The thesis focuses on the design, implementation and validation of a distributed real-time magnetic measurement system for particle accelerators' magnets. In particular the thesis focuses on the study of 4 macro-areas connected to this system.
The first study regards the development and the implementation of a real-time measurement system starting from the system requirements defined by the final users (i.e. LLRF, Power converters and machine operators). To satisfy all the necessities and all the constraints for all the seven particle accelerators that will benefit from this system, it was a crucial point have a measurement system as much flexible as possible to minimize the maintenance costs and to being prone to future requests by the system's users without the necessity of massive changes in the system. In the thesis, a novel system, one of a kind, able to measure and simulate a magnetic field in real-time and provide it over a optic-fiber based network is presented. To satisfy all the requests and to overcome all the constraints it has been necessary to design new custom electronics, both new PCBs modules to be integrated in the framework of the standard CERN electronics, and new FPGA modules described in VHDL. Moreover it has been necessary to design new software modules integrated in the FESA Framework to be compatible with all the CERN infrastructure.
The second study concerns the development of a necessary tool to monitor in real-time the magnetic field provided by the seven installed systems over the optic-fiber. This tool was necessary for three reasons: first, to characterize and calibrate the system looking at the same output that the users will receive. second, to monitor the correct behavior of the systems. The tool was connected to all the seven systems thanks to an optic multiplexer.
The third study regards the DC and dynamic performance evaluation of the presented system. A satisfying agreement with the metrological requirements for the system was found after the fine calibration of the systems.
The fourth study concerns the possibility to use neural networks to predict the magnetic field, and all its non linearity such as eddy currents and hysteresis inside magnets for particle accelerators.
Considering a calibration quadrupole as case study various neural network based architectures have been designed and implemented. The achieved result of this study is a network able to predict the magnetic field leading a percent error below 0.02 %.
A detailed study of these four topics is presented.
All the realized systems and subsystems were benchmarked with simulation and experimental measures performed all around the CERN acceleration complex (i.e.PSB,PS,SPS,AD,ELENA and LEIR). A satisfying agreement respect to the original system's requirements was found in all cases.

Vincenzo Di Capua

Real-time magnetic measurement system for synchrotrons control

2021 - XXXIV

UNIVERSITÀ DEGLI STUDI DI NAPOLI FEDERICO II

Ph.D itee INFORMATION TECHNOLOGY ELECTRICAL ENGINEERING

REAL-TIME MAGNETIC MEASUREMENT SYSTEM FOR SYNCHROTRONS CONTROL

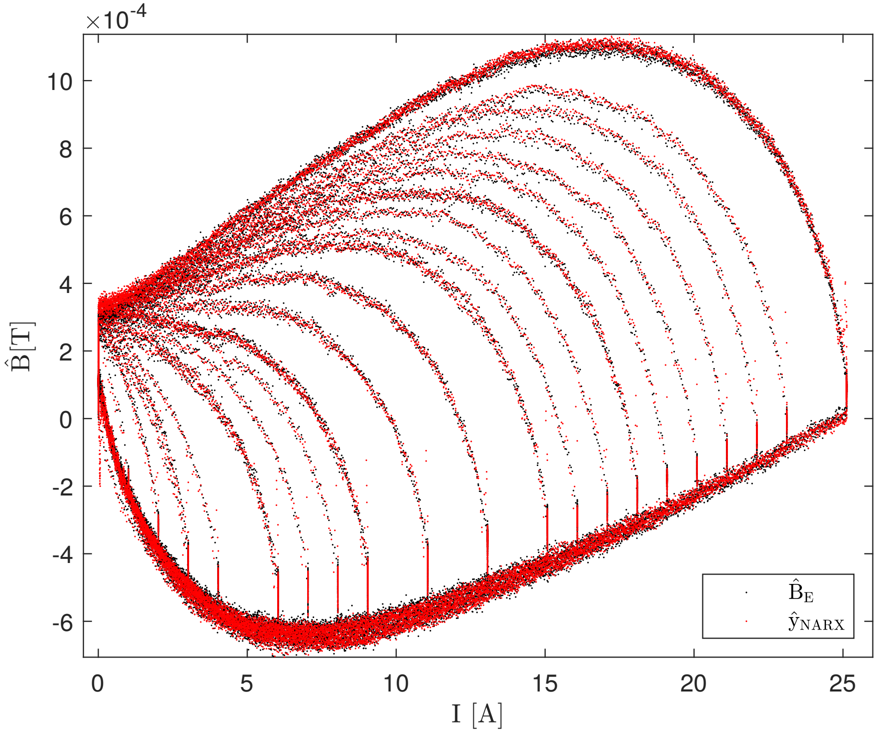

VINCENZO DI CAPUA

PH.D. IN INFORMATION TECHNOLOGY AND ELECTRICAL ENGINEERING

Università degli Studi di Napoli Federico II

Ph.D. THESIS

in

Information Technology and Electrical Engineering

Real-time magnetic measurement system for synchrotrons control

Vincenzo Di Capua

Tutor: Prof. Pasquale Arpaia

Coordinator: Prof. Daniele Riccio

XXXIV Ciclo

Scuola Politecnica e delle Scienze di Base

Dipartimento di Ingegneria Elettrica e Tecnologie dell'Informazione

UNIVERSITY OF NAPLES FEDERICO II

PH.D. THESIS
IN
INFORMATION TECHNOLOGY AND ELECTRICAL ENGINEERING

Real-time magnetic measurement system for synchrotrons control

Supervisors:
Prof. Pasquale Arpaia
Dr. Marco Buzio

Candidate:
Vincenzo Di Capua

2021 - 2022

"Continue what you started and maybe you will reach the top, or at least you will reach a point that you alone will realize that you are not the top."
cit. Lucio Anneo Seneca

To my incredible wife Silvia,
to my amazing daughter Altea,
to my whole Family,
and to my fervent teachers.
You are the motivation behind all the endeavours.

Abstract

The thesis focuses on the design, implementation and validation of a distributed real-time magnetic measurement system tailored for particle accelerators' magnets. In particular the thesis focuses on the study of 4 macro-areas connected to this system.
The first study regards the development and the implementation of a real-time measurement system starting from the system requirements defined by the final users (i.e. low level radio frequency, power converters and machine operators). To satisfy all the necessities and all the constraints for all the seven particle accelerators that will benefit from this system, it was a crucial point have a measurement system as much flexible as possible to minimize the maintenance costs and to being prone to future requests by the system's users without the necessity of massive changes. In the thesis, a novel system, one of a kind, able to measure and simulate a magnetic field in real-time and provide it over a optic-fiber based network is presented. To satisfy all the requests and to overcome all the constraints it was necessary to design new custom electronics, both new PCBs modules to be integrated in the framework of the standard CERN electronics, and new FPGA modules described in VHDL. Moreover it was necessary to design new software modules integrated in the software Framework to be compatible with all the CERN infrastructure.
The second study concerns the development of a necessary tool to monitor in real-time the magnetic field provided by the seven installed systems over the optic-fiber. This tool was necessary for three reasons: first, to characterize and calibrate the system looking at the same output that the users will receive. Second, to monitor the correct behavior of the systems. Third to speed-up the debugging in case of issues. In this layout a novel tool based both on commercial National Instruments and custom hardware is proposed. The tool was connected to all the seven systems thanks to an optic multiplexer.
The third study regards the DC and dynamic performance evaluation of the presented system. A satisfying agreement with the metrological requirements for the system was found after the fine calibration of the systems.
The fourth study concerns the possibility to use neural networks to predict the magnetic field, and all its non linearity such as eddy currents and hysteresis inside magnets for particle accelerators. Considering a calibration quadrupole as case study various neural network based architectures were designed and implemented. The achieved result of this study is a network able to predict the magnetic field leading a percent error below $0.02\ \%$.
A detailed study of these four topics is presented.
All the realized systems and subsystems were benchmarked with simulation and experimental measures performed all around the CERN acceleration complex

(i.e.PSB,PS,SPS,AD,ELENA and LEIR). A satisfying agreement respect to the original system's requirements was found in all cases.

Keywords: Particle accelerators, Magnetic measurements, FPGA, Normal conducting magnets, Neural network, Artificial intelligence, Calibration, Software framework, Magnetic sensors, Linux.

Cover Image: Measured ($\hat{B}_{\mathrm{E}}$, in black) and estimated ($\hat{y}_{\mathrm{NARX}}$ with hyperparameters $\tilde{\theta}_{NARX1}$, in red). The nonlinear component of the field $\hat{B}$ was plotted in function of the current I (hysteresis graph).

Sommario

La tesi si focalizza sul design, sull'implementazione e sulla validazione di un sistema di misura in real-time per il campo magnetico studiato appositamente per i dipoli degli acceleratori di particelle. In particolare questa tesi si focalizza sullo studio di 4 macro-aree che ruotano intorno al sistema sopracitato.
Il primo studio riguarda lo sviluppo e l'implementazione di un sistema di misura in tempo reale a partire dai requisiti di sistema definiti dagli utenti finali (es. radiofrequenza, Convertitori di potenza e operatori di macchina). Allo scopo di tenere in considerazione tutte le necessita' e tutti i vincoli per tutti i sette acceleratori di particelle che beneficeranno di questo sistema e' stato cruciale avere un sistema di misura il più flessibile possibile per ridurre al minimo i costi di manutenzione ed essere proni a futueo richieste da parte degli utenti del sistema senza la necessità di modifiche invasive . Nella tesi viene presentato un nuovo sistema, unico nel suo genere, in grado di misurare e simulare un campo magnetico in tempo reale e fornirlo su una rete basata su fibra ottica. Per soddisfare tutte le richieste e per superare tutti i vincoli è stato necessario avvalersi di elettronica custom, in particolare sia nuovi moduli PCB da integrare nel framework dell' elettronica standard CERN, sia nuovi moduli FPGA descritti in VHDL. Inoltre e' stato necessario sviluppare anche nuovi moduli software integrati nel Framework FESA al fine essere compatibili con tutta l'infrastruttura del CERN.
Il secondo studio riguarda lo sviluppo di uno strumento necessario per monitorare in tempo reale il campo magnetico misurato e distribuito su fibra ottica dai sette sistemi installati. Questo strumento è stato necessario per due ragioni: in primo luogo, per caratterizzare e calibrare il sistema guardando lo stesso output che gli utenti riceveranno. secondo, monitorare il corretto comportamento dei sistemi. In questo layout viene proposto un nuovo strumento basato sia su hardware National Instruments che su hardware custom. Lo strumento è stato collegato a tutti e sette i sistemi grazie a un multiplexer ottico.
Il terzo studio riguarda l' analisi delle performance statiche e dinamiche del sistema presentato. Dopo un accurata calibrazione dei sistemi è stato trovato un accordo soddisfacente con i requisiti metrologici del sistema.
Il quarto studio riguarda la possibilità di utilizzare reti neurali per prevedere il campo magnetico e tutte le sue non linearità come correnti parassite e isteresi all'interno di magneti per acceleratori di particelle. Considerando un quadrupolo di calibrazione come caso di studio, sono state progettate e implementate diverse architetture basate su reti neurali. Il risultato ottenuto da questo studio è una rete in grado di prevedere il campo magnetico con un errore percentuale inferiore al $0.02\ \%$.
Viene presentato uno studio dettagliato di questi quattro argomenti
Tutti i sistemi e i sottosistemi realizzati sono stati confrontati con simulazioni e misure sperimentali eseguite in tutto il complesso di accelerazione del CERN

(cioè PSB,PS,SPS,AD,ELENA and LEIR). In tutti i casi è stato riscontrato un soddisfacente accordo rispetto ai requisiti originari del sistema.

Keywords: Acceleratori di particelle, Misure magnetiche, FPGA, Magneti resistivi, Rete neurale, Intelligenza artificiale, Calibrazione, Framework software, Sensori magnetici, Linux.

Cover Image: Campo magnetico $\hat{B}_{\mathrm{E}}$ misurato (in black) e campo magnetico $\hat{y}_{\mathrm{NARX}}$ stimato con gli iperparametri $\tilde{\theta}_{NARX1}$ (in red). La componente non lineare del campo $\hat{B}$ e' graficata in funzione della corrente I (diagramma di isteresi).

List of publications

Journal publications:

(i) P. Arpaia, M. Buzio, V. Di Capua, S. Grassini, M. Parvis, M. Pentella; *"Drift-Free Integration in Inductive Magnetic Field Measurements Achieved by Kalman Filtering."* **MDPI Sensors** (December, 2021).
doi:10.3390/s22010182

(ii) J. Vella Wallbank, M. Amodeo, A. Beaumont, M. Buzio, V. Di Capua, C. Grech, N. Sammut, D. Giloteaux; *"Development of a Real-Time Magnetic Field Measurement System for Synchrotron Control"*. **MDPI Electronics** 10, no. 17 (September, 2021).
doi:10.3390/electronics10172140

(iii) M. Amodeo, P. Arpaia, M. Buzio, V. Di Capua, F. Donnarumma; *"Hysteresis Modeling in Iron-Dominated Magnets Based on a Multi-Layered NARX Neural Network Approach."* **International Journal of Neural Systems** (July, 2021).
doi:10.1142/S0129065721500337

(iv) C. Grech, M. Amodeo, A. Beaumont, M. Buzio, V. Di Capua, D. Giloteaux, N. Sammut, J. Vella Wallbank; *"Error characterization and calibration of real-time magnetic field measurement systems."* **Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment** (December, 2020).
doi:10.1016/j.nima.2020.164979

(v) P.Arpaia, U.Baracale, F.Corcione, E. De Benedetto, A. Di Bernardo, V. Di Capua, R. Prevete. *"Machine Learning-based assessment of the vascularization quality in laparoscopic colorectal surgery"*. **Nature scientific reports** (submitted (2021)).

(vi) P. Arpaia, G. Annuzzi, E. De Benedetto, V. Di Capua, R. Prevete, E. Vallefuoco. *"Full nutritional factor Neural network based Metabolic predictors"*. **Nature scientific reports** (submitted (2021))

Conference proceedings:

(i) F.M. Velotti, H. Bartosik, M. Buzio, K. Cornelis, V. Di Capua, M.A. Fraser, B. Goddard, V. Kain; *"Characterisation of SPS Slow Extraction Spill Quality Degradation."* **10th Int. Particle Accelerator Conf.(IPAC'19), Melbourne, Australia** (October, 2019).
doi:10.18429/JACoW-IPAC2019-WEPMP034

(ii) P. Arpaia, V. Di Capua, M. Roda, M. Buzio; *"Real-Time Magnetic Measurement Monitoring under cRIO-LabVIEW Based Platform."* **IEEE International Symposium on Precision Clock Synchronization for Measurement, Control and Communication, ISPCS 2019**(December, 2018).

Acknowledgements

For the third time in my life I'm here again to try to express in words the gratefulness for all the people that made this possible. I'm not sure that I'm going to be able to express in this pages the huge feelings I have. But I will try and, I am sure you will understand. Many people indeed deserve to be mentioned here, so it will take a bit. To make it easy, I will mostly write in English, but I will switch to French and Italian for those with a strict preference.
All the research presented in this dissertation was carried out at CERN laboratories, all around the main CERN accelerators in the framework of such an interesting ad challenging project as the B-Train is. If I should describe this project to you in few words I can say that it is a measurement system that plays its part to allow CERN to push the human knowledge further each day. It have been an honour to do my part and it have been an experience I will never forget. I had the oportunity to meet and work amazing people contributing to its operation day by day endure many difficulties, only for the love of their mission: I had the luck of calling these people friends and colleagues. Each one taught me so much from a personal and scientific point of view.
Firstly, I feel deeply grateful to Marco Buzio, my CERN supervisor and B-Train project leader, for always believing in me and pushing me to accept new challenges. Thank you for the constant guidance, excellent advice, confidence, and opportunities for professional growth you have given me. David Giloteaux, thanks for your explanations when I felt lost and stuck. You always have an explanation for something that is not working properly and you have always a brilliant idea to solve the issue. I learned so much from you and I think I will have still a lot to learn. Anthony Beaumont, I admire your tireless work and I would like to thank you for all the good words you always deserved to me. It have been a pleasure to work with you and still now when I have troubles I ask myself what would be your behavior in that situation.
A huge thanks to Alessandro Parrella, Maciej Lipinski, Gianpiero di Giovanni and all the B-train team, for the time spent together and for the helpful and fruitful discussions. You have been an essential piece of my journey, always there to support new ideas and to fight unexpected issues.
I would like to take the time to thank all the collaborators/friends that have always been forthcoming with their knowledge. In particular, I feel proud to mention Mariano Pentella, Stefano Sorti, Lucio Fiscarelli, Carlo Petrone, Melvin Liebsc and all the section TE-MSC-TM. Thank you for your friendly and kind attitude, I learned a lot.
Je suis honoré d'avoir travaillé avec des techniciens d'aussi haut niveau que Maxime, Ricardo, Guy, Regis et toute l'équipe du 311 Vous m'avez toujours fasciné, m'apprenant beaucoup grâce à vos connaissances impressionnantes et votre attitude amicale. Votre professionnalisme et votre précision resteront toujours une grande source d'inspiration. Merci à vous tous.
In particular a great thanks goes to two fiends/colleagues that have indeed greatly contributed to my professional and personal growth: Alessandro Parrella and Mariano Pentella. Alesandro Helped me in so many situations that I can not

even count them. He is a very good friend always ready to help and give his piece of advise. Mariano is a special person and without any doubt one of the best engineers I have ever met in my career. Thank you for all guys, I hope to keep collaborating with you f and maintaining the wonderful relationship we have built.

I am profoundly pleased to thank my mentor, Prof. Pasquale Arpaia, for the invaluable academic supervision, for presenting me 4 years ago the opportunity to start my CERN career, and teaching me to embrace new challenges without fear. As he always says *"no problems, only opportunities"*

Il mio grazie piu grande va pero' alla mia famiglia, a mia moglie Silvia e a mia figlia Altea che ogni giorno con amore e dedizione mi supportano in tutte le mie scelte dalle piu facili a quelle piu difficili e che talvolta comportano dei sacrifici, sapendo sempre come farmi tornare il sorriso nei momenti di sconforto. Con voi so di non essere mai solo e so che su di voi potro' sempre contare. Vi sono eternamente grato.

A mio padre e a mia madre, grazie per avermi permesso di studiare seguendo le mie passioni e di intraprendere questo incredibile viaggio, pur sapendo che mi avrebbe portato se pur solo fisicamente lontano da voi. Non riusciro' mai ad esprimervi tutta la mia gratitudine.

A mia sorella Elvira per aver alleggerito tanti momenti di tensione durante il mio persorso di studio. Ti voglio bene.

A mio Nonno Catello che fin da quando ho memoria ha sempre stimolato in me la curiosita' per il mondo e per il funzionamento delle cose che ci circodnano. Se oggi sono un ingegnere e' anche merito tuo. Grazie infinitamente.

Alle mie nonne Bianca ed Elvira, per aver sempre creduto in me e eper avermi sempre dato buoni consigli sulla vita uno su tutti *"statt'accort!"*

Voglio anche ringraziare mio zio Carmine per aver fatto scoccare in me la passione per l'elettronica. Se oggi questo e' il mio lavoro lo devo sicuramente anche a te.

Grazie Fabio, Tiziana, Nonna Liliana e tutta la famiglia Valeri/Piersanti, mi avete accolto come un figlio e con la vostra gentilezza e allegria mi avete sempre fatto sentire a casa.

Arrivato a questo punto mi rendo conto che non sono capace di scrivere tutto quello che provo per voi quindi e' meglio fermarmi qui con l'augurio di riuscirvelo a dimostrare con i fatti piuttosto che con le parole. GRAZIE GRAZIE GRAZIE.

At this point I realize that I am not able to write everything I feel for you so it is better to stop here with the wish to be able to prove it with facts rather than words. THANKS THANKS THANKS.

À ce stade, je me rends compte que je ne suis pas capable d'écrire tout ce que je ressens pour vous, il vaut donc mieux s'arrêter ici avec le souhait de pouvoir le prouver avec des faits plutôt qu'avec des mots. MERCI MERCI MERCI.

CERN - Geneva, Switzerland
Vincenzo Di Capua

Contents

List of abbreviations

AD Antiproton Decelerator

ADC Analog to Digital Converter

ANN Artificial Neural Network

ANOVA Analysis of Variance

AIC Akaike Information Criterion

BC Bridge Criterion

BIC Bayesian Information Criterion

BPTT Backpropagation Through Time

CERN European Organization for Nuclear Research

CMOS Complementary Metal Oxide Structure

CNAO National Centre of Oncological Hadrontherapy

CTRI Central Timing Receiver

CW Continuous Wave

DAC Digital to Analog Converter

DCCT Direct Current Current Transducer

DDR Double Data Rate

DC Direct Current

DIO Digital Input Output

DMA Direct Memory Access

DNN Deep Neural Network

DSP Digital Signal Processor

ELENA Extra Low Energy Antiproton

FAIR Facility for Antiproton and Ion Research

FD Fourier Descriptor

FE Finite Element

FEC Front End Computer

FESA Front End Software Architecture

FFNN Feed-Forward Neural Network

FFMM Flexible Framework for Magnetic Measurements

FIFO FIrst In First Out

FIRESTORM Field In REal-time STreaming from Online Reference Magnets

FMC FPGA Mezzanine Card

FMR Ferrimagnetic Resonance

FPGA Field Programmable Gate Array

GMT General Machine Timing

GPSDSO GPS Disciplined Oscillator

GSI Helmholtz Center for Heavy Ion Research

HDL Hardware Description language

HDMI High-Definition Multimedia Interface

HIT Heidelberg Ion-Beam Therapy Centre

HP High Priority

ISE Integrated Synthesis Environment

JTAG Joint Test Action Group

LEIR Low Energy Ion Ring

LHC Large Hadron Collider

LINAC Linear Accelerator

LLRF Low Level Radio Frequency

LPC Low Pin Count

LR linear Regression

LSB Least Significant Bit

LSTM Long-Short Memory Network

LVDS Low Voltage Differential Signal

MAE Maximum Absolute Error

MLP Multi Layer Perceptron

MPE Maximum Percentage Error

NARX Non-Linear Autoregressive Exogenous

NMR Nuclear Magnetic Resonance

NN Neural Network

NRMSE Normalized Root Mean Square Error

OHWR Open Hardware Repository

PCIe Peripheral Component Interconnect express

PLL Phase Locked Loop

PPM Pulse to Pulse Modulation

PSB Proton Synchrotron Booster

PS Proton Synchrotron

PTP Precise Time Protocol

PXI PCI Extensions for Instrumentation

RF Radio Frequency

RFQ Radio Frequency Quadrupole

RMSE Root Mean Square Error

SATA Serial Advanced Technology Attachment

SAR Successive Approximation Register

SFP Small Form-factor Pluggable Transceiver

SP Standard Priority

SPEC Simple PCIe Carrier

SPI Serial Periferal Interface

SPS Super Proton Synchrotron

TDNN Time Delayed Neural Network

TTL Transistor Transistor Logic

VHDL Very High Speed Integrated Circuits Hardware Description Language

VME VERSABUS Module Eurocard

WR White Rabbit

Introduction

Introduction

The precise knowledge of the magnetic field produced by dipole magnets is critical not only for the operation of a synchrotron but also in fusion engineering applications, where field measurements are used for diagnostic monitoring of the magnetic field interacting with the plasma. Real-time measurement systems are required, to acquire the magnetic field and feed it back to various subsystems in charge of the control of the magnets, especially in the case of iron-dominated electromagnets with strong non-linear effects such as eddy currents, hysteresis and saturation. Many are the sensors that could be used for magnetic measurements, each kind of sensor has its specific field of application and its peculiarity. For this reason, they have to be carefully chosen and the data coming out need appropriate manipulation to be used. Typically custom hardware [1, 2, 3] and software [4, 5] are needed to elaborate the data and to provide a real-time accurate measure of the magnetic field. In particle accelerators, the beam is accelerated by radio frequency cavities while circulating around a ring made by magnets, which generate a bending field, increasing in proportion to the beam momentum. Accurate knowledge of the magnetic field $B(t)$ at any given time during a magnetic cycle is therefore critical for longitudinal and transverse beam control, power supply control, various beam diagnostics, and qualitative feedback to operators. The required accuracy is typically $0.01\ \%$ [6]. Another aspect to not be neglected is the protocol used to transmit the measurements for long distances, *i.e.* several kilometers, over big laboratories as European Organization for Nuclear Research (CERN) is. For this purpose it is crucial to have a transmission protocol able to handle a large amounts of data streams and to ensure the synchronization of the data coming from different sources at the receiving node. It is also necessary to have a reliable sub-system capable to acquire the data coming out from the magnetic field measurement system with the aim to characterize it and debug it even remotely. Actually, a tailor made system for real-time magnetic measurements has to be characterized to certify the fulfillment of the requirements in terms of Direct Current (DC) and dynamic performance. Moreover, there are cases in which the magnetic field value is still necessary for magnets control but there is no possibility to have an actual measure due to failures in the measurement system or due to the physical impossibility to install sensors in the magnet to be monitored. In these scenarios simulations or predictions of the magnetic field are necessary. The Magnetic field inside electromagnets it is relatively easy to be predicted with closed mathematical models for coil-dominated magnets but this task becomes very difficult for the iron-dominated ones due to the non linearity introduced in the field current relation by eddy currents, hysteresis and saturation effects that characterize ferromagnetic materials. Modelling of quasi-static and dynamic hysteresis loops is one of the most challenging topics in computational magnetism [7, 8, 9, 10, 11, 12, 13]. For example, recent attempts

using the well-known Preisach models[14, 15] could not attain better than $0.2\ \%$ accuracy. Also other classes of methods, such as Jiles-Atherton differential models[16], ultimately turn out to be unsuitable, due to their well-known difficulties in handling minor hysteresis loops. In the thesis, all the aspects described above have been investigated and thoroughly treated. In particular, a novel real-time magnetic measurement system developed to replace the existing systems is presented in the context of a site-wide, consolidation project. The system was designed to cope with the High-Luminosity Large Hadron Collider upgrade, which will require higher beam intensity and improved beam control throughout the injector chain [17]. First, the measurement principle, the general system architecture and the technology employed were discussed focusing in particular on the most critical and specialized components developed, that are, the field marker trigger generator and the magnetic flux integrator. Second, a new monitoring system was developed to provide remote access to the measurement system fiber optic output to have the capabilities to fully characterize the system. The developed hardware and software are presented, together with the results of the validation tests. Third, the results of a detailed metrological characterization of the integrator are discussed, including the aspects of drift estimation and correction, the latency of the whole acquisition chain, as well as absolute gain calibration and frequency response. Finally, a novel approach for magnetic field prediction in electromagnets is proposed, based on tuning a Multi-layered neural network to fit directly the magnet response, by avoiding complementary physical models. Different architectures were considered and selected according to a compromise between the accuracy of the field estimation and the level of complexity of the network. The results of tests carried out on a dedicated experimental setup outperform both traditional and hybrid models, suggesting that this is indeed a very promising approach applicable in a wide range of areas in which the real-time accurate knowledge of the magnetic field in a magnet is required and there are no possibilities for real-time measurements.
The original contributions of this work are the FPGA implementation and integration of the subsystems used to build the presented real-time measurement system, the development from scratch of the tool for the diagnostic and characterization of the main system, the study of drift correction techniques implemented, the DC and dynamic system characterization, the development of NARX neural network approach to model non-linearities in iron-dominated magnets.

This thesis is divided into five parts: part I, Background, part II, Measurement system, part III, a tool for the measurement system characterization, part IV, metrological characterization and, part V Artificial intelligence.

In part I the context and the basic knowledge on particle accelerators needed to fully understand the work is presented, together with the state of the art for all the sub-domains treated in this thesis. In part II, the requirements, the proposal, and the hardware and software implementation for the novel real-time magnetic field measurement system are reported. In part III, a novel design is presented for a custom tool with the aim to fully characterize the system proposed in part II. In part IV, the proposed methodology for the system DC and dynamic characterization and the experimental results are reported. In part V an innovative solution for magnetic field prediction in electromagnets is presented, the requirements, the proposed methodology and the model design are discussed in detail together with the experimental results obtained from a case study carried out on a magnet for particle accelerators.

The structure of the chapters is as follows:

- **Chapter 1: Particle accelerators**. In order to ensure a full understanding of the reasons behind this thesis, an overview of the reasons to accelerate particles is presented. Then, the general architecture of synchrotrons for particle accelerators is depicted. Finally, the CERN accelerator complex is described with a particular focus on the acceleration chain.

- **Chapter 2: State of the art**. A recall of the formalism used in magnetic measurements is presented first. Then the topics magnetic field prediction, artificial intelligence, machine learning and the white rabbit transmission are presented.

- **Chapter 3: System requirements**. The reason behind the novel real-time magnetic field measurement system developed to replace the legacy systems in the acceleration chain is presented. Then requirements for the new system, and the methods used are discussed in detail.

- **Chapter 4: System proposal**. The proposed design for the measurement system is presented. Then the main functionalities of the proposed system are described in detail.

- **Chapter 5:System implementation**. The hardware and software implementation is depicted. All the electronic components used are parented as well as their interconnection and communication protocols. Then all the implemented algorithms together with all the individual modules composing the system are described.

- **Chapter 6: Monitoring system requirements**. The reason behind of the new monitoring system is described, then the main requirements are presented. The proposed architecture for the new monitoring system is described and, an overview of the hardware selected is provided, then the proposed block diagram of the interconnections and functional modules is presented.

- **Chapter 7: Monitoring system implementation**. The implementation of the new monitoring system is presented. First, an overview of the used software tools is provided. Then the implementations of all the Field Programmable Gate Array (FPGA) modules and of the Labview host application are described in detail. Finally, the obtained results proving the fulfillment of all the requirements are discussed.

- **Chapter 8: DC performance**. The results concerning voltage, magnetic flux and integrator drift are discussed separately. The accuracy of the integrator acquisition chain under DC input conditions was evaluated.

- **Chapter 9: Dynamic performance**. The measurements carried out on a test setup of the amplitude transfer function and the latency of the whole acquisition chain are presented.
- **Chapter 10: Hysteresis modelling in iron-dominated magnets based on a multi-layered NARX neural network approach**. Different machine learning approaches for magnetic field predictions are presented, based on tuning a Multi-layered neural network to fit directly the magnet response. First, a description of the problem statement is presented, then the architecture tuning and the model selection phases were described. Finally, a comparison between all the tested architecture is presented, highlighting the one with the best performances.

Part I

Background

Chapter 1

Particle accelerators

In this chapter, to ensure a full understanding of the reasons behind this thesis, an overview of the reasons to accelerate particles is presented. Later, the general architecture of synchrotrons for particle accelerators is depicted. Finally, the CERN accelerator complex is described with a particular focus on the acceleration chain.

1.1 Accelerators for particle physics

Particle physics is the branch of physics that studies the elementary particles that makeup everything that surrounds us such as protons, electrons, quarks, muons, neutrinos and the four forces that govern the interactions between these particles:

- Strong interaction between quarks by gluons;
- Weak interaction;
- Electromagnetic interaction between charged particles;
- Gravitation interaction between everything that has mass (visible only on macroscopic scales).

There are typically two approaches to study particles.

Particles can be studied capturing the ones transported to us by the cosmic rays, nevertheless, most of them have a very short lifetime before they disappear many of them have very low stability and other particles are rarer to be seen in nature.

Particles can also be created and studied artificially thanks to particle accelerators. In this case, they are generated and observed in a closed and more controlled environment than in nature. Accelerators purpose is to provide energy to charged particles to speed them up to $99.9999991\ \%$ of the speed of light (in the Large Hadron Collider (LHC)). When they collide with each other or when they hit a fixed target, their own energy is transformed into new particles (new matter) and vice versa (in accordance with the theory of relativity by Albert Einstein). Different kinds of detectors are placed where these collisions happen. The detectors are used to measure the number of particles, their charge and their mass with the aim to first identify them, and then study their behavior.

1.2 Synchrotrons

Synchrotrons [18] are a specific branch of particle accelerator, in particular, they are circular machines. The main components of a synchrotron are the RF cavities and the magnets and they are disposed in a circular geometry called *ring*

The Radio Frequency (RF) cavities are used to accelerate (or decelerate) the charged particles providing energy to the beam thanks to an electromagnetic field. Magnets are used to steer the beam and keep it tight during its orbit. The particle beam does not circulate in the air but it circulates inside a vacuum pipe, this minimize as much as possible any interaction with spurious surrounding particles present in the air.

A synchrotron however is not composed only of RF cavities and magnets but there are many other devices installed all around the ring to ensure a proper beam quality, the injection, the extraction, the machine protection, and for diagnostic purposes.

Charged particles such as the ones used in synchrotrons respond to the Lorentz law:

$$\boldsymbol{F}(t) = q(\boldsymbol{E}(t) + \boldsymbol{v}(t) \times \boldsymbol{B}(t)), \tag{1.1}$$

where $\boldsymbol{F}$ is the force applied to the particles, t is the time, q is the electric charge, $\boldsymbol{E}$ is the electric field seen by the particles , $\boldsymbol{v}$ is the particle velocity and $\boldsymbol{B}$ is the magnetic flux density seen by the beam refereed later in this work as the magnetic field for simplicity. The electrical and the magnetic field inside synchrotrons must be controlled synchronously to maintain a good quality beam in a stable closed orbit inside the *ring*. In principle looking at Eq.1.1 it seems that the electric field $\boldsymbol{E}$ can be used both to accelerate and steer the beam, but for practical reasons in synchrotrons, the electric field is used only to accelerate the charged particles. There are different kinds of particle accelerators architectures, such as Radio Frequency Quadrupole (RFQ) in Linear Accelerator (LINAC)s and electrostatic septum in the beam injection-extraction areas, in which the electric field is used also to bend the beam. Eq.1.1, states that the force produced by the magnetic field is always orthogonal to $\boldsymbol{v}$ and $\boldsymbol{B}$ due to the vector product, this avoid any longitudinal acceleration by the magnetic field. For a circular orbit, the transverse forces in the act are the centrifugal force produced by the fact that the particles are accelerated on a circular orbit and the force applied to the beam by the magnetic field that has the duty to compensate the centrifugal one in order to maintain a stable closed orbit.

1.3 CERN accelerator complex

The accelerator complex at CERN, represented in Fig.1.1, includes eight accelerators (machines) that accelerate particles in order to increase their energy.

Each machine in the chain increases the energy of the particles by a factor typically around 20 limited by non-linearities, before delivering them to the next more powerful accelerator or to experiments. There are four main experiments

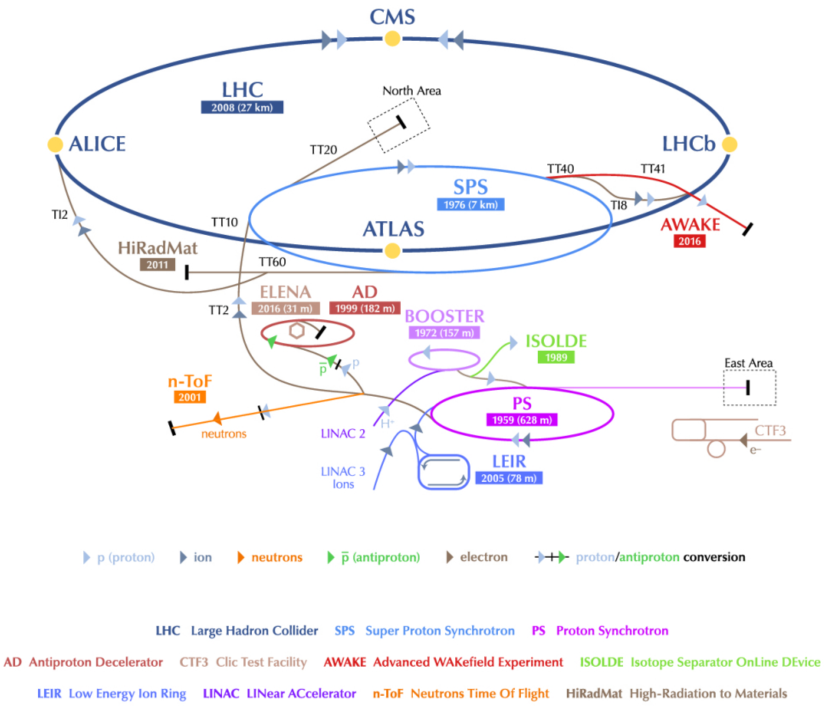

FIGURE 1.1: CERN accelerator complex.

in the LHC (CMS, ATLAS, LHCb and ALICE) and dozen of smaller experiments (LHC or fixed-target); each of them studies particle collisions from a different aspect and with different technologies.

A sophisticated timing system is necessary to synchronize the accelerators. The machines work with cycles of different nature that go to different users. The cycles are organized in supercycles that are repeated during the day, so a supercycle can be seen as a cycle of cycles. Moreover, it has to be considered that each accelerator requires a different time to produce a beam; so the schedule of the beams is a task as complex as essential to operate the entire chain correctly.

In the following paragraphs a quick description of the main accelerators is presented.

1.3.1 Linear Accelerator 2

LINAC2 (Fig.1.2) is the starting point for the protons used in experiments at CERN. Linear accelerators use radiofrequency cavity positive and negative charged alternately. The protons pass through the cavity and they see the electrical field only when it has the right polarity to accelerate them. Small quadrupole magnets (focusing and defocusing) ensure that the protons remain in a tight beam.

The proton source is a bottle of hydrogen gas at one end of LINAC2. The hydrogen is passed through an electric field to strip off its electrons, leaving only protons to enter the accelerator. By the time they reach the other end, the protons have reached an energy of 50 MeV. Then they enter in the Proton Synchrotron Booster (PSB), the next step in CERN's accelerator chain, which takes them to higher energy. LINAC2 started up in 1978 when it replaced LINAC1. It was originally built to allow higher intensity beams for the accelerators that follow it in CERN's accelerator complex. LINAC2 is being replaced by LINAC 4 in 2020-2022.

