D. Glenat and J. Palluel, "LHC WorldFIP Fieldbus Monitoring Project," CERN - AB Department, Geneva, 2004.

[49] R. Denz, "Electronic Systems for the Protection of Superconducting Elements in the LHC," *IEEE Transactions on Applied Superconductivity,* vol. 16, no. 2, pp. 1725-1728, 2006.

[50] F. Bordry and A. Dupaquier, "High Current, Low Voltage Power Converters for LHC Present Development Directions," in *EPAC '96*, Barcelona, Spain, 1996.

[51] Q. King, "Advanced uses of the WorldFIP fieldbus for diverse communications applications within the LHC power converter control system," in *ICALEPCS*, Geneva, 2005.

[52] G. Spiezia, "The LHC Radiation Monitoring System - RadMon," in *10th International Conference on Large Scale Applications and Radiation Hardness of Semiconductor Detectors*, Firenze, Italy, 2011.

[53] T. Wijnands, C. Pignard and R. Tesarek, "An on line radiation monitoring system for the LHC machine and experimental caverns," in *12th Workshop on Electronics For LHC and Future Experiments*, Valencia, Spain, 2006.

[54] P. Lebrun, "Large Cryogenic Helium Refrigeration System for the LHC," in *3rd International Conference on Cryogenics and Refrigeration*, Hangzhou, China, 2003.

[55] P. e. a. Gomes, "The Control System for the Cryogenics in the LHC Tunnel," in *ICMC 2008*, Seoul, Republic of Korea, 2008.

[56] A. J. Burns, "Beam instrumentation," in *10th Workshop on LEP-SPS Performance*, Chamonix, France, 2000.

[57] R. Rausch, "Electronic Components and Systems for the Control of the LHC Machine," in *6th Workshop on Electronics for LHC Experiments*, Krakow, Poland, 2000.

[58] J. G. Proakis and M. Salehi, Digital Communications (5th ed.), McGraw-Hill, 2007.

[59] M. Rice, Digital Communications: A Discrete-Time Approach, Third Edition,, Pearson/Prentice Hall, 2009.

[60] K. Mueller and M. M. , "Timing Recovery in Digital Synchronous Data Receivers," *IEEE Transactions on Communications,* vol. 24, no. 5, pp. 516-531, 1976.

[61] F. Gardner, "A BPSK/QPSK Timing-Error Detector for Sampled Receivers," *IEEE Transactions on Communications,* vol. 34, no. 5, p. 423–429, 1986.

[62] M. Luise and R. Reggiannini, "Carrier Frequency Recovery in All-Digital Modems for Burst-Mode Transmissions," *IEEE Transactions on Communications,* vol. 43, no. 2, p. 1169–1178, 1995.

[63] AMD, "AMD Zynq Ultrascale+ RFSoC," AMD, [Online]. Available: https://www.amd.com/en/products/adaptive-socs-and-fpgas/soc/zynq-ultrascale-plus-rfsoc.html. [Accessed 19 September 2024].

[64] T. H. Lee, The Design of CMOS Radio-Frequency Integrated Circuits, Cambridge, U.K.: Cambridge Univ. Press, 2003.

[65] S. Jagannathan, T. D. Loveless, E. X. Zhang, D. M. Fleetwood, R. D. Schrimpf, T. D. Haeffner, J. S. Kauppila, N. Mahatme, B. L. Bhuva, M. L. Alles, W. T. Holman, A. F. Witulski and L. W. Massengill, "Sensitivity of High-Frequency RF Circuits to Total Ionizing Dose Degradation," *IEEE Transactions on Nuclear Science,* vol. 60, no. 6, pp. 4498-4504, 2013.

[66] W. Chen, V. Pouget, H. Barnaby, J. Cressler, G. Niu, P. Fouillat, Y. Deval and D. Lewis, "Investigation of single-event transients in voltage-controlled oscillators," *IEEE Transactions on Nuclear Science,* vol. 50, no. 6, pp. 2081-2087, 2003.

