UNIVERSITÀ DEGLI STUDI DI TRIESTE

XXXVII CICLO DEL DOTTORATO DI RICERCA IN INGEGNERIA INDUSTRIALE E DELL'INFORMAZIONE

Embedded instrumentation platform on SoC/FPGA for mixed radiation fields

Settore scientifico-disciplinare: **ING-INF/01 Electronics**

DOTTORANDO
IVÁN RENÉ MORALES ARGUETA

COORDINATORE
PROF. FULVIO BABICH

SUPERVISORE DI TESI UNITS
PROF. SERGIO CARRATO

SUPERVISORE DI TESI ICTP
DR. MARIA LIZ CRESPO

ANNO ACCADEMICO 2023/2024

Acknowledgments

I would like to thank my host institutions, the Multidisciplinary Laboratory (MLab) of ICTP, and the University of Trieste for the opportunity to forge my skills as a professional, by providing the right tools, research environment, and the scholarship that permitted the completion of my PhD. To my supervisors, Maria Liz Crespo and Sergio Carrato, for their guidance, patience and support throughout this journey: I am very grateful for their advice and distillation of knowledge. A special acknowledgment to Maria Liz for the moments outside of academia, for believing in me and for motivating me to keep moving forward. Thanks also to Andres Cicuttin for sharing his knowledge and experience.

To the Nuclear Science and Instrumentation Laboratory, IAEA, where most of the experimental sessions of my research were carried out. In addition to their facilities and funding, they also provided key technical and theoretical foundations that brought to life the developments presented in this thesis. Special thanks to Mladen Bogovac and Kalliopi Kanaki.

To my beloved wife Janivelt Oliveros, who originally encouraged me to give the first steps that brought us to Italy. Your support impulsed us to live this adventure together: thanks for your confidence and continuous motivation. You inspire me to keep leaping ahead.

To my family, who have always supported me in many areas of life, even from afar. Most of my academic achievements are also thanks to your aid. Thanks to my dad, uncles, grandmother, grandfather, siblings, cousins, parents- and siblings-in-law.

To my MLab colleagues and friends, with whom, besides carrying out research activities, I lived unforgettable moments: Luis, Romina, Bruno, Werner, Maynor, Charn Loong, and Raul. Definitely this adventure would not have been the same without your friendship and support.

And to my friends from USAC, to whom I have reserve a special place in my heart. Thanks for the countless experiences and your support. Arriving at this goal was also thanks to you all.

Abstract

Instrumentation for nuclear and high-energy physics (HEP) experiments demands continuous improvement of complex custom electronic systems that allow further exploration of the state-of-the-art, fostering new discoveries through faster and more accurate continuous measurements under challenging conditions and harsh environments. These electronic devices usually require a dedicated hardware enabled by multiple firmware developments running atop, leading to embedded systems capable of reliably operate under the accuracy and data rates demanded by novel instrumentation.

Accordingly, with the advent of field-programmable gate arrays (FPGA) and FPGA-based system-on-chips (SoC), embedded systems for nuclear and HEP instrumentation can execute complex algorithms with high-throughput and low-latency requirements in concurrent on-chip deployments, enabling unprecedented sub-microsecond accurate and reproducible results.

This thesis documents a series of developments that resulted on an instrumentation platform for mixed radiation field measurements, enabled by SoC/FPGA technologies. This platform is capable of reliably discriminate online individual events from diverse radiation sources using a single detector, even under challenging distortion conditions caused by high radiation intensity or low signal-to-noise ratio. The deployment of a real-time machine learning model was fundamental to enable efficient event discrimination in the low-power embedded system, exhibiting exceptional performance in a device with portability features.

Many of the developments in the series required experimental data from mixed radiation environments in order to optimize and assess the implementations. In this regard, an open-source remote diagnostics platform was created to control and validate instrument prototypes in nuclear and HEP experiments. This tool permits developers to test in real-time data acquisition systems (DAQs) and detectors located in radiation-controlled areas, while they safely operate the DAQ remotely from a control room. This diagnostics platform enabled agile experiment deployments in constrained time slots, saving valuable time in the data recording sessions during the research. Once tested in the laboratory, a sequence of novel methods has been devised to implement the flexible embedded instrument for mixed radiation fields.

First, a method for fast pulse shape recognition targeting SoC/FPGA was formulated. This technique is based on the Pearson's correlation to create a matching filter, optimized to avoid expensive computational operations like square root. Equivalent detection performance under low signal-to-noise ratio was demonstrated in a SoC/FPGA deployment compared to the Pearson's algorithm, while substantially reducing latency and hardware utilization.

Then, a method based on frequency-domain analysis for gamma/neutron discrimination (FCI) in mixed radiation environments was developed, envisioning embedded real-time applications with single crystal detectors. Superior performance was experimentally validated when compared to pulse-shape discrimination (PSD), especially in the lower energy ranges. The method was optimized to be computationally efficient and was implemented in a SoC/FPGA device, showing excellent resource utilization at sustained operation frequencies compatible with online processing. Besides, a public dataset was released, featuring the first data collection with gamma/neutron labels supported by the neutron reactions of the scintillator. A simulation also proved that FCI is very likely to outperform PSD with organic scintillators, motivating further experiments using diverse detector technologies.

Finally, by gathering several contributions from the aforementioned developments, an embedded system was designed to overcome the limitations of FCI under continuous pile-up distortion. A low-SWaP (size, weight, and power) and high performance event discrimination instrument was created for mixed radiation fields, supported by a real-time machine learning model running on a low-end FPGA. The final system achieved the best performance/SWaP ratio among the latest developments. This optimization milestone enabled top-tier capabilities comparable with benchtop instruments, such as high event-rate continuous pile-up recovery and discrimination, in a flexible embedded platform with outstanding portability features. Experimental results demonstrated such metrics with a commercial integrated monocrystal detector designed for gamma/neutron discrimination.

Contents

List of Figures

List of Tables

Acronyms

ADC Analog to Digital Converter.

ASIC Application-specific integrated circuit.

AUC Area under curve.

BOM Bill of materials.

CLT Cross-level trigger.

CLYC Cs2LiYCl6:Ce scintillator crystal.

COMBLOCK Comunication Block.

ComBlock Communication Block IP Core.

CSI Critical success index.

DAC Digital to Analog Converter.

DAQ Data Acquisition System.

DSP Digital signal processor.

FIR Finite Impulse Response.

FN False negatives.

FP False positives.

FPGA Field Programmable Gate Array.

GUI Graphical user interface.

HB Hardware bridge.

HEPD High Energy Particle Detector.

HLS High-level synthesis.

ILA Integrated Logic Analyzer.

IP Core Intellectual Property Core.

LLD Lower-level discriminator.

MAE Mean absolute error.

MCA Multichannel Analyzer.

NaI Sodium iodide.

NaI(Tl) Sodium iodide doped with thallium.

NMAE Normalized mean absolute error.

PCI Pearson's correlation index.

PL Programmable logic.

PMT Photomultiplier Tube.

PR Precision-recall.

PS7 Processing System.

PSD Pulse Shape Discrimination.

PSNR Peak signal-to-noise ratio.

RTL Register Transfer Level.

SiPM Silicon Photomultiplier.

SOA State-of-the-art.

SoC System On Chip.

SPCI Simplified Pearson's correlation index.

SPI Serial Peripheral Interface.

SWaP Size, Weight, and Power consumption.

TN True negatives.

TP True positives.

UART Universal Asynchronous Receiver-Transmitter.

USB Universal serial bus.

Introduction

Ionizing radiation is daily present in our lives as a central pillar of modern technologies. Either naturally or artificially produced, radiation is utilized in medical diagnostics and treatments [23; 24], energy generation [25; 26], chemical industries [27; 28], agriculture [29; 30], research on clean energy [31; 32], space exploration [33; 34], among others. Thus, accurately measuring and controlling radiation sources is paramount to reliably benefit from such technologies, while keeping users safe from unwanted harmful exposure.

Radiation from other natural origins is among us as well, mainly from cosmic rays coming from outer space and terrestrial sources like soil and atmosphere [35; 36]. Although this background radiation represents a challenge in low-count measurements [37], some applications can take advantage of its existence, for instance to determine soil compositions [38], to accurately date organic samples [39], or to explore massive structures employing highly energetic particles [40].

Furthermore, in the last couple of decades tens of high-energy physics (HEP) experiments and facilities have been built [41], among which many are continuously recording data from highly-energetic ionizing particle interactions. Some of them are also being updated towards increasing their energy limits and event detection rates [42], or substantially increasing their precision [43], aiming at extending their overall discovery potential.

Generally, nuclear and HEP instrumentation rely on detectors that generate pulses

of electric current upon interaction with radiation. These events often last tens of microseconds [44], but in some specific cases they may be as short as a couple of nanoseconds [45]. Such pulses are frequently treated to extend their duration using either dedicated analog shaper circuits [35] or digital signal processing techniques [46], in order to better estimate the energy of the events and to reduce the overall noise. However, there are some applications on which pulse shaping cannot be utilized, due to the inherent distortion that may bury important features required for advanced processing, such as event discrimination in mixed radiation fields.

Commercial systems often exploit a time-domain analysis based on pulse-shape discrimination (PSD) to distinguish among event types in mixed radiation scenarios, estimating the energy deposition rate in the detectors [47]. An improved event discrimination method for mixed radiation detectors is documented in this thesis, leveraging characteristic frequency features of the pulses, while letting off any shaping requirements and increasing the discerning performance in comparison with PSD. Experimental evidence was found using a gamma/neutron (γ/n) Cs2LiYCl6:Ce (CLYC) detector, whereas promising results were obtained with an ad-hoc simulation for fast organic scintillators as well.

Moreover, the presence of noise is unavoidable in these applications, mainly caused by the inherent composition of the detectors and the front-end electronics of the data acquisition systems (DAQ) [35; 48]. Hereof, this thesis also documents how to overcome real-time event detection on severely deteriorated signals due to poor signal-to-noise ratios (SNR), improving performance compared to traditional cross-level trigger implementations generally found in commercial digitizers [49].

Additionally, dealing with event discrimination under intense radiation activity demands advanced reconstruction for piled-up traces, while preserving enough relevant features that enable the distinction among pulse shapes. Thus, combining these requirements in a real-time deployment is a tough task to conceive, but of relevant importance

in nuclear security applications, radiation monitoring, and HEP experiments, which demand high accuracy and low-latency response under high pulse rates [22; 20]. This thesis addresses the aforementioned needs in an embedded instrument featuring low size, weight, and power consumption (SWaP), targeting mixed radiation applications and tailored for efficient γ/n detection under such challenging conditions. The implementation was possible by deploying an optimized machine learning (ML) model in a custom hardware board, carefully tuned to provide on-chip classification in the sub-µs latency range.

Conceiving the aforementioned applications demands high-throughput and low-latency digital signal processing (DSP), on which Field Programmable Gate Arrays (FPGA) and FPGA-based System-on-Chip (SoC) excel. These devices allow flexible algorithm deployments, enabled by their reconfigurable logic elements, high speed digital interfaces as well as their dedicated DSP and memory blocks. However, that flexibility comes along with complex development and testing cycles, causing debugging sessions to take considerable effort. In this regard, a diagnostics platform for DAQ has been developed, on which validating the behavior of systems in radiation-controlled areas is eased by a graphical user interface emulating test equipment, such as oscilloscope and multichannel analyzer. This tool simplifies the deployment of advanced algorithms in custom instrumentation systems based on SoC/FPGA by reducing the effort required for experimental validation.