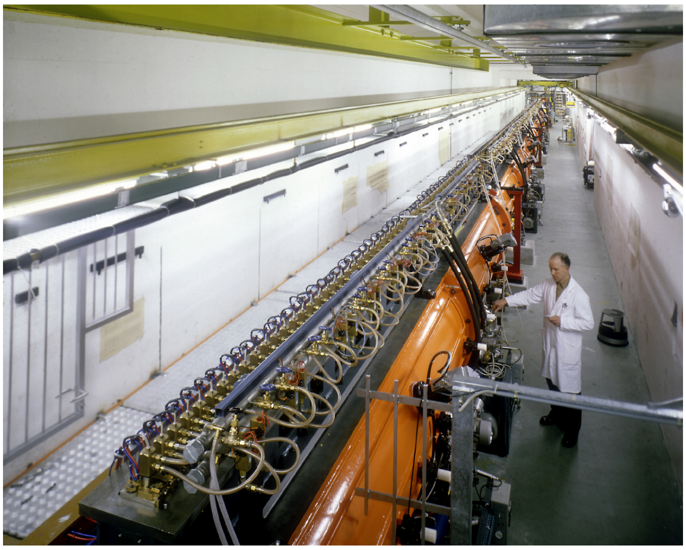

FIGURE 1.2: LINAC2.

1.3.2 Low Energy Iron Ring

Low Energy Ion Ring (LEIR) is the second step in the ion accelerator chain. It receives long pulses of lead ions from Linear accelerator 3 (Linac 3) and transforms them into the short, dense bunches perfect for injection to LHC [19]. LEIR has only four bending dipoles, with a very strong $90°$ curvature, reaching up to 1.15 T (Fig.1.3).

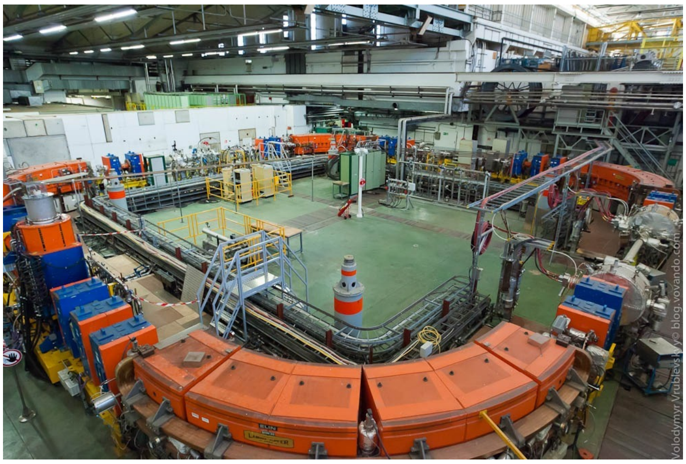

FIGURE 1.3: Low Energy Iron Ring.

1.3.3 Proton Synchrotron Booster

ThePSB (Fig1.4) is made up of four synchrotron rings, one above the other . The Booster receives beams of protons from the linear accelerator Linac 2 at 50 MeV and accelerates them to 1.4 GeV for injection into the Proton Synchrotron (PS). Be-

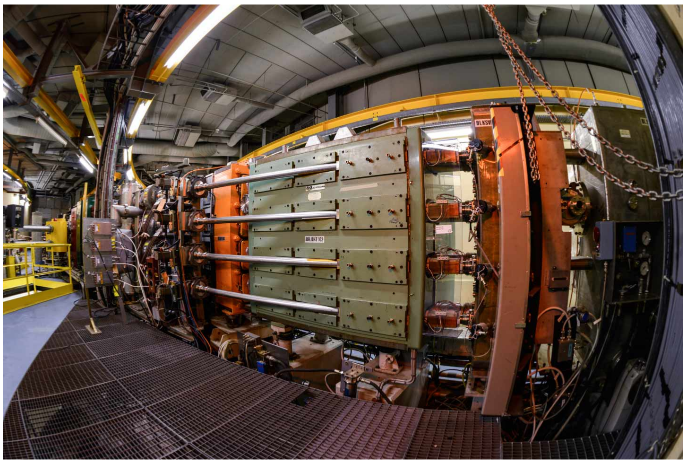

FIGURE 1.4: Proton Synchrotron Booster.

fore the PSB received its first beams protons were injected directly from the linac into the PS, where they were accelerated to 26 GeV. The field level from injection to extraction goes from 0.125 T to 0.861 T, while the duration of the cycle is 1.2 s. There is a limited variety of cycles, which helps the magnetic reproducibility of the machine.

1.3.4 Proton Synchrotron

Proton Synchrotron (PS) typically accelerates either protons delivered by the PSB or heavy ions from the LEIR. In the course of its history, it was used both to fill more powerful accelerators and to perform fixed target experiments The PS (Fig.1.5) was initially CERN's flagship accelerator, but when the laboratory built new accelerators the PS's main role became to supply the beam to the newer machines.

The accelerator operates at up to 25 GeV. The field level from injection to extraction goes from 0.1 T to 1.26 T, while the duration of the cycle can be a multiple of 1.2 s (1.2 s, 2.4 s, 3.6 s).

1.3.5 Super Proton Synchrotron

Super Proton Synchrotron (SPS) (Fig.1.6) is the second-largest machine in CERN's accelerator complex. It's fed by the PS and accelerates the particles to provide beams for the LHC, and to other experiments. The SPS accelerates particles beam up to 450 GeV. It has 1317 conventional (room-temperature) electromagnets, including 744 dipoles to bend the beams around the ring. The accelerator has handled many different kinds of particles: sulfur and oxygen nuclei, electrons,

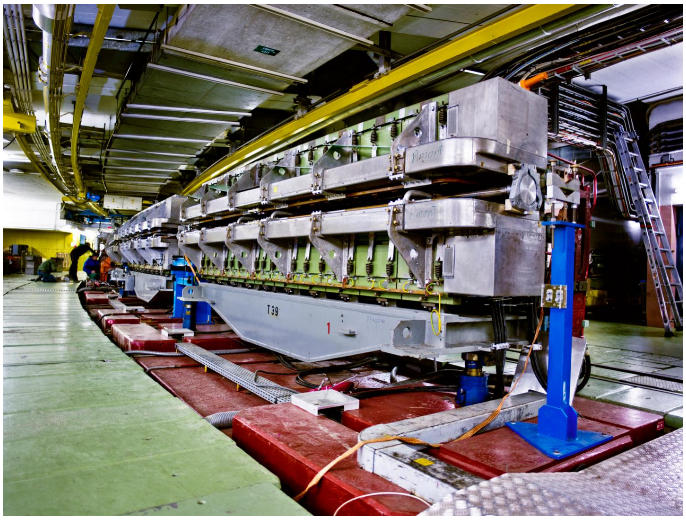

FIGURE 1.5: Proton Synchrotron.

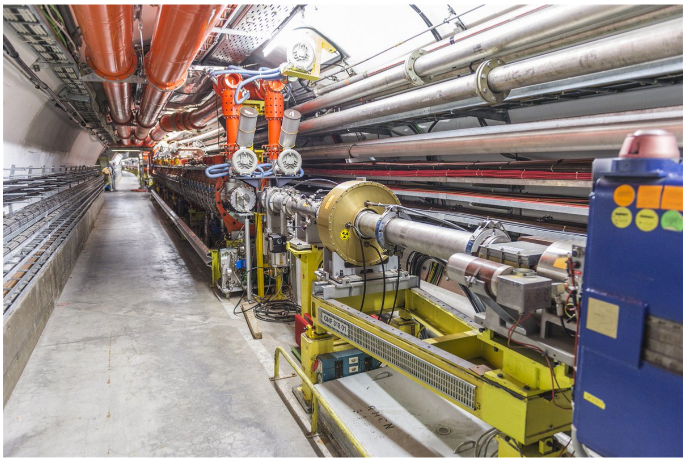

FIGURE 1.6: Super Proton Synchrotron.

positrons, protons and antiprotons. The field in SPS dipoles reaches a high level of 2.02 T, while the cycles are widely different according to the experiments and can be as long as around 30 s in case of slow extraction to certain fixed-target experiments that require a steady particle influx, rather than the more common discontinuous bunches. This process is based on the purposeful excitation of resonant instabilities in the beam.

1.3.6 Large Hadron Collider

LHC is the world's largest and most powerful particle accelerator. It first started up on 10 September 2008 and remains the most powerful machine present in the CERN's accelerator complex. The LHC consists of a 27-kilometer ring of superconducting magnets with a number of accelerating structures to boost the energy of the particles along the way.(Fig.1.7)

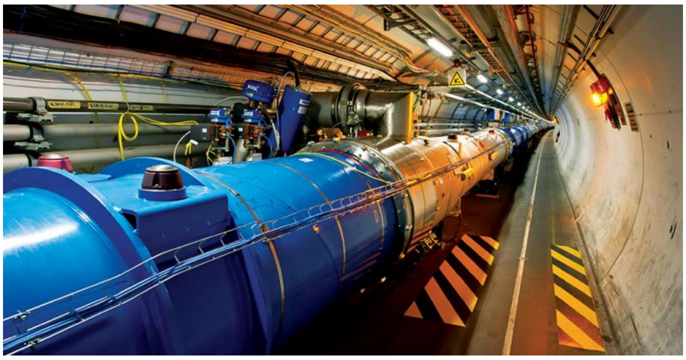

FIGURE 1.7: Large Hadron Collider.

Inside the accelerator, two high-energy particle beams travel at close to the speed of light before they are made to collide. The beams travel in opposite directions in separate beam pipes, two tubes kept at ultrahigh vacuum. They are guided around the accelerator ring by a strong magnetic field maintained by superconducting electromagnets. Much of the accelerator is connected to a distribution system of liquid helium, which cools the magnets, as well as to other supply services. Thousands of magnets of different varieties and sizes are used to direct the beams around the accelerator. These include 1232 dipole magnets 15 meters in length which bend the beams, and 392 quadrupole magnets, each 5–7 meters long, which focus the beams [20].

1.3.7 Antiproton Decelerator

In this accelerators complex, there are also the Antiproton Decelerator (AD) and the Extra Low Energy Antiproton (ELENA), which produce low-energy antiprotons for studies of antimatter, and the Online Isotope Mass Separator (ISOLDE) facility, which is used to produce and study unstable nuclei. The AD (Fig.1.8) is the only machine of its kind. The AD sends the particle that provides to the different experiments around the CERN. The AD's ring is composed of bending and

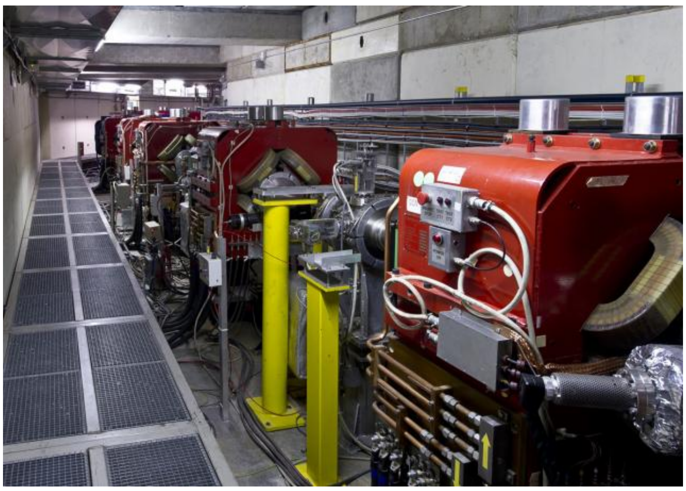

FIGURE 1.8: Aniproton Decelerator.

focussing and defocusing magnets that keep the antiprotons as close as possible, while strong electric fields slow them down to get them usable.

1.3.8 Extra Low Energy Antiproton

ELENA (Fig.1.9) is a compact hexagonal deceleration ring for cooling and further deceleration of 5.3 MeV antiprotons delivered by the CERN AD to an energy of 0.1 MeV. It is based on conventional electromagnets. The field is very low, from 0.05 T to 0.42 T, while the cycle duration is very long, up to two minutes.

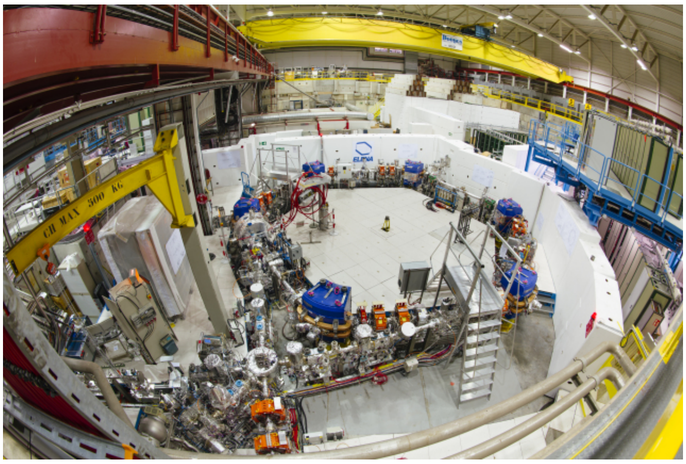

FIGURE 1.9: Extra Low Energy Antiproton.

1.4 Magnets

The charged particle beam is steered by the magnetic field produced by the magnets. There are five types of magnets used in particle accelerators:

- Permanent magnet; the magnetic field is produced by hard ferromagnetic materials.
- Iron-dominated electromagnet; the magnetic field is induced by excitation coils typically made of copper powered by an excitation current. They are called iron-dominated because the magnetic flux is guided by the magnet yoke composed of a soft ferromagnetic material. The saturation field is determined by the yoke.
- Coil-dominated electromagnet; the magnetic field is induced by excitation coils powered by an excitation current. Coil-dominated magnets' coils are generally made of superconducting materials to produce a very high DC field in a large aperture.
- hybrid magnet [21]; the main field is produced by permanent magnets and a correction field produced by an iron dominated electromagnet is added to it

- superferric magnets; composed of a ferromagnetic yoke wounded with superconducting coils [22].

Different magnets' designs exist to produce different magnetic field profiles and each magnetic field profile is used for a precise scope in particle accelerators. The main magnetic field profile used are:

- Dipole field; used to bend the beam and keep it into the desired orbit.
- Quadrupole field; used to *tune* the beam that means focus and defocus the particle beam to keep it tight [23].
- Sextupole field; used to correct the chromatic aberration called also *chromaticity* [24].
- Higher order magnets; used to correct the undesired multipolar fields produced by dipoles and quadrupoles.

In order to take into account all the multipoles of a magnet for particle accelerators the magnetic flux density is expressed as Fourier series expansion,

$$\boldsymbol{B}(\boldsymbol{z}) = B_y + iB_x = \sum_{n=1}^{\infty} (B_n + iA_n)\left(\frac{\boldsymbol{z}}{R_r}\right)^{n-1} \tag{1.2}$$

where $\boldsymbol{B}$ is the field integral parallel to the beam axis, $\boldsymbol{z} = x + iy$, n is the mulipole order (n = 1 dipole, n = 2 quadrupole, etc...), B_n is the normal magnetic field component, A_n is the skew magnetic field component, and R_r is the reference radius of the magnet.

Chapter 2

State of the Art

In this chapter, a recall of the formalism used in magnetic measurements is presented first, which will be largely used in the thesis. Later the topics magnetic field prediction, artificial intelligence, machine learning,and the White Rabbit transmission are presented. In the thesis, a particular focus is devoted to the real-time magnetic field measurements and on the magnetic field prediction using machine learning techniques.

2.1 Magnetic measurements for magnets

Magnetic measurements in magnets for particle accelerators are performed in two scenarios: offline, after the magnet manufacturing and before the magnet installation as part of the quality assurance process; online and in real-time during the machine operations, as necessary feedback for other accelerator's subsystems.

The main kinds of magnetic measurement carried out offline and required to certify the quality of a magnet for particle accelerators are the following:

- The absolute value of the main field component at the peak current, representing the strength of the magnet.
- The field homogeneity inside the magnet aperture.
- The field direction of the main magnet's component.
- the magnetic center for quadrupoles and higher order magnets.
- The eddy current strength and decay time, necessary for pulsed magnets.
- 3D field map for beam tracking.

The real-time magnetic measurements carried out online are necessary since for iron-dominated magnets the field cannot be predicted with sufficient accuracy from mathematical models for magnets operations. Therefore in the CERN synchrotrons: LEIR, PSB, PS, ELENA, and the SPS the real-time feedback becomes necessary.

The magnetic field measured in real-time is generally distributed to three users:

- The Low Level Radio Frequency (LLRF) control system.

- The beam current transformer control system.
- The power converter control system.

The revolution frequency $f_{rev}(t)$ provided to the cavities from the LLRF is calculated from the main bending magnetic field and it is given by

$$f_{rev}(t) = \frac{c}{2\pi R}\sqrt{1 - \frac{1}{1 + \left(\frac{B(t)\rho q}{m_0 c}\right)^2}}, \tag{2.1}$$

where R is the mean orbit radius, $B(t)$is the magnetic field seen by the beam,q is the electric charge of the beam, and m_0 is the particle's rest mass. The revolution frequency is controlled via additional correction and feedback terms that are detailed in the literature [25, 26, 27].

The beam current transformer measures the beam equivalent current, *i.e.* the amount of charge transported by the beam in the time unit [28, 29] to estimate the number of particles that compose the beam $N_p(t)$ given by

$$N_p(t) = \frac{2\pi R}{qc} I_{BCT}(t) \frac{\sqrt{\left(B(t)\right)^2 + \left(\frac{m_0 c}{\rho q}\right)^2}}{B(t)}, \tag{2.2}$$

where $I_{BCT}(t)$ is the measured equivalent current expressed in Ampere. The correction of the particle momentum by the magnetic field is especially relevant for non-relativistic particles circulating in theLEIR, PSB, PS,SPS and ELENA synchrotrons (the AD is a special case, it can work without measurements with a simulated field produced in real-time). Particles circulating in these accelerators are below the relativistic gamma transition.

The power converter controller can use magnetic measurements to regulate the current provided to the main bending magnets. This regulation in the field gives as an advantage the fact that most of the eddy currents, magnetic saturation and hysteresis effects are automatically corrected by the controller. The same does not happen if the regulation of the power converter is done in current. The power converter is voltage controlled by an R-S-T feed-forward regulator with feedback [30, 31].

2.2 Out-of-date B-train system

The Synchrotrons at CERN employ systems so-called *B-train* for determining the dipole field in real-time. The name derives from the discrete positive and negative pulse trains used to distribute incrementally the measured field in out-of-date systems, developed as far back as the 1950s [32].

The out-of-date digital transmission (dating from the early 1960s) uses two coaxial cables to distribute 24 V pulses which indicate a ± 0.1 Gauss increase or decrease of magnetic field, i.e. up and down pulses. These pulses are distributed from the reference magnet to several client applications (Fig.2.1) as described previously.

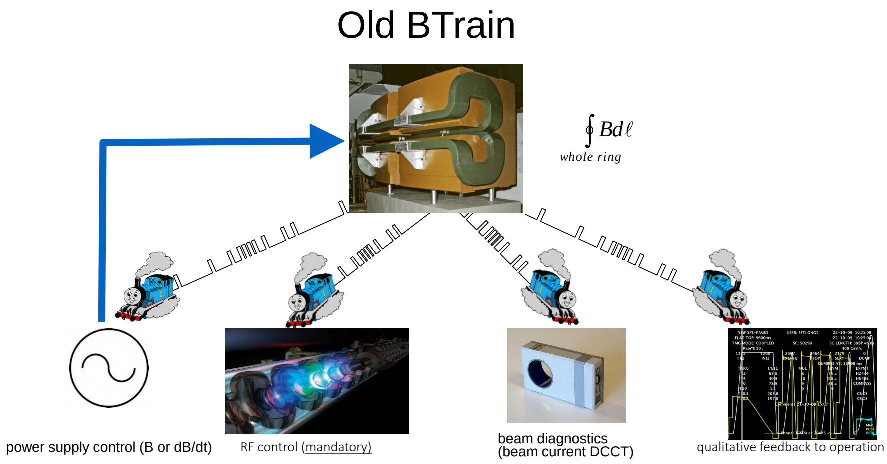

FIGURE 2.1: Old B-train distribution representation.

The B-train measures the average field of a reference magnet and distributes the result in real-time to various other synchrotron sub-systems, as part of a feedback control loop or for diagnostic purposes. The typical range of magnetic field measured goes from about $50\,\mathrm{mT}$ to $2\,\mathrm{T}$. The most critical user of the B-train is the Radio Frequency (RF) subsystem, which uses RF cavities to generate the electric field that accelerates or decelerates the beam. The instantaneous magnetic field must be known with high precision to lock the RF frequency to the particle energy, keeping the beam centered in the vacuum chamber. While the pass-band of the bending magnets does not usually exceed a few hundred hertz, the RF system requires a much faster data rate to ensure smooth feedback, e.g. $250\,\mathrm{kHz}$ in the PS. The use of B-train systems is not unique to CERN: for instance, similar designs are implemented at ion therapy centeres such as the National Centre of Oncological Hadrontherapy (CNAO) [33], MedAustron [34] and the Heidelberg Ion-Beam Therapy Centre (HIT) [35]. For such applications, real-time feedback control of the magnetic field is instrumental, for example, to reduce dead times that would be otherwise spent pre-cycling the magnets to improve their reproducibility.
The out-of-date system has a 0.1 Gauss resolution with long-term stability and reproducibility on the 10^{-4} level. The key specifications satisfy the current request of all users, however, its maintenance is becoming complicated since several components are obsolete and some parameters remain unknown. These are mostly dominated by hardware faults, timing conflicts and electromagnetic interference and have generated 3 to 4 maintenance calls/year in average for all the machines. Moreover, non-desired effects as the drift of the converter or the distribution of negative B-field values at the end of several cycles justified the decision of replacing the 25 years old present B-Train system for a new one. Currently, all

six B-train systems in operation are being upgraded in the frame of a long-term complex-wide consolidation project.

2.3 Magnetic sensors

In this section, an outline of the physical effects behind the magnetic measurement and the most used magnetic sensors and their application are presented.

2.3.1 Induction coil sensor

A generic conductive spire in a magnetic field works according to Faraday's law

$$V_c = \frac{d\Phi}{dt}, \tag{2.3}$$

where V_c is the induced electromotive force, and $\frac{d\Phi}{dt}$ is the magnetic flux variation over time.
The variation of magnetic flux across a winding induces an open-loop voltage to the ends of the winding. The same happens even if instead of a single spire a coil composed by multiple spires is considered and in this case the induced voltage is proportional to the rounds' number. The coil so realized is also called induction coil. Induction coils for accelerator magnets typically consist of multiple loops of single- or multi-filament strands wound around a long, rectangular core [36], where the polarity of the coil's output is chosen in such a way as to be consistent with the sign of the magnetic field. The induction coil provides the dynamic component of the field with intrinsically high linearity and bandwidth. The magnetic flux can change for different reasons: change of magnetic field strength, the relative longitudinal displacement between coil and magnet and coil rotation. Integrating the coil voltage allows obtaining the average magnetic flux from which is possible to obtain the magnetic field $B(t)$ knowing the area of the coil as

$$B(t) = B_0 + \frac{1}{A_c} \int_{t_0}^{t} (V_c(\tau))\, d\tau, \tag{2.4}$$

where B_0 is the magnetic field at the beginning of the integration process, A_c is the effective area of the coil and V_c is the induced voltage. This approach to measure the magnetic field is known as fluxmeter method. The induction coil based magnetic measurement devices such as rotating coils [37, 38], single stretch wire technique [39, 40],static fluxmeter [41], and translating fluxmeter [42]are the most used for their adaptability and for their manufacturability.

Since the main quantity of interest for beam control is the longitudinal integral of the magnetic field, the length of the coil is ideally that of the iron yoke, plus the whole fringe field region (typically, 3 to 4 times the gap height at each end) [43]. In some of the systems at CERN, where this is not possible due to space constraints, a much shorter coil is used instead; this is the case in the PS and LEIR, where the coils used are respectively about $0.3\,\mathrm{m}$ and $0.1\,\mathrm{m}$ long, compared with iron yoke

lengths of approximately $5\,\mathrm{m}$. In such a case, the average field is assumed to be proportional to the local field seen by the coil. The main source of error affecting this technique is given by low-frequency variations of the voltage offset, which is due to various causes including thermocouple voltages along with the cabling and connections, electronic component imbalance and rectification of electrical noise (*i.e.* 1/f noise). This voltage offsets cause, if not corrected, a drift in the integration. Many techniques were studied to compensate and to mitigate this drift [1, 4, 5, 44].
The effective area of the coil can be calibrated within a typical uncertainty of $100\,\mathrm{ppm}$ through the classical flip-coil method, *i.e.* by immersing it in a sufficiently uniform magnetic field B perpendicular to its surface, turning it over by $180°$ and measuring the flux change $\Delta\Phi$, corresponding to twice the magnetic flux through the coil $\Phi = A_\mathrm{c}B$. Alternatively, the flux measured by pulsing the field when the coil is kept fixed can be compared to another reference method, such as a single stretched wire [36]. It is important to stress that, for synchrotron control applications, the absolute calibration of the measurement is of secondary importance with respect to its reproducibility. In fact, during the accelerator setup phase, the relationship between bending field and RF frequency is always adjusted to compensate many sources of systematic errors, including possible errors in the magnetic field measurement itself, provided these remain reasonably small [45].

2.3.2 Nuclear magnetic resonance sensors

Nuclear Magnetic Resonance (NMR), experimentally demonstrated in 1945, is the most accurate method for measuring magnetic field strength, able to reach an accuracy of a few ppm [46, 47]. NMR magnetometers are often used as a reference for calibration due to their high accuracy. The operating principle of an NMR magnetometer is the following: NMR is a physical phenomenon based on the fact that the protons and neutrons possess a magnetic dipole moment, they are equivalent to microscopic bar magnets. This moment, which is positive for protons and negative for neutrons, is analogous to the magnetic moment created by an electric charge rotating around its axis and is described in quantum mechanics by the property called spin. Resonance can occur when the nuclei of certain atoms are immersed in a static magnetic field (B_0) and are exposed to a second oscillating magnetic field B_1 [48, 49]. Some atoms experience this phenomenon, while others never experience it and this depends on the net moment of the nucleus, which is the algebraic sum of the moments of the constituent protons and neutrons (a non-zero moment requires an odd number of protons, or neutrons, or both) Fig.2.2 shows a schematic representation of the NMR working principle. Among the many elements with spin, the most commonly measured nuclei are hydrogen-1 (also chosen as the element with which to produce magnetic resonance imaging since it is the simplest and most abundant element in the human body) and carbon-13.

If subjected to a strong stationary external magnetic field B_0, the proton axis will orient itself along the field itself (previously they were just oriented randomly). This orientation can take place either in the same direction as B_0 (i.e with

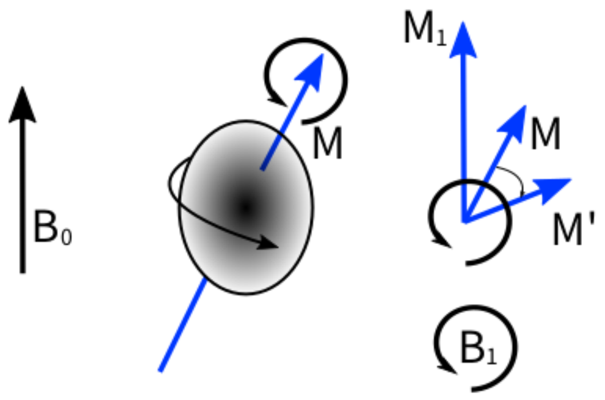

FIGURE 2.2: NMR working principle.

a low energy level) or in the opposite direction (*i.e.* with a high energy level). The difference between these two energy levels depends only on the magnetic moment and the magnetic field. The nuclear magnetic moment is a constant, thus this energy gap is a precise measure of magnetic field strength. The parallel protons are slightly more prevalent than the antiparallel ones. This small prevalence produces a resulting magnetization M, parallel to B_0 and it is measurable. In the classical view, which describes the phenomenon to a very good degree of approximation, due to the effect of B_0, the axis of each proton rotates around the direction of the moment of B_0 (precession). The precession frequency is characteristic of every atomic element and is called Larmor frequency [48]. The resonant frequency is given by the following relation:

$$\nu_{Larmor} = \gamma \cdot B_0, \tag{2.5}$$

where ν_{Larmor} is the Larmor frequency and γ is the gyromagnetic ratio, that is the ratio between its magnetic moment and its angular momentum. The gyromagnetic ratio for an isolated proton is 42.57747892 MHz/T. Consequently, resonant frequencies for a particular substance are directly proportional to the strength of the applied magnetic field. The gyromagnetic ratio is known from fundamental constants with an uncertainty lower than 10^{-8}, which is the reason behind the accuracy of the method. At resonance, the nuclei absorb and release energy while flipping back and forward between the two opposite spin states. The disequilibrium between the two populations gives rise to a measurable signal [47].

2.3.3 Ferrimagnetic Resonance sensors

The working principle of Ferrimagnetic Resonance (FMR) is similar to NMR, FMR instead of protons and neutrons is based on the electron spin, however ferrimagnetic (*i.e.* two unequal population of atoms with anti-parallel magnetic moments) or ferromagnetic (*i.e.* one population of atoms with parallel magnetic moments) materials introduce anisotropy terms into 2.5. Magnetic resonance cannot be established accurately [50] when the material is not fully saturated, therefore the FMR must be operated above this threshold for magnetic measurements. In the FMR sensors gyromagnetic ratio is strongly affected by the temperature and chemical composition of the sample, this is why metrological performance is lower than NMR. The problem with this type of sensor is that it can only function as a dynamic marker since a DC implementation is not yet commercially available. FMR is also widely used in commercial RF devices such as circulators and insulators [51, 52], as well as in tunable filters [53, 54], resonators and oscillators in the spectrum and vector network analyzers. FMR has the relevant advantages of being highly selective (high quality factor, low insertion and return losses) and tunable. The use of FMR in magnetometers is reported in Ref. [55].

2.3.4 Hall effect magnetic sensor

The Hall effect in semiconductors is a particular phenomenon that characterizes the interaction between charged moving particles and magnetic fields. Given a semiconductor strip (type N or type P) and a current flowing into it as described in Fig.2.3, if a magnetic field orthogonal to the strip is applied the moving charges (electrons or holes) are affected by the Lorentz force (Eq.1.1) that drives the moving charges to one end of the strip. This accumulation of charge leads to a charge imbalance that causes a difference of potentials between the two ends of the strip. This voltage can be measured and it is proportional to the strength of the applied magnetic field and in a first order approximation it is correct to assume that

$$V_H = hIBsin(\alpha), \tag{2.6}$$

where V_H is the voltage at the end of the semiconductor strip, I is the current, α is the angle between the magnetic field vector **B** and the Hall plate and its the overall sensitivity which depends on the material, the geometry, and the temperature. This effect can be used to measure the strength of a magnetic field and realize inexpensive Hall effect sensors in Complementary Metal Oxide Structure (CMOS) technology implanting an N well inside substrate P. The main drawback is that this kind of sensor is very temperature sensitive.

2.4 Real-time magnetic field predictions

As aforementioned in sec. 2.1 an accurate knowledge of the magnetic field $B(t)$ at any given time during a magnetic cycle is therefore critical for longitudinal and transversal beam control, power supply control, various beam diagnostics, and qualitative feedback to operators. The required accuracy is typically $0.01\ \%$ [6]. In a restricted number of cases, conventional mathematical models can express the $B(I)$ relationship adequately well. An example is the semi-empirical model of the superconducting bending dipoles of the LHC, which generate a very high field (8.4 T), relatively unaffected by perturbations [56]. In the vastly more common case of iron-dominated magnets, effects due to magnetic saturation, hysteresis,

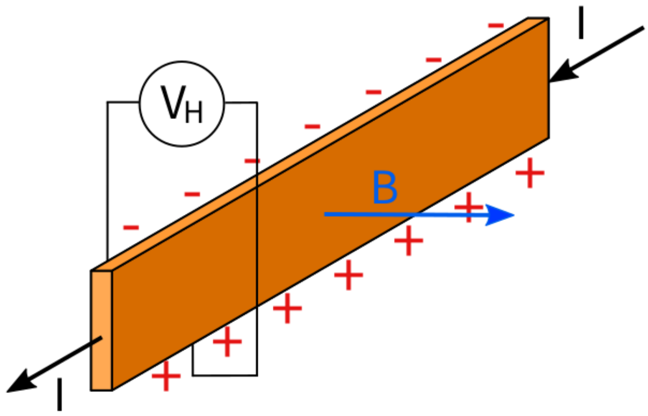

FIGURE 2.3: Hall Effect, B is the applied magnetic field , I is the current and V_H is the voltage at the end of the semiconductor strip.

and eddy currents may be as large as several percent or more, and the problem becomes orders of magnitude more difficult[57, 58]. For example, recent attempts using the well-known Preisach models[14, 15] could not attain better than $0.2\ \%$ accuracy. Also other classes of methods, such as Jiles-Atherton differential models[16], ultimately turn out to be unsuitable, due to their well-known difficulties in handling minor hysteresis loops. As an alternative, real-time measurements can sometimes be carried out in a suitably equipped reference magnet, either a part of the accelerator ring or powered in series with it. At CERN, six of the synchrotrons function thanks to feedback from such measurements. In general, however, this kind of real-time measurement system is complex, expensive, and sometimes very impractical to deploy, for example owing to the lack of space for sensors close to the beam vacuum chamber. As a result, there is a strong incentive to investigate novel kinds of models to complement or even replace measurements. In addition, even where real-time measurement systems are already implemented, such models may serve as a useful complementary role, for example during periods of hardware maintenance. The typical approach followed when real-time feedback is required and when no measurement is available is to use a nominal magnetic cycle corresponding to the excitation current saved in a database. The drawback of this approach consists in the fact that these cycles are fixed and they do not take into account all the non-linearity of the magnet such as saturation, eddy current and hysteresis. Modelling of quasi-static and dynamic hysteresis loops is one of the most challenging topics in computational magnetism, mainly due to the strong non linearity and history dependency shown by ferromagnetic materials[7, 8, 9, 10, 11, 12, 13]. This general problem is commonly addressed in literature in the context of electrical machines, which are excited by sinusoidal current waveforms[59, 60, 61]. Conversely, more complex excitation current waveforms $I(t)$, used in particle accelerators and other magnetic devices operating cyclically, are still an open focus of scientific interest. Such waveforms include quasi-periodic sequences of trapezoidal-shaped pulses, with widely varying slopes (i.e., current or, equivalently, field ramp rates) and flat-top levels. Flat-tops and flat-bottoms correspond to reversal points of the hysteresis loop and their sequence largely determines the relationship $B(I)$ between current and magnetic field. Under these conditions, $B(I)$ becomes much more complex and hard to predict[57, 58].

2.5 Machine learning

Recently, more attention was directed towards Artificial Neural Network (ANN), today used with spectacular results in a variety of domains related to time-series prediction[62, 63, 64, 65, 66, 67, 68, 69, 70, 71], but still relatively unexplored in magnetic applications. In Ref. [72], a hybrid Preisach-Neural Network model is proposed to predict the dynamic hysteresis in ARMCO pure iron, reaching a Normalized Root Mean Square Error (NRMSE) of the order of 0.7 %. ANN techniques are being used more and more often to model magnetic hysteresis in combination with classical approaches like Preisach, Wlodarski, Chua-Stronsmoe and Jiles-Atherton models[59, 61, 73]. In Ref. [59], one of the first attempts at describing the memory mechanism is proposed in systems with rate-independent hysteresis using a combination of Preisach state updating rules and a Feed-Forward Neural Network (FFNN). The architecture consists of two blocks: a Preisach memory and an FFNN bounded to the memoryless relation between the state and the output. The model is identified by tuning the weights of the network architecture through a back-propagation based algorithm.
Developments following the approach in Ref. [59] are also reported in Refs. [60], [74], [75]. In Refs. [76, 77], a different hybrid perspective is proposed by combining an ANN with a Fourier Descriptor (FD) to evaluate dynamic hysteresis loops. The method is suitable when a distorted sinusoidal magnetic field H excites, in steady state, the ferromagnetic core of a device and allows to handle problems that appear in classical approaches when static hysteresis, eddy currents, and anomalous losses should be considered.
An interesting modelling approach is reported in Ref. [78]: an Neural Network (NN) approach for modelling dynamic hysteresis is proposed by combining an array of NNs where each NN is dedicated to a particular fixed portion of the dynamic hysteresis loop. The whole hysteretic path is built by the composition of the evaluations made by different NNs. Moreover, the authors use excitation curves as a sinusoidal waveform.
In Ref. [61], the authors study the hysteresis behavior of a transformer core, determined by a reduced number of real-time measurements of the input currents and the output voltages. The hysteresis loop is computed through a Deep Neural Network (DNN) combined with the Wlodarski magnetic model. The hysteresis loop is then used in a Finite Element (FE)) model to simulate the response of the core to arbitrary excitation waveforms. These hybrid architectures combine a classical approach to model the physics of the problem, together with the ANN prediction of the short-term nonlinear behavior. Their focus is on the prediction of a single hysteresis cycle, and in general, they are not able to take into account the magnetic field response over a long period. Moreover, even in the short term

case, their accuracy does not meet the above accuracy requirements of $0.01\ \%$[15, 59, 61, 76].