[67] D. I. Sotskov, V. V. Elesin, A. G. Kuznetsov, G. N. Nazarova, G. V. Chukov, D. V. Boychenko, V. A. Telets and N. A. Usachev, "Total Ionizing Dose Effects in Phase-Locked Loop ICs and Frequency Synthesizers," in *2015 15th European Conference on Radiation and Its Effects on Components and Systems (RADECS)*, Moscow, 2015.

[68] Z. Chen, D. Ding, Y. Dong, Y. Shan, S. Zhou, Y. Hu, Y. Zheng, C. Peng and R. Chen, "Study of Total-Ionizing-Dose Effects on a Single-Event-Hardened Phase-Locked Loop," *IEEE Transactions on Nuclear Science,* vol. 65, no. 4, pp. 997-1004, 2018.

[69] A. S. Youssouf, M. H. Habaebi and N. F. Hasbullah, "The Radiation Effect on Low Noise Amplifier Implemented in the Space-Aerial–Terrestrial Integrated 5G Networks," *IEEE Access,* vol. 9, pp. 46641-46651, 2021.

[70] T. Turflinger, M. Davey and B. Mappes, "Single event effects in analog-to-digital converters: device performance and system impact," *IEEE Transactions on Nuclear Science,* vol. 41, no. 6, pp. 2187-2194, 1994.

[71] Q. Huang, J. Jiang and Y. Deng, "Comparative evaluation of six wireless sensor devices in a high ionizing radiation environment," *IET Wireless Sensor Systems,* vol. 10, no. 6, p. 276–282, 2020.

[72] J.-M. Armani, J.-F. Gibrat and J. Martinez, "Total ionizing dose tolerance of LoRa wireless devices," in *RADECS 2022 - RADiation and its Effects on Components and systems conference*, Venise, Italy, 2022.

[73] A. Zimmaro, R. Ferraro, J. Boch, F. Saigne, R. G. Alia, A. Masi and S. Danzeca, "Impact of flux selection, pulsed beams and operation mode on system failure observability during radiation qualification," in *2022 22nd European Conference on Radiation and Its Effects on Components and Systems (RADECS)*, Venice, 2022.

[74] J. Budroweit and A. Koelpin, "Design challenges of a highly integrated SDR platform for multiband spacecraft applications in radiation enviroments," in *IEEE Topical Workshop Internet Space (TWIOS)*, 2018.

[75] J. Budroweit and M. Sznajder, "Total ionizing dose effects on a highly integrated RF transceiver for small satellite radio applications in low Earth orbit," in *Proc. IEEE Int. Symp. Phys. Failure Anal. Integr. Circuits (IPFA)*, 2018.

[76] J. Budroweit, S. Mueller, M. Jaksch, R. G. Alía, A. Coronetti and A. Koelpin, "In-situ testing of a multi-band software-defined radio platform in a mixed-field irradiation environment," *Aerospace,* vol. 6, no. 10, p. no. 106, 2019.

[77] J. Budroweit, M. P. Jaksch and M. Sznajder, "Proton induced single event effect characterization on a highly integrated RF-transceiver," *Electronics,* vol. 8, no. 5, p. 519, May 2019.

[78] J. Budroweit, M. Jaksch, R. G. Alía, A. Coronetti and A. Kölpin, "Heavy ion induced single event effects characterization on an RF-agile transceiver for flexible multi-band radio systems in NewSpace avionics," *Aerospace,* vol. 7, no. 2, p. 14, 2020.

[79] D. Chen, T. Mondy and A. Phan, "Heavy ion test report for the AD9364 RF transceiver," NASA, 2017.

[80] Keysight Technol., "Software defined radio measurement solutions," 2017. [Online]. Available: https://www.keysight.com/ch/de/assets/7018-01616/application- notes/5989-6931.pdf.

[81] ETSI, "Evolved universal terrestrial radio access (E-UTRA)," 2008. [Online]. Available: https://www.etsi.org/deliver/etsi_ts/136100_136199/136104/08.03.00_60/ts_136104v080300p.pdf. [Accessed 22 September 2024].