Research objectives

The main goal of this thesis is to develop a flexible embedded instrumentation platform for mixed radiation field measurements, leveraging FPGA and programmable SoC as the supporting hardware technologies. Specific implementations usually required in nuclear and HEP instrumentation are improved respect to the state-of-the-art, seeking the following objectives:

I. Develop an open-source remote diagnostics platform for data acquisition systems in nuclear and HEP applications.

II. Develop an efficient real-time pulse shape recognition method for embedded nuclear and HEP applications.

III. Enhance real-time γ/n discrimination performance for embedded system deployments.

IV. Release a public γ/n dataset using an integrated commercial crystal detector.

V. Create an embedded instrumentation platform for real-time event discrimination under intense activity of mixed radiation fields.

Scientific Contributions

A total of five original contributions have emerged in direct relationship with this thesis, fulfilling each research objective as follows. Conference paper [1] documents the development of an open-source remote diagnostics platform targeting data acquisition systems for nuclear and HEP instrumentation. Real-time measurements and statistics of the devices under inspection are presented through a graphical user interface. Remote access to devices placed in radiation-controlled areas is provided by leveraging existing networking infrastructures such as WiFi and Ethernet.

Journal publication [2] reports a simplified correlation method for real-time pulse shape recognition targeting SoC/FPGA-based applications, useful for event detection in low signal-to-noise ratio scenarios. This approach also has the potential application in nuclear and HEP experiments to distinguish among multiple event shapes, originated by different particles in mixed radiation fields.

In journal publication [3] a method targeting embedded real-time applications for classification of γ/n events in mixed radiation environments is documented, exhibiting

superior discrimination performance compared to traditional mechanisms with CLYC scintillators. The open dataset published in [4] is a result of a data curation process using the system detailed in [3], leading to the first public γ/n tagged dataset with silicon photomultiplier CLYC detectors, supported by the response of the scintillator reactions with neutrons.

The journal publication [5] overcomes the limitation of [3] in high event-rate scenarios, where radiation activity is so high that distortion occurs by pile-up. A machine learning classification model proved remarkable discrimination performance in a low SWaP design, resulting in an efficient real-time embedded mixed radiation instrumentation platform based on FPGA. This deployment combined a high-performance γ/n discrimination system comparable to bench-top implementations, in a portable instrument targeting dosimeters, uncrewed aerial vehicles, among others.

[1] I. R. Morales, M. L. Crespo, S. Carrato, Open Source Remote Diagnostics Platform for Custom Instrumentation in Nuclear Applications, in: Applications in Electronics Pervading Industry, Environment and Society, Springer Nature Switzerland, Cham, 2024, pp. 424–430. doi:10.1007/978-3-031-48121-5_61

[2] A. Cicuttin, I. R. Morales, M. L. Crespo, S. Carrato, L. G. García, R. S. Molina, B. Valinoti, J. Folla Kamdem, A Simplified Correlation Index for Fast Real-Time Pulse Shape Recognition, Sensors 22 (20) (2022). doi:10.3390/s22207697

[3] I. R. Morales, M. L. Crespo, M. Bogovac, A. Cicuttin, K. Kanaki, S. Carrato, Gamma/neutron classification with SiPM CLYC detectors using frequency-domain analysis for embedded real-time applications, Nuclear Engineering and Technology 56 (2024) 745–752. doi:10.1016/j.net.2023.11.013

[4] I. R. Morales Argueta. Gamma and neutron tagged dataset from CLYC SiPM detector [online] (jun 2023). Accessed: 2024-08-12. doi:10.5281/zenodo.8037058

[5] I. R. Morales, R. S. Molina, M. Bogovac, N. Jovalekic, M. L. Crespo, K. Kanaki, G. Ramponi, S. Carrato, Gamma/neutron online discrimination based on machine learning with CLYC detectors, IEEE Transactions on Nuclear Science (2024). doi:10.1109/TNS.2024.3498321

Other relevant publications

Owing to the multidisciplinary nature of the research environment, three relevant contributions related to machine learning applications for SoC/FPGA resulted from the tight collaboration with other colleagues. The conference proceeding [6] documents the deployment of two methods for particle identification in a SoC/FPGA platform: a comparison of a FIR matching filter with a classification based on machine learning is elaborated. In [7], the conference proceeding shows the advantages of compressing machine learning models targeting real-time SoC/FPGA deployments in particle detection systems, reducing the overall resource utilization and inference latency. The journal publication [8] showcases a step-by-step workflow to efficiently deploy several types of machine learning models in SoC/FPGA architectures for classification tasks, yielding a good compromise of memory footprint, inference times, and overall accuracy. The workflow in this publication was also utilized in the creation of the real-time machine learning model developed in [5].

[6] L. G. García Ordóñez, R. S. Molina, I. R. Morales Argueta, M. L. Crespo, A. Cicuttin, S. Carrato, G. Ramponi, H. E. Pérez Figueroa, M. G. Ballina Escobar, Pulse Shape Discrimination for Online Data Acquisition in Water Cherenkov Detectors Based on FPGA/SoC, in: Proceedings of 37th International Cosmic Ray Conference — PoS(ICRC2021), Sissa Medialab, 2021, p. 274. doi:10.22323/1.395.0274

[7] R. S. Molina, L. G. Garcia, I. R. Morales, M. L. Crespo, G. Ramponi, S. Carrato, A. Cicuttin, H. Perez, Compression of NN-based pulse-shape discriminators in

front-end electronics for particle detection, in: International Conference on Applications in Electronics Pervading Industry, Environment and Society, Springer, 2021, pp. 93–99. doi:10.1007/978-3-030-95498-7_13

[8] R. S. Molina, I. R. Morales, M. L. Crespo, V. G. Costa, S. Carrato, G. Ramponi, An End-to-End Workflow to Efficiently Compress and Deploy DNN Classifiers on SoC/FPGA, IEEE Embedded Systems Letters 16 (3) (2024) 255–258. doi:10.1109/LES.2023.3343030

Thesis Outline

The remainder of this thesis is organized as follows. Chapter 1 presents the background and state-of-the-art developments, establishing the foundations of this work. In Chapter 2, a diagnostics platform for embedded nuclear instrumentation systems is described. Chapter 3 dives into an efficient real-time pulse shape recognition method targeting embedded systems. In Chapter 4, a refined γ/n discrimination system is presented, optimized for real-time embedded deployments. Chapter 5 details how the limitations of event discrimination in mixed radiation fields under challenging conditions were overcome in a low-SWaP embedded instrumentation platform running a real-time machine learning model, tuned for γ/n discrimination. Finally, the conclusions and future work are presented in Chapter 6.

Chapter 1

Instrumentation for mixed radiation fields

This thesis studies efficient ways to accurately identify individual events in mixed radiation scenarios using single detectors. In this regard, sensitivity to diverse types of sources is mandatory, enabling information retrieval from mixed fields by diverse signal processing methods. Experiments were carried out to demonstrate the capabilities of the developments presented in the following chapters, mostly relying on a gamma/neutron (γ/n) monocrystal detector. This chapter starts by explaining the working principles of mixed radiation sensing systems, and specifically, the functionality of the commercial detector utilized for the experiments.

The proposed research goals are achieved by deploying a number of optimized state-of-the-art (SOA) algorithms in terms of throughput, latency, and resources, targeting real-time processing in nuclear and HEP instrumentation. As a consequence, the validation in hardware implementations was enforced in order to reveal the actual performance metrics and limitations of the algorithms on SoC/FPGA-based platforms. Hence, novel digital pulse processing methods and data acquisition systems (DAQ) for nuclear and HEP applications are further reported in this chapter, featuring also machine learning advancements for real-time applications.

Moreover, instrumentation demanding efficiency in terms of size, weight, and power consumption (SWaP) imply a critical stage in applications constrained by porta-

bility, autonomy, and reliability. Thus, recent developments featuring low-SWaP are also documented in this chapter, setting a reference for novel implementations requiring such kind of optimizations.

1.1 Mixed radiation detectors

Mixed radiation fields stand for the presence of multiple types of ionizing particles coexisting in a defined location [9]. These particles can be high-energy photons (X-rays, γ rays), particles (electrons, muons, protons, neutrons), or even atom nuclei. In any case, their kinetic energy, measured in electron-volt units (eV), must be high enough to provoke changes in the structure of the surrounding matter to be considered part of a radiation field [35].

Figure 1.1 shows a simulation of the energy spectrum (energy distribution) of a mixed radiation field on a pixel detector, exhibiting the expected behavior of diverse sources in the Large Hadron Collider (LHC) experiment at CERN [9].

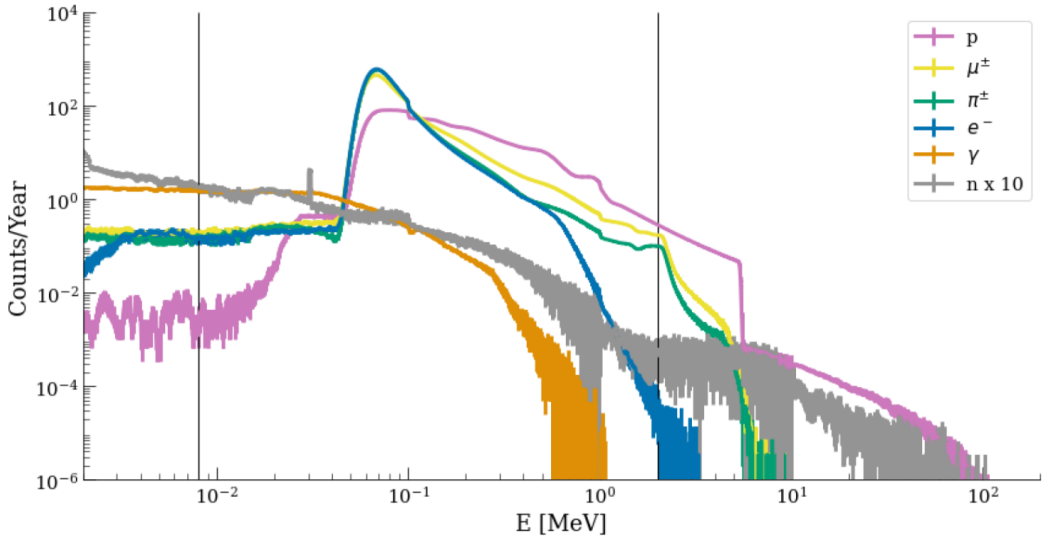

Figure 1.1: Energy spectrum of a simulated mixed radiation field at CERN, featuring several types of particles: protons (p), muons ($\mu^{\pm}$), pions ($\pi^{\pm}$), electrons (e^{-}), gamma rays (γ), and neutrons (n). Figure from [9].

In Figure 1.2, a per-pixel energy deposition measurement of a mixed radiation field from cosmic ray events is shown. Lighter charged particles (like electrons) exhibit a bent track due to a perpendicular external magnetic field. Discrimination among γ rays, electrons, and protons is possible by examining the deposition tracks in these types of pixel detectors [10].