2.5.1 Multi-layered non-linear autoregressive exogenous neural network

Since the proof of universal approximation for Feed forward [79] and recurrent [80] neural networks (with the sufficient condition of one hidden layer) the majority of the NN approaches focused on developing networks in width, almost neglecting the benefits of developing layers in depth. However, an astonishing improvement of NN based system performances was achieved when the possibility of expanding layers in depth became computationally tractable thanks to the development of new smart methods for learning and the increased computational power of computer machines (see Ref. [81] for a comprehensive historical review). Further, when dealing with time series, successful dynamic approaches unfolding the depth of the network through time were proposed, like Long-Short Memory Network (LSTM) and variants [82, 83, 84, 85, 86]. In this scenario, an important role is played by the multi-layer Non-Linear Autoregressive Exogenous (NARX) neural network.
NARX is a popular recurrent neural network architecture having feedback coming only from the output neuron instead of from the hidden neurons. This kind of architecture is typically used for input-output modeling of discrete time nonlinear dynamic systems [87].
The architecture of a NARX Network consists of a Multi Layer Perceptron (MLP) network and two buffers as shown in Fig.2.4. The buffers collect the previous value of inputs and outputs of the network to provide them in input to the MLP network. The size of the two buffers are hyperparameters of the network. If the size of the buffer that collects the output is set to 0 then the NARX is reduced to a Time Delayed Neural Network (TDNN), if also the size of the buffer that collects the previous input is set to 0 then what we obtain is an MLP network that it is the core of the NARX network.
NARX network lies at the edge between a static and a dynamic deep network approach. In general, successful NARX based approaches were extensively studied (see, e.g. Refs. [70] and [71]). NARX models allow a sliding window operation across the feed-forward layers, without relying on recurrent connections.

2.6 White rabbit network

In a modern real-time measurement system it is crucial to have the possibility to transmit the measured data with high accuracy in terms of delay and time synchronization. This is even more critical when the system is spread over several square km, for this we considered the White Rabbit (WR) network. The WR network is a bridge local area network based on existing IEEE standards extended in a backward compatible way in order to meet the CERN's requirements. The used

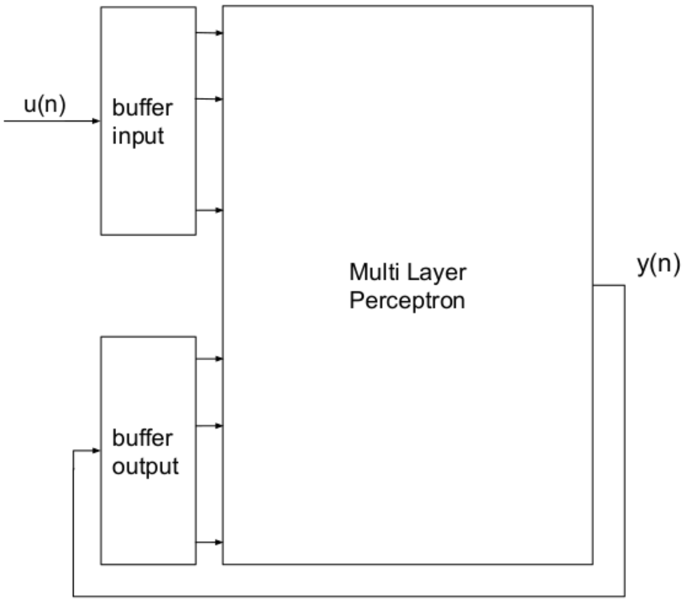

FIGURE 2.4: Basic architecture of a NARX network.

standards are VLANs (IEEE 802.1Q) that use Ethernet (IEEE 802.3) to interconnect switches and nodes as shown in Fig.2.5, and the Precise Time Protocol (PTP) (IEEE 1588-2008) to synchronize them. The main feature of the WR network are:

- Sub-nanosecond accuracy and picosecond precision of synchronization.
- Possibility to connect thousands of nodes.
- Typical distances of 10 km between network elements.
- Gigabit rate of data transfer.
- Fully open hardware, open firmware and open software.
- Hardware commercial availability from many vendors.

2.6.1 History

CERN started to think about a new timing system in 2006 to increase the bidirectionality and the bandwidth to make possible a general timing system for all the accelerators able to auto compensate the cabling delay At the same time, Helmholtz Center for Heavy Ion Research (GSI), start brainstorming about timing system for the Facility for Antiproton and Ion Research (FAIR) facility. Since other collaborations between GSI and CERN were already underway it was natural to try to develop a single timing system that served both sets of requirements. The requirement for a full-duplex high-bandwidth link quickly pushed to Ethernet for the physical layer. Ethernet is not only a high-performance and well studied

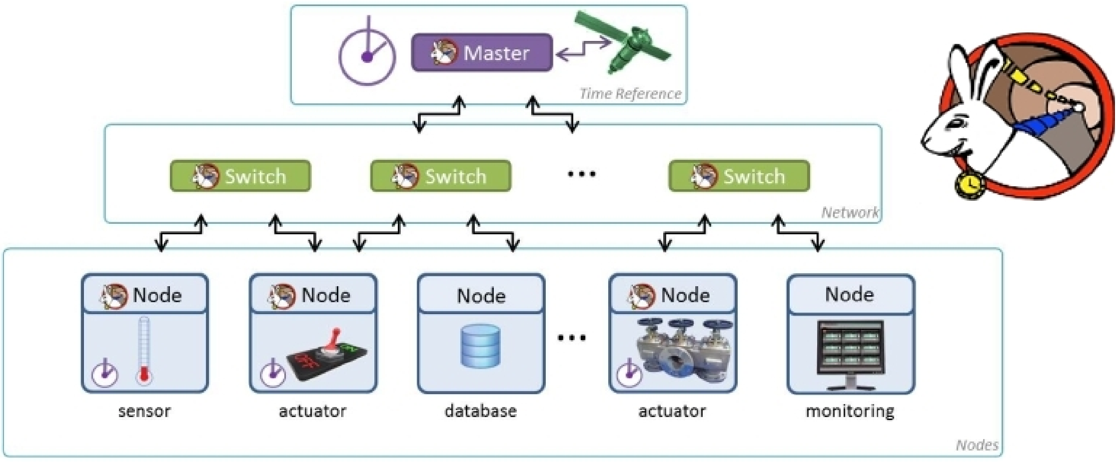

FIGURE 2.5: WR network example.

solution but also long-term support is beyond doubt, and this was important both for CERN and GSI. Synchronous Ethernet defines a clock transmission strategy based on recovering a clock from an Ethernet data stream using a Phase Locked Loop (PLL) [88]. Cabling delay compensation can now be done using the Precise Time Protocol (PTP, IEEE 1588). Mixing these two standards and setting up a strategy to dispatch messages in a deterministic way to all nodes gave birth to the White Rabbit project. The name of the project is referred as the White Rabbit present in Lewis Carroll's novel Alice's Adventures in Wonderland.
Several companies [89] have begun to commercialize White Rabbit for industrial applications by developing their own White Rabbit hardware and software.
The first device on the White Rabbit project was the *white rabbit switch*, financed by the government of Spain and CERN, and produced by Seven Solutions.

2.6.2 Synchronization Scheme

In order to reach the wanted sub-ns accuracy the synchronization of more than 1000 nodes, WR has a timing hierarchy. One of the switch ports is named *uplink* port and all the others are *downlink*. The first switch gets its clock from an external source, *i.e.* a GPS Disciplined Oscillator (GPSDSO) and this clock is used to drive the transmitters inside the *downlink* ports. The *downlink* port are connected or to a final node or to the *uplink* port of another switch leading to a tree of switches in which all the inner clocks are derived from the master's clock.
To compensate the transmission delay we can consider it composed of two components a first one coming from the electronics in the switches and the second one coming from the time of flight in the optical fiber. The first one in a first order approximation can be considered fixed and can be corrected by automatic or manual calibration of the system. The second one can show variations up to 1 ns for fibers lengths about 10 km not buried underground. To solve this problem can be used a two-way scheme like PTP, the drawback is of generating traffic on the network just for synchronization. WR proposes to use instead continuous

measurements of the phase of the bounced-back clocks with respect to the transmit clocks in each one of the switches *downlink* port. A PTP-like exchange can be done initially to figure out a rough estimate of the two-way delay expressed in 125 MHz ticks. From then on, the continuous phase measurement piggybacks on any traffic without perturbing it.

2.6.3 Determinism

Contrary to what happens to the clock, there is no hierarchy for the data traffic. In WR protocol any node can communicate to any node at any time. It is the responsibility of higher layer protocols to keep traffic orderly. For this task, WR provides help in the form of different traffic types in layer 2, with different associated priorities. In order to ensure determinism in the latency of some types of traffic between two nodes, WR specifies different types of traffic. In the event of a High Priority (HP) frame hitting a switch while Standard Priority (SP) frames are waiting for delivery in a pipeline, the HP frame would be output first because of its priority. In order to avoid long SP frames from holding an Ethernet port for too much time, auto-fragmentation of these frames and immediate forwarding of the HP frame is also supported by the WR specification. By automatic fragmentation the SP frame being output would be cut roughly but with a special termination sequence that allows the downstream switch to wait for additional fragments of the SP frame before broadcasting it.

Part II

Measurement system

Chapter 3

System requirements

In this chapter, the reasons behind the novel B-train system called Field In REaltime STreaming from Online Reference Magnets (FIRESTORM) developed to replace the out-of-date systems in the context of a site-wide, long-term consolidation project are presented. Then requirements for the new system, and the methods used are highlighted. FIRESTORM was designed to cope with the High-Luminosity LHC upgrade, which will require higher beam intensity and improved beam control throughout the injector chain [17].

3.1 Measurement goals and method

The main goal of a B-train system is to measure and distribute the average dipole magnetic field, $\bar{B}$, that bends the trajectory of the beam in a synchrotron ring. The magnetic field varies cyclically over time, as illustrated in Fig. 3.1, being proportional to the momentum of the beam particles as they are first injected into the ring, then accelerated and finally ejected. Typical requirements include a measurement uncertainty of $100\ \mathrm{ppm}$ relative to the peak field during a cycle, a bandwidth from DC to $100\ \mathrm{Hz}$, and a maximum latency of $30\ \mu\mathrm{s}$, which is critical especially for the RF subsystem. The measurement is carried out in a suitable reference magnet, which is ideally installed in a dedicated room outside of the synchrotron and is powered in series with the ring magnets. In this case, the absence of a vacuum chamber in the magnet gap leaves the freedom to install sensors along the magnet's longitudinal axis, where the beam is supposed to circulate. When this is not practical, such as in the LEIR bending dipole, sensors must be installed within the accessible fringe field region. Along with $\bar{B}$, the system also distributes the rate of change of the magnetic field, $\dot{\bar{B}}$. This is needed by some machines, such as the PS, which implement multiple excitation circuits that act in parallel on the same magnetic core, in order to compensate induced voltages stemming from inductive coupling effects [90].

3.1.1 Measurement model

The setup for all B-train systems at CERN consists of the combination of two primary sensors [36, 91]: an induction coil to measure the rate of change of the field according to Faraday's law, and a so-called field marker to provide the necessary integration constant B_{m} , according to (3.1-3.3):

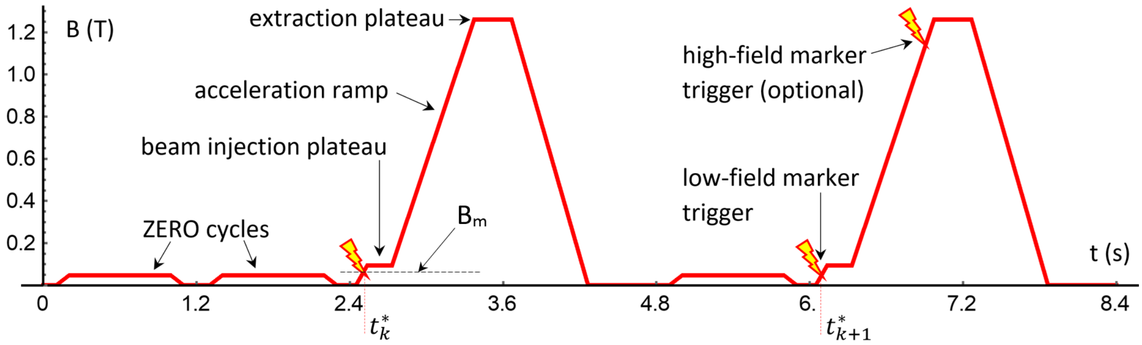

FIGURE 3.1: Schematic example of a sequence of magnetic cycles in the Proton Synchrotron.

$$\Phi = \iint_A N_\mathrm{T} B \, dA = \bar{B} A_\mathrm{c}, \tag{3.1}$$

$$V_\mathrm{c} = -\frac{d\Phi}{dt}, \tag{3.2}$$

$$\bar{B}(t) = B_\mathrm{m} + \Delta\bar{B}(t) = B_\mathrm{m} - \frac{1}{A_\mathrm{c}} \int_{t_k^*}^{t} V_\mathrm{c}(\tau) \, d\tau, \tag{3.3}$$

where Φ is the total magnetic flux linked though the coil, N_T is the number of coil winding turns, A_c is the effective coil area, and V_c is the coil output voltage.

The B-train system must operate uninterrupted over periods that may last for months. The integration process is seamlessly subdivided into a sequence of contiguous integration intervals matching the magnetic cycles $t_k^* \leq t < t_{k+1}^*$, with $k = 1, 2, ...$, where each t_k^* corresponds to a field marker trigger generated when the field crosses the given threshold B_m, practically resetting the process with a new integration constant. By far, the most important error source in (3.3) is the drift of the integral due to a small, but unavoidable voltage offset added to the coil output. The problem of voltage offsets in integrators is well-known not only for induction coils but also in different measurement domains, such as inertial sensor [33, 92, 93]. The offset, which typically ranges in value from a few µV to a few hundred µV, is generated by a number of different mechanisms such as thermoelectric voltages due to temperature gradients along wires; thermocouple voltages at the connections; rectification by non-linear circuit elements of radiated electromagnetic noise; or bias currents due to the imbalance in discrete or integrated components. While some degree of mitigation can be afforded by the thermalization of the whole setup and by careful shielding, the offset can never be eliminated completely [92].

Even if, in principle, the integration constant could be set just once at the beginning of the process, the use of repeated resets effectively prevents the build-up of integrator drift; methods to control the drift within each integration interval are discussed in Section 5.2.3. The value of B_m is often chosen so that a reset happens just before the injection of particles into the synchrotron, when high measurement accuracy is required to capture the incoming beam and preserve its quality. If a reset does occur when the beam is already circulating, B_m must not be assigned to $\bar{B}(t_k)$ abruptly but rather in a gradual manner over a suitable interval, of the order of $10\,\mathrm{ms}$, to prevent any discontinuity that may be harmful. During a magnetic cycle, any given field value B_m will be crossed twice, once on the up-ramp and once on the down-ramp, thus generating two separate field marker triggers. However, in the current implementation of FIRESTORM the field marker signal is gated by a specific time window to avoid spurious noise-induced triggers, thereby generating only one reset trigger per cycle. It is worth noting that, two fully independent markers can be assigned to any integration channel, allowing for the reset timing to be optimised as dictated by the beam quality requirements. A typical application of this feature consists in adding a second marker at high field, just at the beginning of the beam ejection plateau, as depicted in Fig. 3.1.

3.1.2 Field marker

The field marker serves two critical roles in the system: providing the integration constant in (3.3), and periodically resetting the measurement to prevent the accumulation of drift. The marker itself is composed of a magnetic sensor, together with detection electronics (described in Section 5.2.2). It generates a digital trigger pulse whenever the field crosses a pre-set value B_m. As such, this device has inherently a dynamic nature, that is detection can only occur on a field ramp. A variety of different sensors can be used for this purpose. The simplest option is a Hall probe, combined with a voltage comparator. Even though this method is used with success at CNAO [94] and HIT [93], the long-term stability of the offset and gain of Hall probes may be problematic and entail frequent interruptions for recalibration, which are not acceptable for the CERN accelerator chain. The old PS B-train worked satisfactorily over as many as five decades with a so-called peaking strip, described in [95]. However, this marker operates at a very low field of $5\,\mathrm{mT}$ and cannot be scaled easily to higher levels. Instead, the FIRESTORM B-train implements magnetic resonance sensors, which were extensively tested and proven to meet all requirements [96, 97, 98]. These sensors are based on the precession frequency of the elementary magnetic moments (protons or electrons) in a small sample of suitable material. The sample is immersed in a background field B and irradiated with electromagnetic waves at a frequency f, so that at resonance it absorbs and re-emits at the frequency $f = \gamma B$ where γ is the gyromagnetic ratio. In Continuous Wave (CW) mode, the RF excitation frequency is kept fixed and a sharp peak is produced in the output voltage as the background field sweeps through the resonance. Two different types of CW setups were developed (see Fig. 3.2):

- The earliest solution is based on a commercially available instrument, the NMR teslameter Metrolab PT2025 [95, 99, 100]. The probe contains a cylindrical sample with a volume of $226\,\mathrm{mm}^3$ made of a hydrogen-rich substance, water or rubber, where resonance is induced in the hydrogen nuclei based on the gyromagnetic ratio of the proton, $\gamma_\mathrm{p} = 42.577\,478\,92(29)\,\mathrm{MHz\,T^{-1}}$.

This instrument represents a reference standard in magnetometry, as it provides the modulus of the magnetic field with an absolute accuracy of about $5\,\mathrm{ppm}$ provided that field is sufficiently uniform (tolerated relative gradient $\approx 1\ \%/\mathrm{m}$) and stable. NMR probes were used with success as markers in the PSB and SPS systems since the 1980s. In FIRESTORM, a stable excitation frequency is provided by an Aim-TTi Function/Pulse Generator [101], while the teslameter unit demodulates the probe's RF output to obtain its amplitude envelope. Figure 3.2 depicts the output waveform from the teslameter, with the resonance point defined as the first negative peak [100]. Typically, the peak-to-peak amplitude of the signal ranges from $50\,\mathrm{mV}$ up to $1\,\mathrm{V}$, depending on probe type, field uniformity and field ramp rate. The measurement range covers magnetic field levels from $50\,\mathrm{mT}$ to well above $10\,\mathrm{T}$, with field ramp rates up to $0.1\,\mathrm{T\,s^{-1}}$. The effective reproducibility in operation, as derived from jitter measurements at a given field ramp rate under reproducible cycling conditions, is of the order of $5\,\mathrm{\mu T}$.

- The most recent design is represented by FMR devices, based on $\varnothing 0.3\,\mathrm{mm}$ Yittrium-Iron Garnet (YIG) ferrite spheres (for a volume of $0.014\,\mathrm{mm^3}$) as the resonating element [98]. FMR is a form of Electron Paramagnetic Resonance, implying that the precession frequency is three orders of magnitude higher than NMR, about $28\,\mathrm{GHz\,T^{-1}}$. YIG has a narrow resonance peak even at field ramp-rates as high as $5\,\mathrm{T\,s^{-1}}$, with typical quality factors ranging in the hundreds. Another advantage of FMR lies in the small size of the YIG sphere, which makes the resonator compatible with high-gradient fields such as those found in the PS combined-function magnets, where the poles are shaped to add a quadrupole field component with a relative gradient $|\nabla B/B| = 4.6\,\mathrm{m^{-1}}$ [102]. A prototype system based on a commercially available RF filter is installed there since 2012, and a series of tests have shown that the resonance peak remains well defined for absolute gradients up to $12\,\mathrm{T\,m^{-1}}$. On the downside, the anisotropy of conventional monocrystalline YIG introduces a degree of dependence upon temperature and field direction, with equivalent errors up to $40\,\mathrm{\mu T\,{}^{\circ}C^{-1}}$ and $20\,\mathrm{\mu T\,mrad^{-1}}$, respectively. These errors can be reduced by careful mechanical alignment of the YIG sphere and by thermalization of the resonator. Further reduction could be achieved by the use of paramagnetic materials [103], which however have a lower Signal-to-Noise ratio. For FIRESTORM, lumped-element and waveguide resonators were developed in-house and are now being implemented. The existing resonators cover a measurement range from $36\,\mathrm{mT}$ to $110\,\mathrm{mT}$. Overall, the effective reproducibility of the marked field can be as low as $1\,\mathrm{\mu T}$ under optimal conditions [104]. Prototypes of single-chip integrated microwave oscillators [103] working up to $700\,\mathrm{mT}$ are currently being tested; such higher field levels, however, correspond to a frequency range about $20\,\mathrm{GHz}$, which requires complex electronics for the detection.

Figure 3.2 depicts the setup of the two field markers, along with their conditioners and typical examples of output signals. The NMR probe output is demodulated in the teslameter unit, whereas the FMR resonance signal is first amplified and then amplitude-demodulated with a Schottky diode.

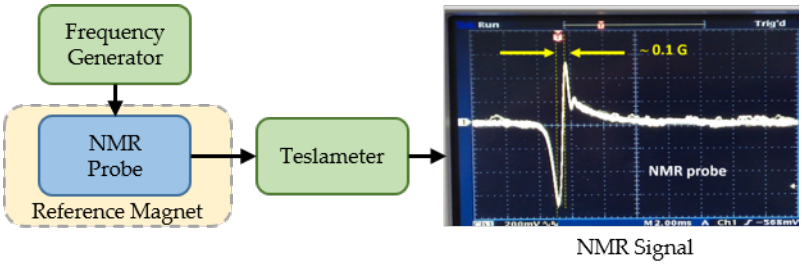

(A) NMR signal conditioning setup

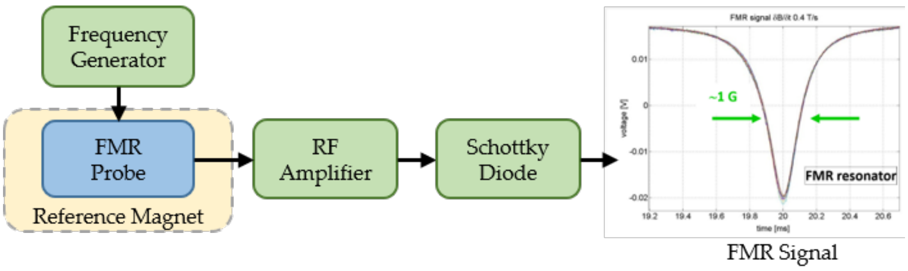

(B) FMR signal conditioning setup

FIGURE 3.2: Field marker block diagram.

3.1.3 Field marker calibration

In all types of field markers described above the sensing element has a very small volume, which represents a problem since the quantity of interest is the average of the field along the whole magnet. As discussed in detail in [105], the ratio of average to local field at the location of the sensor can be considered a constant only within a typical approximation of a few percent, due to magnetic hysteresis and eddy current effects. Since explicit modelling of these non-linear effects is very complex, our calibration procedure takes a different approach, by linking B_{m} directly to the average magnetic field at the time of triggering. In practice, the average field $\bar{B}(t)$ is dynamically measured during any given magnetic cycle waveform, and for a given excitation frequency f of the resonator (i.e. local field value $B = f/\gamma$) $B_{\mathrm{m}} = \bar{B}(t^*)$ is obtained as the average field upon reception of the marker trigger. The dynamic measurement can be performed with a fixed induction coil, provided the initial value of the field $\bar{B}(0)$ is known and integrator drift can be corrected sufficiently well. For example, after a degaussing procedure consisting of low-frequency AC excitation with an exponentially decaying amplitude, the remanent field is of the order of a few µT, and one can safely take $\bar{B}(0) = 0$ [105].

Chapter 4

System Proposal

In this chapter, the proposed design for the measurement system is presented.

4.1 Architecture and functionalities

Figure 4.1 shows the architectural layout of the new system. Unlike previous ones, which contained a large number of custom components with many slight differences, FIRESTORM adopts a modular architecture based upon common hardware and parametric firmware. This implementation allows for each setup to be adapted efficiently based on the different sensor configurations required by the synchrotrons. The simplest case is the SPS B-train, where a long integral coil with one low-field marker provides input to a single integration channel. The ELENA system has also one integral coil but two field markers, each used on different magnetic cycles: a high-field marker for cycles where antiprotons are decelerated, and a low-field marker for special test cycles that accelerate protons and H^- ions (in both cases, the marker is triggered just before injection). Yet another example is given by the combined-function PS magnets; these consist of two halves, one with a focusing and the other with a defocusing gradient, that must in principle be treated as two independent magnets. As such, each half requires a dedicated coil and integration channel. At present, the configuration has only one low-field marker implemented on each channel. Nevertheless, the system allows for the addition of a second pair of high-field markers, that in upcoming operating scenarios will be triggered sequentially on the same magnetic cycle, in order to achieve higher accuracy both at beam injection and extraction.

A flexible architecture is therefore necessary to deal effectively with these different requirements, as well as the adaptations and improvements that could be necessary during the 20- to 30-years lifespan of the system. The modular approach taken by the design, together with the remote configurability and diagnostics capabilities made possible by the tight integration of the software within the site-wide accelerator control system, is expected to improve both the maintainability and longevity of the system.

The key functions of the FIRESTORM B-train are implemented by a set of modules based upon off-the-shelf Simple PCIe Carrier (SPEC) hosted in an industrial Front End Computer (FEC). Each SPEC card hosts a custom-made FPGA Mezzanine Card (FMC) that implements analogue and digital I/O. This architecture allows splitting out the different functions with a fine level of granularity,

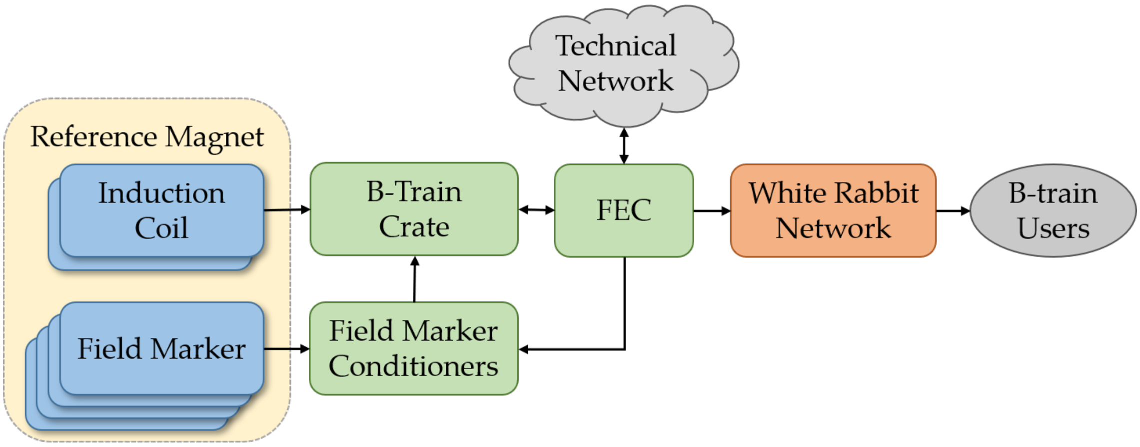

FIGURE 4.1: Block diagram of the main functions composing the FIRESTORM system.

improving both the flexibility and the maintainability of the final system. All design elements, including PCB layouts and firmware, are released on the Open Hardware Repository (OHWR) [106], a CERN initiative aimed primarily at the High Energy Physics community to stimulate collaboration, as well as the commercialization of the designs by industrial partners. These cards are linked to the magnetic sensors through the B-train crate, acting as a central hub. The final design element is a fiber-optic Ethernet-based network WR [107], used to distribute the measured field with high speed and noise immunity [108]. The use of an Ethernet frame allows for the transmission of the measured magnetic field alongside various ancillary signals, metadata and, crucially, three other versions of the field itself. These are: the field measured by the out-of-date system, where available; a copy of the nominal field obtained from the magnetic cycle database ("simulated field", see Section 5.2.4); and a mathematical model of the field based on the magnet excitation current ("predicted field", see Section 5.2.5. This is currently only at the prototype stage and is not implemented in the deployed systems). Access to these high-resolution, synchronized versions of the magnetic field is expected to greatly facilitate system diagnostics and to enhance operational flexibility in certain situations, e.g. when the measured field is not the most appropriate feedback source for the users (see Section 5.2.4).

Chapter 5

System Implementation

In this chapter, the hardware and software implementation are described. First, all the electronic component used are parented as well as their interconnection and communication protocols. Later all the implemented algorithms together with all the individual modules composing the system are presented.

5.1 Hardware architecture of the FIRESTORM system

The core of the FIRESTORM system is the FEC, an industrial diskless rack-mounted PC hosting the main electronic components [109]. About 2000 FECs are deployed throughout CERN for interfacing with devices that are involved in synchrotron control, such as RF and vacuum control systems, beam diagnostics instrumentation as well as the B-train systems. The current generation of FEC is the Siemens SIMATIC IPC847E with up to 11 free Peripheral Component Interconnect express (PCIe) slots. The operating system is 64-bit CentOS7 Linux [110] and the software is based on a distributed, real-time C++ class framework called Front End Software Architecture (FESA), which is at the heart of accelerator controls at CERN and GSI [111]. FESA abstracts the interface between the high-level accelerator control infrastructure and the local hardware, which is accessed via user-written device drivers. Tools are provided to help with the generation and debugging of C++ code. Automatic mechanisms are provided to store and retrieve class properties representing configuration parameters from a common database, as well as broadcasting measurement and diagnostic data vectors across the complex in quasi-real-time, i.e. with a latency of the order of a full magnetic cycle, which is adequate for many non-critical tasks. All communication happens on the Technical Network, a segregated Ethernet network secured against intrusion that is used to control and monitor all accelerator systems.

FESA is tightly integrated with the hardware timing system, used to synchronize the accelerators and their subsystems within a few microseconds [112]. This system consists physically in a network of coaxial cables, independent for each accelerator, distributing several hundreds of trigger pulses representing timing events relevant for beam or equipment monitoring and control. A separated serial channel for each accelerator (the so-called Machine Telegram) distributes information including the type of magnetic cycle being run, the destination of the beam and the type of the next synchrotron cycle that will be run. The framework

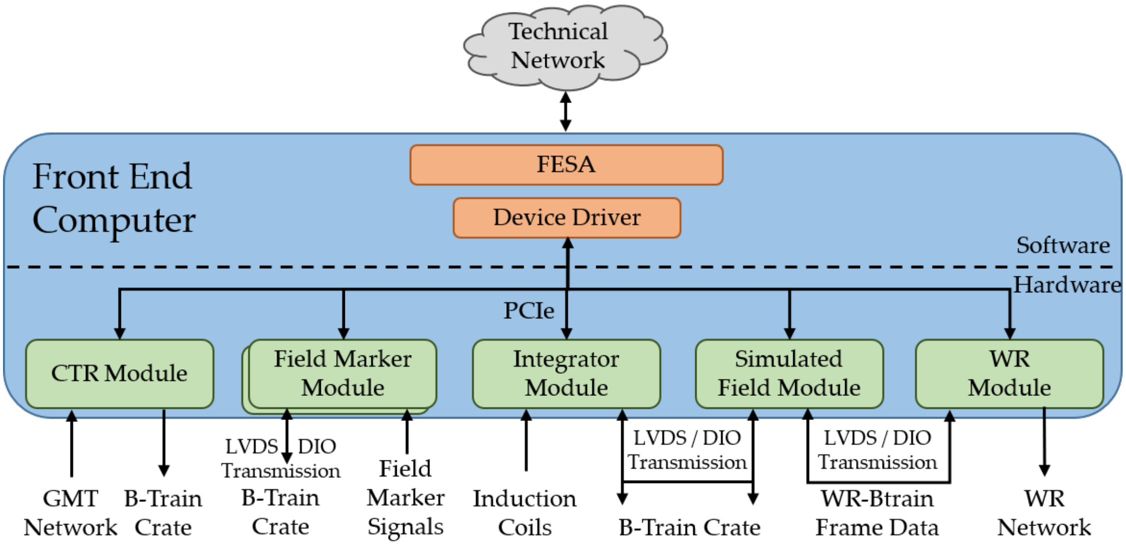

FIGURE 5.1: Architecture of the FEC, depicting the software hierarchy as well the flow of data through the FEC.

was recently fully endowed with so-called Pulse to Pulse Modulation (PPM) capabilities, which enable or disable specific actions such as class property setting and broadcasting, according to cycle type. PPM is a novel, crucial functionality that allows for the automatic adaptation of sensor calibration parameters to the magnetic characteristics of the synchrotron cycle; in particular, this applies to the field marker level B_{m}, which for the best accuracy should be calibrated independently for each cycle type. Compared to the manual updating carried out in the older B-train systems, this mechanism improves dramatically the flexibility and reliability of the configuration process. At present, the FIRESTORM FESA software comprises four different classes: the B-train class, which interfaces with all sub-systems that produce the measured field $\bar{B}$; the FSBT_BTG class, specifically for controlling the simulated field; the CosmosCheckWRS, which monitors the status of White Rabbit network; and the Comet_EVM, for environmental monitoring. Altogether, the FESA framework allows for the adjustment of more than 200 different configuration variables, inherent to the operation of the B-train, as well as access to over a 100 acquisition parameters, including internal registers and measurement values.

5.1.1 SPEC - Simple PCIe Carriers

The SPEC (see Fig. 5.2) is a general-purpose FMC carrier card with ready-made drivers ("spec-sw" on OHWR) that encapsulate the complexity of the host bus communication protocol, thus greatly simplifying the whole development cycle [113]. Several bus variants are available on the market, including VERSABUS Module Eurocard (VME), PCI Extensions for Instrumentation (PXI)E as well as PCIe interfaces, along with various types of FPGA modules.

The 4-lane PCIe version implemented in the FIRESTORM B-train is based on a $250\,\mathrm{MHz}$ Xilinx Spartan-6 LXT FPGA [114], offering 101,261 logic cells and

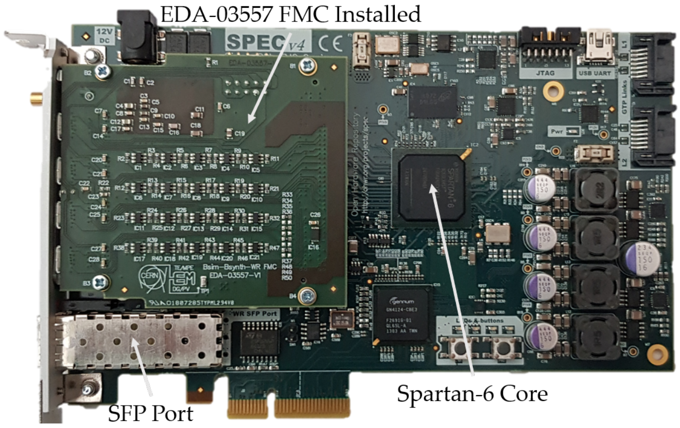

FIGURE 5.2: SPEC Card with EDA-03557 FMC card installed.

4.824 kbit memory. Currently, at most 50 % of the gate resources are used up in any module, which leaves considerable room for future improvements. The FPGA implements a finite state machine that defines the logical behavior of each component, and a number of Digital Signal Processor (DSP) cores that carry out the real-time signal processing tasks in 32-bit fixed-point representation (matching the WR distribution data format). Additional connectivity features include a Small Form-factor Pluggable Transceiver (SFP) fiber optic port which can be used for WR distribution (see Section 5.2.6), a Low Pin Count (LPC) connector as the FMC interface, plus standard Serial Advanced Technology Attachment (SATA), mini-USB and Joint Test Action Group (JTAG) (for FPGA programming) connectors.

The internal memory structure of the FPGA and its external interface are defined with the help of Wishbone, an open-source core-to-core logic bus [115]. A tool ("wbgen2" on OHWR [116]) is available to generate semi-automatically Very High Speed Integrated Circuits Hardware Description Language (VHDL) or Verilog cores that implement registers, memory blocks, FIrst In First Out (FIFO) registers and interrupts, along with the corresponding C header files. In this way, FESA software components can easily access and manipulate related variables and data structures. In particular, the transfer of large memory blocks representing the waveforms of various acquired or processed data are transferred via Direct Memory Access (DMA) through the PCIe bus, to be broadcast across the network. It is possible to note that, in the current implementation of the system, only one SPEC card at at a time is allowed to have DMA, to ensure stability. All the basic functionalities of the SPEC, including configuration, initialization etc. are managed automatically by the FESA framework.

5.1.2 FMC - FPGA Mezzanine Cards

The FIRESTORM B-train design includes four different types of FMC conforming to the FPGA Mezzanine Card ANSI/VITA 57 standard [117], which decouples I/O functions from the FPGA and allows simpler, modular designs. Except for the CERN-standard Central Timing Receiver (CTRI) card, that realizes the interface to the timing system, the other three cards were developed for the specific functions of the field marker trigger generator, magnetic flux integrator, White Rabbit I/O, simulated and predicted field features, all described in detail in Section 5.2. The last three functions require no specialized hardware, so they share the same FMC card design. The FMC designs are based on a small form factor that connects to the FPGA via a 160-pin LPC interface, which allows a theoretical bandwidth up to $40\,\mathrm{GB\,s^{-1}}$ with negligible latency and no protocol overhead. The major drawback of this choice is the difficulty of transmitting clock signals from the carrier to the mezzanine, thus preventing true hardware synchronization between the different cards. At present this does not represent a limitation, as the overall latency meets the requirements (see Section 9.2). Both the integrator and trigger generator FMC cards implement small-footprint, ultra-low phase noise Crystek CCHD-575 oscillators to generate locally a $80\,\mathrm{MHz}$ clock with $\pm 20\,\mathrm{ppm}$ worst-case frequency stability.

Communication between the FMC cards, the B-train crate and the WR transmitter is provided by a daisy chain of standard HDMI cables with 19-pin mini-HDMI connectors, chosen for their small size and robustness. Each HDMI cable carries a $250\,\mathrm{MHz}$ Low Voltage Differential Signal (LVDS) link, which allows bi-directional, self-clocked Manchester-encoded serial transmission with a theoretical $250\,\mathrm{Mbit\,s^{-1}}$ throughput. Transmission latency is typically less than $1\,\mu\mathrm{s}$, mainly due to the serialization/de-serialization steps. In parallel, eight conductors are dedicated to differential $2.5\,\mathrm{V}$ logic DIO channels, which allow the relaying of various kinds of trigger pulses with no protocol overhead.

An important goal of the LVDS links is to convey the different versions of the magnetic field to the White Rabbit frame assembler (see, Section 5.2.6), bypassing the PCIe bus with its associated programming complexity and uncontrolled latency. The daisy chain starts with the integrator module (which was the first component to be designed during the development phase), proceeds through the Simulated Field module and terminates at the WR transmitter. As discussed in Section 9.2, this arrangement results in a slight increase in the overall measurement latency, while still remaining acceptable. A separate LVDS link allows the direct exchange of data between the integrator and the trigger generator card.