[82] A. Ildefonso, I. Song, G. N. Tzintzarov, Z. E. Fleetwood, N. E. Lourenco, M. T. Wachter and J. D. Cressler, "Modeling Single-Event Transient Propagation in a SiGe BiCMOS Direct-Conversion Receiver," *IEEE Transactions on Nuclear Science,* vol. 64, no. 8, pp. 2079-2088, 2017.

[83] P. K. C. Mishu, M.-K. Cho, A. Khachatrian, S. P. Buchner, D. McMorrow, P. Paki, J. D. Cressler and I. Song, "Voltage-Controlled Oscillator Utilizing Inverse-Mode SiGe-HBT Biasing Circuit for the Mitigation of Single-Event Effects," *IEEE Transactions on Nuclear Science,* vol. 69, no. 6, pp. 1242-1248, 2022.

[84] A. Ildefonso, G. N. Tzintzarov, N. E. Lourenco, Z. E. Fleetwood, A. Khachatrian, S. P. Buchner, D. McMorrow, J. H. Warner, M. Kaynak and J. D. Cressler, "Tradeoffs Between RF Performance and SET Robustness in Low-Noise Amplifiers in a Complementary SiGe BiCMOS Platform," *IEEE Transactions on Nuclear Science,* vol. 67, no. 7, pp. 1521-1529, 2020.

[85] Analog Devices, "AD9361, AD9364 and AD9363," 2018. [Online]. Available: https://wiki.analog.com/resources/eval/user-guides/ad9361.

[86] G. C. J. Clark and J. Bibb Cain, Error-Correction Coding for Digital Communications, New York: Plenum Press, 1981.

[87] S. B. Weinstein, "The history of orthogonal frequency-division multiplexing [history of communications]," *IEEE Commun. Mag., ,* vol. 47, no. 11, p. 26–35, 2009.

[88] J. J. v. d. Beek, M. Sandell and P. O. Börjesson, "ML estimation of time and frequency offset in OFDM systems," *IEEE Trans. Signal Process.,* vol. 45, no. 7, p. 1800–1805, 1997.

[89] European Space Agency, "Single event effects test method and guidelines—Basic specification," 1995. [Online]. Available: https://escies.org/download/webDocumentFile?id=62690.

[90] R. A. Shafik, M. S. Rahman and A. R. Islam, "On the extended relationships among EVM, BER and SNR as performance metrics," in *Proc. Int. Conf. Electr. Comput. Eng.*, Dec. 2006.

[91] T. P. a. M. Moeneclaey, "The effect of carrier frequency offset on the performance of band limited single carrier and OFDM signals," in *Proc. IEEE Global Telecommun. Conf. (GLOBECOM)*, “The effect of carrier frequency offset on the performance of band limited single carrier and OFDM sig- nals,” in Proc. IEEE Global Telecommun. Conf. (GLOBECOM), vol. 1, Nov. 1996, pp. 719–723., 1996.

[92] C. L. a. T.-S. Ng, "Pilot-based carrier frequency offset estimation in OFDM systems," in *Proc. 4th Int. Conf. Inf., Commun. Signal Process*, 2003.

[93] J. Budroweit, "Development of a highly integrated and radiation-tolerant software-defined radio platform for multi-band radio applications in space systems," 2021, Dissertation. doi: 10.15480/882.3779..

[94] J.-P. Thomesse, "Fieldbus Technology in Industrial Automation," *Proceedings of the IEEE,* vol. 93, no. 6, pp. 1073-1101, 2005.

[95] T. Rappaport, "Indoor radio communications for factories of the future," *IEEE Communications Magazine,* vol. 27, no. 5, pp. 15-24, 1989.

[96] R. T.S, Wireless Communications—Principles and Practice, Upper Saddle River, NJ: Prentice-Hall, 2002.