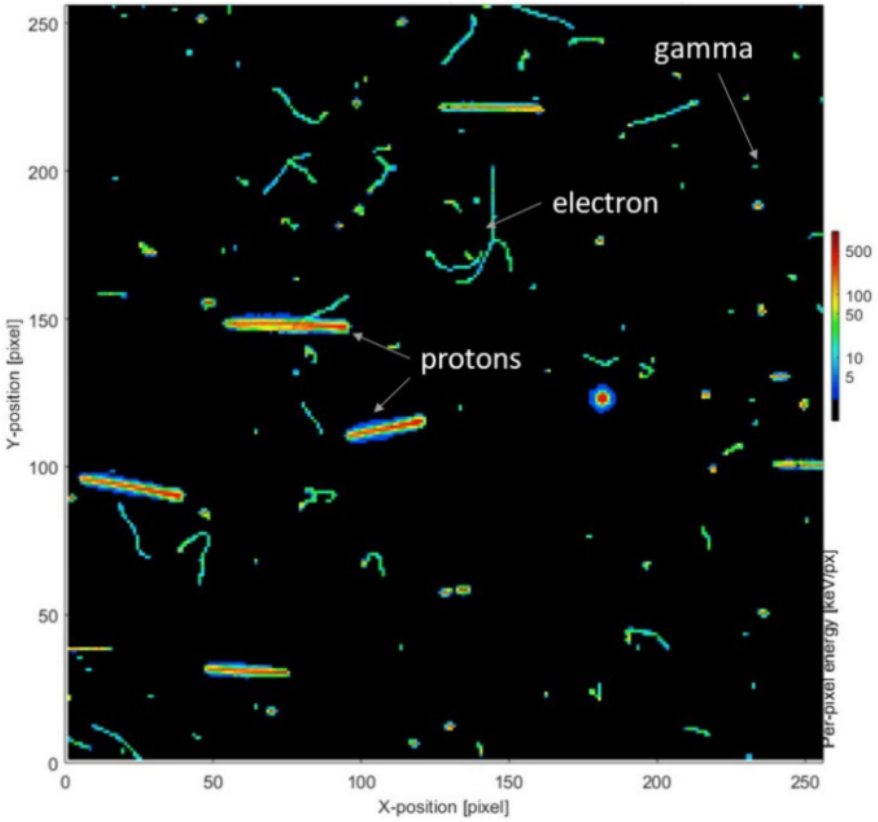

Figure 1.2: Pixel detector showing distinctive particle tracks of a mixed radiation field caused by cosmic ray events. Figure from [10].

Alternatively, event discrimination is mandatory to account for specific features of individual radiation sources in single (non-pixel) detectors, often carried out with pulse-shape discrimination (PSD). PSD is conducted by analyzing specific features of the produced waveforms (detailed in Section 1.3.2), such as rise-time [6], or leading

edge decay rate [50]. The performance metric is further quantified by a Figure-of-Merit (FoM) [51], regardless of the type of detector, with the highest values representing the best discrimination figures.

Although the discrimination features vary according to each application, this process relies on analyzing current pulses generated by detectors in HEP and nuclear applications. These events normally surge from the interaction of the ionizing particles in a homogeneous scintillator material, further captured by single or arrays of light sensors, such as photomultiplier tubes (PMT) or silicon photomultipliers (SiPM) [35]. Examples of such detectors capable of event discrimination are: water Cherenkov tanks in cosmic rays for muon/electron (μ/e) [52; 53], alpha/beta (α/β) systems for water quality assessment in distribution plants [54], proton/electron (p/e) detectors in dark matter studies [55], and gamma/neutron (γ/n) scintillators in space exploration and fusion research [56; 32].

Compound scintillators have also been developed, originally conceived as "phosphor sandwich" (phoswich) [57]. Nowadays, phoswich detectors are manufactured as a combination of scintillators optically coupled to exploit their individual discrimination capabilities in a single device, resulting in triple- or quadruple-purpose systems [58]. These phoswiches are usually attached to fast PMTs to extract most of the event features out of each pulse [59; 60], as shown in Figure 1.3. Despite the evident advantage, besides increased production costs, some technical challenges derived from stacking scintillators must also be considered, such as self-attenuation [11] and cross-talk [61].

Other mixed radiation detector technologies are based on direct current sensing, such as micromegas [62] and silicon carbide (SiC) [45]. Rather than using a light conversion through a scintillation material coupled to a light sensor, micromegas take advantage of gas ionization under strong electric fields, forcing free electrons into a defined path that further produce a measurable avalanche. These detectors are useful for large areas and high count rates, such as beam monitoring in boron

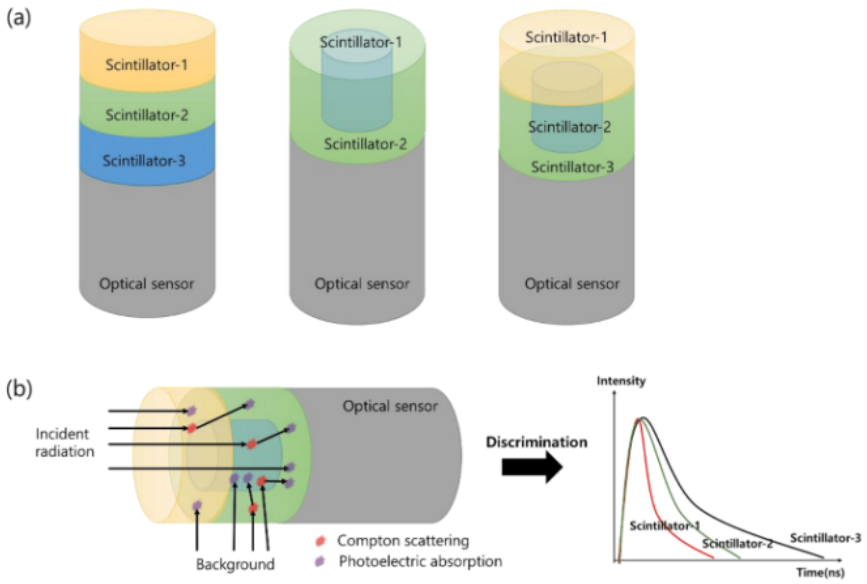

Figure 1.3: (a) Typical phoswich detector composition, featuring an array of optically-coupled scintillators attached to an optical sensor. (b) Event discrimination from different radiation types in a phoswich, yielding distinct pulse shapes. Figure from [11].

neutron capture therapy, demanding high spatial resolution. However, an important trade-off in terms of poor energy resolution should not be disregarded when opting for this technology [63]. The solid-state SiC have demonstrated outstanding detection performance in harsh environments, where elevated temperatures and considerably high mixed radiation doses (neutrons, gamma rays, and charged particles) would impair other detectors, such as germanium or cadmium-zinc telluride. They represent an active research area, due to their resilience and low-noise features specially useful in rough environments found in fission power and fusion experimental reactors [64; 65].

Typically, the complexity of discrimination systems scale up exponentially in HEP experiments where hundreds or thousands of detectors work simultaneously. However, the foundations of signal processing techniques for each element remain similar to those of the single-detector developments, further building up transmission and compression schemes atop to handle the corresponding huge amounts of data [66].

1.1.1 Mixed γ/n detectors

One of the most efficient options for thermal neutron (n_t) detectors at the time reckons the ^{3}He isotope [36], taking advantage of its large neutron capture cross section and low sensitivity to γ rays. However, due to the shortage of such material, alternatives have been developed [16; 67; 68]. For example, ^{6}LiF and BF3 proportional neutron counters offer stable operation over time and are designed to be also insensitive to γ events [69; 70].

Other alternatives rely on γ/n scintillators with discrimination capabilities, being two technologies that prevail in single-detector solutions: organic and crystal. The former is characterized by its fast response, enabling high event rates and precise time-stamping for time-of-flight (ToF) applications. The latter generally offers higher stopping power, exhibiting improved energy resolution for γ and n spectroscopy applications, resulting in better detection efficiency [71].

Fast organic scintillators such as liquid NE213 [72] have a renowned legacy for fast neutron (n_f) detection with excellent discrimination from γ rays. However, the toxicity, flammability, and handling considerations may complicate the usability out of controlled laboratory environments nowadays. Novel plastics like EJ-276 have been developed to overcome these issues, which, despite a slight decrease in the discrimination capabilities, are more suitable for field deployments such as beamline applications, homeland security and source identification [73]. In any case, thermal neutron sensitivity on these detectors is negligible, turning crystal scintillators into a more effective solution.

Single-crystal scintillators for mixed γ/n fields are commercially available either for double- and triple-function detectors. The traditional NaI(Tl) γ has recently shown excellent PSD performance when co-doped with ^{6}Li (NaIL) [74; 75], resulting in a low-cost reliable crystal for n_t detection and γ spectroscopy. However, neutron detection efficiency is not as high as other alternatives [76; 77].

Developments of elpasolite single-crystals have gained recent traction, mainly sparked thanks to their great proportionality and good light yield (LY) for gamma rays, while providing sensitivity to n_t and n_f spectroscopy capabilities. In this regard, CLLB (Cs2LiLaBr6) has shown better γ energy resolution and higher LY compared with NaI(Tl), while exhibiting a decent γ/n discrimination score (FoM) with n_t. Besides, CLLC (Cs2LiLaCl6) offers triple-functionality, being also sensitive to n_f, at the cost of reduced LY and a slight increase in FoM [78].

Both TLYC (Tl2LiYCl6) and CLYC (Cs2LiYCl6) exhibit similar overall characteristics, including n_f spectroscopy. Specifically, TLYC has shorter decay time profiles, allowing reliable detection at higher event rates. However, the n_t LY is almost half compared with CLYC, reducing the total efficiency when paired to similar photodetectors. Furthermore, γ/n discrimination FoM in CLYC is superior, enabling more reliable event separation at lower energy ranges [79].

Apart from a n_t capture event cluster from the ^{6}Li atoms in the n_f energy range, the TLYC presents a strong resonance from neutrons with energy around 244 keV [80], severely hindering the neutron spectrum deconvolution procedure. More details about the n_t capture cluster can be found in Section 1.1.2.

In general terms, CLYC has the best γ/n discrimination FoM among the evaluated crystals [78], and exceeds the NaI(Tl) γ resolution [12; 81; 82]. Although the γ LY is lower than its counterparts, CLYC offers the best overall features compared with other elpasolites recently developed, making it the ideal choice to exploit its capabilities and confidently explore novel techniques in several scenarios. Consequently, this thesis addresses most of the studies using a commercial mixed radiation detector based on the triple-purpose CLYC scintillator.

1.1.2 CLYC scintillator

The CLYC scintillator is designed to be mostly sensitive to n_t promoted by their capture in the Lithium nuclei within the crystalline array, enriched with at least 95% of ^{6}Li isotopes (CLYC-6) [81]. The nuclear reaction ([$n + {}^6\text{Li} \rightarrow T + \alpha$ (Q = 4.78 MeV)]) produces indirect scintillation in the Ce^{3+} ions with a Q-value (Q) proportional to the energy released during this interaction, which further manifests in the expelled tritium (T) and α particle pairs. In this regard, CLYC results in an excellent alternative for n_t detection, yielding higher efficiency per volume compared with ^{3}He probes [83; 84; 85; 86; 87; 88].

Moreover, γ rays are detected in CLYC through both direct electron-hole capture and core-to-valence luminescence (CVL) [89]. The CVL process is responsible for the differences in the rise and fall times of γ and neutron events in the produced light, enabling the γ/n discrimination capabilities of the crystal [90].