5.1.3 B-train crate

The B-train crate, shown in Fig. 5.3, is the external interface of the FIRESTORM system, working as a hub for routing internal and external signals. The crate includes analogue and digital interface modules that allow local diagnostic access to all sensor outputs, the field marker triggers as well as the distributed magnetic field. In particular, a module is designed to accept as input the incremental pulse distribution of the out-of-date B-trains, based on two parallel 24 V pulse trains which represent respectively $\pm 10\,\mu\mathrm{T}$ field increments, accumulate the field value and send it via LVDS to the Integrator module for inclusion in the output stream.

The analogue outputs are duplicated on the back-plane of the crate in order to feed OASIS, a distributed acquisition system that allows for the monitoring, with some bandwidth and resolution limitations, of all of CERN's operation-related equipment signals [118]. The front panel hosts a set of High-Definition Multimedia Interface (HDMI) connectors that allow to make the links between the FMC cards, or to break them to access individual signals for diagnostics. Finally, an LCD multi-screen panel is provided to display real-time status information such as the measured field $\bar{B}$, sensor calibration parameters and other FPGA registers.

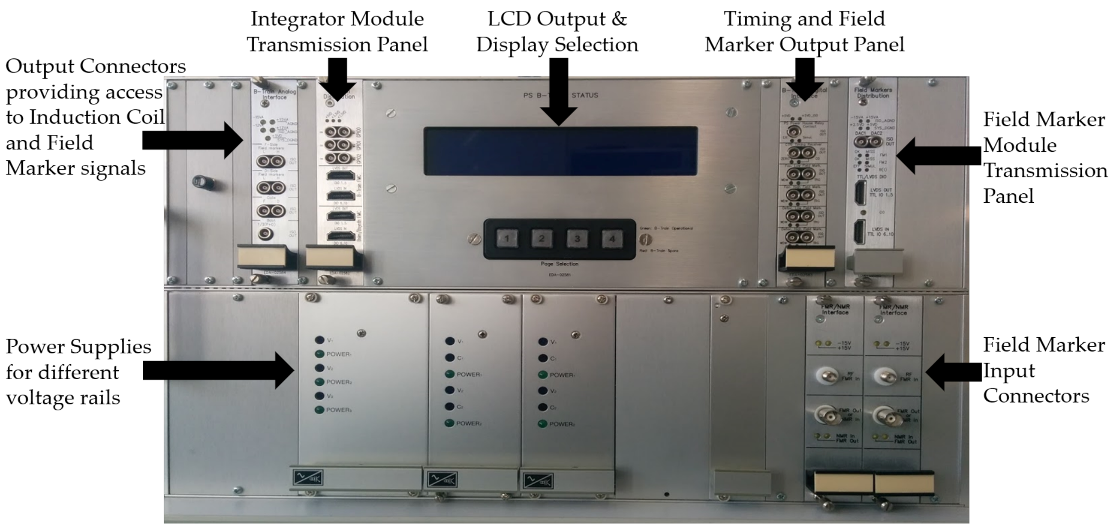

FIGURE 5.3: The front side of the B-train crate, showing the various 3U Eurocard front panels associated with the FMC boards.

5.2 Functional SPEC modules

In this section, the design and functions of the six kinds of SPEC/FMC boards used in the FIRESTORM B-train is described in detail.

5.2.1 Central Timing Receiver

The Central Timing Receiver card is a CERN-standard component installed in all FECs, where it receives and decodesGeneral Machine Timing (GMT) events that contain information on the cycle being performed in each accelerator [119]. The FIRESTORM system utilises the CTRI for generating two critical local timing Transistor Transistor Logic (TTL) triggers:

- the "C0" trigger, which signals the start of a new accelerator cycle and is used as an internal reference for various time related functions, such as the integrator calibration procedures and the field marker gating function described below. Optionally, C0 can be used to enforce a restart of the flux integration process to a given preset value. This is useful, for example, when a field marker malfunction is suspected.

- the “ZERO” cycle trigger, which signals the start of special cycles where no beam is circulating and the magnet excitation current is kept at a low (or zero) level. ZERO cycles are run from time to time in some (but not all) of CERN synchrotrons, either as low-power fillers in the machine schedule, or to allow capacitive-discharge magnet power supplies time to recharge. Whenever available, ZERO cycles are used for self-calibration of the integrators as described in Section 5.2.3.

These triggers are distributed through standard coaxial cables to the B-train crate, where they are first converted to $2.5\,\mathrm{V}$ pulses and then relayed to the integrator and the other FMCS via the HDMI DIO lines.

5.2.2 Field marker trigger generator module

The Field Marker Trigger Generator module has the goal of detecting the resonance peak in the output V_{m} of an NMR or FMR resonator, and to generate a TTL trigger pulse accordingly. The module has a dual-channel design which allows, for example, to have a high-field and low-field marker acting at different times on the same integration channel (as in ELENA [104]), or two markers acting in parallel on two separate integration channels (as in the PS). It is useful to recall that the B-train crate has a number of connectors sufficient to handle up to four field marker signals, corresponding to up to two SPEC/FMC cards operating in parallel in the same FEC. The field marker output is initially routed through a signal conditioner board in the B-train crate. This removes the DC component, allowing for the subsequent comparison to a known threshold, and then optionally amplifies the signal (this is necessary only for the FMR sensor, not the NMR) [98]. In the following, the hardware on the Field Marker FMC and the peak detection algorithm implemented in the FPGA is described in detail .

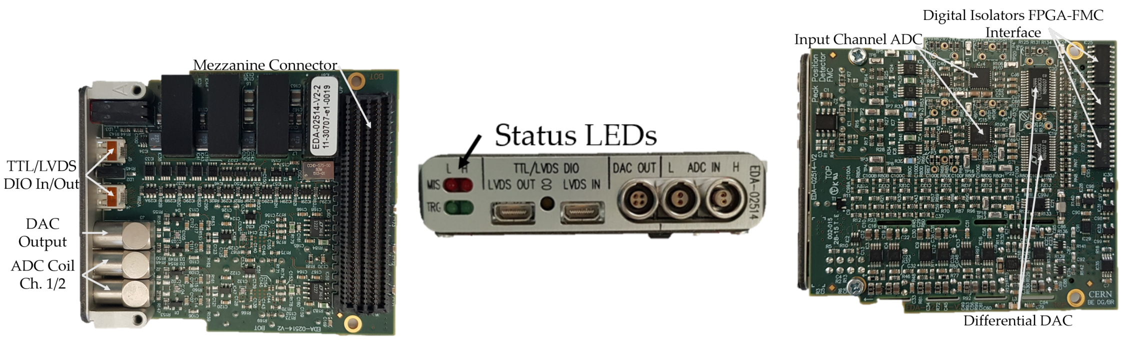

(A) Top View (B) Front View (C) Bottom View

FIGURE 5.4: Field marker trigger generator FMC.

Field Marker FMC

The Field Marker FMC (EDA-02514, see Fig. 5.4) includes two fully independent, parallel acquisition channels based on a 16-bit, $10\ \mathrm{MSamples}/s$ AD7626 Analog to Digital Converter (ADC) with a $\pm 4\,\mathrm{V}$ differential input range. This ADC was chosen for it's resolution and sampling rate because the resonance peak is a very fast event that has to be accurately recognized for the marker being effective. The analogue input stage includes a low-pass filter that is essential in removing the noise generated by the sensor or picked up along the way, thus avoiding spurious triggers. The cut-off frequency is usually set around $2\,\mathrm{kHz}$, which corresponds to a detection delay of the order of $100\,\mu\mathrm{s}$; being systematic, this has negligible impact on the calibration of B_{m}, at least as long as the field ramp rate at the time of triggering is constant [100]. No additional anti-aliasing filter is necessary because the low-pass filter added to remove the noise generated by the sensor or picked up along the way serve also the anti-aliasing function.

I/O connectors include dual LEMO inputs for analogue field marker signals, a LEMO analogue output for the on-board Digital to Analog Converter (DAC), and two mini-HDMI sockets for the input and output LVDS DIO. The generated $1\,\mathrm{ms}$ trigger pulse is transmitted on the LVDS output to the Integrator module, from where it is propagated down to the WR module to be written in to the distributed WR Ethernet frame (see 5.2.6). In addition, four diagnostic status LEDs that signal are present, on any given machine cycle, the detection of the high- and low-field marker triggers, or the lack thereof within the allowed time window.

Peak detection algorithm

The resonance peak is defined as the first zero-crossing of the derivative of V_{m} and the corresponding time $t^* = t_j$ is defined by:

$$j = min(i) : \begin{cases} t_1 \leq t_i \leq t_2 \\ |V_{\mathrm{m},i}| \geq \overline{V} \\ sign(\dot{V}_{\mathrm{m},i}) \neq sign(\dot{V}_{\mathrm{m},i-1}) \end{cases} , \tag{5.1}$$

where i is the running index of the waveform samples; $[t_1, t_2]$ is a pre-defined gating window, typically $20\,\mathrm{ms}$ long, that prevents spurious triggers to happen too far from the expected time during a cycle; $\dot{V}_{\mathrm{m}}$ is the time derivative of the sensor output, calculated with a seven-point finite-difference scheme; and $\overline{V}$ represents a voltage threshold, set independently for each system above the residual noise level after filtering. As the zero-crossing can happen anywhere within the $[t_{i-1}, t_i]$ interval, this simple algorithm considering an ADC sample rate equal to $10\ \mathrm{MSamples}/s$ has an uncertainty of $\pm 50\,\mathrm{ns}$, which is negligible for our application because the time duration of the peak depends from the field ramp rate and it is in the range of few ms. It is useful to recall that all the algorithm's parameters are stored in FPGA registers loaded at run time by the FESA software, which can be adapted automatically to the type of cycle being run via the PPM mechanism.

5.2.3 B-train integrator

The dual-channel B-train Integrator module has the primary role to determine the value of the average field $\bar{B}$ that is distributed to the B-train users. In addition, it has the capability to accept as input the incremental pulse distribution of the out-of-date B-trains, based on two parallel 24 V pulse trains which represent respectively $\pm 10\,\mu\mathrm{T}$ field increments, to accumulate it and to distribute the result alongside the FIRESTORM measurement. The main issue affecting this measurement is the drift caused by a voltage offset δV superposed to the coil output V_c. The offset has a spectrum akin to $^1/_f$ pink noise, with a slowly drifting, almost systematic component superposed to random fluctuations with periods of the order of a few seconds to a few minutes, comparable with the duration of most accelerator cycles [120]. Such an offset can be mitigated, for example, by choosing high-quality discrete components causing imbalances in the analog input stages, by reducing thermal gradients leading to thermoelectric voltages and, in general, by ensuring long-term thermal stability via adequate ventilation. Respect to other voltage integrators described in the literature [33][93][121], the specificity of the presented design lies in the method used to estimate it and correct in real-time. For simplicity, It could be assumed that throughout each integration interval (or, equivalently, accelerator cycle) δV is a constant, and re-evaluate it periodically.

The application of this specific module is not limited to the presented application but it could in principle be extended at all the industrial or scientific applications in which a fast digital integration is required.

In the following subsections the hardware of the mezzanine card, the integration and error correction algorithms are discussed in detail .

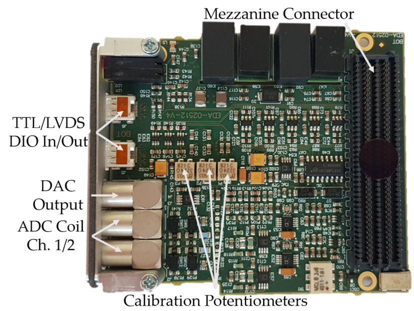

(A) Top View.

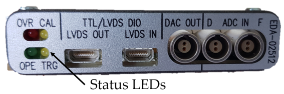

(B) Front View.

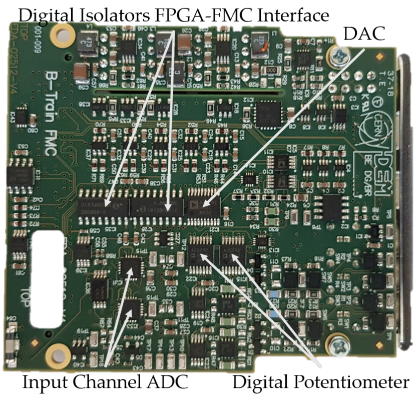

(C) Bottom View.

FIGURE 5.5: Integrator module FMC

Integrator FMC

The Integrator FMC (EDA-02512) is shown in Fig. 5.5. The core of each FMC integration channel is a high-linearity 18-bit, $2\ \mathrm{MSamples}/s$ AD7986 Successive Approximation Register (SAR) ADC with a $0\,\mathrm{V}$-$5\,\mathrm{V}$ differential input range. This ADC was chosen mainly for it's resolution in order to have the field resolution within specification. Each channel includes the following conditioning stages:

- A three-way selection switch with a $200\,\mu\mathrm{s}$ settling time for the auto-calibrating function, as explained below.
- An input buffer with a $27\,\mathrm{MHz}$ bandwidth and a R_{in} $=2\,\mathrm{M\Omega}$ impedance. The impedance stems from a compromise between the need to limit signal attenuation for high-resistance input loads, and the need to limit the offset voltages arising due to input bias currents. For a typical measurement coil resistance on the order of R_{c} $=1\,\mathrm{k\Omega}$, the low-frequency attenuation can be easily calculated from $\frac{R_{\mathrm{c}}}{R_{\mathrm{c}}+R_{\mathrm{in}}} \approx 500\ \mathrm{ppm}$, and corrected in the post-processing stage. As the specified input signal bandwidth is just $100\,\mathrm{Hz}$, a more rigorous dynamic study of the parasitic capacitive effects was not considered a priority at this stage.
- A two-stage pre-amplifier that scales the nominal $\pm 10\,\mathrm{V}$ induction coil signal to the $\pm 5\,\mathrm{V}$ differential input range of the ADC. First, a voltage divider realized with high-precision discrete resistors attenuates the signal by a factor $5/8$; then, a fully differential funnel amplifier AD8475 with attenuation factor $4/5$ and nominal passband $15\,\mathrm{MHz}$ prepares the signal for the ADC, while ensuring that the total attenuation factor is $1/2$.
- An AD5291 digital potentiometer used in a voltage divider to provide a programmable voltage source with $1\,\mathrm{mV}$ range and about $1\,\mu\mathrm{V}$ resolution, injected between the two attenuation stages and used for fine offset compensation.
- A simple first-order RC anti-aliasing filter with a $1\,\mathrm{MHz}$ cutoff frequency, which gives a nominal 100 ppm maximum error at the upper end of the $100\,\mathrm{Hz}$ signal bandwidth.

The board includes also a multi-purpose AD5791 DAC with a $\sim 1\,\mu\mathrm{s}$ settling time, whose output can be applied to the integrator input by switching the input selector to the position 2, as shown in Fig. 5.6. This is used both for the periodic gain self-calibration, and to generate various kinds of analogue output signals as may be needed for diagnostics (e.g. an image of the measured field $\bar{B}$ to be visualized on the spot with an oscilloscope) or for special purposes (e.g. an image of the field derivative $\dot{\bar{B}}$ that is used to compensate eddy current effects in the PS magnets). Three mechanical potentiometers are also included to adjust manually the offset, positive and negative range of the DAC as needed for gain calibration, as explained below. I/O connectors include, beside the dual ADC input and the DAC output, two TTL/LVDS DIO connectors for the daisy chaining and diagnostics of the card's output. Finally, four diagnostic status LEDs signal, on any given machine cycle, the reception or lack thereof of a high- or low-field marker trigger.

Integration algorithm

The integrator implements two identical acquisition and computation chains in parallel, which are combined linearly to provide the final output:

$$\bar{B} = k_1 \bar{B}_1 + k_2 \bar{B}_2. \tag{5.2}$$

This implementation provides the flexibility to use only one set of sensors, or to mix two sets according to the circumstances, as is required for example in the PS B-train system. In the following, the operation of a single channel, dropping for simplicity the index from all related variables is described in detail. The data-flow is represented schematically in Fig. 5.6, where the analogue pre-processing and signal digitization performed by the FMC is on the left, while the numerical processing carried out by the FPGA is on the right.

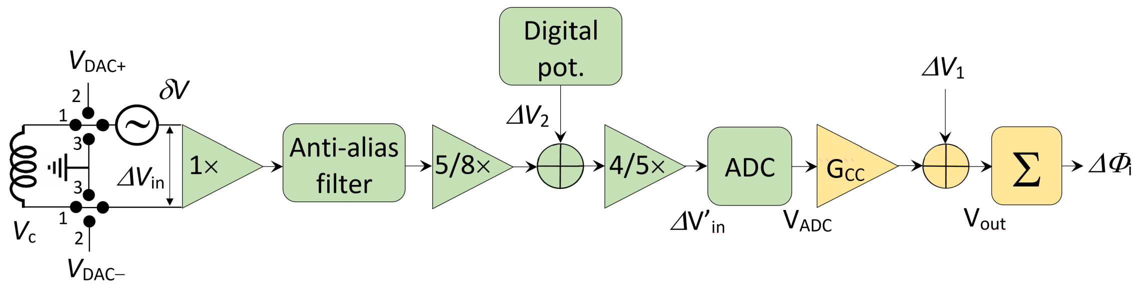

FIGURE 5.6: Schematic flowchart of the integrator, including offset and gain correction. The green blocks denote analogue processing steps, while the yellow ones, processing in the digital domain.

It could be assumed that the differential voltage ΔV_{in} at the input of the conditioning stage is the sum of the coil output voltage V_{c}, and the offset δV:

$$\Delta V_{\text{in}} = V_{\text{c}} + \delta V. \tag{5.3}$$

In other words, only the sources of offset internal to the card (e.g. due to discrete component imbalances) can be considered and the external ones can be neglected, such as thermoelectric gradients on the cabling between the induction coil and the card. This approximation is usually sufficient to obtain good results, as shown in Section 8.3.
The differential voltage $\Delta V'_{\text{in}}$ at the input of the ADC can be expressed as:

$$\Delta V'_{\text{in}} = 0.8\,(0.625 \Delta V_{\text{in}} + \Delta V_2) = \frac{1}{2} \Delta V_{\text{in}} + \frac{4}{5} \Delta V_2, \tag{5.4}$$

where ΔV_2 is the programmable offset added by the voltage divider thanks to the digital potentiometer. The sampled voltage is first corrected according to (5.5):

$$V_{\text{out}} = G_{\text{cc}} V_{\text{ADC}} + \Delta V_1, \tag{5.5}$$

where $G_{\mathrm{cc}} \approx 1$ is the internal gain correction factor and ΔV_1 represents the coarse offset correction. To remain within the FPGA resource limits with a reasonable margin, all variables in (5.5) are represented in 18-bits, with an effective resolution of $1\ \mathrm{LSB} \approx 76\,\mu\mathrm{V}$. The change in magnetic flux $\Delta\Phi$ is integrated according to:

$$\Delta\Phi_i = \tau_{\mathrm{s}} \sum_{j=i_k^*}^{i} V_{\mathrm{out},j}, \tag{5.6}$$

where j is a running sample index, i_k^* marks the start of the current integration interval upon reception of the $k-th$ field marker trigger, and $\tau_{\mathrm{s}} = 500\,\mathrm{ns}$ is the sampling time. The calculations in (5.5) and (5.6) are carried out with a 56-bit depth to avoid overflow, and the flux change $\Delta\Phi_i$ is represented with a depth of 32-bits ($1\ \mathrm{LSB} \approx 5\,\mathrm{nV\,s}$) to match the format of the final output. Finally, the average magnetic field is computed according to the model (5.7):

$$\bar{B}_i = \gamma(B_{\mathrm{m}} - \alpha\frac{\Delta\Phi_i}{A_{\mathrm{c}}}), \tag{5.7}$$

where the non-dimensional coefficients γ and α represent correction factors, accounting respectively for the difference between the reference magnet and the average of those in the accelerator, and any error in the effective area or the position of the coil, as discussed in [105].

Drift correction

Drift correction relies on the availability of beam-less ZERO cycles during which the integrator input can be safely short-circuited (position 3 of the input switch in Fig. 5.6), and the observed drift can be attributed entirely to the voltage offset δV. Since sometimes the accelerators operate with many, closely spaced ZERO cycles, a dead time of 5 minutes between corrections is imposed, which in practice was found to avoid possible instabilities. During the correction process, the distributed field values will be of course meaningless and must be disregarded by the users; in particular, the power converters feeding PS and PSB magnets must open their control loops to avoid runaway instability.

The estimation and compensation of the voltage offset are carried out in two stages. The first stage is purely numerical and occurs in the FPGA, where the coarse offset correction ΔV_1 is derived by averaging the voltage during a portion of a ZERO cycle:

$$\Delta V_1 = -\delta V = -\frac{\Delta\Phi_0}{n_0\tau_s} = -\frac{1}{n_0}\sum_{i=i_0}^{i_0+n_0-1} V_{\mathrm{out},i}, \tag{5.8}$$

where the index i_0 marks the start of the short-circuit measurement, n_0 represents its duration in samples and $\Delta\Phi_0$ is the measured flux drift. The duration of this measurement should be as long as possible to improve the accuracy of the computed average, which scales as $n_0^{-1/2}$; however, it is important to leave some margin at the start of the cycle for the control loop of the power converters to be opened. As an example, in the PS system i_0 = 400 kSamples or, equivalently,

$200\,\mathrm{ms}$ after C0 and $n_0 = 200$ kSamples, corresponding to a $100\,\mathrm{ms}$ duration was set.

Since the resolution of ΔV_1 is limited to 1 Least Significant Bit (LSB) = $76\,\mu\mathrm{V}$, It was decided to implement an additional correction stage, adding a much finer offset ΔV_2 to the signal in the analogue input stage. This offset can be set with $1\,\mu\mathrm{V}$ resolution over a range of $\pm 500\,\mu\mathrm{V}$. Different strategies are currently being evaluated to set optimally ΔV_2, including differentiation followed by low-pass filtering of the measured flux, or an iterative binary search strategy that aims at zeroing the measured drift. As this feature is still at the prototype stage, all the results reported in Section 8.3 were obtained by setting $\Delta V_2 = 0$.

Gain correction

The linear gain correction procedure is also performed during a ZERO cycle, immediately after the offset calibration described above, except that the input is switched on the position 2 of Fig. 5.6. This applies to the input of the acquisition chain the output of the high-precision DAC, used as a voltage reference in the range between $\pm V_{\mathrm{ref}}$, with $V_{\mathrm{ref}} = 8.75\,\mathrm{V}$. In this range, which covers the majority of cases, the DAC shows a very good linearity; moreover, V_{ref} has an exact hexadecimal representation in the VHDL code, which improves the accuracy of the gain correction. The DAC itself is calibrated manually at least once, as part of the production tests, with an external Agilent 34401A multimeter [122]. The three mechanical potentiometers installed on the FMC are used respectively to remove first any offset at $0\,\mathrm{V}$, and then to adjust the values of $\pm V_{\mathrm{ref}}$. This calibration procedure can be repeated during operation if deemed necessary. The gain correction procedure consists of applying to the input first $+V_{\mathrm{ref}}$ and then $-V_{\mathrm{ref}}$, over two sequences of n_1 samples each, during which the FPGA computes the average of the sampled voltage. Taking into account the scaling done by the conditioning module (5.4), the gain correction factor is then computed as:

$$G_{\mathrm{cc}} = \frac{1}{2} \frac{V_{\mathrm{ref}}}{\frac{1}{n_1}\sum_{i=i_1}^{i_1+n_1-1} V_i - \frac{1}{n_1}\sum_{i=i_2}^{i_2+n_1-1} V_i} \approx 1, \tag{5.9}$$

where $i_1 = i_0 + n_0 + \Delta n$ is the starting sample of the $+V_{\mathrm{ref}}$ acquisition , $\Delta n = 1kS$ is an interval of $0.5\,\mathrm{ms}$ introduced to give the input time to stabilize, $n_1 = 300$ kSamples corresponds to the duration of the acquisition of$150\,\mathrm{ms}$, and $i_2 = i_1 + n_1 + \Delta n$ is the starting sample of the $-V_{\mathrm{ref}}$ acquisition .

5.2.4 Field Simulation

The simulated field module, schematically represented in Fig. 5.8, generates in real-time a high-resolution image of the nominal, pre-programmed magnetic cycle as it is stored on the centralized LHC Software Architecture (LSA) database [123]. The role of this feature is twofold:

- as a normal part of the accelerator restart sequence, when the accelerating RF cavity control system needs a realistic value of the field to be input via the B-train for its own frequency program, even when no beam is circulating yet and the magnets are not powered.

- under certain special circumstances, when magnetic field measurement feedback is not the best option. For instance, machine operators may want to replace the measured field temporarily with the simulated one as a beam diagnostic tool. As another example, in the case of a power converter trip, the value fed back to the RF cavities must switch automatically from the measured to the simulated field, in order to avoid large, potentially harmful discontinuities. Even more crucially, in the specific case of the AD, the simulated field is always preferred because the machine is magnetically very reproducible, and the RF system is adversely affected by the noise inherent in the measured field.

Vector cycle representation

The image of each cycle is stored in the LSA as a two-column vector table, where the first column represents time and the second, in general, the magnetic field. One exception to this rule is provided once more by the AD, where the LSA image contains the magnet excitation current waveform, and the B-train software must apply a given analytical relationship to derive the magnetic field. (This involved procedure is not necessarily more precise than a measurement, but the accelerator was finely adjusted accordingly in the 1990's, and today there is hardly reason to change.)

The vector cycle representation is very compact, most cycles being described accurately by a few dozens to a few hundred vectors. A finer resolution is generally required at high field, where the current-to-field relationship is non-linear due to iron saturation, or to smooth out discontinuities at the junction of current ramps and plateaus. A maximum number of 7025 vectors can be accommodated in the SPEC's on-board RAM, which is more than enough for any present or anticipated need. A small memory footprint is also critical to pre-fetch quickly from LSA the table for the next cycle while the current cycle is still running. The telegram provides an advance of at least $1.2\,\mathrm{s}$, i.e. one basic accelerator period, which is largely sufficient for the FESA software running on the FEC to interrogate the database, download the data via the Technical Network and transfer it onto the Simulated Field card[1]. This new strategy, unlike the out-of-date B-train systems which kept in memory the full high-resolution time series corresponding to a few cycle types, is way more efficient and general as it can adapt transparently to any of the thousands of cycles already stored, or expected in future.

[1]By default, all external data are transferred onto the Integrator module via PCIe DMA, and from there they get handed to the other modules down the LVDS daisy chain. The only exception to this is the AD system, where there is no Integrator module and DMA is implemented directly in the Simulate Field module FPGA

Magnetic cycle interpolation

The table of vectors is interpolated to the desired resolution (by default, 4 µs) in real-time in the FPGA, using Bresenham's line algorithm. Practical details are different, according to the accelerator. Implementation is simpler for the accelerators in the PS complex (LEIR, PSB and PS), where the cycle length can only be one, two or three basic 1.2 s periods. Conversely, in the antiproton decelerators (AD and ELENA) cycle vectors are not necessarily known *a priori*, but are defined at run time by specific start and stop timing events, triggered manually by operators in the Control Room. This mechanism provides the possibility to prolong a plateau for an arbitrary duration, up to a couple of minutes, as required for beam electron cooling or to accumulate antimatter for various experiments. During these pauses, the interpolation is temporarily stopped and the B-train outputs a constant value. The possibility of pausing a cycle on the flat-top, by means of a specific set of timing events, is also implemented in the SPS, where it used to adjust beam extraction for ion-beam momentum slip stacking [[124]].

Simulated/Predicted/WR FMC

The FMC of the Simulated Field module (EDA-03557), which is physically the same for the Predicted Field and WR modules, is shown in Fig. 5.7. there are two input and two output mini-HDMI connectors for LVDS DIO, one SFP optical port for WR and one coaxial output for the on-board AD5791 DAC.

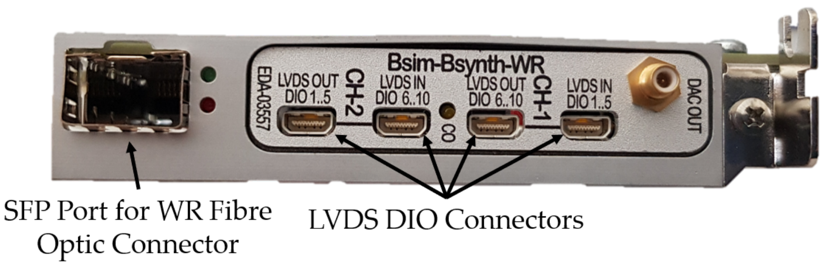

FIGURE 5.7: EDA-03557 front view mounted on a SPEC card.

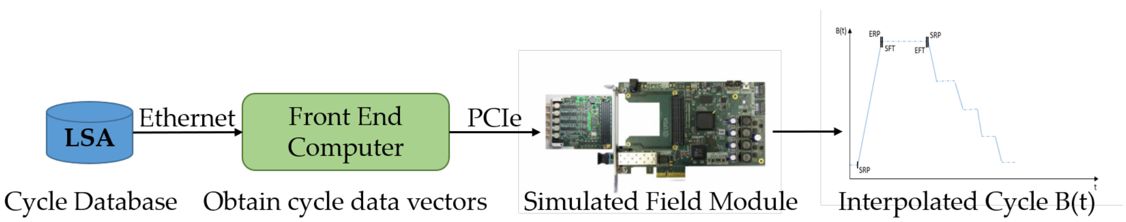

FIGURE 5.8: Schematic representation of the Simulated Field dataflow.

As already said for the integrator module, also the simulated module combined with a SPEC card is not limited to be used in this application but it can be more widely used in industrial and scientific application where it is mandatory to transmit data at a very high rate between devices even 10 km far from each other. Indeed the WR is spreading more and more in the industrial and scientific world for these kind of applications[88].

5.2.5 Field Prediction

The Predicted Field module, which at the time of writing is at an early prototype stage, implements a mathematical model to derive in real-time the magnetic field from the excitation current. This can be useful in a variety of scenarios, which overlap with the Simulated Field use-cases. For example, machine operators might want to switch temporarily from the measured to the predicted field as a diagnostic measure, whenever they suspect a sensor malfunction; or, the difference between measurement and expectation can be continuously monitored, as a powerful real-time diagnostic tool. In the long-term, if proven to be sufficiently accurate, the prediction might replace the measurement altogether, drastically cutting the cost and complexity of future B-train systems. For the time being this is not yet possible and further studies are required to achieve this result in operation

Different categories of mathematical models are being considered as candidates for this functionality. At CERN, a semi-empirical analytical model called FIDEL [125] is used with success since 2007 to derive offline the inverse field-to-current relationship of the superconducting LHC magnets, as needed for open-loop control with $100\,\mathrm{ppm}$ accuracy. This is possible thanks to the coil-dominated character of these magnets and to the very high field they reach, well above $8\,\mathrm{T}$, which minimize the impact of non-linear effects such as iron saturation, hysteresis and eddy or persistent currents. In iron-dominated magnets below $2\,\mathrm{T}$, such as those found in all other accelerators at CERN, these effects can affect the magnetic field much more severely, which makes the task more challenging.

This is especially true for history-dependent features, such as the remanent field or the response to minor hysteresis loops, as illustrated by the failure of an early linear dynamic model tested in the PSB [126]. Different classes of hysteresis models are discussed widely in the literature, including closed-form or differential analytic expressions, operator-based and neural network formulations, which have had some success such as those in [127]. At present, these are still being evaluated to identify the most suitable one for real-time FPGA implementation.

FMC

The Predicted Field FMC carries no specialized hardware and is indeed the same as the one for the Simulated Field described above. The FMC can accept as an input the current measured with a high-precision Direct Current to Current Transformer (DCCT) and distributed by the controller of the power converters over a dedicated WR network, using a specific definition of Ethernet frame that includes additionally the status of the power supply and its output voltage. The frame rate in this case is much lower than for the B-train WR, i.e. usually $10\,\mathrm{kfps}$, which corresponds to the operational frequency of the digital controller of the power converter.

5.2.6 White Rabbit interface

White Rabbit, adopted as the IEEE 1588-2019 standard, is an Ethernet-based network for real-time, large-scale distributed control systems, featuring deterministic data delivery at gigabit speed with sub-nanosecond synchronization over multiple kilometers of optical fiber [108]. Originally developed at CERN, it is now openly accessible under the GNU general public license in the OHWR and is supported by National Instruments and many other vendors. Given the tight timing constraints and the requirement to distribute serially the actual value of the magnetic field, rather than just the increments as in the old systems, WR was a rather natural choice [128]. The added value of this solution resides in the improved maintainability of the network based on commercially available routers and switches. These components can be remotely configured and upgraded with a full suite of powerful remote debugging and diagnostic facilities that allow for measurement of the transmission latency, warn of packet loss and more.

The measured, out-of-date, simulated and predicted fields are distributed through the LVDS daisy chain connections to the WR FMC card, which is physically identical to those of the simulated and predicted modules. The FMC offers physically the possibility of being used either as a receiver or a transmitter, the actual functionality being implemented in the SPEC gateware. In the FIRESTORM B-train, the FPGA handles the task to assemble the data in the so-called WR-Btrain frame, illustrated in Fig. 5.9. The B-train frame is 26 bytes long and includes, in this order:

- A 16-bit frame control header, consisting of an 8-bit frame type ID and various status and error flag bits, essential for remote diagnostics and calibration. In particular, there is a flag bit that signals if the current cycle is a ZERO cycle; and two additional flag bits, which are set to 1 to indicate the reception respectively of the C0 (cycle start) and field marker trigger pulses. Since the pulses have a nominal width of $1\,\mathrm{ms}$, each one of these two flags will normally be set on 250 consecutive frames.

- The first part of the payload, comprising two 32-bit slots for the active $\bar{B}$ and $\dot{\bar{B}}$ values. The active version of the magnetic field is selected among the four possibilities and is positioned in the first slot, to ensure that all users read by default the same value. In signed fixed-point representation, the distributed field has a resolution (1 LSB) of $10\,\mathrm{nT}$ and a range of $\pm 20\,\mathrm{T}$, which is more than enough for all foreseeable applications. The second slot contains the numerical time derivative of the first one, with a resolution of $1\,\mathrm{\mu T\,s^{-1}}$ and a range of $\pm 2\,\mathrm{kT/s}$, also exceeding all foreseeable demands.

- The second part of the payload, which includes four 32-bit slots for the measured, out-of-date, simulated and predicted field, with the same resolution as the active field slot. Repetition of the active field slot implies a redundancy of 4 bytes in the payload, which has negligible impact; in return, this guarantees that all users and diagnostic tools may conveniently access at any time the four versions of the field, at fixed slots within the frame.

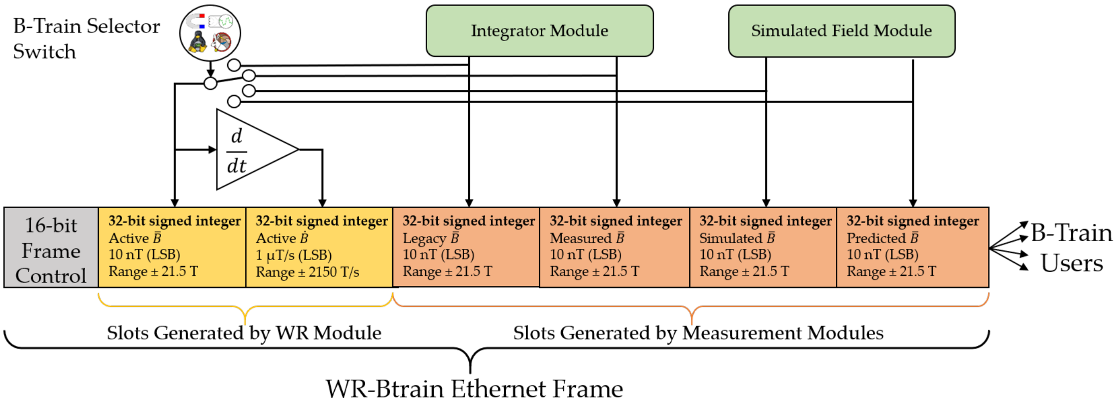

FIGURE 5.9: Functional diagram of the broadcast White Rabbit Ethernet frame.

The WR-Btrain frame is distributed as part of the larger streamer frame consisting of a 46-bytes payload, the minimum-length for an Ethernet frame. As padding is implemented to achieve the required payload size, this arrangement leaves five additional 32-bit slots free for future expansion [129]. Based on the frame size, this corresponds to a maximum theoretical rate of 1.4 Mfps over a gigabit optical link. At present, a steady transmission rate of 250 kfps is achieved reliably, while tests are ongoing to establish a practical upper limit.

Part III

A tool for the measurement system characterization

Chapter 6

Monitoring system requirements and proposal

In this chapter, the reason behind the new monitoring system is described, later the requirements are illustrated. Than the proposed architecture for the new monitoring system is described. First, an overview of the hardware selected is presented, later the proposed block diagram of the interconnections and functional blocks is illustrated.

6.1 System requirements

Monitoring the WR payloads at full speed and for extended periods of time is an essential part of both, the commissioning and operation phases. Initially, system debugging requires that the measurements distributed via WR be compared systematically to old measurements and the signals as received by the numerous end users. Noise and glitches in the field measurement can affect destructively the beams and require full bandwidth to be evaluated correctly. In operation, continuous comparison of twin parallel acquisition chains to each other and to expectations (simplified mathematical models) will be needed to diagnose instrumentation errors such as integrator drift, identify fault conditions to raise the appropriate alarms and evaluate measurement uncertainty estimates.
So far, a Python-based acquisition system was used to intercept the WR stream and save it to file. This implements a DMA to retrieve upcoming WR data previously stored on a Double Data Rate (DDR) memory. There are some drawbacks to this solution, such as: not easy customization; not easy installation procedure; impossibility to check simultaneously multiple measurement chains; no remote access, and impossibility to select remotely the system to be monitored and debugged.