[97] Telecommunications and Information Exchange Between Systems—Local and Metropolitan Area Net- works—Specific Requirements—Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications,, "Telecommunications and Information Exchange Between Systems—Local and Metropolitan Area Net- works—Specific Requirements—Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications,," 1999.

[98] A. Willig and A. Wolisz, "Ring Stability of the PROFIBUS Token-passing Protocol over Error-prone Links," *IEEE Transactions on Industrial Electronics,* vol. 48, no. 5, pp. 1025-1033, 2001.

[99] Union Technique de l’Electricite., "General Purpose Field Communication System, EN 50170, Volume 2: PROFIBUS," 1996.

[100] A. Willig, "Analysis of the PROFIBUS Token Passing Protocol over Wireless Links," *Industrial Electronics, 2002. ISIE 2002. Proceedings of the 2002 IEEE International Symposium on,* pp. 56-60, 2002.

[101] A. Willig, M. Kubisch, C. Hoene and A. Wolisz, "Measurements of a Wireless Link in an Industrial Environment using an IEEE 802.11-compliant Physical Layer," *IEEE Transactions on Industrial Electronics,* vol. 49, no. 6, pp. 1265-1282, 2002.

[102] A. Willig, "Polling-based MAC protocols for Improving Real-time Performance in a Wireless PROFIBUS," *IEEE Transactions on Industrial Electronics,* vol. 50, no. 4, pp. 806-817, 2003.

[103] J. a. R. L. Haehniche, "Radio Communication in Automation Systems: the R-fieldbus Approach," in *2000 IEEE International Workshop on Factory Communication Systems. Proceedings (Cat. No.00TH8531)*, 2000.

[104] L. Rauchhaupt, "System and device architecture of a radio based fieldbus-the RFieldbus system," in *4th IEEE International Workshop on Factory Communication Systems*, 2002.

[105] F. Pacheco, E. Tovar, A. Kalogeras and N. Pereira, "Supporting Internet Protocols in Master-Slave Fieldbus Networks," in *Proceedings of 4th IFAC FET Conference*, 2001.

[106] M. Alves, E. Tovar, F. Vasques, G. Hammer and K. Rother, "Real-time communications over hybrid wired/wireless PROFIBUS-based networks," in *Proceedings 14th Euromicro Conference on Real-Time Systems. Euromicro RTS 2002*, 2002.

[107] F. Pacheco, N. Pereira, B. Marques, S. Machado, L. Marques, L. M. Pinho and E. Tovar, "Industrial Multimedia put into Practice," [Online]. Available: http://www.cister.isep.ipp.pt/docs/industrial_multimedia_put_into_practice/145/view.pdf. [Accessed 10 October 2024].

[108] A. Kutlu, H. Ekiz and E. Powner, "Performance analysis of MAC protocols for wireless control area network," in *Proceedings Second International Symposium on Parallel Architectures, Algorithms, and Networks (I-SPAN'96)*, 1996.

[109] A. Kutlu, H. Ekiz and E. Powner, "Wireless control area network," in *IEE Colloquium on Networking Aspects of Radio Communication Systems, pp. 3/1 –3/4*, 1996.

[110] W. Ng, C. Ng, B. Mohd. Ali and N. Noordin, "Wireless Controller Area Network Using Token Frame Scheme," in *International Conference on Informatics Engineering and Information Science, pp 544–556*, Kuala Lumpur, Malaysia, 2011.

[111] P. Morel, A. Croisier and J.-D. Decotignie, "Requirements for Wireless Extensions of a FIP Fieldbus," in *Proceedings 1996 IEEE Conference on Emerging Technologies and Factory Automation. ETFA '96*, 1996.

[112] P. Morel and A. Croisier, "A Wireless Gateway for Fieldbus," in *Proceedings of 6th International Symposium on Personal, Indoor and Mobile Radio Communications*, 1995.

[113] A. Willig, K. Matheus and A. Wolisz, "Wireless Technology in Industrial Networks," *Proceedings of the IEEE,* vol. 93, no. 6, pp. 1130-1151, 2005.