The detection and spectroscopy of n_f was further discovered [91], and is mostly caused by the [$n + {}^{35}\text{Cl} \rightarrow {}^{35}\text{S} + p$ (Q = 616 keV)] scattering reaction. These features render the CLYC crystal a viable alternative to n_t detectors based on He-3 [83] with an added value of n_f and γ spectroscopy.

In general, neutron interactions with the crystal provoke indirect scintillation events with Q corresponding to the reaction, observed as pulses on the γ energy-equivalent scale (keVee or MeVee). Quenching also occurs when the α and p (charged particles) byproducts of the n_t and n_f events travel through the crystalline network, being this effect more significant for heavier particles [89]. Thus, following the Briks' Law, the α and T mostly generated by n_t captures in the ^{6}Li will suffer a larger degradation in the total energy deposition observed as the scintillation light, compared to that of the protons mainly caused by the n_f reactions [81].

The indirect effect of n_t capture yields scintillation photons with an energy-equivalent magnitude of Q (4.78 MeV) from the exothermic reaction multiplied by the average

0.67 quenching factor for α and T, leading to pulses around 3.2 MeVee in the energy spectrum. In the case of n_f, the events are detected as the incoming neutron energy plus Q of the reaction proton (Q = 616 keV) multiplied by the quenching factor of 0.84 [92]. This mechanism also allows the computation of the minimum energy at which n_f can be detected as a product of scintillation, around 616 keV $\times$ 0.84 $=$ 517 keVee. Figure 1.4 shows in the upper lobe both the n_t cluster and a continuum illustrating the n_f spectrum. The vertical axis represents a pulse-shape discrimination (PSD) ratio that permits the distinction between γ and neutron events (further discussed in Section 1.3.2), whereas the horizontal axis shows the energy spectrum.

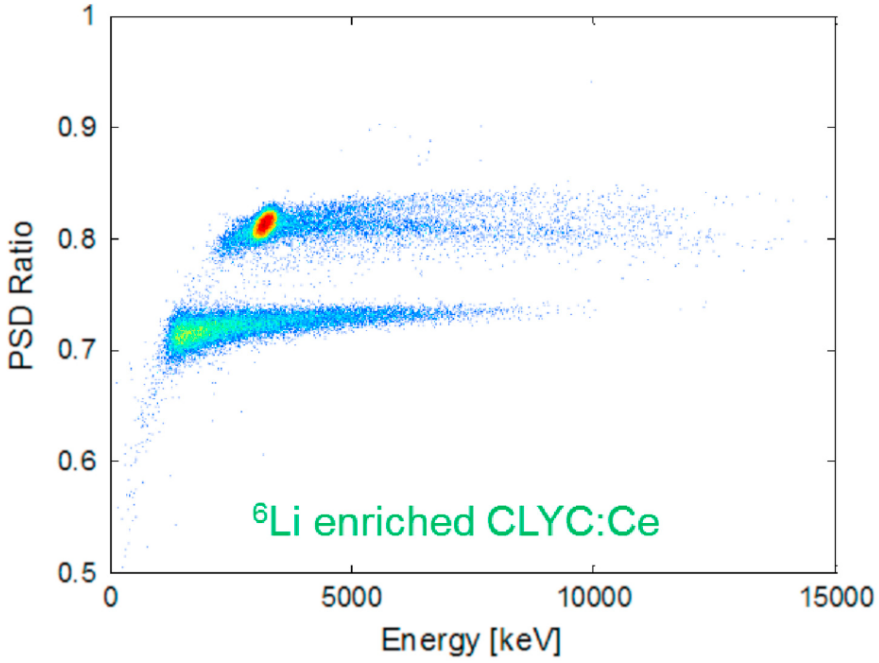

Figure 1.4: Thermal neutron capture cluster at 3200 keVee in CLYC represented by a 2D histogram (close to PSD Ratio = 0.8). Vertical axis represents a pulse-shape discrimination (PSD) ratio that distinguishes between γ and neutron events, whereas the horizontal axis shows the energy spectrum. Figure from [12].

The neutron detection efficiency in CLYC scintillators for n_t and n_f mostly depends on the atom count in the crystal and their neutron reaction cross sections. Thus, the

pulses corresponding to n_t detected through the ^{6}Li capture reaction are very likely to form a visible cluster of events close to 3.2 MeVee in the energy spectrum, surpassing by 3 orders of magnitude the amount of detected n_f [93]. This collection of n_t events is superimposed with the n_f spectrum, hindering reliable estimation of their energy deposition around that value. In this thesis, the 3.2 MeVee cluster is utilized for energy calibration purposes (further detailed in Chapter 4), leveraging the substantial amount of events evidencing a sharp distribution in the γ energy-equivalent scale.

Nonetheless, some applications require a continuous n_f spectrum reconstruction without the presence of the 3.2 MeVee cluster from the n_t events. In this regard, CLYC crystals enriched with 99% of ^{7}Li have been studied (CLYC-7), significantly reducing the ^{6}Li count, and naturally suppressing the n_t cluster [84]. From now on, CLYC will refer to the ^{6}Li-enriched scintillator (or CLYC-6) for sake of simplicity, whereas CLYC-7 will be explicitly addressed.

1.1.3 Integrated CLYC detector

Most of the results presented in this thesis were experimentally validated using an off-the-shelf CLYC probe manufactured by Scionix B.V. (model V12.7B30/SIP-E3-CLYC-X) [94], aiming at assessing the developments towards event discrimination in mixed radiation fields. This detector features a small cylindrical CLYC crystal with dimensions (ϕ=12.7 mm, h=30 mm), optically coupled to a 2x2 silicon photomultiplier (SiPM) array model ArrayC-60035-4P. A temperature-compensated bias power supply is integrated within the probe casing, along with a 50 Ω output preamplifier [94]. This setup, as shown in Figure 1.5, allows for easy deployment without any external circuitry for the SiPM polarization, temperature compensation, or impedance matching, rendering an excellent alternative for low-SWaP applications.

A CLYC detector coupled to a SiPM provides advantages for applications where coexistence with strong magnetic fields is mandatory [32; 95] or linearity comparable to

Figure 1.5: Integrated CLYC detector (Scionix V12.7B30/SIP-E3-CLYC-X). It includes a small CLYC scintillator optically coupled to a SiPM array, a temperature-compensated bias supply, and an output preamplifier.

NaI(Tl) for spectroscopy is required [96; 97]. Moreover, the integrated bias supply and preamplifier within the casing allow for a lightweight, smaller [32] and less complex setup (compared to the PMT counterparts) [98; 99] entitled for applications such as remote mixed radiation monitoring with uncrewed aerial vehicles (UAV) [100] and space exploration dosimetry [34; 56]. The main drawback of such integration level is the effect of pulse-shape degradation caused by the SiPM inherent capacitance [96; 101; 102], imposing a challenging task in terms of reliable event discrimination based on pulse-shape analysis, motivating the development of sophisticated methods for online processing in data acquisition systems (DAQ).

1.2 Embedded systems for nuclear and HEP instrumentation

An embedded system can be defined as a self-contained electronic device enabled by dedicated software (firmware) to accomplish its functionality, devised to reliably execute a set of specific tasks. The correctness of the system is dictated by an harmonious coexistence between hardware and firmware, providing time-accurate execution to fulfill the design requirements [103].

Usually, an embedded system is expected to continuously accomplish the designated duties without modifications through its entire lifespan. However, with the recent advent of reconfigurable digital devices, such as FPGA and SoC, the functionality may be easily upgraded, providing unprecedented flexibility with few or no changes in the hardware design [104].

Regarding nuclear and HEP instrumentation, mixed-signal electronic systems are mandatory, featuring analog interfaces to interact with sensors and advanced detectors, coupled to a digital domain through analog-to-digital and digital-to-analog converters (ADC and DAC). Some examples of embedded developments for these applications are portable radiation dosimeters [105], high-performance event discrimination instruments [20], space radiation monitors [106], and planetary exploration devices [107]. Figure 1.6 depicts a ThermoFischer active γ/n dosimeter, a great example of a portable embedded system for nuclear instrumentation.

Figure 1.6: ThermoFischer active γ/n dosimeter at NSIL, IAEA. This is an example of an embedded system targeting nuclear instrumentation: a self-contained portable device adapted to reliably accomplish a specific task.

This thesis presents a series of developments that document an embedded platform for reliable event discrimination in mixed radiation fields, featuring efficient operation towards portable applications operating under challenging conditions. The following chapters will guide the reader through the development stages, culminating with the embedded system in Chapter 5.

1.2.1 FPGA and SoC

FPGA and FPGA-based SoC play a crucial role in nuclear [108; 109; 110] and HEP [66; 111; 112] instrumentation. The flexibility of FPGAs to implement advanced digital designs, along with dedicated computation blocks, and a plethora of input/output ports, permit the deployment of custom algorithms to target specific requirements. The ability to design hardware subsystems and interconnect them using description languages enables the concurrent execution of operations with arbitrary computational precision, surpassing the capabilities of any general-purpose microprocessor or microcontroller for custom real-time requirements. These features enable combined low-latency and high-throughput in orders of magnitude only possible previously with application-specific integrated circuits (ASIC) [113].

FPGA-based SoC are an evolution of FPGA, where a general-purpose microprocessor and several peripherals (such as communication blocks, timers, and memory controllers) are included in the same chip with a reconfigurable FPGA fabric. This integration permits not only fast communication between the microprocessor and the FPGA with the on-chip network, but a reduced bill of materials (BOM) in complex board designs. Increased flexibility to co-partition the computational tasks is also possible, generally performing the high-performance real-time calculations in the FPGA fabric, whereas the non-critical operations are carried out by the microprocessor firmware [114].

Commercial products leverage the potential of SoC/FPGA technologies to produce nuclear instrumentation systems. For example, CAEN SpA. has recently released a high-level configuration software to allow users to deploy pre-defined processing and communication blocks in their SoC-based systems, enabling custom high-performance designs for demanding applications [115].

Regarding custom applications, the new electronics for the Pierre Auger Observatory features multiple detectors, including several PMT and SiPM channels as well as

a radio-frequency input. This embedded system is enabled by the firmware developments in a AMD Zynq-7000 SoC, aiming at capturing events of interest of extensive air showers induced by ultra-high energy cosmic rays [13]. The processing system in the SoC performs high-level tasks (such as communications) using a Linux distribution. The high-performance processing is carried out in the FPGA fabric, such as ADC readout, trigger generation, fast interfaces, and time synchronization. Everything is optimized for low power consumption, due to the restrictions provided by the solar power supply. Figure 1.7 shows the board developed for this purpose, featuring multiple SMA connectors for the PMT and SiPM inputs, coupled to multiple ADCs, several power domains for biasing and internal regulation, the SoC/FPGA, among others.