A more flexible solution was needed, the requirements of the new system were the following:

- Be able to catch all the frames sent (at 250 kHz) on the optical fiber with the WR protocol.
- Be able to check the quality of the transmission to see directly what the users see

- Be able to watch at the same time both chains one and two for each machine to compare them.
- Select remotely the machine to monitor via the fiber optical switch.
- A simple graphical user interface.
- An easy way to log locally the waveform acquired.
- A portable solution, remote access to the device from everywhere inside the CERN network.

6.2 System proposal

6.2.1 Hardware selection

The strong requirements in terms of flexibility and remote access led the development to use the CompactRIO (cRIO) platform from National Instruments (NI). The optical fibers, chosen physical link to broadcast the measurement signals, arrive in two optical fiber switches from the operational (OP) and spare (SP) chains. To decode the WR B-train data, a cRIO White Rabbit module (CRIO-WR) [130] was used. CRIO-WR is a standalone WR node implementation on a PCB with a form factor for NI cRIO crates. The board is originally derived from and keeps maximum firmware compatibly with the established boards SPEC [131] and CUTE-WR [132]. To control the optical fiber switch, a DIO module NI-9401 connected to a Switch Controller was used.
In order to host and control the CRIO-WR and the DIO modules, a cRIO crate controller NI-cRIO-9040 is used. The processor runs a real-time software target used to access the device and therefore the B-train data from any PC on the same network with an Ethernet connection. A simple overview of the proposed architecture is depicted in Fig. 7.3

NI-9040 Embedded Controller

The cRIO-9040 (Fig.6.1) is a four slots high performance embedded controller from NI. Within the controller, there is a Xilinx FPGA Kintex-7 series and a dual core Intel Atom running NI Linux Real-Time at 1.3 GHz. It manages the connection between the controller and a host PC over Ethernet. This device communicates with the cRIO modules collecting the data coming from the CRIO-WR node and addressing the optical fiber switch through the cRIO-DIO.

NI-9401 DIO module

The NI-9401 (Fig.6.2) is a configurable digital I/O module working between 0 V and 5 V TTL. In this application it controls the optical fiber switch controller, selecting the machine (measurement chain) to be debugged and monitored. The switch is a multiplexer, three digital lines are necessary to select one of the 8 machines remotely and a fourth line to select between OP and SP.

FIGURE 6.1: NI-9040 Embedded Controller.

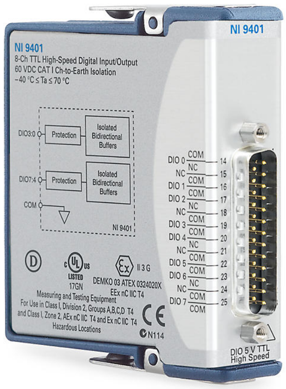

FIGURE 6.2: NI-9401 DIO module.

cRIO-WR module

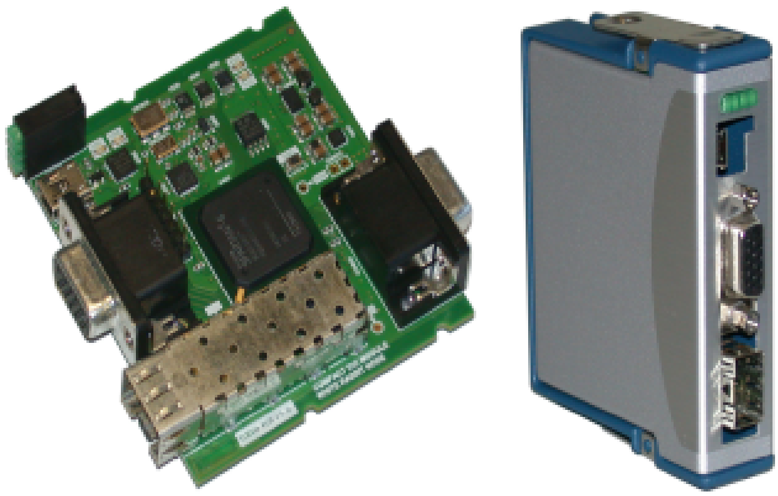

FIGURE 6.3: cRIO-WR module.

CRIO-WR [130] is a custom board with native WR support based on an XC6SLX45T Spartan-6 FPGA (Fig.6.3) The board is originally derived from and keeps maximum firmware compatibly with the established boards SPEC [131] and CUTE-WR [132]. The FPGA gateware includes the WR PTP Core [133][134], the WR streamers [135][136], the needed logic to extract and decode the B-train frames over WR and to manage the cRIO Serial Periferal Interface (SPI) interface. In addition, some glue logic is used to allow a good interface between the FPGA modules.

6.2.2 Block Diagram

Optical fibers coming from the various B-Train and connected to the optical fiber switch, as presented in Fig. 6.4, are plugged into the CRIO-WR module. The WR PTP Core and the streamers extract the Ethernet payload sent via WR and pass it to the BTrain receiver module decoding the raw signals forming the B-frame. In the current implementation, a new valid frame arrives every 4 μs (250 kHz).
In this proof-of-concept, header, B, Bdot, and oldB are sent. This module takes care of sending the data to the controller's FPGA via cRIO interface.
The structured data arrive in the FPGA target, in which, a comparison is performed to avoid the process of the same frame twice. The data pass then through a FIFO to cross the two clock domains present inside the FPGA target, 40 MHz and 5 MHz. The raw values are pulled from the structure payload and filled into a FIFO, one for each slot. Data is then retrieved by the host application through the Ethernet interface. Thanks to the real-time target, received and decoded data are sent to the network and to the host PC on the same subnetwork of the chassis. On the host side, data are collected and elaborated to be plotted and logged into .csv files.

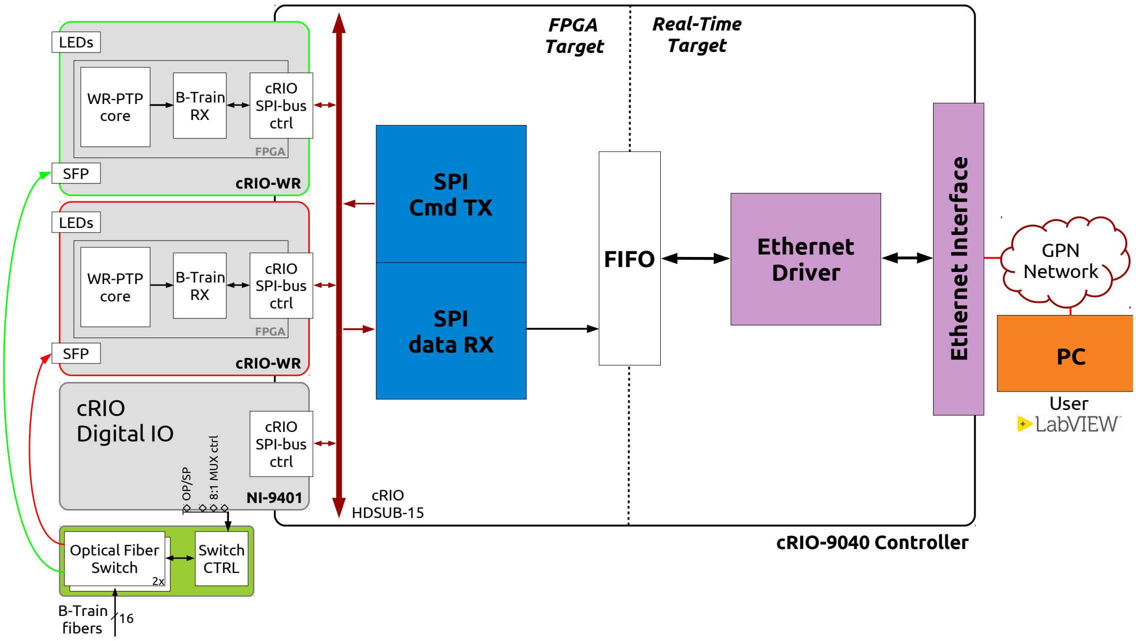

FIGURE 6.4: Block diagram of the B-train monitoring system.

Chapter 7

Monitoring system implementation

In this chapter, the implementation of the new monitoring system is described. First, an overview of the used software tools is illustrated. Later, the implementations of all the FPGA modules and of the Labview host application are described in detail. Finally, the obtained results proving the fulfillment of all the requirements are fully discussed.

7.1 Used software

- LabVIEW is a systems engineering software for applications that require test, measurement, and control with rapid access to hardware and data insights. LabVIEW offers a graphical programming approach that helps you visualize every aspect of your application, including hardware configuration. This visualization makes it simple to develop data analysis algorithms, and design custom user interfaces. LabVIEW2017 was used to develop the software running on the host computer to interface with the controller and to have the graphical user interface of the system.

- LabVIEW FPGA is an additional module of the LabVIEW platform. It allows to program FPGAs present in some National Instruments chassis directly in LabVIEW but with the possibility to import and reuse existing HDL (VHDL, Verilog) code with the IP Integration Node. LabVIEW FPGA is also fully equipped with built-in simulation capabilities and debugging tools. This tool was used to program the FPGA embedded in the cRIO-9040 controller.

- Xilinx Integrated Synthesis Environment (ISE) is a software produced by Xilinx for synthesis and analysis of Hardware Description language (HDL) designs. The version 14.7 was used to synthesize of the code for the Spartan-6 FPGA and to program the flash memory present on the cRIO-WR board to allow the self programming of the FPGA at each reboot.

- EDA Playground is an online tool that allows simulating HDL circuits with commercial simulators. EDA Playground is specifically designed for small prototypes and examples and for this purpose it was used. This tool was used to test the developed sub-modules to certify their functionality before to integrate them into the main project.

- ModelSim is a multi-language HDL simulation environment by Mentor Graphics, for simulation of hardware description languages such as VHDL, Verilog, and SystemC, and includes a built-in C debugger. ModelSim can be used independently, or in conjunction with Intel Quartus Prime, Xilinx ISE, or Xilinx Vivado.

7.2 VHDL and Software Development

7.2.1 cRIO-WR FPGA Gateware

The FPGA gateware on the CRIO-WR module was developed starting from an example project available on the CERN Open Hardware Repository [130]. The existing demo project contains only the WR-PTP core and it sends a time-stamp of a few bytes. To reach the desired goals, some blocks were added: Streamer module to extract the Ethernet payload out of the WR-PTP core; BTrain Transceiver module to decode the user specific B Frame [137]; and cRIO SPI-bus controller to send/receive data to/from the cRIO embedded controller. B frames are identified through the header and latched out when a new valid frame arrives. The cRIO SPI-bus controller block is then transmitting the data into the cRIO controller.

7.2.2 FPGA Kintex-7 70 T

The cRIO controller's embedded FPGA collects the data from the CRIO-WR modules and provides the decoded signals to the host user application. On the FPGA there are two different clock domains, 40 MHz for the SPI interface with the CRIO-WR modules and 5 MHz for the host application and DIO module interface.

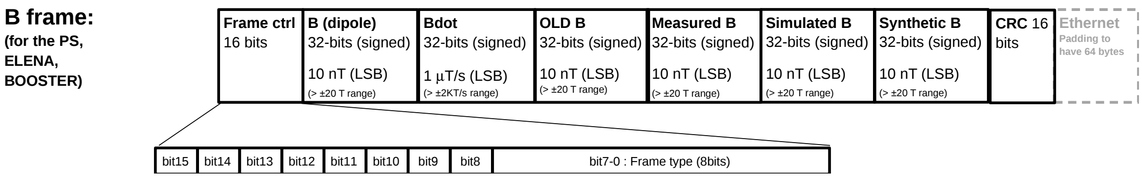

FIGURE 7.1: White Rabbit B-frame.

The data, at the controller's FPGA target, is then connected to the cross clock domain FIFO (2k samples of 14 bytes), in which its output is split into 5 different signals, the WR B-frame slots. FIFO registers were needed to have all slots synchronized and aligned for plotting on the host application.

The used FIFOs have a capacity of 2024 samples of 32 bits on the target side and a capacity of 250k samples of 32 bits on the host side. These values were calculated in order to avoid timeouts caused by FIFO overflow.

The DIO module is also controlled by the controller's FPGA. It drives the switch controller and multiplexes the 16 optical fiber switch inputs to its 2 outputs, OP and SP chains. The fibers are then connected to the CRIO-WR and processed resulting in 16 different monitored systems with only two fiber SFP port inputs.

7.2.3 LabVIEW Host Application

On the host side, a dedicated LabVIEW user application was developed. At each start of its VI, the depth of the FIFO registers is initialized (5x 250k sample FIFOs of 32 bits each).
The main while loop reads out and extracts the data from the 32 bits five slots. A scaling factor is then applied and samples are decimated. The decimation is performed to shrink the decoded samples for plotting purposes.
As shown in Fig. 7.2, there are two graphs, one for B, old B, and C0, and a separate one for Bdot. Each waveform on the graphs can be enabled or disabled depending on the analysis to be performed. In a second loop, based on events, the multiplexer is controlled to select which measurement chain to be plotted. It is also possible to log the signal waveforms into .tdms and/or .csv files for further post processing with other tools, assuring a correct time synchronization due to the White Rabbit protocol implementation.

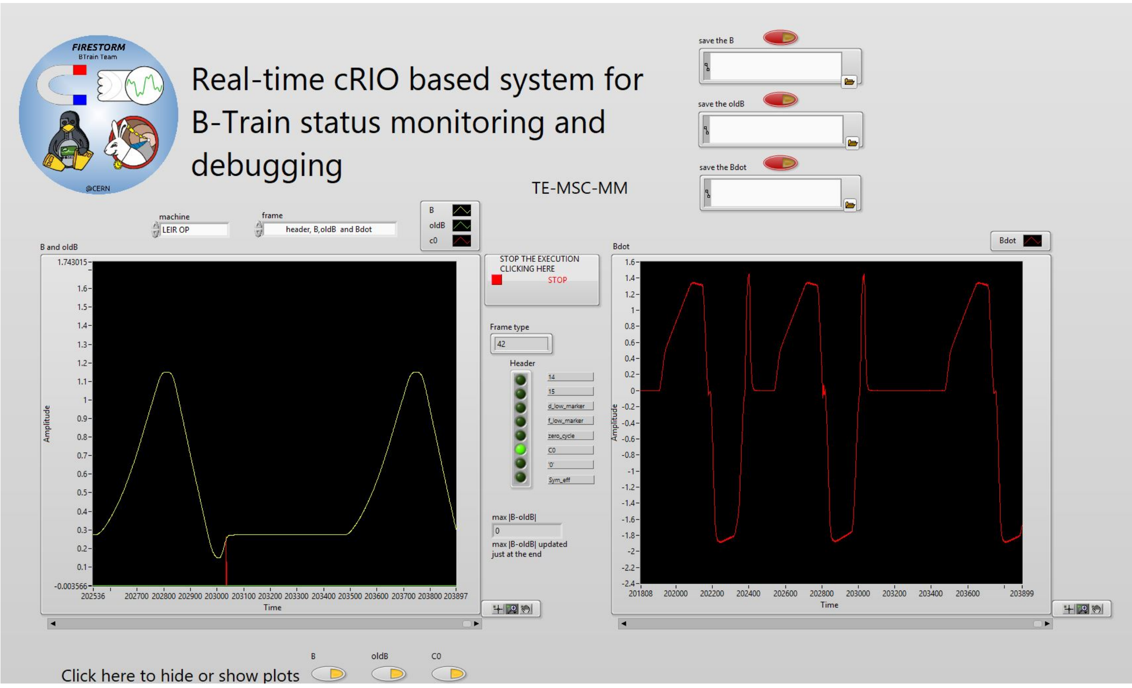

FIGURE 7.2: B-train monitoring proposed architecture.

7.3 Results

The whole full system was assembled and tested in the B-Train lab at CERN. The tests were carried out following this list of requirement.

- System stability
- Correct reception of the WR Frame and correct decoding of the B_Train frame
- Satisfaction of real-time requirements (reduced random delay)
- Correct switching between the machines
- Correct frame switching
- Correct file logging

In the following paragraphs the chosen method for the most critical test where described more in detail.

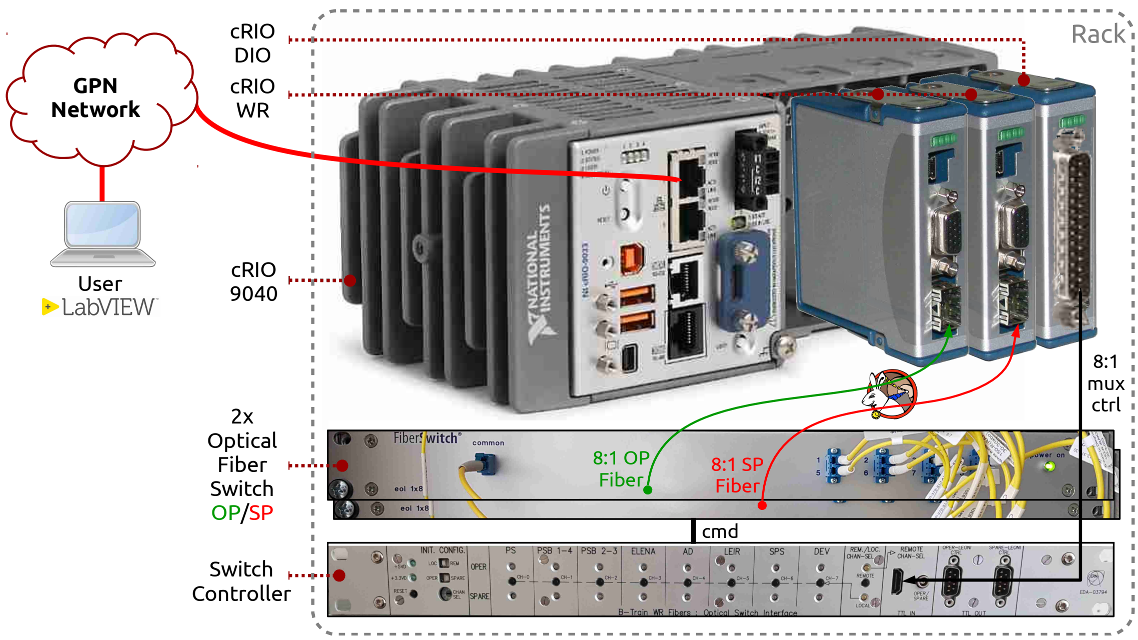

FIGURE 7.3: B-train monitoring proposed architecture.

The legacy Python application and the FESA navigator were used as a reference to prove the correct decoding of the WR frames. Starting from a known point of a known cycle the data were saved and compared with the same cycle acquired with the other two monitoring tools during this test no problems were discovered and the data were always correctly decoded. The real-time feature was tested looking at the debugging info provided by the WR-switch and collected in a Grafana based user interface. The WR switch reported no frame lost in all the observation periods, also the delay was within the specification. These tests reported in Fig.7.4 and in Fig.7.5 validated the system real-time capabilities.

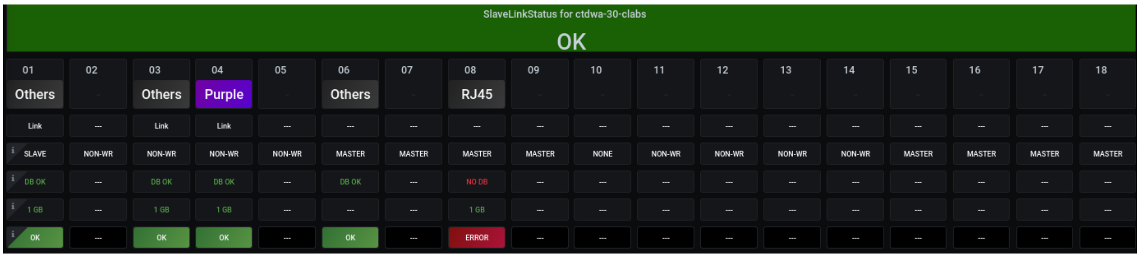

FIGURE 7.4: WR witch status connected to the monitoring system in a period of a month. the only error is related to an unplugged RJ45 debug adapter not related to this work.

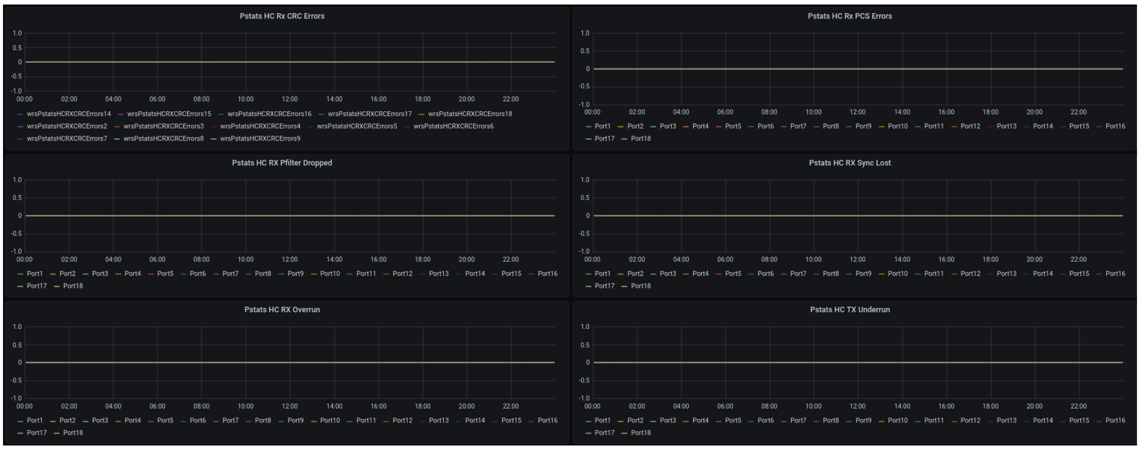

FIGURE 7.5: TX and RX frames flow through the switch connected to the monitoring tool in a period of a month.

The complete system stability was tested letting the system run for a month. The only reboot of the system was caused by an accidental power cut in the lab. The presented device is still running to be used on the field and to collect data about its stability to improve it even further.

Part IV

Metrological characterization and calibration

Chapter 8

DC performance

In this chapter, the results concerning voltage, magnetic flux and integrator drift are discussed separately. The accuracy of the integrator acquisition chain under DC input conditions was evaluated by applying a known reference voltage to V_{in}, and comparing it to the measurement results. The voltage source used was a Data-Precision-8200 multifunction calibrator, which is characterised by a nominal resolution of $10\,\mu\mathrm{V}$ and an RMS output noise level of about $300\,\mu\mathrm{V}$. A total of 20 values in the range between $\pm 8.75\,\mathrm{V}$ were applied sequentially, while at the same time the DAC output was measured with an Agilent 34401A multimeter [122][138].

8.1 Voltage measurement

The accuracy of the voltage measurement, performed with the ADC 7986 providing a Voltage Reference (VREF) equal to $5.0V$, including the in-built gain and the coarse offset corrections, was determined by comparing the known input with the values of V_{ADC} and V_{out} taken from the internal FPGA registers. During the tests the fine offset correction $\Delta \mathrm{V}_2$ was set to zero, and the results are plotted in Fig. 8.1. The error bars represent the repeatability $\sigma_{\mathrm{V}} = 400\,\mu\mathrm{V}$, obtained from the standard deviation over 150 measurements, corresponding to about $5\,\mathrm{LSB}$. For the most part, this scatter is due to the external source, as confirmed by performing the measurements with the input shorted, in which the intrinsic noise of the acquisition chain is about $100\,\mu\mathrm{V}$, slightly more than $1\,\mathrm{LSB}$.

The difference between V_{ADC} and V_{in} provides an indication of the error introduced by the conditioning and digitisation stages, which from Fig. 8.1 is approximately linear. It can be seen in Table 8.1 that the slope and offset, as obtained from an off-line linear regression absed calibration, are very close to the parameters determined by the in-built correction algorithms this prove the effectiveness of the in-built correction method. Since these are more than one order of magnitude above the nominal ratings of the ADC, the error must be ascribed essentially to the analogue conditioning stage. From the difference between V_{out} and V_{in}, it is possible to determine the residual error following the in-built correction, which is random across the whole input range and has an RMS average of $135\,\mu\mathrm{V}$, i.e. a little less than $2\,\mathrm{LSB}$.

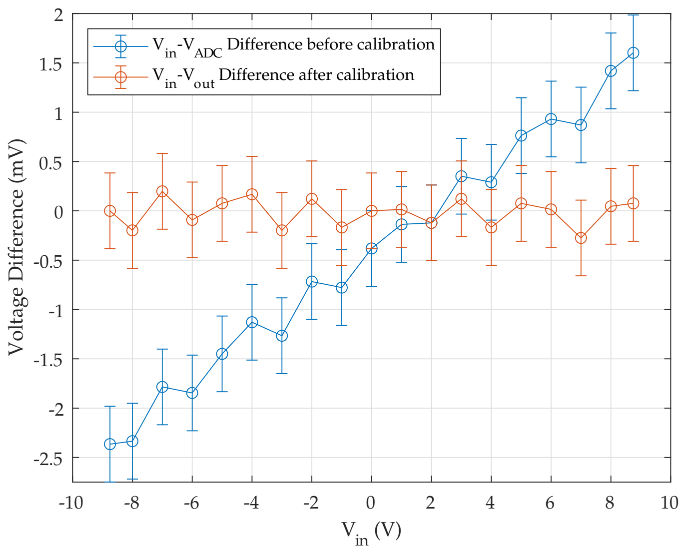

FIGURE 8.1: Error of FIRESTORM digital voltages with respect to measured input value.

Calibration Method	ΔV_1 (μV)	$G_{cc} - 1$ (ppm)
In-built Correction	381	221
Manual Linear Least-Squares Fit	423	220

TABLE 8.1: Comparison between calibration methods that demonstrate the effectiveness of the in-built correction.

8.2 Flux measurement

The accuracy of the flux measurement was determined by integrating a constant input voltage V_{in} in the range $\pm V_{\text{ref}}$ over a precisely set duration of $\Delta t = 1\,\text{s}$, taking a zero integration constant consisting in integrating the measured values when the input of the system is short circuited, and comparing the output $\Delta\Phi$, as read from the FPGA register, to the expected value $V_{\text{in}}\Delta t$. The results, expressed in terms of the equivalent voltage difference:

$$\Delta V_{\text{e}} = \Delta\Phi/\Delta t \; - \; V_{\text{in}}, \tag{8.1}$$

are plotted in Fig. 8.2. It can be seen that all measurements lie within the expected range of $\pm 1/2$ LSB from V_{out}, with the error bars representing the standard deviation over 1000 consecutive repetitions. On RMS average, the repeatability thus evaluated as the standard deviation of the data reported in Fig. 8.2 is about $3\,\mu\text{V}$ or, equivalently, $0.04\,\text{LSB}$. The improvement compared to the voltage noise level can be attributed to the numerical integration suppressing high-frequency noise components. The RMS average of the mean errors across the whole input range is $141\,\mu\text{V}$, corresponding to about $2\,\text{LSB}$. This is consistent with the residual error of V_{out} reported in Section 8.1.

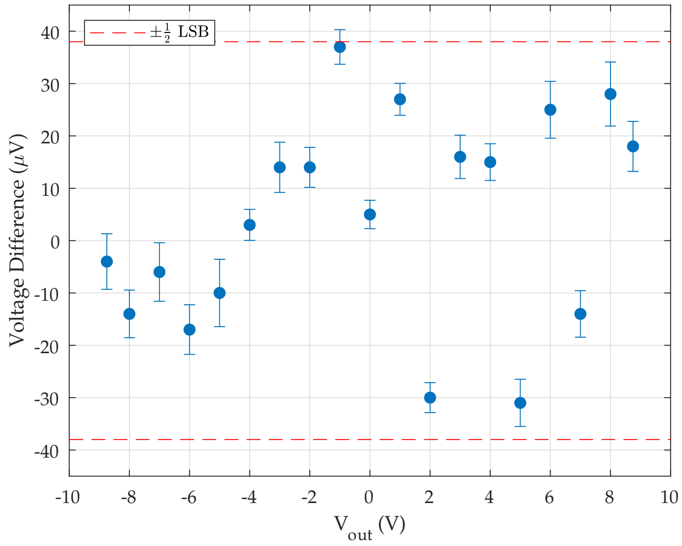

FIGURE 8.2: Mean and standard deviation of ΔV_{e} (8.1) over the full input dynamic range.

8.3 Integrator drift

The short- and long-term performances of the in-built offset voltage compensation method were evaluated by first shorting the input, then retrieving the measured flux waveforms over integration periods of duration $\Delta t = 1\,\mathrm{s}$ and $\Delta t = 120\,\mathrm{s}$. The shortest duration is representative of the typical cycle lengths in most of CERN accelerators; at the other extreme, two minutes is the longest expected in ELENA antiproton cycles. Measurements were repeated respectively 1000 and 8 times and the results are plotted in Fig. 8.3a and 8.3b. The average and standard deviation of the equivalent voltage offset $\delta V = \Delta\Phi/\Delta t$ are given in Table 8.2 for each set of measurements.

The overall RMS voltage offset after the in-built correction, calculated as the average of the equivalent voltage offset over 1000, and reported in Table 8.2 is about $8\,\mu\mathrm{V}$, which under typical operating conditions i.e. using a coil of area $A_{\mathrm{c}} \approx 1\,\mathrm{m}^2$, is equivalent to a measured field drift of the order of $8\,\mu\mathrm{T\,s^{-1}}$. Such an error is usually acceptable for the shorter cycles, but may not be so for longer ones; to cover for this case, a specific novel correction strategy is currently under development [120]. To conclude, the curves in 8.3b illustrate the time evolution of the offset, which is observed to fluctuate with a scatter as large as $\sim$50 % of its mean value on a time scale of a few seconds. This result allows to assume that the offset is constant during the whole magnetic cycle, considering also that the standard magnetic cycle duration is always under 3 s.

TABLE 8.2: Linearly estimated offset voltage

Integration Period (s)	Mean (μV)	σ (μV)
1	7.7	3.5
120	6.5	3.6

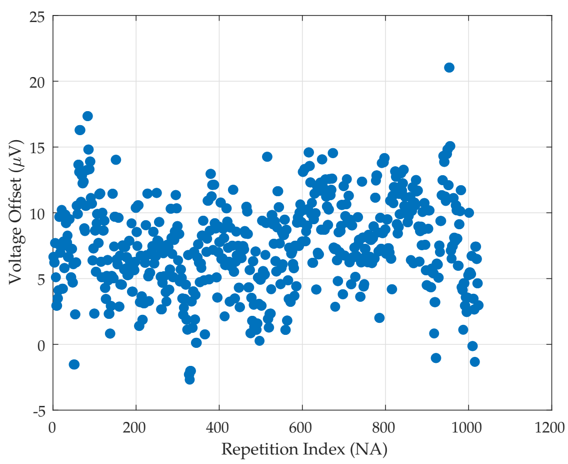

(A) Linearly estimated offset over 1 second plotted as a pseudo time-series to highlight the its evolution over the repetitions.

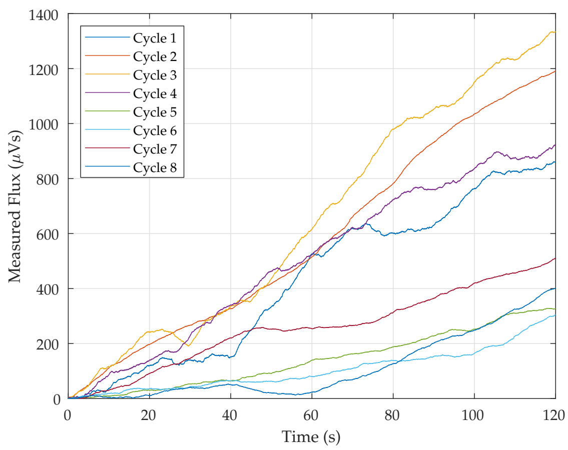

(B) Integrator drift over 120 second.

FIGURE 8.3: Results of integrator drift tests.

Chapter 9

Dynamic performance

In this chapter, the measurements of amplitude transfer function and the latency of the whole acquisition chain are presented. the measurement were performed on a the test setup represented schematically in Fig. 9.1, including: a signal generator, a complete FEC with Integrator, Simulated Field and WR output modules, an external WR switch (simulating the distribution network) and a WR receiver (simulating an end user).

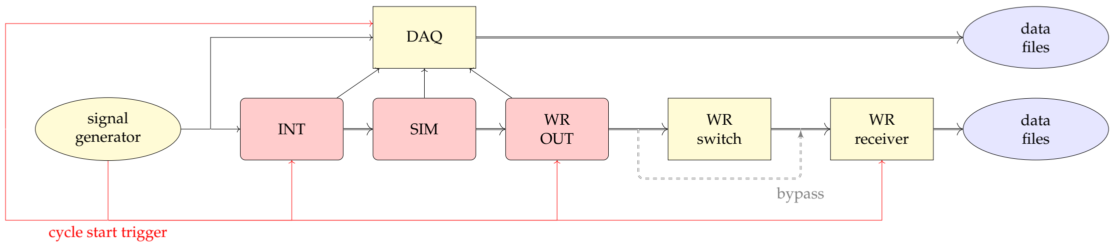

FIGURE 9.1: Schematic layout of the dynamic performance test setup.

9.1 AC amplitude transfer function

The gain response was measured as a function of frequency by injecting into the integrator for $1\,\mathrm{s}$ sine waves of varying amplitude i.e. 1, 5 and 10 $\mathrm{V_{pp}}$, and then retrieving the $\Delta\Phi(t)$ waveform from the White Rabbit stream at the receiver's end. For this test, the receiver used was the WR monitoring system presented in chapter 6, which is able to stream continuously multiple WR channels to disk at a peak aggregated rate of about 100 kfps. This entails an uncertainty of $\pm\ 10\,\mu\mathrm{s}$ on any individual timing measurement, which can however be improved by averaging over a sufficiently high number of repetitions. The peak-to-peak amplitude of the response was then derived as the mean difference between successive maxima and minima, thereby canceling out the effect of integrator drift. The amplitude response ratio is shown in Fig. 9.2a along with the $-20\,\mathrm{dB/decade}$ slope of an ideal integrator, while the difference between the two is magnified in Figure 9.2b. Below $100\,\mathrm{Hz}$, mean errors are within the target tolerance of $100\,\mathrm{ppm}$; whereas, starting from $1\,\mathrm{kHz}$, the effect of the anti-aliasing filters starts to manifest, and the error increases by several hundred ppm. For frequencies below $100\,\mathrm{Hz}$ the scatter of the results, about $200\,\mathrm{ppm}$, is comparable to the errors measured during the DC tests. At 1 kHz the scatter is about one order of magnitude higher, which could be ascribed to the decreasing number of samples per period.

The standard deviation of the integrated voltage was the calculated to be similar, approximately 100 µV, and is attributed to the noise of the signal generator. At higher frequencies, the sampling rate limits the drift correction that can be performed thereby increasing the standard deviation of the measurement results, hence the larger error bars at 1 kHz.

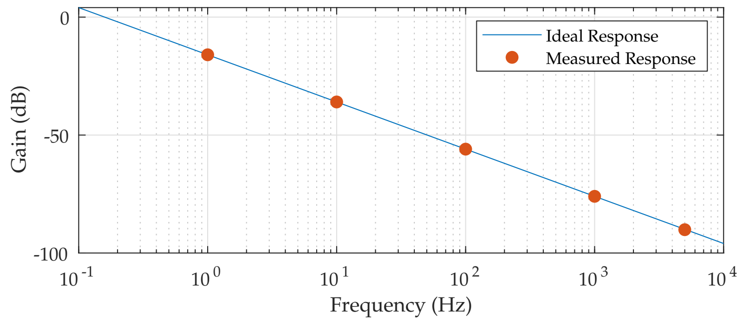

(A) Amplitude frequency response.

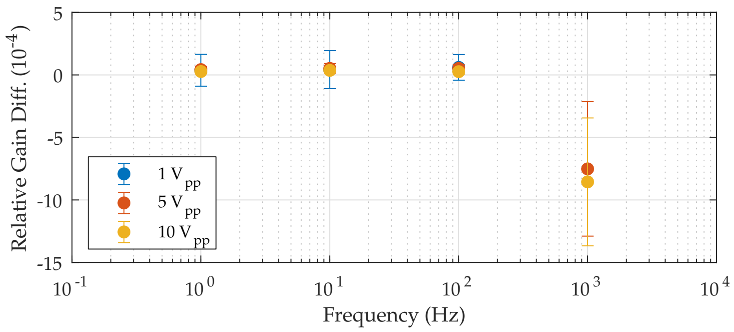

(B) Relative difference between measured and ideal gain (detail).

FIGURE 9.2: Dynamic gain response results.

9.2 Latency measurements

The latency of the whole measurement chain is of fundamental importance for the qualification of the system as a feedback source for the user control loops, and in particular for the RF subsystem. The propagation of a step change in the input along the chain was analyzed, to determine the contribution of each processing and transmission stage. The schematic layout of the test setup is shown in Fig. 9.3. A constant voltage was applied to the integrator input, to generate a constant-slope ramp in the measured field . The initial time reference was given by the C0 cycle start trigger provided by a Central Timing Card, via the B-train crate. (All triggers and DIO connections are done via the crate, which does not appear in the layout for the sake of simplicity.) The C0 trigger was used in place of a field marker to reset to zero the integration, thus providing an easy-to-identify falling edge propagating through the chain.

At the end of the chain a WR receiver was set up, which for this test consisted of an additional SPEC/FMC module, specifically designed for diagnostics. A single WR switch was included to reproduce the configuration common to most deployed systems.

To measure the propagation delay between components, two different, complementary methods were used: 1) the time interval between TTL triggers and/or the diagnostic analogue output of the modules was measured with a NI 6366 USB DAQ, sampling at $10\ \mathrm{MSamples}/s$; 2) the timing information built in the WR stream was retrieved by means of the standard WR diagnostic facilities.