[114] J. Song, S. Han, A. Mok, D. Chen, M. Lucas, M. Nixon and W. Pratt, "WirelessHART: Applying Wireless Technology in Real-Time Industrial Process Control," *2008 IEEE Real-Time and Embedded Technology and Applications Symposium,* pp. 377-386, 2008.

[115] ISA, "Wireless Systems for Industrial Automation: Process Control and Related Applications," 2009.

[116] Microchip, "PolarFire FPGA and PolarFire SoC FPGA Programming User Guide," [Online]. Available: https://ww1.microchip.com/downloads/aemDocuments/documents/FPGA/ProductDocuments/UserGuides/PolarFire_FPGA_and_PolarFire_SoC_FPGA_Programming_User_Guide_VB.pdf. [Accessed 2 11 2024].

[117] A. Scialdone, R. Ferraro, R. G. Alía, L. Sterpone, S. Danzeca and A. Masi, "FPGA Qualification and Failure Rate Estimation Methodology for LHC Environments Using Benchmarks Test Circuits," *IEEE Transactions on Nuclear Science,* vol. 69, no. 7, pp. 1633-1641, 2022.

[118] R. Seelam, "I/O design flexibility with the FPGA mezzanine card (FMC)," *Xilinx White Paper WP315,* 2009.

[119] F. Vasey, D. Hall, T. Huffman, S. Kwan, A. Prosser, C. Soos, J. Troska, T. Weidberg, A. Xiang and J. Ye, "The Versatile Link common project: feasibility report," *Journal of Instrumentation,* vol. 7, no. 1, p. C01075, 2012.

[120] P. Moreira, S. Baron, S. Bonacini, O. Cobanoglu, F. Faccio, S. Feger, R. Francisco, P. Gui, J. Li, A. Marchioro, C. Paillard, D. Porret and K. Wyllie, "The GBT-SerDes ASIC prototype," *Journal of Instrumentation,* vol. 5, no. 11, p. C11022, 2010.

[121] P. e. a. Moreira, "The GBT Project," 2009. [Online]. Available: https://cds.cern.ch/record/1235836.

[122] M. Marin, A. Boccardi, C. Godichal, J. Gonzalez, T. Lefevre and T. Levens et al., "The Giga Bit Transceiver based Expandable Front-End (GEFE) - a new radiation tolerant acquisition system for beam instrumentation," *Journal of Instrumentation,* vol. 11, no. 2, p. C02062, 2016.

[123] F. Faccio, S. Michelis, G. Blanchot, G. Ripamonti and A. Cristiano, "The bPOL12V DCDC converter for HL-LHC trackers: towards production readiness," in *PoS TWEPP 2019*, 2020 070.

[124] A. Caratelli, S. Bonacini, K. Kloukinas, A. Marchioro, P. Moreira and R. O. e. al., "The GBT-SCA, a radiation-tolerant ASIC for detector control and monitoring applications in HEP experiments," *Journal of Instrumentation,* vol. 10, p. C03034, 2015.

[125] J. S. e. al., "WirelessHART: Applying Wireless Technology in Real-Time Industrial Process Control," in *2008 IEEE Real-Time and Embedded Technology and Applications Symposium*, St. Louis, MO, USA, 2008, pp. 377-386.

[126] M. Luvisotto, Z. Pang and D. Dzung, "High-Performance Wireless Networks for Industrial Control Applications: New Targets and Feasibility," *Proceedings of the IEEE,* vol. 107, no. 6, pp. 1074-1093, 2019.

[127] Microchip, "PolarFire® FPGA and PolarFire SoC FPGA Power-Up and Resets User Guide," Aug 2023. [Online]. Available: https://ww1.microchip.com/downloads/en/DeviceDoc/PolarFire_FPGA_and_PolarFire_SoC_FPGA_Power_Up_and_Reset_User_Guide_60001676.pdf. [Accessed 2025].

[128] R. H. Barke, "Group Synchronizing of Binary Digital Sequences," in *Communication Theory*, Butterworth, 1953.