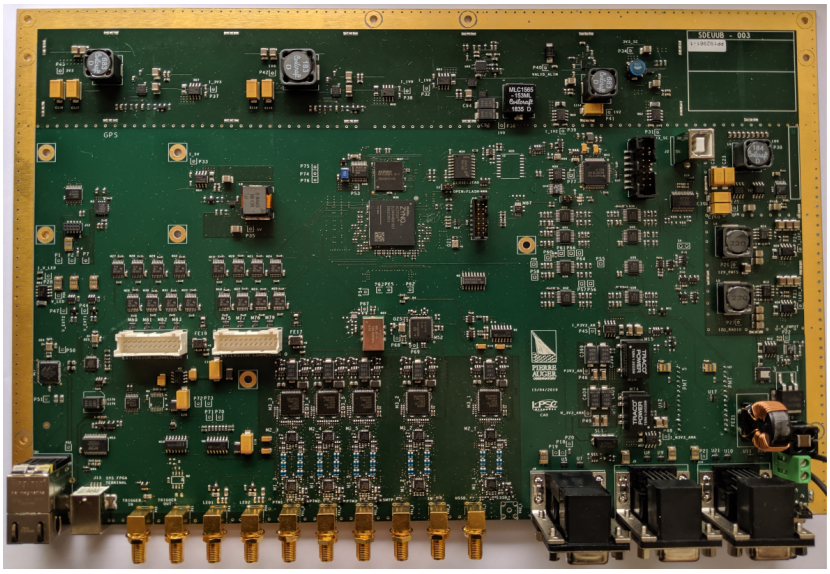

Figure 1.7: Upgraded Unified Board from the Pierre Auger Observatory. It features a SoC, several ADC channels for PMT and SiPM inputs, interface for GPS, among others. Figure from [13].

Also, Figure 1.8 shows a hardware board designed for a multichannel DAQ in a HEP experiment at CERN. The MSADC contains two interleaved ADCs and a AMD Virtex-4 FPGA, whereas the SoM is based on a newer AMD Zynq-UltraScale+ SoC. The embedded system relies on three firmware domains to enable a continuous feature

extraction device: FPGA design in the Virtex-4, FPGA fabric design in the SoC, and real-time operating system in the SoC multi-core processor [14]. This example highlights the relevance of a firmware working in symbiosis atop of a custom hardware board, on which a huge effort is devoted to obtain a embedded system platform [116].

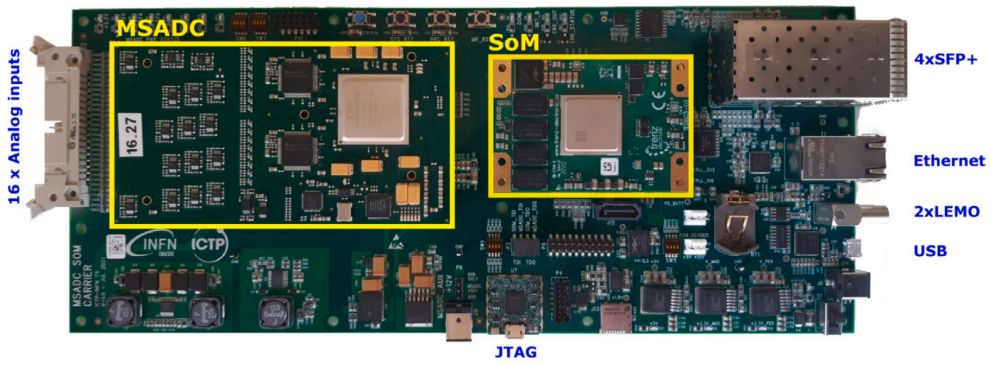

Figure 1.8: The FPGA Frontend Carrier Card is hardware board featuring a complex design with several analog interfaces, high-speed communication links, one FPGA, and one multicore SoC. This board is enabled by multiple firmware developments: FPGA design in MSADC, FPGA fabric design in SoC, and real-time operating system in SoC hard processor. The result is an embedded system for a multichannel feature extraction device in a HEP experiment at CERN. Figure from [14].

Throughout the development of this thesis, a similar approach was carried out in terms of firmware, targeting diverse hardware boards. FPGA firmware in different languages was coded for various target parts, along with C/C++ firmware for embedded processors. Details on this regard will be provided in the next chapters.

1.2.2 Analog interface

Detectors in nuclear and HEP applications provide signals whenever an event or set of events interact with them. Usually, a current pulse is created, which in some cases exhibits an amplitude proportional to the the deposited energy in the detector. Such pulses tend to be in the amplitude range of µV or mV after preamplification, and represent slight variations in the signal. Furthermore, this output may be either referenced to ground, or superimposed on a high-voltage potential in the orders of

thousands of volts. These events in the signal must be translated to the digital domain to be processed by the SoC/FPGA in the embedded system, often carried out by an ADC with specific characteristics that maximize the information retrieval of each pulse. However, the ADC inputs are usually coupled to a matched impedance of 50 Ω, compatible with single-ended or differential voltage signals with maximum amplitude ranging from 1 V to 2 V [13; 117].

In order to accommodate the incoming pulses from a detector output to the ADC input, an analog front-end (AFE) is required. The AFE consists of a set of passive and active analog devices capable of transforming the tiny current pulses from a detector into voltage signals, compatible with the input range of the ADC, while, in some cases, also isolating high potentials from the rest of the system [118]. A low-power AFE for SiPM-based detectors is documented in [119], potentially beneficial for portable and battery-operated instruments, further elaborated in Section 1.5.

The analog interface often includes a bias supply as well, featuring an isolated high-voltage output to power either SiPM or PMT sensors [18; 19]. Moreover, a temperature probe may be attached to the system, enabling automatic temperature compensation to accomplish stabilized measurements without further offline analysis [120]. Some applications also require advanced multichannel high-voltage control, demanding more complex digital closed-loop systems [121].

1.2.3 Custom DAQ hardware

Two hardware boards were utilized to validate the developments presented in this thesis. First, a commercial ZedBoard board served as a SoC-based system comprising an AMD Zynq-7000, external DRAM memory, Ethernet interface, and several input-output (I/O) ports [122]. The ZedBoard was mostly useful to deploy large computational models as means of comparison with the optimized techniques presented in this thesis, detailed in Chapter 3.

However, most of the research was conducted on a custom board featuring a more power-efficient FPGA, developed by the Nuclear Science and Instrumentation Laboratory (NSIL), International Atomic Energy Agency (IAEA). It was devised with multi-functionality in mind, i.e., its deployment in different field applications, targeting low-power operation [120]. It accommodates various sub-assemblies, including: (*i*) input power monitoring and management sub-system, (*ii*) digital interface controller, (*iii*) digital signal processing block, (*iv*) data storage, and (*v*) analog interface domain.

The backbone of the input power monitoring and management sub-system is based on a switched-mode Li-Poly battery charger with integrated system power path management. It performs several safeguard tasks such as over-voltage, over-current, and short-circuit protection.

Regarding low-power features, the board integrates two additional regulators and four load switches. These load switches can be used to shut down each sub-assembly providing a configurable system power consumption management.

The digital interface controller is based on a Cortex-M7 MCU, used to provide connectivity towards different peripherals through a USB port. In addition, tight control over different power domains is achieved via the input/output pins, managing the platform's low-power operation modes by enabling/disabling individual load switches. It is worth mentioning that the MCU lays in a separate power domain, automatically switched on when the external power supply is connected.

This hardware offers three additional serial ports to connect other peripherals, such as altimeter, global navigation satellite system receiver, and short range radio for wireless communication. The modular architecture expands its potential use in different applications, such as deployment on UAV remote radiation monitors and portable dosimeters.

The main component in the digital signal processing block is a low-end Artix-7 Series FPGA from AMD. This component does not contain any hardcore processors,

demanding significantly less power consumption compared to similar SoC solutions. Additionally, this approach does not require a heat sink, relaxing the platform enclosure constraints. Therefore, the FPGA-based design is crucial for the SWaP reduction requirements of the system. Still, a MicroBlaze soft-core processor was instantiated in all the designs to facilitate communication with external interfaces.

The system also allows storing data onboard using an 8 GB eMMC memory. This increases the field deployment reliability when packet loss occurs in the wireless data link, recording measurements for further offline readout.

The analog interface domain encompasses the AFE, digitally controlled bias power supply for SiPM, and analog interface for the temperature sensor attached to the detector. The AFE enables direct coupling of the SiPM output using an impedance-matched input. These signals are further converted to the digital domain using a 14-bit 250 Msps analog-to-digital converter (ADC) with differential and integral nonlinearities of 1.7 and 4.5 LSB, respectively. The AFE features a resolution of 122 μV/LSB with a 2 V_{pk-pk} analog input range, and a measured baseline noise RMS of 1.78 LSB (217 μV).

Temperature compensation is also performed to ensure stable event energy measurements, achieved by dynamically steering the voltage of the programmable bias supply.

The 2D render of the hardware board is shown in Figure 1.9, emphasizing the most relevant sub-assemblies.

In the context of this thesis, custom firmware was written from ground-up for the FPGA fabric and for the embedded processors (Arm Cortex-A9 and AMD MicroBlaze) targeting both boards, as individually demanded by the applications detailed in the following chapters. The firmware running in the Arm Cortex-M7 MCU for the NSIL's board was also developed, based on the original bring-up code. The MCU executes the peripheral and AFE setup, clocking system management, and low-power optimizations.

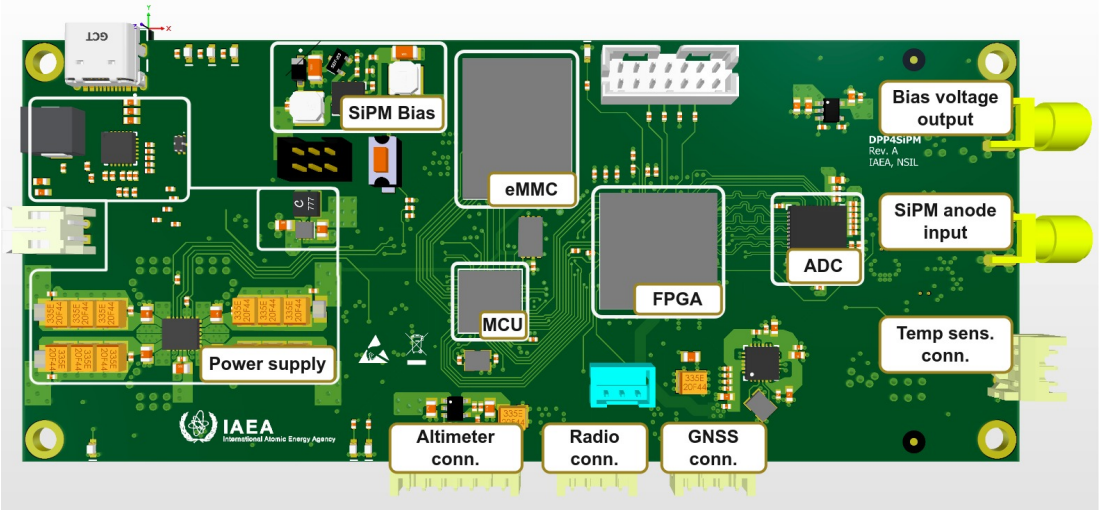

Figure 1.9: 2D render of NSIL's hardware board with marked sub-assemblies. Picture taken from [5].

1.3 Digital pulse processing

Early pulse processing systems were implemented in hardware, where analog systems were used to extract relevant information from each pulse, such as energy and timestamp, dealing with reduced SNR and high event rates. Nowadays, with the vast availability of fast analog-to-digital converters (ADC) and SoC/FPGAs the pulse processing stages have been pushed towards the digital domain, leaving just the essential pre-amplifying step to the analog counterpart [35]. This change of paradigm permits repeatability and significantly improves the flexibility to adapt existing designs to diverse detectors and applications. Tasks like tuning the shaping constants, swapping shaping functions, adjusting pile-up rejection parameters, and setting calibration gains can be achieved by modifying the firmware design in the FPGA. These advantages not only avoid hardware reworking, but allow to finely tune the system parameters dynamically, achieving optimal working conditions regardless of changes in temperature and analog component tolerances.