The tests were done at two different WR frame rates i.e. $250\ \mathrm{kfps}$, which is currently the default, and $100\ \mathrm{kfps}$, which is under consideration to match the control loop requirements of the RF subsystem in the AD, ELENA, LEIR and PSB accelerators. The output data rate of the integrator remained fixed at $250\ \mathrm{kfps}$ in both cases.

The results obtained are summarized in Fig. 9.3 and listed in detail in Table 9.1, 9.2 and 9.3, where their extrema, mean and standard deviation (jitter) over 10,000 repetitions are reported.

Below the different quantities that were measured are listed, and the details of the procedure are discussed:

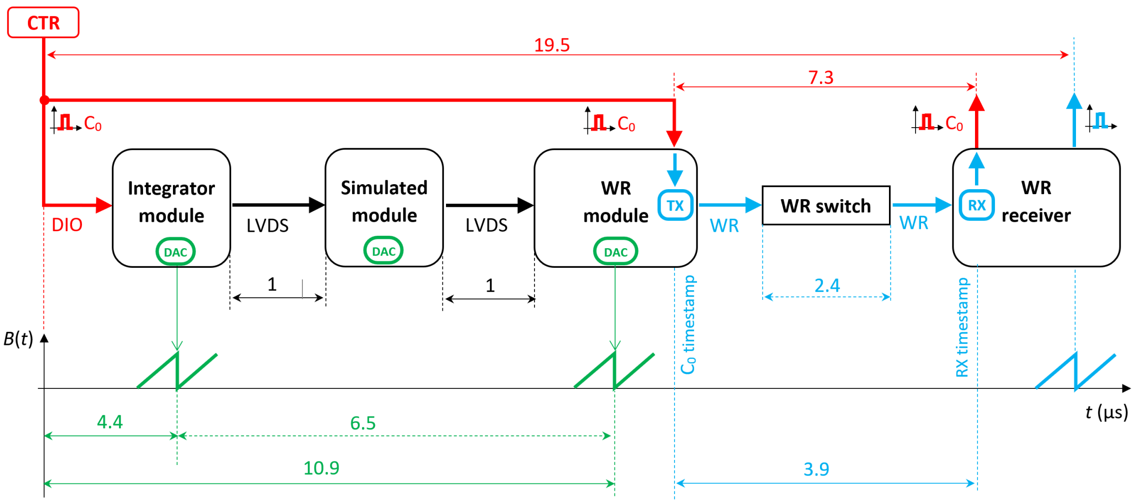

FIGURE 9.3: Latency test of the different components of the acquisition and transmission chain @ 250 $kfps$. In red, TTL trigger pulses associated with the cycle start C_0; in green, the analogue outputs of the SPEC modules; and in cyan, quantities derived from the WR frame. The horizontal axis is used to represent time differences (not to scale).

- The overall mean propagation delay, $19.5\ \pm 2.3\ \mu\mathrm{s}$ ($2\ \sigma$) at $250\ \mathrm{kfps}$, was obtained from the time difference between C0 and a TTL trigger pulse emitted by the WR receiver, as soon as it detected the field step change. The full statistics are reported in Table 9.1. This result is the most important, as it qualifies the whole acquisition chain. The accuracy of a single measurement is dominated by the frame period length in the WR stream, but due to the high number of repetitions the uncertainty of the mean is as low as $0.02\,\mu s$.

- The WR delay, from transmitter to receiver, was measured as the time difference between C0 and a second TTL trigger pulse emitted by the WR receiver. For this, a C0 was directly injected in the WR module via one of the two DIO connections available. There, the WR core set to 1 the start cycle flag in the header of the frames transmitted throughout the duration of the pulse. At the other end, the WR receiver generated a pulse upon arrival of the first active flag. The average delay thus measured is $7.3 \ \pm 1.2\ \mu s$ @ 250 kfps (Table 9.1). The setup included a total fiber length of about $5\,m$ which adds a negligible delay, taking into account a typical propagation speed of $5\,\mu s\,km^{-1}$.
- The WR delay was also cross-checked by subtracting the high-resolution, GPS-synchronized timestamps injected in every frame by the WR transmitter, from the timestamp available at the receiver end. Again, the specific frame used was the first one having its field marker flag set in the Frame Control header, thereby representing the rising edge of the field marker trigger pulse. The result obtained with this method is $3.9 \ \pm 1.2\ \mu s$ @ 250 kfps (Table 9.2), which is significantly lower than the previous result. This difference can be ascribed to the functional sequence of the operations executed in the modules, since timestamping is the last one before transmission, and the first one upon reception. The uncertainty of each single timestamp difference is equal to the WR frame period, i.e. $4\,\mu s$.
- The delay due to the WR switch alone, i.e. $2.4 \ \pm 0.1\ \mu s$ @ 250 kfps, was measured by repeating the previous test while bypassing the switch.
- The details of the propagation through the FEC were measured by a different method, based on generating an analog output image of the integrated field via the DAC built in the FMCs of the Integrator, the Simulated Field and the WR modules. [1] The mean delays in the Integrator alone, and in the whole chain up to the WR module are 4.4 and $10.9\ \mu s$ respectively, with a single-take uncertainty of $1\,\mu s$ due to the sampling rate of the DAC.

The overall delay is well below the specified tolerance of $30\,\mu s$, even taking into account the possibility of multiple switches in the WR network, and of physical fiber lengths up to $100\,m$ as it applies to most installations. [2]

When decreasing the WR frame rate to 100 kfps, the overall latency increases by about $3\,\mu s$, which is half than what could have been expected from the $6\,\mu s$

[1]Since the Simulated Field module is by far the less computationally intensive of the three, its latency is very low and the relative results were unstable, which is the reason why they are not reported here.

[2]In the SPS, the B-train measurement is transmitted about $3\,km$ away to the beam dump subsystem, where it is used by the safety interlock PLC. As this subsystem has a high tolerance, about $1\,mT$, it is unaffected by the additional $15\,\mu s$ delay, even during fast field ramps.

frame period increase. However, this is consistent with the frame period itself being only a small component of the delay, which is dominated by the processing taking place in the three FEC modules in series.

TABLE 9.1: Overall FIRESTORM latency at different transmission speeds

	Input to Output at 250 kHz (µs)	**Transmission at 250 kHz** (µs)	**Input to Output at 100 kHz** (µs)	**Transmission at 100 kHz** (µs)
Average	19.5	7.3	22.6	8.3
Minimum	15.6	4.9	15.7	5.0
Maximum	23.8	9.9	29.7	11.7
2σ	2.3	1.2	3.6	1.9

TABLE 9.2: Delay across the White Rabbit network @ 250 $kfps$

	Without Switch (µs)	**With Switch** (µs)
Average	1.53	3.93
Minimum	1.52	3.86
Maximum	4.32	7.2
2σ	0.13	0.13

TABLE 9.3: Internal FEC propagation time @ 250 $kfps$

	C_0 **to Integrator DAC** (µs)	C_0 **to WR DAC** (µs)
Average	4.4	10.9
Minimum	1.8	8.4
Maximum	6.8	13.3
2σ	1.2	1.7

Part V

Machine learning

Chapter 10

Hysteresis modelling in iron-dominated magnets based on a multi-layered NARX neural network approach

In this chapter, different machine learning approaches for magnetic field predictions are presented, based on tuning a Multi-layered neural network to fit directly the magnet response, by avoiding complementary physical models. Different architectures are considered and selected according to a compromise between the accuracy of the field estimation and the level of complexity of the network. In the optics of a future real-time implementation of the network in FPGA the complexity of the network is relevant not only in the training phases bu also in the operational ones because all the weighs of the neurons have to be locally stored. First a description of the problem statement is presented. Then an overview of the measurement setup and the dataset preparation is depicted. Later, the architecture tuning and the model selection phases were described. Finally, a comparison between all the tested architecture is presented, highlighting the one with the best performances.

10.1 Problem statement and architecture proposal

The problem of real-time magnetic field prediction can be formally expressed as the estimation of a generic unknown output y as a function f of H previous outputs, the input u and K previous inputs:

$$y(n) = f(y(n-1), ..., y(n-H), u(n), u(n-1), ..., u(n-K)), \tag{10.1}$$

where y is the estimated magnetic field, u is the excitation current and n is a discrete time index. The model relies on two buffers, one of network outputs and another of past observations of the input current.
The network's predictive capability is enhanced by endowing it with two buffers taking into account the previous input and output of the system. These are expected to model dynamic features that are either time-dependent, such as eddy current decay transients, or history-dependent, such as magnetic hysteresis. In practice, it is proposed to approximate Eq. 10.1 with a Multi-layered NARX model.

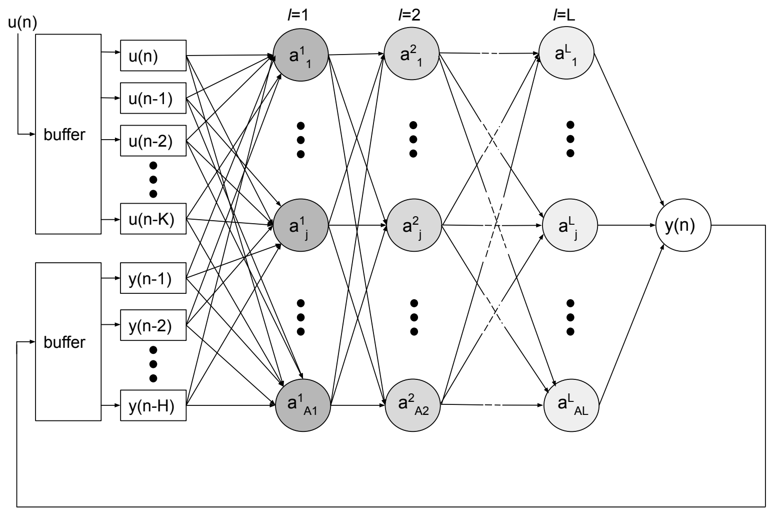

FIGURE 10.1: Multi-layered NARX scheme. The computation of the states a_j^l of the network is forward, starting from first layer (Eq.10.2), then updating internal layers (Eq.10.3), and finally computing the output $y(n)$ (Eq.10.4).

Given a NARX network with L layers (Fig.10.1), the input can be propagated forward through the layers in the following way. The output values a_j^1 of the A_1 neurons belonging to the first layer $l = 1$ of the network are computed by means of the equations:

$$a_j^1(n) = \phi^{A_1}\left(\sum_{h=1}^{H} W_{jh}^{O} y(n-h) + \sum_{k=0}^{K} W_{jk}^{I} u(n-k)\right) \tag{10.2}$$

where ϕ^{A_1} is the activation function of the neurons of the layer 1, the weights W_{jh}^{O} control the strength of the connections from the h-th output to the neuron j of the layer 1 and the weights W_{jk}^{I} control the strength of the connection from the k-th input to the neuron j of the layer 1. Similarly, for each successive layer it is possible to compute the output values of each neuron by the equations:

$$a_j^l(n) = \phi^{A_l}\left(\sum_{i=0}^{A_l} W_{ji}^{l} a_i^{l-1}(n)\right) \tag{10.3}$$

where $l \in \{2, \ldots, L\}$ is the layer index, ϕ^{A_l} is the activation function of the neurons of the l-th layer, and W_{ji}^{l} are the weights that control the strength of the connection from the neuron i of the layer $l-1$ to the neuron j of the layer l. the conventional formulation is used in which $a_0^{l-1} = 1$ are set in order to simulate the bias term W_{j0}^{l}. The value of the output neuron y is computed as:

$$y(n) = \sum_{i=0}^{A_{L-1}} W_i^E a_i^L(n) \tag{10.4}$$

where the identity as activation function is used, while the weights W_i^E control the strength of the connection coming from the neuron i of the layer L to the output neuron y. In the simulation the activation function $\phi^{A_l}(x) = \tanh(x)$ is set for all the layers $l \in \{1, L\}$. In summary, the architecture is defined by an array of hyperparameters $\theta = (L, \mathbf{A}, K, H)$.

The presented modellization takes inspiration from recent 'Jordan' NARX neural network models [139, 140] and it is augmented with internal static layers. Although in those models the output layer is sent back to the input layers (that is why this kind of modelization is also referred to as recurrent), it is worth noting that the computation is completely forward. Consequently, this approach avoids the shortcomings of Backpropagation Through Time (BPTT) learning [141], which requires unfolding the network through time for as many timesteps as there are in the sequence, which significantly slows down learning and/or causes large memory consumption. Note that the presented formalizations collapse to the definition of Deep MLP (when $H = K = 0$) and Deep TDNN (when $H = 0$). Selecting a model architecture is therefore equivalent to assigning a set of integer values to the hyperparameters $\boldsymbol{\theta}$.

10.2 Measurement Setup

An extensive measurement campaign was performed as a case study on a spare reference quadrupole available at CERN (Fig. 10.2). The measurement setup is represented schematically in Fig. 10.3. The magnet was fed by an A&D AG BIP1540 power supply, capable of providing an output current up to 40 A and an output voltage of 10 V. To control the power supply, an NI PXI 4461 card is used driven by custom C++ software based on the Flexible Framework for Magnetic Measurements (FFMM) [142]. The current was measured with a LEM IT60 ULTRASTAB Direct Current Current Transducer (DCCT) having an accuracy of 3 ppm.

The magnet was excited with ten different cycle sequences. In order to enforce a specific initial condition (for H and B) it was decided to perform a degaussing of the magnet before starting the acquisitions [143]. In this way, the initial conditions with H=0 A/m and B=0 T are ensured. The described shape of the waveforms was designed to use curves similar to the ones actually used in magnets for particle accelerators [144].

Each sequence includes seven trapezoidal cycles about 4 to 5 seconds long, starting from $I = 0$ and reaching increasingly higher flat-top values, designed to scan the whole interior area of a major hysteresis loop. The major loop corresponds to the maximum applied current of 25 A and each flat-top level represents an inversion point in the magnetic history, which determines the subsequent branch of the hysteresis loop. The flat-top levels in the different sequences were

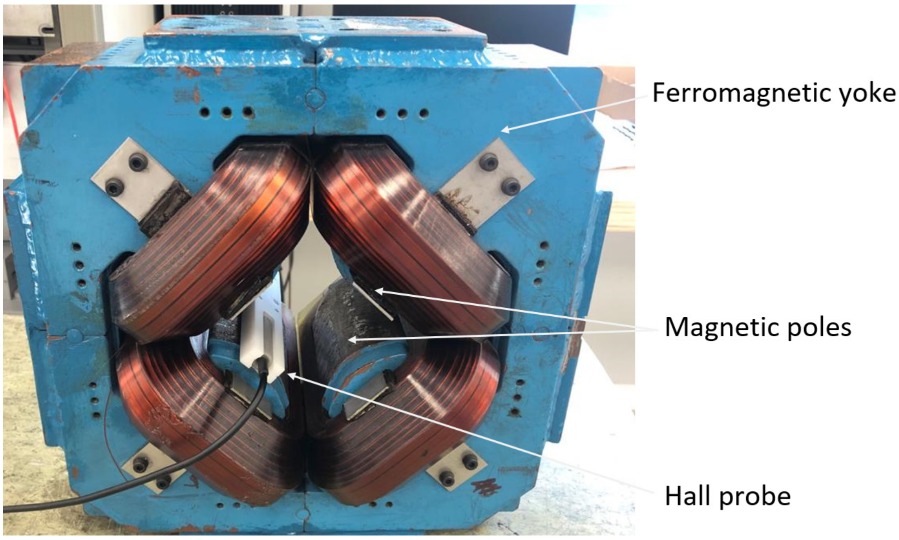

FIGURE 10.2: Reference quadrupole used for the case study. At the tip of the pole, the Hall probe based FM302 teslameter provided by Projekt Elektronic GmbH was used to measure the magnetic field.

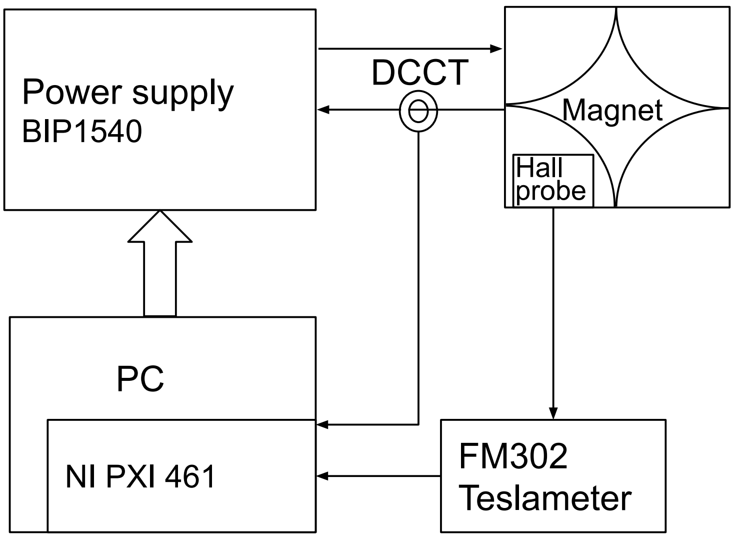

FIGURE 10.3: Schematics of measurement setup.

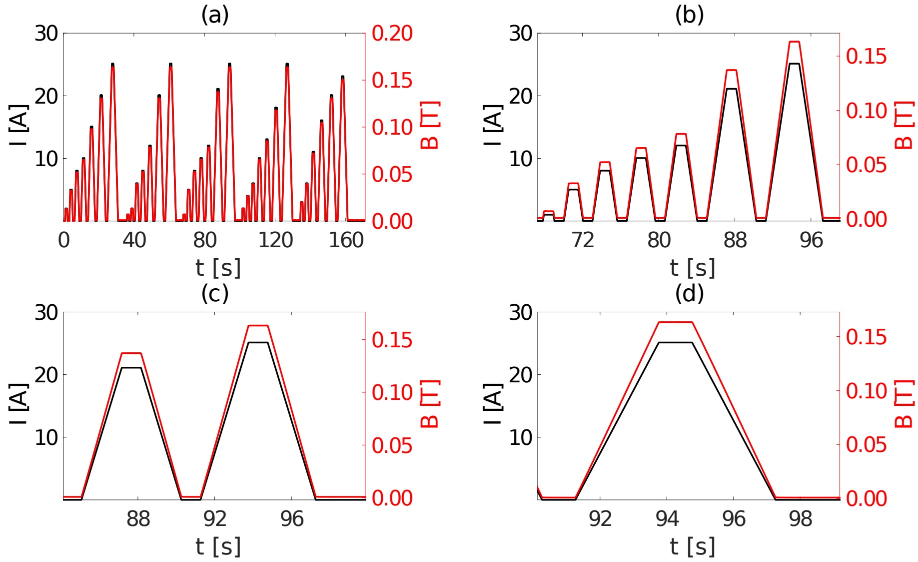

FIGURE 10.4: Sample Plot of the current $I(t)$ (in black) and the magnetic field, $B(t)$ (in red) from the collected data. The abscissa represents the time in seconds (obtained knowing the sample rate of the acquisitions of 2.5 kS/s - corresponding to a sampling interval of 0.4 ms). The ordinate has a double scale: on the left for the excitation current $I(t)$ and on the right for the magnetic field $B(t)$. Panel (a) shows the form of a whole sequence of data collected. Panel (b) shows a selection focusing on the first set of cycles, containing seven magnetic cycles; Panels (c) and (d) represent a further zoom on the last two and one magnetic cycle of the set of cycles respectively.

all different, in order to test the interpolating capability of the network. All different levels are listed in Tab. 10.1, while a sample subset is plotted in Fig. 10.4. The field ramp-rate, defined as the slope of the excitation current during the ramps, is kept constant for all training and test cycles, so as to minimize the impact of this variable on the predictive performance of the neural network. Even if the ramp-rate is assumed to be constant, the magnetic field will still be affected by hysteresis due to the different amplitude of the plateau and eddy currents. This is because the main objective of this paper is modelling the hysteretic part of the response, rather than different dynamics, so the aim was to eliminate as much as possible any confounding variable.

The magnetic field was measured with a Hall probe-based FM302 teslameter, from Projekt Elektronic GmbH, with a sensitivity of 1 V/T. The output voltage was acquired with the same PXI card used to control the power supply, with 24 bits of resolution at the raw sampling rate of 2.5 kS/s. The probe was placed at the tip of one of the magnet poles to measure the peak field in the gap (proportional to the quadrupolar field gradient acting as a magnetic lens on the particle beam). The maximum field measured was $B_{\max} = 0.16$ T. The noise level of the measurement, estimated from its standard deviation on the current plateaus, is approximately $13\ \mu$T, i.e. $8.1 \cdot 10^{-5}$ relative to the maximum.

10.3 Magnet response

The measured relationship g between the current I and the magnetic field B represents the so called *hysteresis graph* of the magnet:

$$B(I) = g(I) \tag{10.5}$$

The relationship appears to be essentially linear, although a zoom-in reveals that the field follows a different path when the current is reduced. The area of the hysteresis loop is indicative of the losses, which include a quasi-static contribution intrinsic to the material, plus a dynamic component due to the eddy currents, which increases with the ramp rate. The maximum width of the loop, relative to the full range of the field, is approximately $1\ \%$ in the region between 7 and 17 A.

It is possible to get an insight into the magnetic response, using a Linear Model [145] and in first approximation neglecting its nonlinear part:

$$B(I) = B_0 + G \cdot I + \hat{B}(I) \tag{10.6}$$

Consequently, it was possible to compute $B_0 = 7.79 \cdot 10^{-4}$ T and $G = 6.50 \cdot 10^{-3}$ T/A, respectively the offset and gain of the least-square linear regression shown in Fig. 10.5.

$\hat{B}(I)$ is the residual of the regression and contains the nonlinear component. This decomposition is crucial in the construction of datasets used to learn and test models on the nonlinear part of the signal.

A different representation of the magnetic behavior as a function of the current is given in Fig. 10.6, in terms of the so-called transfer function T_{f}, defined by the field-to-current ratio:

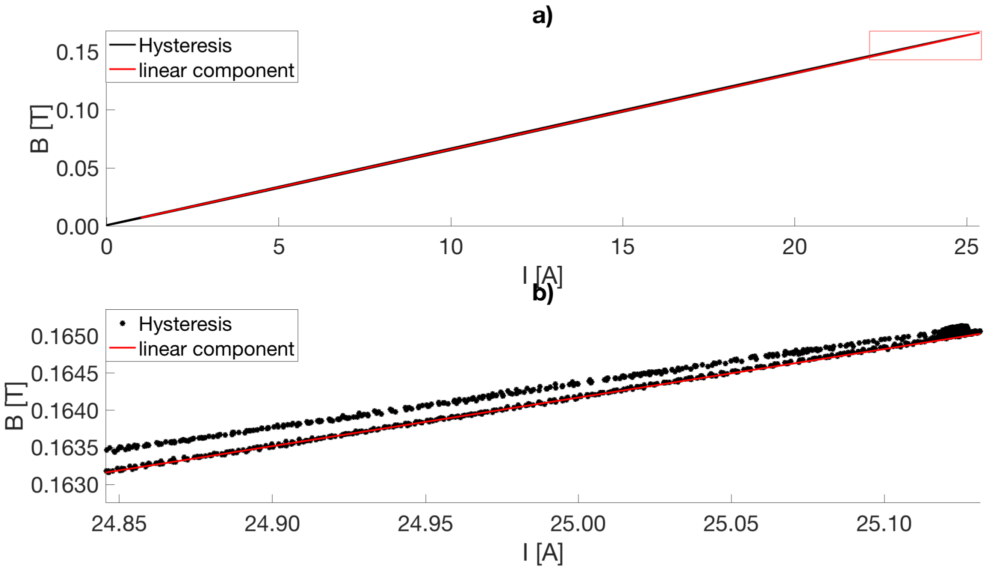

FIGURE 10.5: Hysteresis graph. (a): the field as a function of the current, together with the least-square regression. (b): corresponding plot magnified in a sample rectangle.

$$T_{\mathrm{f}}(I) = \frac{B(I)}{I} = \frac{B_0}{I} + G + \frac{\hat{B}(I)}{I} \tag{10.7}$$

In this figure one can more clearly see how the response switches discontinuously from the lower to the upper branch of the hysteresis loop, as the current starts to decrease at the end of each flat-top. The flatness of the lower branch on the up-ramp between about 5 and 25 A corresponds to an almost constant transfer function, i.e. the desired linear behavior under typical operating conditions. Due to the field level being very low, there is no visible saturation leading to a reduction of the transfer function at high current. The nonlinear component $\hat{B}$ vanishes at the high and low reversal points of the hysteresis loops. As a result, the vertical asymptote for $I \to 0$ can be entirely attributed to the remanent field B_0. The transfer function B/I is equivalent to the H-B graph and it contains the same amount of information, since in a magnetic circuit the field H is proportional to the excitation current I. In the context of this application, the transfer function is the preferred representation because it allows an operator to visualize more readily the degree of linearity corresponding to a given level of excitation, perfect linearity corresponding of course to a constant transfer function [146].

10.4 Dataset Preparation

The raw dataset $\bar{D}$ is composed of the excitation current and field waveforms of the 10 sequences of 7 magnetic cycles, acquired at 2.5 kS/s (sampling time 400 μs) for a total of 871 440 samples. During the network architecture selection and training phases were carried out on a reduced subset D including 87 144

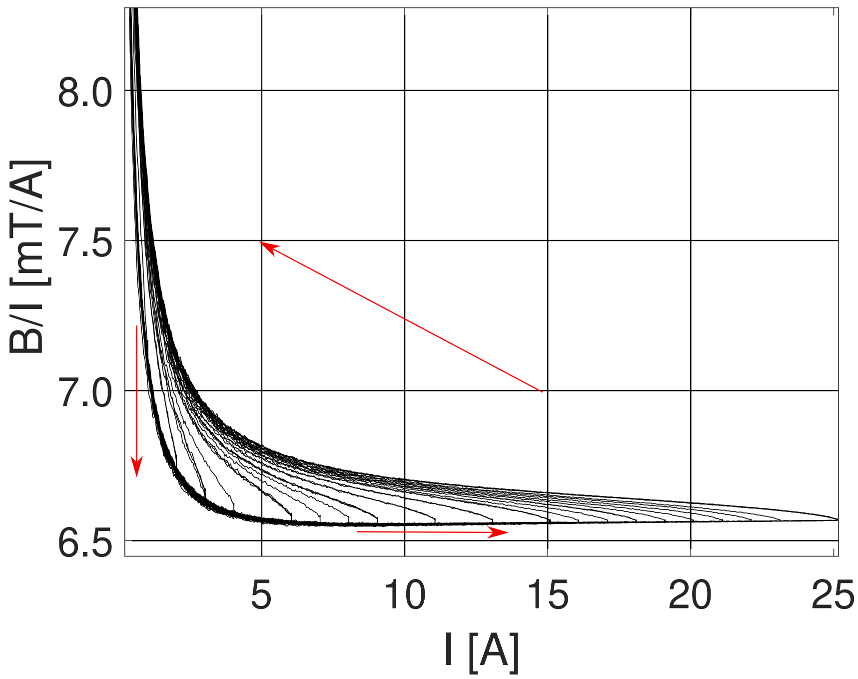

FIGURE 10.6: Transfer function $T_{\mathrm{f}}(I)$ of the magnet, defined by the ratio between magnetic field and excitation current as in Eq. (10.7).

samples at the actual sample rate of 250 S/s (corresponding to a sampling interval of 4 ms.).

One half of the subsampled dataset D, D_{L} was used for the architecture learning, and the remaining half, D_{E} was used for the test and statistical error evaluation.

The data arrays were organized in pairs $[I_{\mathrm{L}}(n), B_{\mathrm{L}}(n)]$ and $[I_{\mathrm{E}}(n), B_{\mathrm{E}}(n)]$ including the measured current and field. The data in D_{L} was further split into training $D_{\mathrm{L}}^{\mathrm{train}}$, validation $D_{\mathrm{L}}^{\mathrm{val}}$ and test $D_{\mathrm{L}}^{\mathrm{test}}$ datasets with a 60:20:20 ratio. The splitting of the dataset is detailed in Tab. 10.1, where the sequences of flat-top currents are also listed. It should be noted that the training subset includes only a few of the possible transitions between different successive flat-top levels. Each different combination is associated with a different branch of the magnetic hysteresis loop, and the accuracy of the inference made on the test combinations gives a measure of the interpolating power of the trained network. In addition, a second version of the datasets, $\hat{D}_{\mathrm{L}} = [\hat{I}_{\mathrm{L}}(n), \hat{B}_{\mathrm{L}}(n)]$ and $\hat{D}_{\mathrm{E}} = [\hat{I}_{\mathrm{E}}(n), \hat{B}_{\mathrm{E}}(n)]$, was created by replacing the measured field with its nonlinear component, as derived from Eq. (10.6), $\hat{B}(I) = B(I) - B_0 - G \cdot I$. In this case, the magnetic field can be computed by adding back the NN output $\hat{y}$ to the previously subtracted linear regression. The rationale of this decomposition is to isolate the physically interesting part of the magnet's response, focusing the training process on a dataset having a much smaller dynamic range.
Architectures were simulated within the Neural Network Toolbox in Matlab 2018b. The training and simulations were performed on a computer equipped with an Intel Core i5, Clock 3.2GHz, Ram 8 GB.

TABLE 10.1: Current cycle flat-top values of the 10 cycle sequences tested. The role of each dataset, be it training, validation or test is given in the second column.

Dataset		Cycle Index							Unit
Index	Type	1	2	3	4	5	6	7	
1	$D_{\mathrm{L}}^{\mathrm{train}}$	2	5	8	10	15	20	25	A
2	$D_{\mathrm{L}}^{\mathrm{train}}$	1	2	6	8	12	20	25	A
3	$D_{\mathrm{L}}^{\mathrm{train}}$	1	5	8	10	12	21	25	A
4	$D_{\mathrm{L}}^{\mathrm{val}}$	2	4	6	10	13	18	25	A
5	$D_{\mathrm{L}}^{\mathrm{test}}$	3	6	11	16	20	23	25	A
6	D_{E}, $\bar{D}_{\mathrm{E}}$	3	6	11	16	20	23	25	A
7	D_{E}, $\bar{D}_{\mathrm{E}}$	2	6	9	13	17	22	25	A
8	D_{E}, $\bar{D}_{\mathrm{E}}$	3	6	9	15	18	21	25	A
9	D_{E}, $\bar{D}_{\mathrm{E}}$	1	3	7	9	13	18	25	A
10	D_{E}, $\bar{D}_{\mathrm{E}}$	2	4	8	11	15	19	25	A

TABLE 10.2: Parameters in input to Algorithm 1

params	Value	Description
epochs	200	Max. training epochs
maxFail	6	Max. validation failures
minGrad	$1 \cdot 10^{-7}$	Min. gradient
muMax	$1 \cdot 10^{10}$	Max. μ value
R	200	Learning repetitions

10.5 Architecture Tuning

10.5.1 Model Selection and Evaluation

An incremental approach was adopted, by increasing the complexity of the model progressively. First, a static structure without feedback was considered, the number of layers L and then the number of neurons on each layer A_l were independently optimized. Next, a feedback was added first on the input, optimizing K, and then on the output, optimizing H to finally achieve a NARX structure (see section 10.5.3). The process was carried out on D_L and also on $\hat{D}_L$ as defined in Section 10.4. This is because, besides testing the capability of the network architecture of reconstructing the magnetic field ($B \approx y$), the capability of obtaining B when focusing only on the learning of its nonlinear component and reconstructing it by adding the linear component was also tested (i.e., $B \approx B_0 + Gu + \hat{y}$). The optimal hyperparameters are given in Tab. 10.3 and were later used also to evaluate the models trained on the full dataset D_{L}.

The pseudo-code representing each step of the selection process is listed in Algorithm 1, where the input dataset D represents either the full or the nonlinear component only version. The algorithm includes three main loops. The first

Algorithm 1 *Model Evaluation and Selection* $(\mathcal{M}, params, D)$

Require: set $\mathcal{M}$ of hyperparameters θ_j for each model to evaluate (see Table 10.3)
set $params$ of simulation parameters (see Table 10.2)
the dataset D
Ensure: set $Score$ of evaluations for each model in $\mathcal{M}$

$[D_{\mathrm{L}}^{\mathrm{train}}, D_{\mathrm{L}}^{\mathrm{val}}, D_{\mathrm{L}}^{\mathrm{test}}] = Split(D_{\mathrm{L}}, 60:20:20)$
for $\theta_j \in \mathcal{M}$ **do** ▷ $Loop1$: iterate over the set of different models
  for $i = 1:R$ **do** ▷ $Loop2$: repeat learning R times
    repeat
      $W = train(\theta_j, params, D_{\mathrm{L}}^{\mathrm{train}}, D_{\mathrm{L}}^{\mathrm{val}})$
    until *term_condition (params)* ▷ $Loop3$: actual execution of a training instance
    $[B_{\mathrm{L}}^{\mathrm{test}}, I_{\mathrm{L}}^{\mathrm{test}}] \leftarrow D_{\mathrm{L}}^{\mathrm{test}}$
    $y_i \leftarrow y(\theta_j, W; I_{\mathrm{L}}^{\mathrm{test}})$
    $Error(i) \leftarrow RMSE(y_i, B^{\mathrm{test}})$ ▷ as defined in Eq. 10.8
  end for
  $Score(j) = model\_evaluate(Error, \theta_j)$ ▷ compute BC scores, Eq. 10.9
end for

loop, $Loop1$, iterates over a set of models $\mathcal{M}$, each one defined by its own hyperparameter vector, θ_j. The second loop, $Loop2$, trains the network and estimates its prediction error R times on the test dataset $D_{\mathrm{L}}^{\mathrm{test}}$, in order to improve the statistical significance of the results. The third loop, $Loop3$, is an actual instance of training performed according to the training parameters given in Tab. 10.2, and the training and validation datasets, $D_{\mathrm{L}}^{\mathrm{train}}$ and $D_{\mathrm{L}}^{\mathrm{val}}$. During the training procedure the hyperparameters of the model θ_j are kept fixed, while the connection weights W are updated iteratively with the objective to minimize the output reconstruction error with the Levenberg-Marquardt (LM) method [147], until one of the termination criteria is met. These are contained in the function $term\_condition$ and include:

- the validation error fails to decrease for $maxFail$ iterations
- the maximum number of epochs for the training ($epochs$) is reached
- the LM damping factor μ exceeds its maximum acceptable value ($muMax$). In the LM algorithm, the factor μ switches continuously from a Newton-like ($\mu \approx 0$) to a steepest gradient descent ($\mu \gg 0$) optimization. Too large values of μ imply that one is too far from a minimum and the search has failed.

At the end of the learning phase, the validation dataset $D_{\mathrm{L}}^{\mathrm{val}}$ is used to optimize the network generalization. The test dataset $D_{\mathrm{L}}^{\mathrm{test}}$ is used to evaluate the prediction performance of the network after the training in terms of the Root Mean Square Error (RMSE), computed for each iteration of $Loop2$ as:

$$RMSE(y_i, D^{\mathrm{test}}) = \sqrt{\frac{\sum\limits_{n \in N} (y_i(n) - B^{\mathrm{test}}(n))^2}{|N|}} \tag{10.8}$$

where $y_i = y(\theta_j, W; I_{\mathrm{L}}^{\mathrm{test}})$ is estimated according to Eq. (10.1) with the current set of hyperparameters and weights, and $|N|$ is the number of samples of the test dataset. To choose the best architecture it is possible to perform a statistical model selection: the model evidence $P(D|\boldsymbol{\theta}_j)$ was maximized, i.e. a probability term that expresses the preference shown by the data for the $j-th$ model of hyperparameters $\boldsymbol{\theta}_j$. In general, the computation of this term is analytically intractable, thus different approximations of this term were proposed in literature [145]. Popular approximations rely on different penalization terms of the model complexity computed as the number of weights $|W|$ determined by specific choice of hyperparameters θ_j, like Akaike Information Criterion (AIC) or Bayesian Information Criterion (BIC).
A recently proposed information criterion, named Bridge Criterion (BC), was used. BC aims at bridging the advantages of both AIC and BIC in the asymptotic regime. [148] To this end, the function $model\_evaluate$ assigns to each hyperparameter θ_j a score based on the BC term:

$$Score(j) = |N| \ln \left(\frac{1}{R} \sum_{i=1}^{R} Error^2(i) \right) + |N|^{2/3} \cdot (1 + 1/2 + ... + 1/|W|) \quad (10.9)$$

Thus, the minimization of the BC term corresponds to the maximization of the *evidence* of the j-th model, ensuring a balanced model fit as the first term in Eq. (10.9) weighs the reconstruction error, while the second term penalizes the number of weights.Moreover, a smaller parameter space allows us to find more stable solutions during the training phase.

10.5.2 Static Network Structures

The model selection started by evaluating the performance of a static network, without any feedback. This is the case of a Deep MLP Neural Network, defined by hyperparameters θ_{MLP} that include the number of hidden layers L, and the number of neurons in each of them, A_l.

An example of static architecture is shown in Fig. 10.8, with two hidden layers characterized by a $tanh$ activation function, while the output layer is linear.
Following Algorithm 1, increasingly complex structures were evaluated. Fig. 10.7 shows the overall model selection guided by BC scores, in four steps: first two steps for selecting the static structure, last two steps for selecting the input-output buffer size. In the first step, an overall comparison was performed over structures with a fixed number of nodes per layer (10) but a different number of hidden layers (up to 15) , aimed at determining the optimal depth for the neural network. In Fig. 10.7a, BC scores are shown as a function of the number of layers. It is possible to appreciate that networks with a number of layers under 8 gave the best BC scores. After 8 layers the BC scores increase, meaning that the performance of the network models is in general poorer. In Fig. 10.7b, the second step of the selection procedure is shown, in which the structures with 2 and 4 layers were selected as suitable candidates for the next selection step, since they provide a good

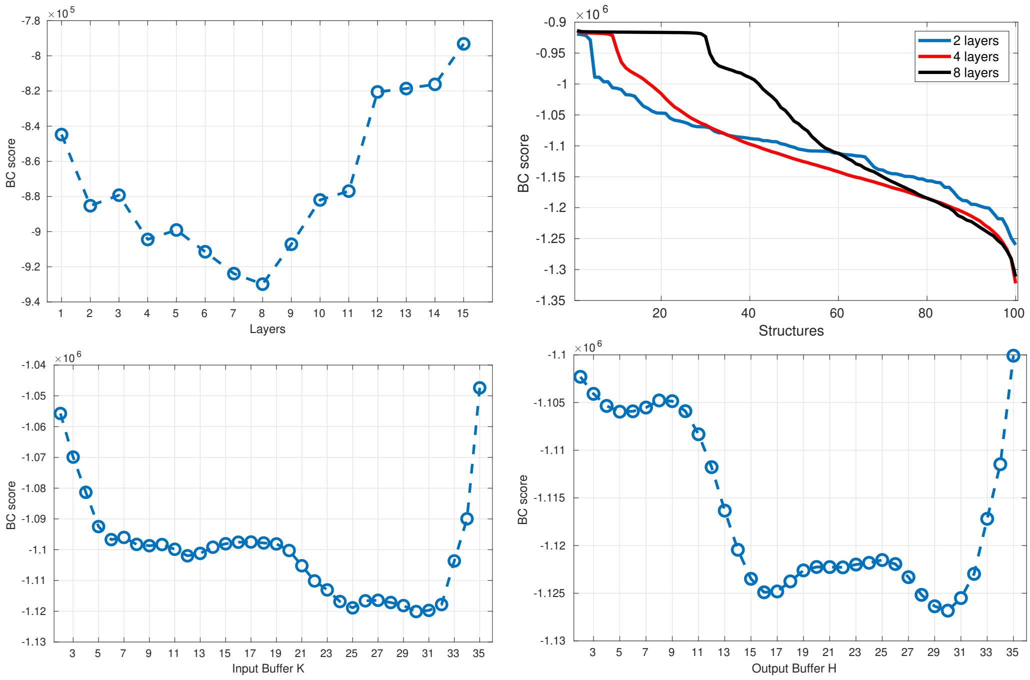

(C) (D)

FIGURE 10.7: The process of model selection in four steps. On the top the static, and on the bottom the dynamic hyperparameter selection, respectively. Panel (a): firstly, once fixed the number neurons to 10 per layer, it was let the number of levels vary and compute BC scores in order to select best promising structures. Panel (b): secondly, fixing the number of the layers, the number of neurons for each layer was varied from 1 to the maximum of 10 neurons, selecting the best 100 winning structures. Panel (c): thirdly, the dimension of the input buffer varid. Panel (d): finally, the dimension of the output buffer were varied while all other hyperparameters remain fixed.

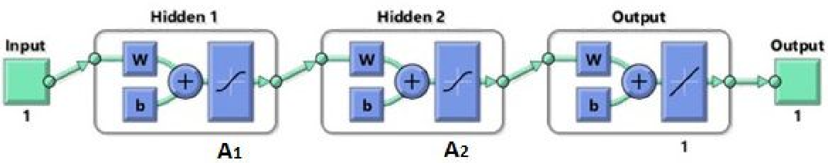

FIGURE 10.8: Generic scheme of a MLP neural network. The first green block represents the input layer, where the input is the excitation current *I(n)*; the second and the third blocks represent the two hidden layers with A_1 and A_2 neurons respectively. In each hidden layer, there are two blue boxes representing the weights, *W*, and the bias, *b*; the two hidden layers use a *tanh* activation function. The fourth block represents the output layer, which is linear. The output of the neural network is the predicted value $y(n)$ estimating the magnetic field $B(n)$ and corresponding to the last green block. Note that this scheme is obtained as outputs of the Matlab Toolbox "nntraintool", in which the parameters used for the simulations were inserted.

TABLE 10.3: Hyperparameters definition for the model selection.

Neural Network Model	Multilayer Perception	Time Delay	Autoregressive Exogenous
Hyperparameter	θ_{MLP}	θ_{TDNN}	θ_{NARX}
L	$\{1,\ldots,15\}$	$\tilde{L}_{MLP}$	$\tilde{L}_{MLP}$
$\mathbf{A}$	$\{1,\ldots,10\}^L$	$\tilde{\mathbf{A}}_{MLP}$	$\tilde{\mathbf{A}}_{MLP}$
K	0	$\{1,\ldots,35\}$	$\tilde{K}_{TDNN}$
H	0	0	$\{1,\ldots,35\}$

compromise between performance and complexity, and once verified that even adding 8 layers do not significantly improve BC scores. With these two steps the procedure helps us select static MLP structures with hyperparameters L,**A**, shown in Tab. 10.4.

10.5.3 Dynamic Network Structures

The static models was enhanced with temporal feedback. The strategy consisted in adding feedback to hidden layers, starting from the best static network structures selected previously.

When buffering past observation in input, the model architecture reduces to a deep TDNN with hyperparameters θ_{TDNN}.

During this step of model selection, the best hyperparameters found in the previous section for MLP networks ($\tilde{\theta}_{MLP}$) were used, while the impact of the input delay buffer length K were explored as an additional hyperparameter. Fig. 10.7c and Fig. 10.7d show the BC scores as a function of K and H in the range $\{1,\ldots,35\}$ for the best two- and four-layer structures selected in section 10.5.2. BC remains low for K between 5 and 33, with a minimum range between 25 and 32. For higher buffer lengths the RMSE increases rather steeply, which indicates instability. Similar behavior can be appreciate for K in the range between 13 and 32, after that performance for higher buffer lengths rapidly degrades. The resulting architecture is a NARX with hyperparameters θ_{NARX}, including the length H of the output buffer. Also in this case, the model was fine tuned with an incremental approach: the set of optimal model hyperparameters previously selected (L, $\mathbf{A}$ and K) was fixed and only H was varied in the range $\{1,\ldots,35\}$. It was found that the best performances are associated with long buffers: in Tab. 10.4 the best combinations of hyperparameters $\tilde{\theta}$ selected by Alg. 1 are shown.

10.6 Results and Discussion

The performance of these models was evaluated on a completely new collected dataset D_{E}, which include different sequences of hysteresis inversion points with respect to the learning phase, in order to stress the interpolating power of the networks.

TABLE 10.4: Hyperparameters $\tilde{\theta}$ selected by Alg. 1 and used in tests shown in Section 10.6. In the last column, the corresponding number of neuron connections are listed $|W|$.

Hyperparameters	L	$\mathbf{A}$	K	H	$\|W\|$
$\tilde{\theta}_{MLP1}$	2	$(10, 9)$	0	0	129
$\tilde{\theta}_{MLP2}$	4	$(1, 1, 1, 10)$	0	0	37
$\tilde{\theta}_{TDNN1}$	2	$(10, 9)$	26	0	379
$\tilde{\theta}_{TDNN2}$	4	$(1, 1, 1, 10)$	31	0	67
$\tilde{\theta}_{NARX1}$	2	$(10, 9)$	26	31	689
$\tilde{\theta}_{NARX2}$	2	$(7, 8)$	26	17	381
$\tilde{\theta}_{NARX3}$	4	$(1, 1, 1, 10)$	31	31	98
$\tilde{\theta}_{NARX4}$	4	$(1, 8, 5, 4)$	31	31	153

First, for each model the estimation $y = y(\theta, W; I_{\mathrm{E}})$ was computed given by the network on the new dataset. Then, the $RMSE(y, D_{\mathrm{E}})$ was calculated according to Eq. 10.8. In order to facilitate comparison to the requirements, the RMSE were normalized with respect to the maximum measured field:

$$NRMSE(y, D_{\mathrm{E}}) = \frac{RMSE(y, D_{\mathrm{E}})}{B_{\max}} \cdot 100 \tag{10.10}$$

Two other measures of performance were used, i.e. the MAE:

$$MAE(y, D_{\mathrm{E}}) = \max\left\{|y(n) - B_{\mathrm{E}}(n)|\right\}_{n \in N} \tag{10.11}$$

and the MPE, normalized with respect to the maximum measured field:

$$MPE(y, D_{\mathrm{E}}) = \frac{MAE(y, D_{\mathrm{E}})}{B_{\max}} \cdot 100 \tag{10.12}$$

The results obtained are summarized in Tab. 10.5 and Tab. 10.6. In Tab. 10.5 the test dataset D_{E} is considered with samples at the same sample rate of the learning dataset 250 S/s, while in Tab. 10.6 the $\bar{D}_{\mathrm{E}}$ dataset used for the final testing was collected at the raw sample rate of 2.5 kS/s, ten time faster than the learning dataset. In both tables, for each type of model, the reference to the optimal hyperparameter vector along with the corresponding error norms are listed. The optimal hyperparameters are listed separately in Tab. 10.3. The linear regression alone gives a relative error of the order of the percent, which corresponds to the relative width of the hysteresis loop. Such an error, which in other contexts might be taken as an indication of good linearity of the magnet tested, is unacceptable for the considered application. Next, let us consider the results of the networks trained on the dataset D_{L}, which are given in the upper half of Tables 10.5 and 10.6. Both the static (MLP) and the dynamic networks with input feedback (TDNN) perform as the linear regression alone.

The NARX networks, instead, are two orders of magnitude better, achieving a best-case NRMSE of 0.006 %. The results evaluated on the reduced dataset are about a factor of two worse, i.e. 0.01 %.

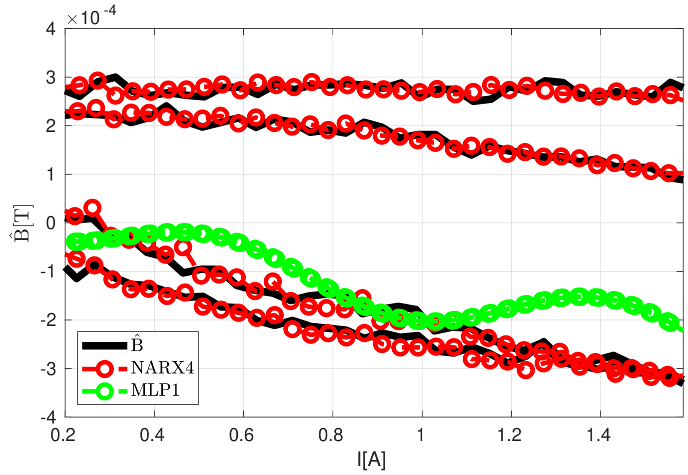

FIGURE 10.9: Approximation of four cycles in a sample interval $[0.2, 1.6]$ of the current I. Black lines show the measured nonlinear magnetic response $\hat{\mathrm{B}}$, along with a NARX4 (red) and MLP1 (green).

Fig. 10.9 gives a qualitative insight on solutions explored by the different models. It is shown within four different cycles the nonlinear magnetic response $\hat{\mathrm{B}}$ (black lines) along with a NARX network (NARX4) and an MLP network (MLP1) response in function of a sample interval of input current I. It is possible to appreciate that across different cycles, the best that an MLP can do is to perform a weighted mean of different magnetic responses, a typical known behavior in machine learning because the network is trying to approximate an input-output signal response that is not a mathematical function, resulting in an ill-posed problem [145].

On the other hand, the addition of the finite-temporal information in input lets the NARX disambiguate the different cycles and gives a fine approximation of the different magnetic responses. This further gives an insight into what the used incremental model selection approach does: the first steps of structural model selection found the best structures that better approximate mean values of the nonlinear signal through different cycles, then next steps of model selection learn the dimension of the buffer that carries out the sufficient dynamical information in order to disambiguate each cycle.

Thanks to having memory of past outputs, the NARXs are shown to be able to reproduce the dynamics of the magnet very accurately.

The maximum length of the output buffer, $K = 35$, corresponds during the training phase to a total duration of 140 ms, much shorter than the time span necessary to cover even just two consecutive inversion points of the magnetic cycle. While in classical approaches, such as the Preisach models, the complete sequence of inversion points is a necessary input to reconstruct accurately the magnetic history, here it was found instead that such information appears to be encoded implicitly by the network, despite the shortness of the output delay buffer. The role of the buffers might therefore be limited to the modelling of short-term dynamics, such as the decay of eddy currents or the ripple of the power supply. This hypothesis seems to be confirmed by the improved performance at the higher sampling rate,

corresponding to a buffer duration of 14 ms, which allows a finer modelling on an even shorter time-scale. Tab. 10.7 shows a comparison of the obtained results with respect to those of the state of the art in literature facing similar reconstruction problems. In Ref. [61], the authors used a Deep Neural Network to model the magnetization curve, achieving an RMSE of 0.13 %. This result is comparable with the obtained result for an MLP architecture ($8.70 \cdot 10^{-4}$ T), but it is higher than the RMSE obtained from the NARX architecture ($2.12 \cdot 10^{-5}$ T). In Ref. [59], the authors used a Preisach memory block and a feed-forward Neural Network to magnetic hysteresis modelling. The maximum prediction error achieved is around 13 % that is higher than the achieved MAE reported in Tab. 10.5. In fact, the MAE for a NARX network is of the order of 10^{-4}. In Ref. [15], the author used Preisach to model the hysteretic behavior of a combined magnet, reaching a relative error in the order of 0.2 %. For the presented case study, the achieved MPE with a NARX architecture is of about 0.2 %.
In Ref. [72], the authors proposed a Preisach-recurrent neural network model to predict the dynamic hysteresis in ARMCO® Pure Iron. The proposed model is able to predict the magnetic flux density of ARMCO® Pure Iron with a NRMSE of about 0.7 %. Comparing their result with the ones in Tab. 10.5, it is possibe to see that the NRMSE obtained from a NARX architecture is of the order of 10^{-2}. In Ref. [73], the authors presented a neural network model of nonlinear hysteretic inductors, achieving a relative error less than 8 %. In Tab. 10.5, the MPE resulting from a NARX architecture is about $2 \cdot 10^{-1}$ %. Finally, in Ref. [76], the authors proposed a combined approach (Genetic Algorithm and Neural Network) to modelling dynamic hysteresis. This approach allows them to achieve a Mean Square Error less than 5 %. For the presented case study, in Tab. 10.5 the RMSE for the various architectures is shown. In particular, for the NARX architecture, an RMSE of the order of 10^{-5} was achieved, giving therefore a better result compared to the literature.
An example of the nonlinear field component $\hat{B}_E$ (see Section 10.2) is plotted as a function of the time or the current, along with the corresponding reconstruction by a NARX network, in Fig. 10.10. These plots allow to appreciate visually the high quality of the reconstruction, which matches the measured field closely and consistently. Let us now consider the results of the networks trained on the nonlinear component dataset, $\hat{D}_{\mathrm{L}}$, (LR+$*$ group) which are given in the bottom halves of Tables 10.5 and 10.6.
This kind of reconstruction can be considered as an hybrid modelling of the magnetic field: a first module corresponding to the linear regression module is coupled with a network which models the nonlinear part only of the signal. In this case, the result is qualitatively different, since all tested architectures learned on $\hat{D}_L$ perform almost equally and with performance comparable to the MLP and TDNN networks alone, and unable to reach the better performance of NARX learned on the full signal in D_L. These results confirm that avoiding pre-processing at the same time relying complex delay structured in NARX networks is a successfull choice to capture the full dynamic of the magnetic field.

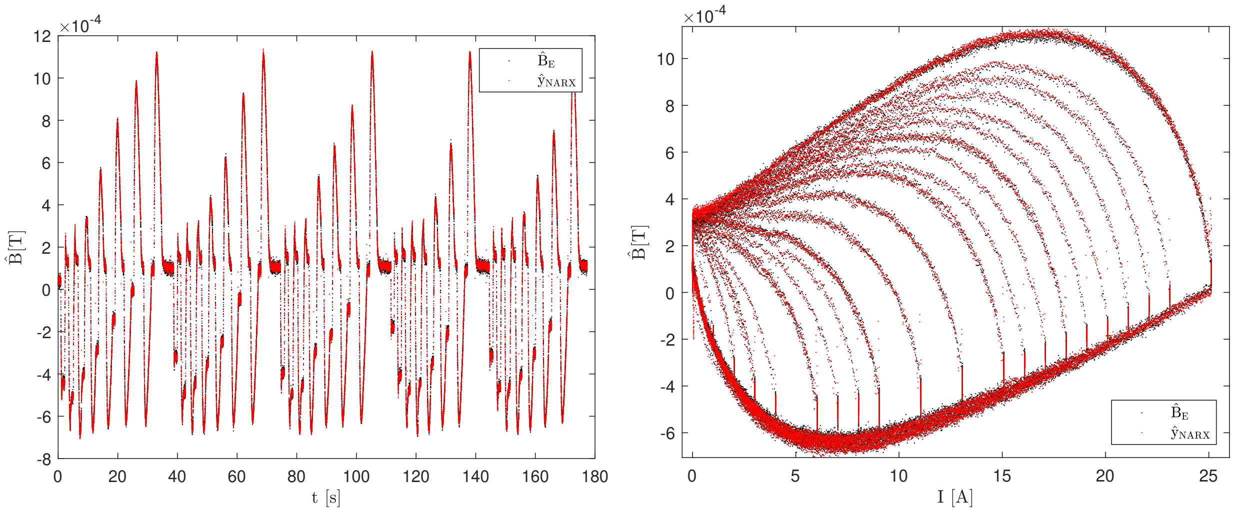

(A) Measured vs Estimated Field

(B) Hysteresis graph

FIGURE 10.10: Measured ($\hat{B}_{\mathrm{E}}$, in black) and estimated ($\hat{y}_{\mathrm{NARX}}$ with hyperparameters $\tilde{\theta}_{NARX1}$, in red) nonlinear component of the field $\hat{B}$ in function of time t (Panel (a)) and in function of the current I (Hysteresis graph - Panel (b)).

To perform a statistical evaluation among the models, a one-way Analysis of Variance (ANOVA) was performed for absolute errors in the reconstruction respectively of dataset D_{E} reported in Tab. 10.5 and dataset $\bar{D}_{\mathrm{E}}$ reported in Tab. 10.6. A first ANOVA on D_{E} reconstruction revealed a significant statistical effect on groups, $F[16, 754 \cdot 10^3] = 13 \cdot 10^3$, $p < 10^7$. A Post hoc analyses (Tukey's test) revealed that all models are significantly different from Linear Regression (null hypothesis). From this analysis, it was found that TDNN2 (best performing model excluding NARX models) is not statistically different from MLPs, TDNN1 and LN+TDNN2 ($p > 0.01$). On the other hand, all NARXs are statistically different from all the other models ($p < 10^{-5}$) and are not statistically different from each other ($p > 0.01$). The second ANOVA on $\bar{D}_{\mathrm{E}}$ revealed again a a significant statistical effect on groups, $F(16, 754 \cdot 10^3) = 14 \cdot 10^3$, $p < 10^{-7}$. Post hoc analyses revealed that models augmented with linear regression (LN+$*$) were not significantly different from each other ($p > 0.01$). Moreover, this LR+$*$ group was not statistically different from at least one of the MLP group ($p > 0.01$). This means that at their best they could at most replicate the performance of MLP models and this confirms that the learning restricted to the nonlinear part only does cut off important information of the original signal.
The statistical results on TDNN models are even more interesting: TDNN2 is not statistically different from MLP1 ($p > 0.01$) and TDNN1 is different from all the other models being the worst one, failing to generalize the case of faster sample rate. This is explained by observing that changing the buffer on the input is not sufficient to the TDNN model to let it adapt to signal when different input rates are given. On the other hand, in the case of NARX the buffer on input delay allows the nets to adapt very smoothly to the new faster rate, thanks to the buffer on internal outputs. Thus, all models in NARX group perform better and are statistically different from other groups ($p < 10.6 \cdot 10^{-6}$) and interestingly the

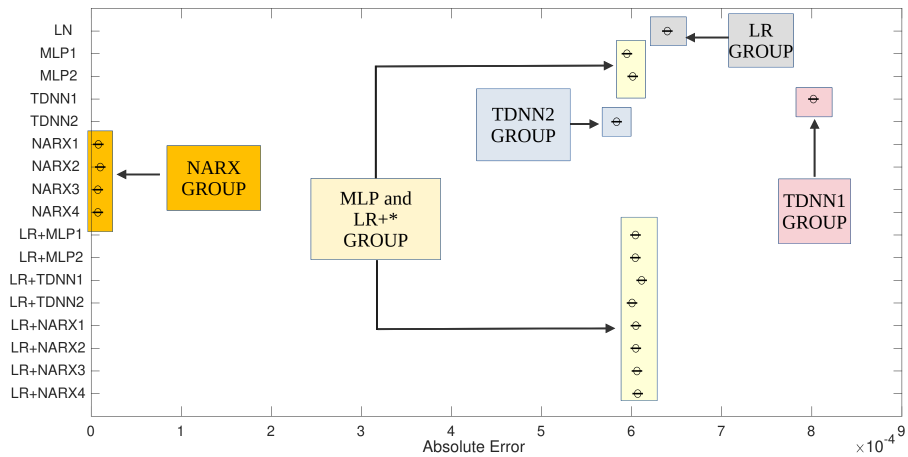

FIGURE 10.11: ANOVA results on Absolute Errors computed for the competing models on Dataset $\bar{D}_{\mathrm{E}}$. The horizontal lines are the 95 % confidence interval for each model. Five groups are highlighted: LR group (gray), MLP group and LR+$*$ group (yellow), TDNN1 (pink), TDNN2 (in blue) and the NARX group (orange).

different NARXs with different selected hyperparameters are not statistically different ($p > 0.1$). It is possible to visualize the result of this second ANOVA in Fig. 10.11: here the absolute error with 95 % confidence interval (straight line) is shown for each model. The models can be partitioned into four groups: linear regression is the baseline (in gray). The only model performing worst of this baseline is TDNN1 (pink group) that completely fails to adapt to the new sample rate of $\bar{D}_{\mathrm{E}}$. Then, the behavior of the models learned on linear models (LN+$*$ group) is equiparable to the MLP networks, performing better than LN alone (yellow group). On the other hand, TDNN2 is slightly better than this group (blue group). And finally, all NARXs (orange group) are successful in adapting to the new rate and significantly outperform the performances of all the other groups. A final test is made computing training and simulation time of execution of the winning architectures and results are shown in Tab. 10.8. The first column reports the architectures on which the training and simulation times are evaluated. In particular, for the training/simulation time evaluation the neural networks trained on the dataset D_{L} was considered. The second column contains the training time for each architecture. The training time refers to the time needed for the function Train (see line 5 in Alg. 1). The third and the fourth columns contain the simulation time computed on the test dataset D_{E} at the decimated data rate of $250\ S/s$ (see Tab. 10.5) and the dataset $\bar{D}_{\mathrm{E}}$) at full data rate of $2.5\ kS/s$ (see Tab. 10.6), respectively.

The simulation time refers to the time needed for the evaluation of the predictions on all the points of the dataset (see line 8 in Alg. 1).
It is possible to appreciate that this result nicely fits the complexity measure in Tab. 10.4. The more complex a model is, the more execution time is required to complete the different steps of the Alg. 1. In accordance with the complexity measure, it is worth mentioning that deeper models of NARX perform better than

NARXs with fewer layers. On the other hand, it should be noted that computation occurring on the same layer can be further optimized by processing them in parallel, thus a trade-off can be achieved between those constraints in order to reach the best performance in time execution.

TABLE 10.5: Performance comparison among the different architectures, computed on the test datasets D_{E} at the decimated data rate of 250 S/s. RMSE, NMRSE, Maximum Absolute Error (MAE) and Maximum Percentage Error (MPE) are shown for the linear Regression (LR) alone , the ANNs trained on the full dataset D_{L} and the hybrid models LR+*, combining LR plus the ANN trained on the nonlinear component only $\hat{D}_{\mathrm{L}}$. The reconstructed magnetic field for models LR+* is computed by Eq.10.6, in which the linear component ($B_0+G{\cdot}I$) is computed by means of LR coefficients, while the nonlinear component ($\hat{B}$) is approximated by the corresponding network output. Corresponding hyperparameters are given in Tab. 10.4.

Architecture	Test Dataset	Hyper- -parameters	RMSE [T]	NMRSE (%)	MAE [T]	MPE (%)
Linear Regression		G, B_0	$9.07 \cdot 10^{-04}$	$5.67 \cdot 10^{-01}$	$2.60 \cdot 10^{-03}$	1.61
MLP1	D_{E}	$\tilde{\theta}_{MLP1}$	$8.70 \cdot 10^{-04}$	$5.44 \cdot 10^{-01}$	$2.60 \cdot 10^{-03}$	1.64
MLP2	D_{E}	$\tilde{\theta}_{MLP2}$	$8.83 \cdot 10^{-04}$	$5.52 \cdot 10^{-01}$	$2.50 \cdot 10^{-03}$	1.55
TDNN1	D_{E}	$\tilde{\theta}_{TDNN1}$	$7.95 \cdot 10^{-04}$	$4.97 \cdot 10^{-01}$	$1.90 \cdot 10^{-03}$	1.21
TDNN2	D_{E}	$\tilde{\theta}_{TDNN2}$	$8.05 \cdot 10^{-04}$	$5.03 \cdot 10^{-01}$	$2.20 \cdot 10^{-03}$	1.39
NARX1	D_{E}	$\tilde{\theta}_{NARX1}$	$2.12 \cdot 10^{-05}$	$1.32 \cdot 10^{-02}$	$3.36 \cdot 10^{-04}$	$2.10 \cdot 10^{-01}$
NARX2	D_{E}	$\tilde{\theta}_{NARX2}$	$2.13 \cdot 10^{-05}$	$1.33 \cdot 10^{-02}$	$3.17 \cdot 10^{-04}$	$1.98 \cdot 10^{-01}$
NARX3	D_{E}	$\tilde{\theta}_{NARX3}$	$2.05 \cdot 10^{-05}$	$1.28 \cdot 10^{-02}$	$3.11 \cdot 10^{-04}$	$1.95 \cdot 10^{-01}$
NARX4	D_{E}	$\tilde{\theta}_{NARX4}$	$2.05 \cdot 10^{-05}$	$1.28 \cdot 10^{-02}$	$3.12 \cdot 10^{-04}$	$1.95 \cdot 10^{-01}$
LR+MLP1	D_{E}	G, B_0, $\tilde{\theta}_{MLP1}$	$8.00 \cdot 10^{-04}$	$5.00 \cdot 10^{-01}$	$1.90 \cdot 10^{-03}$	1.16
LR+MLP2	D_{E}	G, B_0, $\tilde{\theta}_{MLP2}$	$8.00 \cdot 10^{-04}$	$5.00 \cdot 10^{-01}$	$1.90 \cdot 10^{-03}$	1.17
LR+TDNN1	D_{E}	G, B_0, $\tilde{\theta}_{TDNN1}$	$8.05 \cdot 10^{-04}$	$5.03 \cdot 10^{-01}$	$1.90 \cdot 10^{-03}$	1.21
LR+TDNN2	D_{E}	G, B_0, $\tilde{\theta}_{TDNN2}$	$7.97 \cdot 10^{-04}$	$4.98 \cdot 10^{-01}$	$1.90 \cdot 10^{-03}$	1.20
LR+NARX1	D_{E}	G, B_0, $\tilde{\theta}_{NARX1}$	$7.99 \cdot 10^{-04}$	$4.99 \cdot 10^{-01}$	$1.90 \cdot 10^{-03}$	1.19
LR+NARX3	D_{E}	G, B_0, $\tilde{\theta}_{NARX2}$	$7.99 \cdot 10^{-04}$	$5.00 \cdot 10^{-01}$	$1.90 \cdot 10^{-03}$	1.19
LR+NARX3	D_{E}	G, B_0, $\tilde{\theta}_{NARX3}$	$8.00 \cdot 10^{-04}$	$5.00 \cdot 10^{-01}$	$1.90 \cdot 10^{-03}$	1.20
LR+NARX4	D_{E}	G, B_0, $\tilde{\theta}_{NARX4}$	$8.01 \cdot 10^{-04}$	$5.01 \cdot 10^{-01}$	$1.90 \cdot 10^{-03}$	1.19

TABLE 10.6: Performance comparison among the different architectures, computed on the test dataset $\bar{D}_{\mathrm{E}}$ at the full data rate of 2.5 kS/s. RMSE, NRMSE, MAE and MPE are shown for the LR alone, the ANNs trained on the full dataset D_{L} and the hybrid models combining LR plus the ANN trained on the nonlinear component only $\hat{D}_{\mathrm{L}}$. The values of the corresponding hyperparameters are given in Tab. 10.4

Architecture	Test Dataset	Hyper- -parameters	RMSE [T]	NMRSE (%)	MAE [T]	MPE (%)
Linear Regression		G, B_0	$9.07 \cdot 10^{-04}$	$5.67 \cdot 10^{-01}$	$2.60 \cdot 10^{-03}$	1.61
MLP1	$\bar{D}_{\mathrm{E}}$	$\tilde{\theta}_{MLP1}$	$8.70 \cdot 10^{-04}$	$5.44 \cdot 10^{-01}$	$2.60 \cdot 10^{-03}$	1.64
MLP2	$\bar{D}_{\mathrm{E}}$	$\tilde{\theta}_{MLP2}$	$8.83 \cdot 10^{-04}$	$5.52 \cdot 10^{-01}$	$2.50 \cdot 10^{-03}$	1.55
TDNN1	$\bar{D}_{\mathrm{E}}$	$\tilde{\theta}_{TDNN1}$	$1.10 \cdot 10^{-03}$	$6.66 \cdot 10^{-01}$	$2.70 \cdot 10^{-03}$	1.67
TDNN2	$\bar{D}_{\mathrm{E}}$	$\tilde{\theta}_{TDNN2}$	$8.52 \cdot 10^{-04}$	$5.32 \cdot 10^{-01}$	$2.50 \cdot 10^{-03}$	1.56
NARX1	$\bar{D}_{\mathrm{E}}$	$\tilde{\theta}_{NARX1}$	$9.92 \cdot 10^{-06}$	$6.20 \cdot 10^{-03}$	$4.63 \cdot 10^{-05}$	$2.89 \cdot 10^{-02}$
NARX2	$\bar{D}_{\mathrm{E}}$	$\tilde{\theta}_{NARX2}$	$1.27 \cdot 10^{-05}$	$8.00 \cdot 10^{-03}$	$6.90 \cdot 10^{-05}$	$4.31 \cdot 10^{-02}$
NARX3	$\bar{D}_{\mathrm{E}}$	$\tilde{\theta}_{NARX3}$	$9.22 \cdot 10^{-06}$	$5.80 \cdot 10^{-03}$	$3.98 \cdot 10^{-05}$	$2.49 \cdot 10^{-02}$
NARX4	$\bar{D}_{\mathrm{E}}$	$\tilde{\theta}_{NARX4}$	$9.28 \cdot 10^{-06}$	$5.80 \cdot 10^{-03}$	$4.06 \cdot 10^{-05}$	$2.54 \cdot 10^{-02}$
LR+MLP1	$\bar{D}_{\mathrm{E}}$	G, B_0, $\tilde{\theta}_{MLP1}$	$8.00 \cdot 10^{-04}$	$5.00 \cdot 10^{-01}$	$1.90 \cdot 10^{-03}$	1.16
LR+MLP2	$\bar{D}_{\mathrm{E}}$	G, B_0, $\tilde{\theta}_{MLP2}$	$8.00 \cdot 10^{-04}$	$5.00 \cdot 10^{-01}$	$1.90 \cdot 10^{-03}$	1.17
LR+TDNN1	$\bar{D}_{\mathrm{E}}$	G, B_0, $\tilde{\theta}_{TDNN1}$	$8.05 \cdot 10^{-04}$	$5.03 \cdot 10^{-01}$	$1.90 \cdot 10^{-03}$	1.18
LR+TDNN2	$\bar{D}_{\mathrm{E}}$	G, B_0, $\tilde{\theta}_{TDNN2}$	$7.97 \cdot 10^{-04}$	$4.98 \cdot 10^{-01}$	$1.90 \cdot 10^{-03}$	1.18
LR+NARX1	$\bar{D}_{\mathrm{E}}$	G, B_0, $\tilde{\theta}_{NARX1}$	$7.98 \cdot 10^{-04}$	$4.99 \cdot 10^{-01}$	$1.90 \cdot 10^{-03}$	1.17
LR+NARX2	$\bar{D}_{\mathrm{E}}$	G, B_0, $\tilde{\theta}_{NARX2}$	$7.98 \cdot 10^{-04}$	$4.99 \cdot 10^{-01}$	$1.90 \cdot 10^{-03}$	1.17
LR+NARX3	$\bar{D}_{\mathrm{E}}$	G, B_0, $\tilde{\theta}_{NARX3}$	$7.99 \cdot 10^{-04}$	$5.00 \cdot 10^{-01}$	$1.90 \cdot 10^{-03}$	1.17
LR+NARX4	$\bar{D}_{\mathrm{E}}$	G, B_0, $\tilde{\theta}_{NARX4}$	$8.00 \cdot 10^{-04}$	$5.00 \cdot 10^{-01}$	$1.90 \cdot 10^{-03}$	1.17

TABLE 10.7: Related results

Architecture	Metric	Value
Deep Neural Network (MLP with two hidden layers), Ref. [61]	Root Mean Square Error	0.13 %
Preisach + Feed-forward neural network (one hidden layer), Ref. [59]	Maximum Absolute Error	13 %
Preisach, Ref. [15]	Relative Error	0.2 %
Preisach + Recurrent Neural Network, Ref. [72]	Normalized Root Mean Square Error	0.7 %
Neural Network, Ref. [73]	Relative Error	< 8 %
Genetic Algorithm + Neural Network, Ref. [76]	Mean Square Error	< 5 %
Proposed architecture NARX4	Normalized Root Mean Square Error	$5.80 \cdot 10^{-3}$ %

TABLE 10.8: Training and simulation times.

Architecture	Training time [s]	Simulation time (D_{E}) [s]	Simulation time ($\bar{D}_{\mathrm{E}}$) [s]
MLP1	16.88	0.40	1.08
MLP2	14.18	0.44	0.88
TDNN1	46.30	1.73	12.50
TDNN2	22.13	1.50	16.16
NARX1	125.52	2.66	25.23
NARX2	42.96	2.30	22.23
NARX3	19.91	2.91	27.98
NARX4	15.89	2.84	28.67

Conclusions

Conclusions

In this thesis, the concepts behind the B-train systems at CERN were introduced, beginning from their role in a synchrotron, their method of operation, as well as the different sensors required for their implementation. The design of the systems was discussed, describing in detail the functions linked to magnetic flux integration, self-calibration and distribution over a White Rabbit network. A new real-time WR monitoring and debugging system for the B-train was presented, to improve flexibility and remote access to the signals of the new B-Trains at CERN. A working prototype was successfully obtained under a platform never used previously in this context, demonstrating that the requirements in terms of both performance and operational flexibility can be met. This work showcases the power and possibilities offered by WR, as a new and flourishing standard in the general context of distributed acquisition chains and control systems, especially when accurate timing and synchronization are crucial. The presented monitoring device covered a key role in the final implementation of the new B-train system at CERN, as it will allow real-time logging, monitoring and visualization of multiple data streams to an unprecedented level of accuracy and time resolution. The Monitoring system was also profitably used for the system characterization and debugging.

First, the DC response of the system was experimentally characterized, showing that after calibration the voltage acquisition error across the $\pm 10\,\mathrm{V}$ input range has arithmetic and an RMS mean of respectively 3 and $135\,\mu\mathrm{V}$ (8.1). then the equivalent voltage offset error was evaluated integrating the same, known constant voltage inputs over one second, obtaining closely comparable results, thus proving the soundness of the integration stage (8.2). Next, integration tests were carried out with the input shorted for up to $120\,\mathrm{s}$, measuring on RMS average an equivalent voltage offset of $8\,\mu\mathrm{V}$. Under typical operational conditions, these errors lead to a measurement drift of the order of $8\,\mu\mathrm{T\,s}$. Such offset, while remaining high in absolute terms, is much lower than in the legacy systems, and is well within the tolerance of all magnetic cycles being run.

The response of the system to time-changing inputs was characterized as well. First, the measurement gain as a function of frequency was evaluated for sine-wave inputs of varying amplitude, showing that errors remain below the specified $100\,\mathrm{ppm}$ tolerance up to the required $100\,\mathrm{Hz}$ bandwidth. Then, the contribution of all system components to the measurement latency was quantified, finding that overall it remains always well below the specified $30\,\mu\mathrm{s}$ tolerance. Therefore, it is possible to conclude that the new CERN B-train electronic acquisition system meets all its requirements, as it was also independently verified by running an extensive series of operational tests, as reported in detail in [104].