Still, digital pulse processing (DPP) is based on the same models implemented

on the analog counterparts, adapted to discrete-time domain. A typical DPP design for spectroscopy (event energy statistics) is composed by a pulse shaper, followed by a pile-up rejector, a baseline restorer, and a trigger mechanism, ending with a peak detector [35]. These stages permit the estimation of the energy for each event, by further feeding the multichannel analyzer (MCA) also deployed within the DPP chain. Figure 1.10 depicts a block diagram of such deployment in the FPGA firmware on a DAQ board.

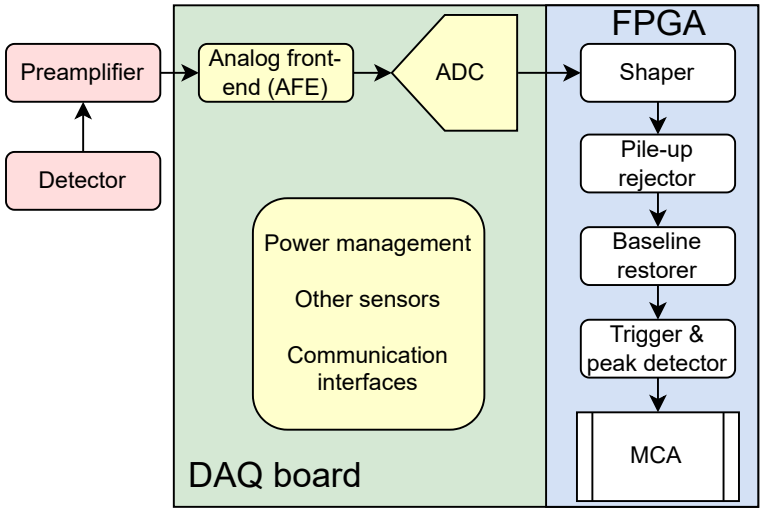

Figure 1.10: Typical embedded system setup for single-detector spectroscopy in high event rates. The DPP is implemented in the SoC/FPGA firmware (white background blocks). The FPGA and the components shaded in yellow represent physical devices onboard the DAQ, whereas blocks shaded in red are external to the board (detector and preamplifier), still part of the embedded system.

The pulse shaper is devoted to transform the digitized signal from the preamplifier to produce a train of short pulses with amplitude proportional to the deposited energy in the detector. Usually a band-pass filter time constants tuned to the expected input pulse shape is designed, taking care of avoiding overshooting and maximizing SNR. An efficient methodology to deploy such shaping filters in digital domain was proposed

by Jordanov et al., and remains as one of the optimal mechanisms to translate shaping circuits into DSP hardware implementations [123].

Pile-up rejection permits reliable energy accounting of events under intense radiation activity. This method is intended to drastically reduce spectral distortion caused by pile-up, at the cost of losing the actual event count [35]. Some systems, as the one presented in Chapter 5, use advanced techniques to recover information from piled-up traces, incurring in much lower pulse count losses, while enabling low spectral distortion.

Baseline represents the average digital output value of a detector in idle state, i.e., with no events present in the trace. However, in severe pile-up scenarios, baseline tends to be shifted towards higher values, increasing the uncertainty of the measured energy for spectroscopy. Long-term baseline shifts also occur, mainly caused by variations in temperature and pressure, or self-degradation of the detector in prolonged periods. The baseline restorer prevents this slow variations by serving as a high-pass filter featuring a long constant, gated (disabled) when an event arrival is detected. This mechanisms permits not only stable spectra, but allows to broaden the dynamic range of the system by utilizing the whole input range of the ADC [35].

Trigger allows to detect the presence of an event of interest, by analyzing a continuous stream sampled by the ADC of the system. The simplest and widely spread mechanism is based on leading-edge detection (also called cross-level trigger): a threshold is statically set just above the baseline noise level, signaling the presence of an event whenever the threshold is crossed by [35]. Although easy to implement in real-time hardware systems, scenarios with high dynamic range or low SNR may be drastically affected, limiting the detection capabilities and causing missed events. In systems demanding precise pulse timing, significant jitter caused by noise or high dynamic range may require more sophisticated approaches. In this regard, an advanced pulse detection mechanism targeting low SNR for precise detection regardless of the pulse

amplitude is presented in Chapter 3.

Once the arrival of an event has been signaled by the trigger, the peak detector finds the highest amplitude of the shaped pulse within a defined time period [35]. It can be also defined as a peak holder, on which the maximum value is attained for a defined time limit. Its output corresponds to a single digital value proportional to the deposited energy of the event.

Finally, the MCA consists of a histogram builder block that efficiently updates the bins according to each value provided from the peak detector. The output of this element is defined as the energy spectrum, revealing important features of the detected events. This information is often utilized for isotope identification [35], to reveal elemental compositions in X-ray fluorescence [124], to account for the presence of muons in cosmic-ray experiments [125], among others.

Precise pulse time-stamping and accurate counting may also take advantage of the DPP elements utilized for the MCA [35]. More importantly, thanks to the parallelism and flexibility provided by SoC/FPGA technologies, these DPP systems can be easily integrated within the same design, allowing for a multi-purpose reconfigurable instrument. Moreover, advanced DPP systems may support much complex methods to precisely extract the most relevant features in HEP experiments, significantly reducing the data transmission throughput requirements in multi-detector deployments [14].

1.3.1 Advanced pulse detection

Event detection is usually performed through the leading edge trigger mechanism with a threshold set just above the noise level. However, spurious events may seldom occur [53], creating false positives that contribute to higher uncertainty in event counts and distortion in the reconstruction of the energy spectrum. Tackling this issue, some methods have been designed to take out most of the events a system can detect, even in the lowest energy ranges, where the SNR is below the expected in a typical trigger

system. Some of these novel implementations usually try to emulate the expected pulse shape of the detector output with a matching filter, obtaining a measure based on the Pearson's correlation between the trace and the pattern filter.

Huang et al. [126] introduced a method for comparing the likeliness of a processed nonstationary signal to the expected output from a simulation. The discrete signal under test is synthesized by summing a known noiseless trace and random noise. This signal is then compared (after being pre-processed) with the expected trace using the Pearson's correlation coefficient, which sets the threshold for the detection capability of the system.

Pani et al. [127] developed a real-time neural signal decoding system based on a digital signal processor (DSP). Their algorithm also uses the Pearson's correlation to match a triggered signal against a known template to recognize the type of signal they capture. The fixed templates are normalized using z-score to improve the processing throughput. Iniguez-Lomeli et al. [128] implemented a FPGA-based real-time detection and sorting system, specifically designed for bio-signals classification. Their classifier relies on a pulse-shape recognition algorithm based on correlation. A voting stage assigns the type of signal that was detected, by choosing the highest correlation value between the signal and multiple pre-recorded patterns. Wang et al. [129] developed a discrimination algorithm based on Pearson's correlation, aiming at discerning the signal likelihood among multiple triggered events in a multi-channel neural processing system. Their method drops the events that do not match the pulse shape of their neighbors, avoiding spurious events being recorded and further improving the signal-to-noise ratio. Their tests were carried out using a SoC/FPGA-based device, capable of replicating their correlation algorithm in multiple channels. Although not directly related to nuclear instrumentation, the low-SNR scenarios and the continuous trace analysis in these works can be relevant for the applications concerning this research.

Another use of z-score normalization is described by Pollastrone et al. [130], where

a matching pattern or template is compared against a triggered signal event. However, owing to the method of measuring the similarity between the template and the triggered pulse (mean-squared error), precise triggering timing is required to align the incoming pulse with the pattern.

Moore C. and Lin W. [131] recently exposed the growing demand of optimized methods to compute correlation algorithms in embedded devices. Their goal was to develop a fast and accurate solution to correlate two signals in real-time. Their approach used a low-level hardware design to carefully leverage the available resources in FPGA devices. This thesis addresses the same challenge in Chapter 3, by proposing a simplified correlation index based on mean average deviation.

Aiming at quantifying the performance of a pulse detection system, Glenn et al. [132] used the precision–recall metric to assess the capabilities of their single-event counting system, complying with their imbalanced event distribution. They also used z-score normalization as the preprocessing stage in their algorithm.

Alharbi [133] compared a city-block algorithm (based on the absolute difference between two terms) with the Euclidean distance method, which is based on the square root of the difference between two squared terms. The research found that by separating two types of pulses using both algorithms, the classification accuracy was very similar, whereas the city-block algorithm resulted as more efficient in terms of computational complexity. This idea is exploited in Chapter 3 to significantly reduce the resource utilization in an SoC/FPGA implementation.

1.3.2 Pulse shape discrimination

The pulse shape of individual events derived from a scintillator are often relevant in mixed radiation fields, where a detector produces pulses with distinctive features that enable the discrimination among diverse type of radiation interactions [51]. These traces are usually captured through a leading edge trigger mechanism directly from

the digitized preamplifier output, disregarding the shaping stage that may wipe out the discrimination features of the pulses [50; 47]. The type of event is determined using a PSD factor, estimated as the ratio between the collected charge in the trailing edge (short window) and the total charge of the pulse (long window). Let x be the area under the trailing edge of the pulse and y the total pulse area (being the area proportional to the deposited energy of the event), the PSD factor can be defined as

$$PSD = x/y \tag{1.1}$$

Since x is a fraction of y and both are positive areas, PSD is dimensionless and expected to be in the range [0,1]. Figure 1.11 shows a simulation of two different captured pulses analyzed through this method: event type 1 (ET1) exhibits a faster decay rate, permitting its discrimination from event type 2 (ET2).

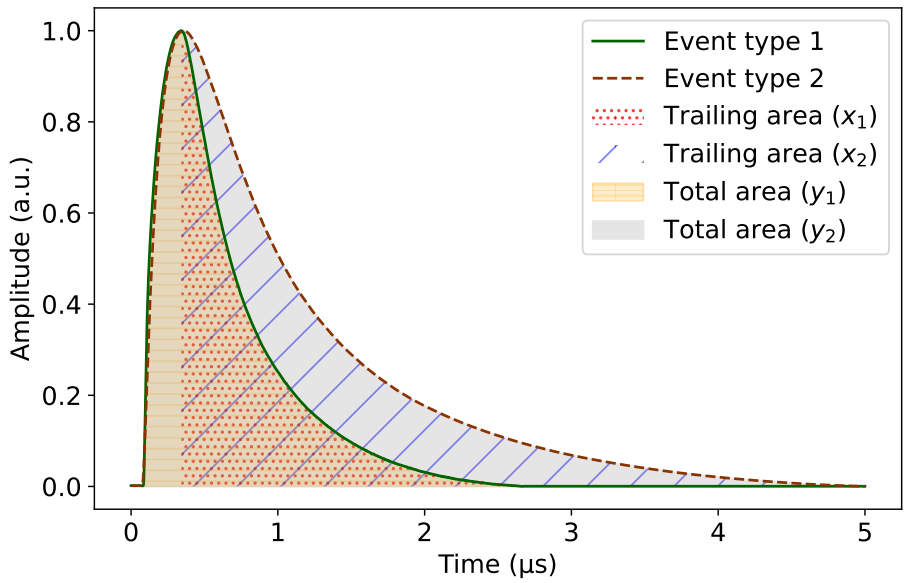

Figure 1.11: Pulse shape discrimination featuring two event types from a single detector simulation output. Short and long integration windows are marked as trailing (x) and total (y) areas, respectively.