An incremental method to select an optimal DNN architecture was developed to predict the field generated by a magnet when excited by a sequence of cyclic

excitation current waveforms. the method was experimentally validated on a case study in conditions representative of those found in particle accelerators and similar, pulsed-mode machines. The response of the magnet tested is linear within about $1.5\ \%$, and the work was focused essentially on predicting its residual nonlinear component, which is dominated by ferromagnetic hysteresis. The presented networks were trained and tested directly on the raw datasets, founding that NARX networks achieve in general the required level of performance *i.e.* an NRMSE better than $0.01\ \%$, while simpler architectures with buffers only on the input (TDNN) or no buffers at all (MLP) do not. Interestingly, it was discovered that by isolating the nonlinear residual component of the measured magnetic response (datasets $\hat{D}_{\mathrm{L}}$ and $\hat{D}_{\mathrm{E}}$), the performance of all the networks improves. In particular, the simple MLP architecture improves by two orders of magnitude and can achieve an NRSME as low as $5.9 \cdot 10^{-5}$ with only 19 neurons on two hidden layers.
Such excellent performance is well within the initial requirements and paves a very promising way for future applications in this context. It was observed that the prediction accuracy generally improves when the network is trained on low data rate (250 S/s) signals and tested at a higher data rate (2.5 kS/s). This may be linked to the fact that the reduced dataset is less affected by noise, and therefore allows the network to better focus on capturing the underlying dynamics. The presented method was based on an incremental approach that firstly optimizes the static structure parameters ($L, \mathbf{A}$) and then the time buffers (*i.e.*, K, H). While this approach does not ensure the finding of an optimal solution, It was proved that it constitutes an efficient heuristic able to computationally minimize the model selection procedure. Future work will also focus on extending and refining the model selection by including smart procedures for parameter grid search, trainable activation functions and sparse structure learning (see e.g., Refs. [149, 150, 151, 152, 153]) that would allow deeper structures to be better managed. This work opens the door to further investigation on this aspect by decreasing further the data rate of the training dataset, while at the same time increasing the interval between the samples in the output buffer so that it may cover the period of two or more cycles. Train and testing NARX networks on a wider variety of excitation waveforms, such as sequences of cycles with flat-tops increasing or decreasing randomly, which are representative of the most challenging actual operating conditions of accelerator magnets. In addition, in future experimental campaigns, the range of the tested currents will be extended, to ensure the introduction of relevant levels of saturation, as well as the range of current ramp rates to deal with different levels of eddy current-related effects. Moreover, a further improvement will be to expand the framework to include classification capabilities [69, 154] to identify different branches of the hysteresis cycles in real-time. Overall, in this framework, the analysis of the simpler network solutions found by the procedure could also open to the possibility of producing even more efficient solutions, by substituting blocks of NN operations with equivalent mathematical equations or equivalent smaller networks.

Finally, as part of the renovation of the real-time magnetic measurement systems currently ongoing at CERN, the obtained result open exciting possibilities on a real-time version of the NARX networks implemented in FPGA hardware that will be able to carry out a continuous field prediction, in parallel to the measurement. This facility will provide the opportunity to gather huge amounts of data concerning thousands of different sequences of cycles, covering all relevant dynamic scenarios. This will ultimately allow to fine-tune the parameters of the networks and estimate their robustness and performance in the long term with high statistical significance.

References

[1] Donglai Zhang, Xiaolan Yan, Enchao Zhang, and Shimin Pan. "A long time low drift integrator with temperature control". In: *Review of Scientific Instruments* 87.10 (2016), p. 105119.

[2] DM Liu et al. "An analog-digital integrator for EAST long pulse discharge". In: *Fusion Engineering and Design* 165 (2021), p. 112255.

[3] Antonio JN Batista et al. "F4E prototype of a chopper digital integrator for the ITER magnetics". In: *Fusion Engineering and Design* 123 (2017), pp. 1025–1028.

[4] Maria Amodeo, Pasquale Arpaia, and Marco Buzio. "Integrator drift compensation of magnetic flux transducers by feed-forward correction". In: *Sensors* 19.24 (2019), p. 5455.

[5] P Arpaia, P Cimmino, L De Vito, and L Fiscarelli. "Offset correction in a digital integrator for rotating coil measurements". In: *IMEKO World Congress" Metrology for green growth"*. 2012.

[6] Eike Feldmeier, Th Haberer, Andreas Peters, C Schoemers, Rudolf Steiner, et al. "Magnetic field correction in normal conducting synchrotrons". In: *Proc. IPAC* (2010).

[7] Augusto Visintin. "Mathematical models of hysteresis". In: *Modelling and Optimization of Distributed Parameter Systems Applications to Engineering*. Springer, 1996, pp. 71–80.

[8] Vahid Hassani, Tegoeh Tjahjowidodo, and Thanh Nho Do. "A survey on hysteresis modeling, identification and control". In: *Mechanical systems and signal processing* 49.1-2 (2014), pp. 209–233.

[9] Isaak D Mayergoyz. *Mathematical models of hysteresis and their applications*. Academic Press, 2003.

[10] Edward Della Torre. "Hysteresis modeling". In: *COMPEL-The international journal for computation and mathematics in electrical and electronic engineering* (1998).

[11] J Takács. "A phenomenological mathematical model of hysteresis". In: *COMPEL-The international journal for computation and mathematics in electrical and electronic engineering* (2001).

[12] Monia Ferjani Jaafar. "Magnetic hysteresis modeling and numerical simulation for ferromagnetic materials". In: *2013 International Conference on Control, Decision and Information Technologies (CoDIT)*. IEEE. 2013, pp. 516–523.

[13] David C Jiles and David L Atherton. "Theory of ferromagnetic hysteresis". In: *Journal of applied physics* 55.6 (1984), pp. 2115–2120.

[14] Ferenc Preisach. "Über die magnetische Nachwirkung". In: *Zeitschrift für physik* 94.5-6 (1935), pp. 277–302.

[15] Valentin Pricop. "EFECTELE HISTEREZISULUI DIN MATERIALELE FOLOSITE PENTRU CIRCUITELE MAGNETICE ALE ACCELERATOARELOR DE PARTICULE, HYSTERESIS EFFECTS IN THE CORES OF PARTICLE ACCELERATOR MAGNETS". PhD thesis. Brașov, Romania: Școala Doctoral Interdisciplinară, Departament: Inginerie Electrică și Fizică Aplicată, Transilvania U., 2016.

[16] David C Jiles and David L Atherton. "Theory of ferromagnetic hysteresis". In: *Journal of magnetism and magnetic materials* 61.1-2 (1986), pp. 48–60.

[17] Klaus Hanke et al. "The LHC Injectors Upgrade (LIU) Project at CERN: Proton Injector Chain". en-us. In: *Proceedings of the 8th Int. Particle Accelerator Conf.* IPAC2017 (2017), Denmark–. DOI: 10.18429/jacow-ipac2017-wepva036.

[18] Edmund J N Wilson. "Fifty years of synchrotrons". In: (1996). URL: https://cds.cern.ch/record/339572.

[19] C Carli et al. "Commissioning and Performance of LEIR". In: *COOL*. Vol. 7. 2007, p. 134.

[20] Stephen Myers. *LHC commissioning and first operation*. Tech. rep. 2010.

[21] M Modena et al. "Design, assembly and first measurements of a short model for CLIC final focus hybrid quadrupole QD0". In: *Conf. Proc.* Vol. 1205201. CERN-ATS-2012-200. 2012, THPPD010.

[22] Hanno Leibrock et al. "Prototype of the Superferric Dipoles for the Super-FRS of the FAIR-Project". In: *IEEE Transactions on Applied Superconductivity* 20.3 (2010), pp. 188–191. DOI: 10.1109/TASC.2010.2040167.

[23] Helmut Wiedemann. "Periodic Focusing Systems". In: *Particle Accelerator Physics*. Springer, 2003, pp. 237–280.

[24] Helmut Wiedemann. "Perturbations in Beam Dynamics". In: *Particle Accelerator Physics*. Springer, 2015, pp. 477–538.

[25] Claude Bovet, Robert Gouiran, Igor Gumowski, and Karl Helmut Reich. *A selection of formulae and data useful for the design of A.G. synchrotrons; rev. version*. Geneva: CERN, 1970. URL: https://cds.cern.ch/record/104153.

[26] P. Baudrenghien. "Low-level RF - Part I: Longitudinal dynamics and beam-based loops in synchrotrons". In: (Jan. 2012). Comments: 27 pages, contribution to the CAS - CERN Accelerator School: Specialised Course on RF for Accelerators; 8 - 17 Jun 2010, Ebeltoft, Denmark, 27 p. DOI: 10.5170/CERN-2011-007.341. arXiv: 1201.2597. URL: https://cds.cern.ch/record/1415913.

[27] Maria Elena Angoletta et al. "Initial Beam Results of CERN ELENA's Digital Low-Level RF System". In: (2017), THPAB142. 4 p. DOI: 10.18429/JACoW-IPAC2017-THPAB142. URL: https://cds.cern.ch/record/2289141.

[28] J C Denard. "Beam current monitors". In: (2009). DOI: 10.5170/CERN-2009-005.141. URL: https://cds.cern.ch/record/1213275.

[29] Robert C Webber. "Charged particle beam current monitoring tutorial". In: *AIP Conference Proceedings CONF- 9410219*. Vol. 333. 1. American Institute of Physics. 1995, pp. 3–23.

[30] Q King, ST Page, H Thiesen, and M Veenstra. "Function generation and regulation libraries and their application to the control of the new main power converter (POPS) at the CERN CPS". In: *Proc. 13th Int. Conf. on Accelerator and Large Experimental Physics Control Systems (ICALEPCS'11)*. 2011, pp. 886–889.

[31] Fulvio Boattini, Jean-Paul Burnet, and Gregory Skawinski. "POPS: The 60MW power converter for the PS accelerator: Control strategy and performances". In: *2015 17th European Conference on Power Electronics and Applications (EPE'15 ECCE-Europe)*. IEEE. 2015, pp. 1–10.

[32] Marco Buzio et al. "Development of Upgraded Magnetic Instrumentation for CERN's Real-time Reference Field Measurement Systems". In: *Proceedings, 1st International Particle Accelerator Conference (IPAC'10): Kyoto, Japan, May 23-28, 2010*. Vol. 100523. Geneva, Switzerland: JACoW, May 2010, MOPEB016. URL: http://accelconf.web.cern.ch/AccelConf/IPAC10/papers/MOPEB016.pdf.

[33] G. Franzini et al. "Final Design and Features of the B-train System of CNAO". In: *Proceedings, 1st International Particle Accelerator Conference (IPAC'10): Kyoto, Japan, May 23-28, 2010*. Vol. 100523. Geneva, Switzerland: JACoW, May 2010, WEPEB079. URL: http://accelconf.web.cern.ch/AccelConf/IPAC10/papers/wepeb079.pdf.

[34] Philippe Fraboulet, Kristian Ambrosch, Ivan De Cesaris, and Andre Beuret. "The current and field regulation system of the MedAustron Ion-Beam accelerator". In: *IECON 2013 - 39th Annual Conference of the IEEE Industrial Electronics Society*. IEEE, nov 2013. DOI: 10.1109/iecon.2013.6700312.

[35] E. Feldmeier, T. Haberer, A. Pters, C. Schoemers, and R. Steiner. "Magnetic Field Correction in Normal Conducting Synchrotrons". In: *Proceedings, 1st International Particle Accelerator Conference (IPAC'10): Kyoto, Japan, May 23-28, 2010*. Vol. 100523. Geneva, Switzerland: JACoW, May 2010, MOPD004. URL: http://accelconf.web.cern.ch/AccelConf/IPAC10/papers/mopd004.pdf.

[36] M. Buzio. "Fabrication and calibration of search coils". In: *CAS-CERN Accelerator School: Specialised course on Magnets*. Bruges, Belgium, Apr. 5, 2011, pp. 387–421. arXiv: 1104.0803v1 [physics.acc-ph]. URL: https://arxiv.org/ftp/arxiv/papers/1104/1104.0803.pdf.

[37] J Billan et al. "Twin rotating coils for cold magnetic measurements of 15 m long LHC dipoles". In: *IEEE Transactions on Applied Superconductivity* 10.1 (2000), pp. 1422–1426.

[38] P Arpaia, E De Matteis, and R Schiano Lo Moriello. "Unscented transform-based uncertainty analysis of rotating coil transducers for field mapping". In: *Review of Scientific Instruments* 87.3 (2016), p. 035004.

[39] J. DiMarco et al. "Field alignment of quadrupole magnets for the LHC interaction regions". In: *IEEE Transactions on Applied Superconductivity* 10.1 (2000), pp. 127–130. DOI: 10.1109/77.828192.

[40] Carlo Petrone. "Wire methods for measuring field harmonics, gradients and magnetic axes in accelerator magnets". Presented 2013. 2013. URL: https://cds.cern.ch/record/1601973.

[41] G Golluccio et al. "PCB coil array for measuring curved accelerator dipoles: two case studies on the MedAustron accelerator". In: *proceedings of IMEKO 2014* (2014).

[42] Stephan Russenschuck. "Rotating-and translating-coil magnetometers for extracting pseudo-multipoles in accelerator magnets". In: *COMPEL-The international journal for computation and mathematics in electrical and electronic engineering* (2017).

[43] Louis Walckiers. "Magnetic measurement with coils and wires". eng. In: *CAS - CERN Accelerator School: Specialised course on Magnets*. CERN, June 2009. DOI: 10.5170/CERN-2010-004.357. URL: https://cds.cern.ch/record/1345967.

[44] JA Garcia and M Rivas. "A quasi-static magnetic hysteresis loop measurement system with drift correction". In: *IEEE transactions on magnetics* 42.1 (2005), pp. 15–17.

[45] S. Hancock. *A fit-based frequency programme for the PS*. Tech. rep. CERN-AB-Note-2007-036. Geneva: CERN, Dec. 2007. URL: http://cds.cern.ch/record/1072486.

[46] J Taylor, H du T Mouton, and T Jones. "High-bandwidth NMR teslameter". In: *AFRICON 2007*. IEEE. 2007, pp. 1–7.

[47] Christian Grech et al. "Metrological characterization of nuclear magnetic resonance markers for real-time field control of the CERN ELENA ring dipoles". In: *IEEE Sensors Journal* 18.14 (2018), pp. 5826–5833.

[48] Richard R Ernst, Geoffrey Bodenhausen, Alexander Wokaun, et al. *Principles of nuclear magnetic resonance in one and two dimensions*. Vol. 14. Clarendon press Oxford, 1987.

[49] Peter J Hore. *Nuclear magnetic resonance*. Oxford University Press, USA, 2015.

[50] J Nicolas. "Microwave ferrites". In: *Handbook of Ferromagnetic Materials* 2 (1980), pp. 243–296.

[51] Wee Tee Soh, Chin Yaw Tan, and Chong Kim Ong. "Novel Coplanar Waveguide Circulator Using Metallized Yttrium Iron Garnet Disk". In: *Microwave and Optical Technology Letters* 55.9 (2013), pp. 2143–2145.

[52] Jarikasem Warit, T Ueda, and Makoto Tsutsumi. "Microwave circulator using yttrium iron garnet film". In: *2000 Asia-Pacific Microwave Conference. Proceedings (Cat. No. 00TH8522)*. IEEE. 2000, pp. 1212–1213.

[53] PETER Roschmann. "Compact YIG bandpass filter with finite-pole frequencies for applications in microwave integrated circuits (short papers)". In: *IEEE Transactions on Microwave Theory and Techniques* 21.1 (1973), pp. 52–57.

[54] Jerzy Krupka, Bartlomiej Salski, Pawel Kopyt, and Wojciech Gwarek. "Electrodynamic study of YIG filters and resonators". In: *Scientific reports* 6.1 (2016), pp. 1–9.

[55] Kenneth H Carpenter and Marcus K daSilva. "Phase-locked yttrium iron garnet magnetometer for remote measurement of small field changes in a fluctuating background". In: *Review of Scientific Instruments* 53.9 (1982), pp. 1414–1417.

[56] M. Di Castro et al. "Parametric field modeling for the lhc main magnets in operating conditions". In: *2007 IEEE Particle Accelerator Conference (PAC)*. 2007, pp. 1586–1588.

[57] N Sammut, L Bottura, and J Micallef. "Mathematical formulation to predict the harmonics of the superconducting large hadron collider magnets". In: *Phys. Rev. Accel. Beams* 9.1 (2006), p. 012402.

[58] P Arpaia, M Buzio, F Caspers, G Golluccio, and C Petrone. "Static metrological characterization of a ferrimagnetic resonance transducer for real-time magnetic field markers in particle accelerators". In: *2011 IEEE International Instrumentation and Measurement Technology Conference*. IEEE. 2011, pp. 1–4.

[59] Claudio Serpico and Ciro Visone. "Magnetic hysteresis modeling via feed-forward neural networks". In: *IEEE Transactions on Magnetics* 34.3 (1998), pp. 623–628.

[60] Dimitre Makaveev, Luc Dupré, Marc De Wulf, and Jan Melkebeek. "Modeling of quasistatic magnetic hysteresis with feed-forward neural networks". In: *Journal of Applied Physics* 89.11 (2001), pp. 6737–6739.

[61] A Du, J Pan, and A Guzzomi. "Modeling the Hysteresis Characteristics of Transformer Core under Various Excitation Level via On-Line Measurements". In: *Electronics* 7.390 (2018).

[62] Jorge Garcıa-González, Juan M Ortiz-de-Lazcano-Lobato, Rafael M Luque-Baena, and Ezequiel López-Rubio. "Background subtraction by probabilistic modeling of patch features learned by deep autoencoders". In: *Integrated Computer-Aided Engineering* Preprint (2020), pp. 1–13.

[63] Safa Hamreras, Bachir Boucheham, Miguel A Molina-Cabello, Rafaela Benitez-Rochel, and Ezequiel Lopez-Rubio. "Content based image retrieval by ensembles of deep learning object classifiers". In: *Integrated Computer-Aided Engineering* Preprint (2020), pp. 1–15.

[64] David Simões, Nuno Lau, and Luis Paulo Reis. "Exploring communication protocols and centralized critics in multi-agent deep learning". In: *Integrated Computer-Aided Engineering* Preprint (2020), pp. 1–19.

[65] David Diaz-Vico, Jesus Prada, Adil Omari, and José Dorronsoro. "Deep support vector neural networks". In: *Integrated Computer-Aided Engineering* Preprint (2020), pp. 1–14.

[66] U Rajendra Acharya et al. "Automated EEG-based screening of depression using deep convolutional neural network". In: *Computer methods and programs in biomedicine* 161 (2018), pp. 103–113.

[67] Guilherme Brandão Martins, João Paulo Papa, and Hojjat Adeli. "Deep learning techniques for recommender systems based on collaborative filtering". In: *Expert Systems* 37.6 (2020), e12647.

[68] Hidir Selcuk Nogay and Hojjat Adeli. "Detection of Epileptic Seizure Using Pretrained Deep Convolutional Neural Network and Transfer Learning". In: *European Neurology* 83.6 (2020), pp. 602–614.

[69] Mehran Ahmadlou and Hojjat Adeli. "Enhanced probabilistic neural network with local decision circles: A robust classifier". In: *Integrated Computer-Aided Engineering* 17.3 (2010), pp. 197–210.

[70] Tsungnan Lin, Bill G Horne, Peter Tino, and C Lee Giles. "Learning long-term dependencies in NARX recurrent neural networks". In: *IEEE Transactions on Neural Networks* 7.6 (1996), pp. 1329–1338.

[71] Hava T Siegelmann, Bill G Horne, and C Lee Giles. "Computational capabilities of recurrent NARX neural networks". In: *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* 27.2 (1997), pp. 208–215.

[72] Christian Grech, Marco Buzio, Mariano Pentella, and Nicholas Sammut. "Dynamic Ferromagnetic Hysteresis Modelling Using a Preisach-Recurrent Neural Network Model". In: *Materials* 13.11 (2020), p. 2561.

[73] Silvano Cincotti, Michele Marchesi, and Antonino Serri. "A neural network model of parametric nonlinear hysteretic inductors". In: *IEEE transactions on magnetics* 34.5 (1998), pp. 3040–3043.

[74] AA Adly and SK Abd-El-Hafiz. "Using neural networks in the identification of Preisach-type hysteresis models". In: *IEEE Transactions on Magnetics* 34.3 (1998), pp. 629–635.

[75] Maurizio Cirrincione, Rosario Miceli, Giuseppe Ricco Galluzzo, and Marco Trapanese. "Preisach function identification by neural networks". In: *IEEE Transactions on Magnetics* 38.5 (2002), pp. 2421–2423.

[76] Alessandro Salvini, F Riganti Fulginei, and Christian Coltelli. "A neuro-genetic and time-frequency approach to macromodeling dynamic hysteresis in the harmonic regime". In: *IEEE Transactions on Magnetics* 39.3 (2003), pp. 1401–1404.

[77] Alessandro Salvini, Francesco Riganti Fulginei, and Giuseppe Pucacco. "Generalization of the static Preisach model for dynamic hysteresis by a genetic approach". In: *IEEE transactions on Magnetics* 39.3 (2003), pp. 1353–1356.

[78] Francesco Riganti Fulginei and Alessandro Salvini. "Neural network approach for modelling hysteretic magnetic materials under distorted excitations". In: *IEEE Transactions on Magnetics* 48.2 (2012), pp. 307–310.

[79] Kurt Hornik, Maxwell Stinchcombe, and Halbert White. "Multilayer feedforward networks are universal approximators". In: *Neural Networks* 2.5 (1989), pp. 359–366. ISSN: 0893-6080.

[80] Kenichi Funahashi and Yuichi Nakamura. "Approximation of dynamical systems by continuous time recurrent neural networks". In: *Neural Networks* 6.6 (1993), pp. 801–806. ISSN: 0893-6080.

[81] Jürgen Schmidhuber. "Deep learning in neural networks: An overview". In: *Neural Networks* 61 (2015), pp. 85–117. ISSN: 0893-6080.

[82] Sepp Hochreiter and Jürgen Schmidhuber. "Long short-term memory". In: *Neural computation* 9.8 (1997), pp. 1735–1780.

[83] Juan Antonio Pérez-Ortiz, Felix A Gers, Douglas Eck, and Juergen Schmidhuber. "Kalman filters improve LSTM network performance in problems unsolvable by traditional recurrent nets". In: *Neural Networks* 16.2 (2003), pp. 241–250.

[84] Francesco Donnarumma, Roberto Prevete, Andrea de Giorgio, Guglielmo Montone, and Giovanni Pezzulo. "Learning programs is better than learning dynamics: A programmable neural network hierarchical architecture in a multi-task scenario". In: *Adaptive Behavior* 24.1 (2016), pp. 27–51.

[85] Francesco Donnarumma, Roberto Prevete, Fabian Chersi, and Giovanni Pezzulo. "A Programmer–Interpreter Neural Network Architecture for Prefrontal Cognitive Control". In: *International Journal of Neural Systems* 25.06 (2015), p. 1550017.

[86] Yang Li et al. "Automatic Seizure Detection using Fully Convolutional Nested LSTM". In: *International Journal of Neural Systems* 30.04 (2020), p. 2050019.

[87] Hang Xie, Hao Tang, and Yu-He Liao. "Time series prediction based on NARX neural networks: An advanced approach". In: *2009 International Conference on Machine Learning and Cybernetics*. Vol. 3. 2009, pp. 1275–1279. DOI: 10.1109/ICMLC.2009.5212326.

[88] Maciej Lipiński, Tomasz Włostowski, Javier Serrano, and Pablo Alvarez. "White rabbit: A PTP application for robust sub-nanosecond synchronization". In: *2011 IEEE International Symposium on Precision Clock Synchronization for Measurement, Control and Communication*. IEEE. 2011, pp. 25–30.

[89] *White Rabbit Companies*. https://www.ohwr.org/companies. Accessed: 26/07/2018.

[90] Jean Paul Burnet et al. *CERN Proton Synchrotron working point control using an improved version of the pole-face-windings and figure-of-eight loop powering*. Tech. rep. revised version submitted on 2006-08-28 16:05:02. Jul 2006, 4 p. URL: https://cds.cern.ch/record/972423.

[91] M. Benedikt, F. Caspers, and M. Lindroos. "Application of magnetic markers for precise measurement of magnetic fields in ramped accelerators". In: *Part. Accel.* 63 (1999), pp. 57–73. URL: http://cds.cern.ch/record/379810/files/p57.pdf.

[92] Seong-Heon Seo, Andreas Werner, and M. Marquardt. "Development of a digital integrator for the KSTAR device". In: *Review of Scientific Instruments* 81.12 (Dec. 2010), p. 123507. DOI: 10.1063/1.3519303.

[93] Eike Feldmeier and T. Haberer. "Development Of A High Precision Integrator For Analog Signals To Measure Magnetic Fields In Real-time". In: *IPAC*. May 2013. URL: http://accelconf.web.cern.ch/accelconf/ipac2013/papers/mopwa001.pdf.

[94] M. Pezzetta et al. "B-train Performances at CNAO". In: *IPAC2011*. San Sebastian, Spain, Sept. 2011. URL: https://accelconf.web.cern.ch/accelconf/IPAC2011/papers/mopo039.pdf.

[95] Anthony Beaumont. "NMR magnetic field markers of the B-Train for the Proton Synchrotron accelerator". EDMS:993482. MA thesis. Belfort, France: UTBM, Mar. 2009. URL: https://edms.cern.ch/document/993482/1.

[96] M. Benedikt, F. Caspers, and T. Salvermoser. *Absolute Calibration of the CERN PS-Booster Bending Field with Beam and NMR*. Tech. rep. PS-OP-Note-98-23. Geneva: CERN, 1998.

[97] P. Di Cesare, C. Reymond, H. Rottstock, and P. Sommer. *SPS magnetic field cycle measurement system: Investigation on Nuclear Magnetic Resonnance (NMR) system to get a faster reading of field values*. Tech. rep. CERN-SPS-PCO-Note 89-9. Geneva: CERN, Apr. 1989. URL: https://cds.cern.ch/record/2158783.

[98] Anthony Beaumont, Marco Buzio, and Giovanni Boero. "Ferrimagnetic Resonance Field Sensors for Particle Accelerators". In: *Review of Scientific Instruments* 90.6 (June 2019). DOI: 10.1063/1.5097508.

[99] METROLAB. *PT2025 NMR Teslameter User Manual*. 2.0. METROLAB Instruments SA. Geneva, Switzerland, Sept. 2003. URL: https://www.metrolab.com/wp-content/uploads/2015/07/PT2025_user_manual.pdf.

[100] Christian Grech, Anthony Beaumont, Marco Buzio, and Nicholas Sammut. "Performance comparison of Nuclear Magnetic Resonance and FerriMagnetic Resonance field markers for the control of low-energy Synchrotrons". In: *Journal of Physics: Conference Series* 1065.5 (2018), p. 052022. DOI: 10.1088/1742-6596/1065/5/052022.

[101] aimTTi. *Aim-TTi TF930*. Available Online. AIM & Thandar Thurlby Instruments. Cambridgeshire, United Kingdom. URL: http://docs-europe.electrocomponents.com/webdocs/0d08/0900766b80d084a0.pdf.

[102] S. Gilardoni and Ed. D. Manglunki. *Fifty years of the CERN Proton Synchrotron:Volume 1*. Tech. rep. CERN-2011-004. Geneva, 2011. URL: https://cds.cern.ch/record/1359959.

[103] Anthony Beaumont, David Giloteaux, Alessandro V. Matheoud, Marco Buzio, and Giovanni Boero. "Electron paramagnetic resonance magnetic field sensors for particle accelerators". In: *Review of Scientific Instruments* 91.10 (2020), p. 105001. DOI: 10.1063/5.0004128.

[104] Christian Grech et al. "Error characterization and calibration of real-time magnetic field measurement systems". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 990 (Feb. 2021), p. 164979. ISSN: 0168-9002. DOI: 10.1016/j.nima.2020.164979. URL: https://www.sciencedirect.com/science/article/pii/S0168900220313760.

[105] Christian Grech, Marco Buzio, and Nicolas Sammut. "A Magnetic Measurement Model for Real-Time Control of Synchrotrons". In: *IEEE Transactions on Instrumentation and Measurement* (apr 2019). DOI: 10.1109/TIM.2019.2904073. URL: https://ieeexplore.ieee.org/document/8680651.

[106] *Open Hardware Repository*. https://ohwr.org/. Accessed July 08, 2020. URL: https://ohwr.org/.

[107] Patrick Loschmidt, Georg Gaderer, Natasa Simanic, Arshad Hussain, and Pedro Moreira. "White rabbit - sensor/actuator protocol for the CERN LHC particle accelerator". In: *2009 IEEE Sensors*. IEEE, 2009. DOI: 10.1109/icsens.2009.5398529.

[108] P. Jansweijer, H. Peek, and E. de Wolf. "White Rabbit: Sub-nanosecond timing over Ethernet". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 725 (2013), pp. 187–190. ISSN: 0168-9002. DOI: 10.1016/j.nima.2012.12.096. URL: http://www.sciencedirect.com/science/article/pii/S016890021201652X.

[109] Siemens. *SIMATIC IPC847E*. First. A5E43920357-AA. Nov. 2018. URL: https://support.industry.siemens.com/cs/document/109763148/simatic-ipc847e?dti=0&lc=en-WW.

[110] *About CentOS*. https://centos.org/about/. Accessed July 08, 2020. CERN. URL: https://www.ohwr.org/project/spec/wikis/home.

[111] T. Hoffmann. *FESA Primer*. Tech. rep. https://wiki.gsi.de/foswiki/pub/Personalpages/FESA/FESAPrimer_TH.pdf. Geneva: GSI, 2007. URL: https://wiki.gsi.de/foswiki/pub/Personalpages/FESA/FESAPrimer_TH.pdf.

[112] P. Alvarez-Sanchez, J. C. Bau, D. Domínguez, J. Lewis, and J. Serrano. "The Evolution of the CERN SPS Timing System for the LHC Era". In: CERN-AB-2003-103-CO (Nov. 2003), 4 p. URL: https://cds.cern.ch/record/693149.

[113] *Simple PCIe FMC carrier (SPEC)*. https://www.ohwr.org/project/spec/wikis/home. Accessed July 08, 2020. CERN. URL: https://www.ohwr.org/project/spec/wikis/home.

[114] Xilinx. *Spartan 6 FPGA*. https://www.xilinx.com/support/documentation/data_sheets/ds160.pdf. Xilinx. Oct. 2011. URL: https://www.xilinx.com/support/documentation/data_sheets/ds160.pdf.

[115] *WISHBONE*. https://cdn.opencores.org/downloads/wbspec_b4.pdf. Revision B.4 Specification, OpenCores, Accessed July 08, 2020. URL: https://cdn.opencores.org/downloads/wbspec_b4.pdf.

[116] *Wishbone slave generator*. https://ohwr.org/project/wishbone-gen/. Wishbone tool developed by OHWR, Accessed July 08, 2020. URL: https://ohwr.org/project/wishbone-gen.

[117] VITA. *FPGA Mezzanine Card (FMC) standard*. Approved in 2008, revised in 2010. Fountain Hills, AZ: VMEbus International Trade Association, 2008. URL: https://cds.cern.ch/record/1172409.

[118] L. Bojtar. *OASIS User Manual*. Feb. 2009. URL: https://cds.cern.ch/record/1164782.

[119] Pablo Alvarez. *CTR User Guide*. Tech. rep. CERN, 2007. URL: https://edms.cern.ch/document/598907.

[120] Maria Amodeo, Pasquale Arpaia, and Marco Buzio. "Integrator Drift Compensation of Magnetic Flux Transducers by Feed-Forward Correction". In: *Sensors* 19.24 (Dec. 2019), p. 5455. DOI: 10.3390/s19245455.

[121] P. Arpaia, L. Bottura, L. Fiscarelli, and L. Walckiers. "Performance of a fast digital integrator in on-field magnetic measurements for particle accelerators". In: *Review of Scientific Instruments* 83.2 (Feb. 2012). EDMS n. 024702, p. 024702. DOI: 10.1063/1.3673000.

[122] Keysight. *Agilent Technologies 34401A Digital Multimeter Manual*. Available: http://literature.cdn.keysight.com/litweb/pdf/34401-90013.pdf. Keysight Technologies. Feb. 2015. URL: https://docs-apac.rs-online.com/webdocs/0b40/0900766b80b40128.pdf.

[123] D. Jacquet, R. Gorbonosov, and G. Kruk. "LSA - the High Level Application Software of the LHC - and Its Performance During the First Three Years of Operation". In: (Mar. 2014). URL: https://cds.cern.ch/record/1696961.

[124] Thomas Bohl et al. "Momentum slip-stacking in CERN SPS for the ion beams". In: *Proceedings, IPAC 2019: Melbourne, Australia, May 19-24, 2019*. 2019.

[125] E. Todesco et al. *The Magnetic Model of the LHC in the Early Phase of Beam Commissioning*. Tech. rep. CERN-ATS-2010-154. CERN, June 2010, 4 p. URL: https://cds.cern.ch/record/1277918.

[126] Friedhelm Caspers, W. Heinze, J. Lewis, M. Lindroos, and T. Salvermoser. "An Alternative to Classical Real-time Magnetic Field Measurements using a Magnet Model". In: *International Conference on Accelerator and Large Experimental Physics Control Systems (ICALEPCS 97)*. CERN-PS-97-063-CO-OP. Nov. 1997. URL: https://cds.cern.ch/record/338904.

[127] Maria Amodeo, Pasquale Arpaia, Marco Buzio, Vincenzo Di Capua, and Francesco Donnarumma. "Hysteresis Modeling in Iron-Dominated Magnets Based on a Multi-Layered NARX Neural Network Approach". In: *International Journal of Neural Systems* 31.09 (2021), p. 2150033.

[128] Damien Perrelet et al. "White-Rabbit Based Revolution Frequency Program for the Longitudinal Beam Control of the CERN PS". In: *Proceedings, 15th International Conference on Accelerator and Large Experimental Physics Control Systems (ICALEPCS 2015): Melbourne, Australia, October 17-23, 2015*. 2015, MOPGF091. DOI: 10.18429/JACoW-ICALEPCS2015-MOPGF091. URL: http://jacow.org/icalepcs2015/papers/mopgf091.pdf.

[129] Maciej Marek Lipinski. *B-train over White Rabbit Wiki*. https://gitlab.cern.ch/BTrain-TEAM/Btrain-over-WhiteRabbit/wikis/home. URL: https://gitlab.cern.ch/BTrain-TEAM/Btrain-over-WhiteRabbit/wikis/home.

[130] *CRIO-WR wiki project*. https://www.ohwr.org/projects/crio-wr/wiki/wiki. Accessed: 19/07/2018.

[131] *White Rabbit Project*. www.ohwr.org/projects/white-rabbit. Accessed: 19/07/2018.

[132] *CUTE-WR wiki project*. https://www.ohwr.org/projects/cute-wr-dp/wiki. Accessed: 19/07/2018.

[133] M. Lipinski et al. ""White Rabbit Application and Enhancements"". In: ().

[134] M. Lipinski. ""Real-Time distribution of magnetic field values using White Rabbit the FIRESTORM project"". In: ().

[135] *White Rabbit Streamers IP Core*. www.ohwr.org/projects/wrcores/wiki/wr-streamers. Accessed: 19/07/2018.

[136] *SPEC wiki project*. https://www.ohwr.org/projects/spec?jump=wiki-wiki. Accessed: 19/07/2018.

[137] *B-Train over WR frame*. https://gitlab.cern.ch/BTrain-TEAM/Btrain-over-WhiteRabbit/wikis/wr-btrain-frame-format. Accessed: 19/07/2018.

[138] Analogic. *Data Precision 8200 User Manual*. URL: https://testequipmentconnection.com/specs/DATA_PRECISION_8200.PDF.

[139] Han Wang and Gangbing Song. "Innovative NARX recurrent neural network model for ultra-thin shape memory alloy wire". In: *Neurocomputing* 134 (2014), pp. 289–295.

[140] Molla S Hossain Lipu et al. “State of charge estimation for lithium-ion battery using recurrent NARX neural network model based lighting search algorithm”. In: *IEEE access* 6 (2018), pp. 28150–28161.

[141] P.R. Vlachas et al. “Backpropagation algorithms and Reservoir Computing in Recurrent Neural Networks for the forecasting of complex spatiotemporal dynamics”. In: *Neural Networks* 126 (2020), pp. 191–217. ISSN: 0893-6080.

[142] P Arpaia et al. “A software framework for flexible magnetic measurements at CERN”. In: *2007 IEEE Instrumentation & Measurement Technology Conference IMTC 2007* (2007).

[143] Marco S Fabus and Stephen Peggs. *Hysteresis and degaussing of H1 dipole magnets*. Tech. rep. Brookhaven National Lab.(BNL), Upton, NY (United States), 2019.

[144] Simone Gilardoni et al. “Fifty years of the CERN Proton Synchrotron: Volume 2”. In: *arXiv preprint arXiv:1309.6923* (2013).

[145] Christopher M Bishop. *Pattern recognition and machine learning*. springer, 2006.

[146] W. Venturini Desolaro, L. Bottura, Y. Chaudhari, M. Karppinen, and N. Sammut. “MEASUREMENT AND MODELING OF MAGNETIC HYSTERESIS IN THE LHC SUPERCONDUCTING CORRECTORS”. In: *Tenth European Particle Accelerator Conference (EPAC’06)*. 2006, pp. 26–30.

[147] Sam Roweis. “Levenberg-marquardt optimization”. In: *Notes, University Of Toronto* (1996).

[148] Jie Ding, Vahid Tarokh, and Yuhong Yang. “Model selection techniques: An overview”. In: *IEEE Signal Processing Magazine* 35.6 (2018), pp. 16–34.

[149] Pushpalatha C. Bhat, Harrison B. Prosper, Sezen Sekmen, and Chip Stewart. “Optimizing event selection with the random grid search”. In: *Computer Physics Communications* 228 (2018), pp. 245–257. ISSN: 0010-4655.

[150] Kazi Md Rokibul Alam, Nazmul Siddique, and Hojjat Adeli. “A dynamic ensemble learning algorithm for neural networks”. In: *Neural Computing and Applications* 32.12 (2020), pp. 8675–8690.

[151] Francesco Donnarumma et al. “A framework to identify structured behavioral patterns within rodent spatial trajectories”. In: *Scientific reports* 11.1 (2021), pp. 1–20.

[152] Pasquale Arpaia, Francesco Donnarumma, Antonio Esposito, and Marco Parvis. “Channel selection for optimal EEG measurement in motor imagery-based brain-computer interfaces”. In: *International Journal of Neural Systems* 31.3 (2021), p. 2150003.

[153] A. Apicella, F. Isgrò, R. Prevete, and G. Tamburrini. “Middle-Level Features for the Explanation of Classification Systems by Sparse Dictionary Methods”. In: *International Journal of Neural Systems* 30.08 (2020), p. 2050040.

[154] Mohammad Hossein Rafiei and Hojjat Adeli. "A new neural dynamic classification algorithm". In: *IEEE transactions on neural networks and learning systems* 28.12 (2017), pp. 3074–3083.

List of Figures

List of Tables