In this case, ET1 will result in a smaller PSD factor, which in an experimental environment would probably represent lighter charged particles (such as electrons) or photons interacting in the scintillator; whereas the slowly ET2 decay time will constitute a higher PSD value, commonly associated to heavier particles, like muons or atom nuclei [125; 134; 135]. In this simulation, the PSD value for ET1 resulted in 0.71, whereas for ET2 the factor was 0.84. Discrimination concludes by assigning a label to each event based on the PSD factor, usually separated by a constant classification threshold, as conducted in Chapter 4.

Although the implementation methodology may vary in diverse real-time instrumentation systems, the definition in Equation 1.1 remains unchanged. The most common mechanism consists on delaying the incoming trace of the detector (by less than the duration of a pulse) and summing up the area of both the original and delayed traces within a determined time window. The PSD factor is then computed as the ratio of collected charges between the windows [35]. The window limit definitions may also vary according to each experiment requirements and might be defined by other regions in the pulse, which is the case of the analysis further detailed in Chapter 4.

Due to the nature of this approach, other nomenclature may be found in literature, like prompt and delayed gates, or short and long integration windows [136]. Moreover, the discrimination technique can also be referenced as PSD with charge-comparison method, or CCM [137].

1.3.3 Other pulse discrimination methods

Although PSD is widely used, novel developments have demonstrated to outperform its capabilities in diverse aspects, matching requirements for specific use-cases. Nonetheless, PSD has the advantage of easy deployment in hardware instruments demanding online processing, imposing an important starting point for cutting-edge systems. In this regard, Johnson et al. [138] developed a real-time method based on

digital filters for γ/n discrimination under elevated radiation activity scenarios. They computed a smoothed discrete-time derivative filter on the detector traces, obtaining a sharp pulse to signal the time of arrival of each event. These pulses were further re-shaped by a pass-band filter to enhance the differences between the gamma and neutron events. Although their solution provided a straightforward online method for FPGA implementation, their proposal required a long buffering stage, causing a dead-time close to 100 milliseconds. Unfortunately, the accuracy decreased at higher count-rates and the algorithm proof-of-concept was tested offline: no FPGA deployment was performed.

Garcia et al. [6] developed a discrimination method for cosmic ray detectors based on a finite impulse response (FIR) matching filter. They tuned the z-score-normalized filter coefficients such that the signal is correlated with the pulse shape of the expected pulses. The method was implemented on a SoC/FPGA platform for real-time event classification. Blair et al. [139] used a normalized cross-correlation method to distinguish between two types of pulse shapes. Their algorithm is based on the Pearson's correlation, similarly to those presented in Section 1.3.1.

Wen and Enqvist [15] used a fine-tuned triangular filter matched to the pulse shape of a CLYC detector applying the recursive implementation from Jordanov et al. [123]. As in [138], they measured the timing between zero-crossings at the filter output to determine the event type under high rates, achieving a constant dead-time of 2.67 µs. Although the measured neutron discrimination sensitivity was fairly low (28.9%) at the maximum count-rate (375 kcps), the detection rate increased to 100% at 25 kcps or lower. No real-time implementation was further developed, leaving the FPGA deployment as a future work. These performance figures are of relevance when compared to the development presented in Chapter 5. Figure 1.12 shows an experimental trace from their detector output with severe pile-up distortion, along with the filtered output using the proposed method.

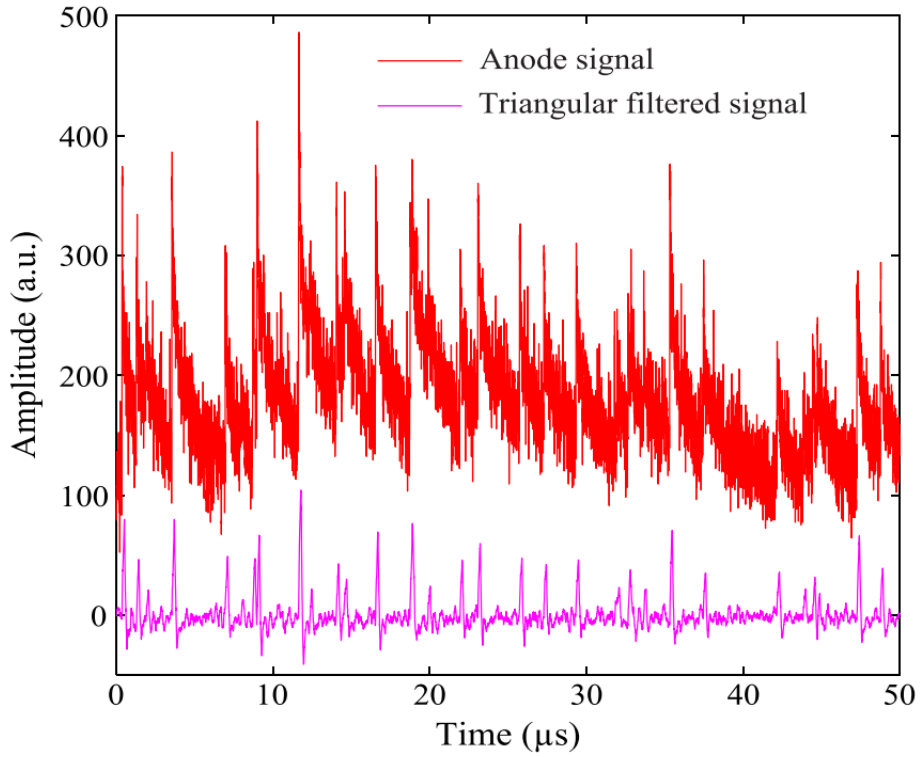

Figure 1.12: Experimental trace of a CLYC detector with a fast PMT under intense radiation activity, evidencing strong pile-up distortion. The filter proposed by Wen and Enqvist reveals a cleaner output (lower part) of the input trace (upper part of the plot). Figure from [15].

PSD with CCM is an analysis purely conducted in the time domain. However, methods based on frequency-domain (FD) have demonstrated their advantages in specific applications. For instance, Safari et al. [140] proposed a method for pulse identification in mixed γ/n radiation environments. Their main goal was to obtain a new classification metric based on a nonlinear weighting method. It was also demonstrated that a special case of the discrete Fourier transform (DFT) defines a new PSD parameter used in their method, described by a piecewise function that includes the zero-frequency (DC) component. The dataset was recorded at 5 GHz with an 8-bit amplitude resolution. They set a lower energy boundary to evaluate the performance, discarding low-energy events. Their FoM was computed using only events higher than

200 keVee and was reported to be $(1.48 \pm 1.28 \times 10^{-2})$. In general, the FoM represents a single value to score the discrimination performance, and will be relevant when compared to the method detailed in Chapter 4.

Dutta et al. [16] proposed analyzing the pulse decay times and the energy spectral density (ESD) of long trace windows that span through multiple events, recorded from a CLYC-based detector at 250 MHz with 12-bits resolution ADC. Their method aims at distinguishing gamma and neutron events in high count-rate scenarios with a fast PMT optically coupled to the crystal. Pulse alignment is mandatory to achieve a precise pulse-decay time estimation. By using Parseval's theorem, they demonstrated how the ESD analysis (computed with DFT) correlates with the deposited energy of gamma or neutron (equivalent) events in the crystal. Their method uses a pulse-detector algorithm that executes an offline search for each recorded long trace (sample-by-sample) after a low-pass FIR filter is applied. The best FoM value at the lowest measured event rate was 1.35. Figure 1.13 shows the offline pulse seeking methodology, accurately detecting events under high event rates.

Nakhosti [141] carried out an individual pulse analysis of triggered traces recorded at 4 GHz, aiming at considerably reducing the required sampling rate while still providing an acceptable FoM for γ/n discrimination. The proposed method is based on the Fast Fourier Transform (FFT), on which the individual pulse shape information containing the features is preserved for signals recorded from liquid scintillators. After subsampling the raw data from 4 GHz down to 32 MHz, two energy cuts were analyzed individually: the first one for events between 50 keVee and 200 keVee, while the other accounted for the remaining pulses with energies higher than 200 keVee. The FoM for the reported energy ranges were (0.62 ± 0.06) and (1.31 ± 0.04), respectively.

Polack et al. [142] stated that dealing with γ/n discrimination with energies lower than 1 MeV (low SNR) is challenging. In addition, all the authors in [34; 50; 141; 16; 140] applied energy cuts to disregard low-energy events when computing the FoM, even

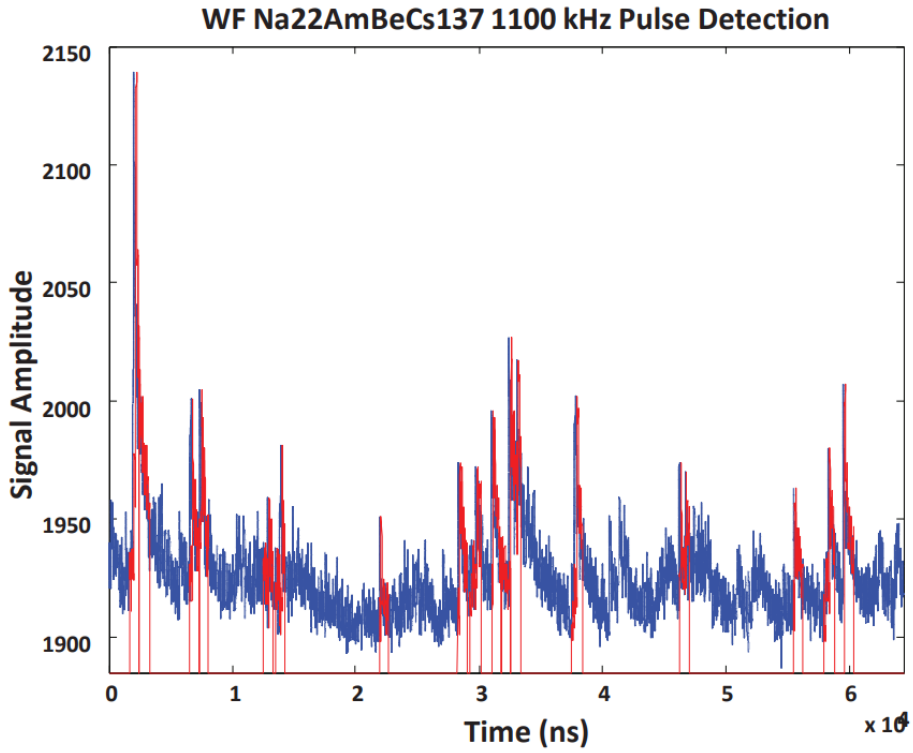

Figure 1.13: Offline seeking algorithm implemented by Dutta et al. for γ/n discrimination using a CLYC detector. The blue signal represents the digitized output from the PMT and red peaks reveal the detection of an event in the piled-up trace. Figure from [16].

with different types of scintillators (liquid, plastic, and CLYC) and light sensors (PMT and SiPM). Consequently, performing discrimination in the full detector energy range is still an issue to work with. Chapter 4 explores a way to overcome this issue, exhibiting higher FoM than PSD with CCM and low-SNR reliable γ/n discrimination in CLYC detectors.

1.4 Machine learning for HEP and nuclear instrumentation

Pulse shape discrimination has demonstrated to perform reliably in several deployments, including commercial devices for general-purpose DPP [47]. However, diverse

scenarios hinder the PSD capabilities, such as reduced SNR or intense radiation activity. Although frequency-based analysis proved better performance in low SNR [3], pile-up distortion is still an issue to deal with [143]. Aiming at overcoming such limitations, ML has been extensively utilized as an advanced technique to classify (or reject) individual event types, to reconstruct the trajectories of particles, and to estimate their energy in detector arrays.

Regarding HEP experiments, ML showed superior performance in a real-time muon/electron discrimination system in a water Cherenkov detector for a cosmic rays experiment, providing an overall accuracy higher than 96% with comparable latency to that of a PSD implementation [6]. Also, the muon telescope (MuTE) experiment has been recently exploring supervised and unsupervised ML techniques to reject the muon background from the vertical muon events [144]. In terms of ultra-high energy neutrino discovery, the IceCube experiment has developed a method based on a set of deep neural networks (DNN) to estimate the energy of individual events and reconstruct the direction of incidence with outstanding precision. These ML models demonstrated to successfully detect rare events from a set of traces with low SNR in an array of detectors, that otherwise may have been lost in the background [145].

CERN is also pioneering in the field of open-source ML developments, targeting both massive offline analyses and fast inference on FPGAs for real-time triggers [146; 147]. Of relevance is their *hls4ml* package, which allows to systematically deploy several types of ML models on SoC/FPGA devices. This package permits energy-efficient, low-latency, and high-throughput parallelized designs, systematically exploiting the available hardware resources (such as DSP blocks and BRAMs), tailoring the designs to specific characteristics of the target device [17].

It is worth noting that, following the in-house open-source workflow presented in [8], this research took advantage of *hls4ml* to produce a hardware block abstraction (IP core) out of an optimized ML model, further implemented on a low-end FPGA for event

discrimination in mixed radiation fields, as detailed in Chapter 5. Figure 1.14 depicts the end-to-end workflow to efficiently deploy a ML model in SoC/FPGA, whereas Figure 1.15 shows a typical workflow to translate an optimized ML model into a high-level synthesis project using *hls4ml*. In this case, the *hls4ml* workflow corresponds to the "Hardware synthesis tool for ML" intermediate step from Figure 1.14, as tested in [8].

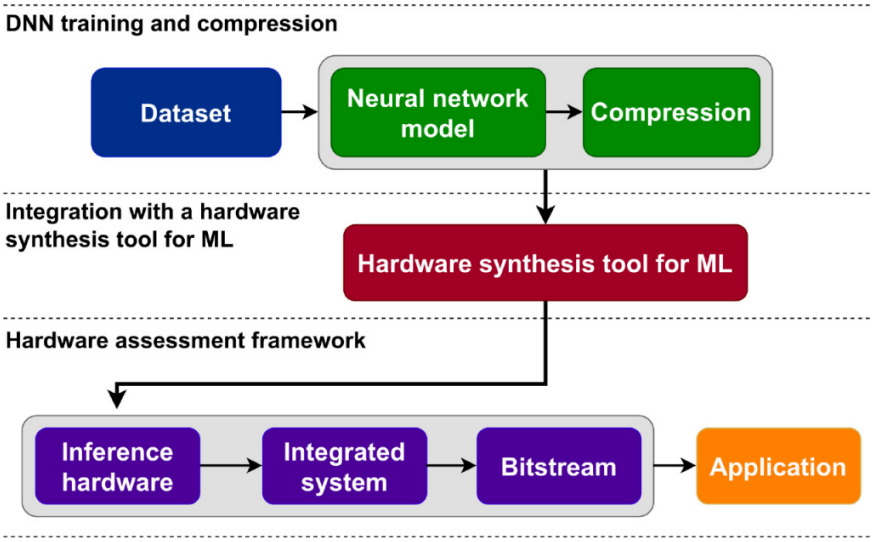

Figure 1.14: End-to-end workflow to efficiently deploy ML models in SoC/FPGA. The hardware synthesis tool for ML corresponds to *hls4ml* shown in Figure 1.15. Figure from [8].

In terms of nuclear instrumentation, many of the HEP principles can be applied for real-time data acquisition as well as for signal conditioning and processing. Since alike issues such as low SNR and pile-up persist in this field, ML turns to be a good alternative to overcome the limitations. As a reference, Yoon et al. [137] were able to improve the PSD performance on a γ/n discrimination system in the low-energy region, requiring pulses of low amplitude to be reliably binary classified either as gamma ray or neutron. They demonstrated how their trained convolutional neural network (CNN) architecture extracted the most relevant features of the pulses, increasing the sensitivity in the detection of fast neutrons in organic scintillators.

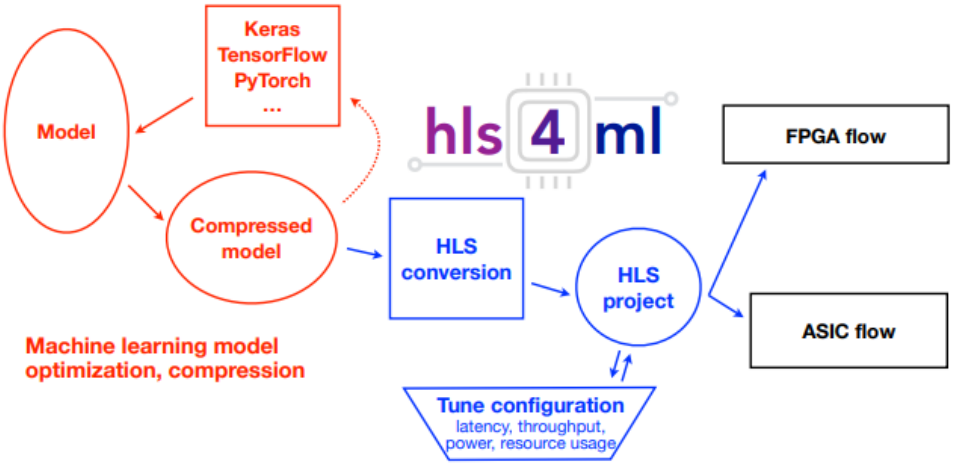

Figure 1.15: Workflow based on *hls4ml* to translate a ML model into a high-level synthesis project for SoC/FPGA deployments. Figure from [17].

1.4.1 Online ML discrimination

This research targets event discrimination with single detectors in mixed radiation fields, optimized for real-time execution in SoC/FPGA systems. In this regard, the most relevant SOA developments based on supervised and unsupervised online ML models are described next. Specific performance parameters, such as dead-time, count-rate, discrimination accuracy, and pile-up recovery are of relevant importance to contrast with the reported metrics from the system presented in Chapter 5.

Simms et al. [148] designed a system to detect the presence of neutron sources using liquid scintillators in mixed radiation fields, aiming at minimizing the contamination of false gamma counts. Their system relied on an unsupervised gamma mixture model executed on a FPGA. Although the measured latency was remarkable (3 µs), offline pre-processing was required. Moreover, no pile-up recovery nor rejection strategies were implemented, rendering the method unsuitable for high count rates.

Qi et al. [149] proposed a set of solutions based on unsupervised ML algorithms for γ/n discrimination running on FPGA with a CLYC detector, seeking to minimize

the required time to detect the presence of a neutron source with a confidence higher or equal than 95%. No single-event detection dead-time nor count-rate metrics were reported. However, they claim to be capable of identifying the presence of a neutron source in less than one second with a confidence of 99%.

Astrain et al. [22] implemented a real-time γ/n discrimination system based on a 1D convolutional neural network. Their deployment achieved a maximum count-rate of 79 kcps and a maximum dead-time of 50 µs with a steady overall classification accuracy of 98%. The system is reported to effectively perform single-event γ/n identification and pile-up recovery of up to two simultaneous events. Of relevant interest is that they designed the neural network using experimental data from the tokamak fusion reactor at the Joint European Torus.

The work recently presented by Michels et al. [20] documents a complete solution for real-time γ/n discrimination based on ML using a set of MLP models. Classification with an overall accuracy of 98.2% and pile-up recovery was tested at count-rates as high as 1.11 Mcps with a maximum dead-time of 7.7 µs. As in [22], up to two piled-up events could be identified using a captured (triggered) trace within a time window. Triple and quadruple pile-ups could be detected but not discriminated, while more than four piled-up events were not tested. A methodology to guarantee the ground-truth of the event types for the ML training was not conducted in the real-time validation. These high-performance metrics were achieved using a top-tier benchtop FPGA system, unsuitable for portable applications. An example of two piled-up pulses is shown in Figure 1.16, depicting a distortion that is very likely to be recoverable as indicated by [20; 22].

1.5 Low-SWaP instrumentation

When it comes to instruments meant for deployment in constrained scenarios, such as portable devices, aerial platforms, battery-operated unsupervised systems, or space-

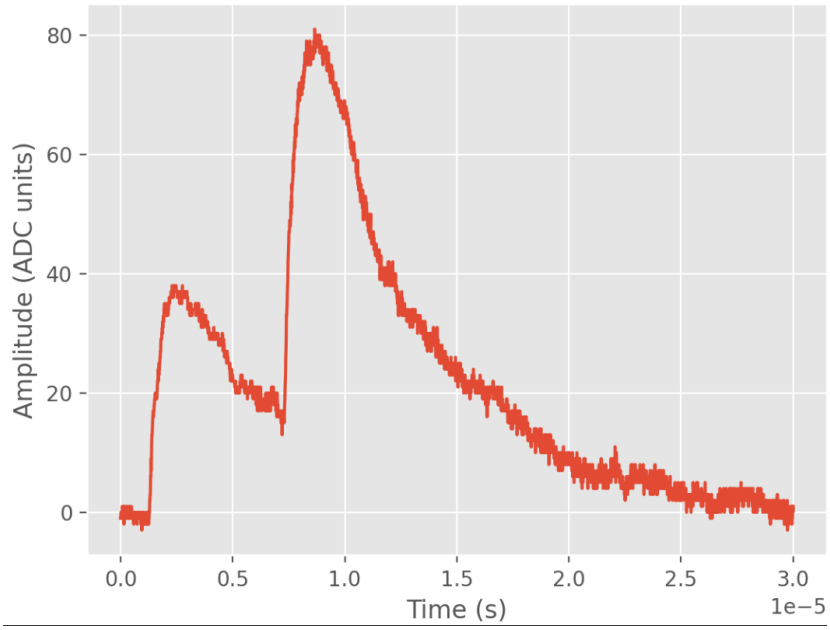

Figure 1.16: Piled-up trace from an integrated CLYC detector, caused by two independent events. ML algorithms have demonstrated the potential to extract accurate energy and timing information from this kind of distorted events.

related applications, optimizing critical parameters to ensure reliable autonomous operation is mandatory. In this aspect, it is widely accepted to tackle three specific variables: size (or volume), weight, and power consumption (SWaP). The objective is to reduce these parameters as much as possible, resulting in low-SWaP devices that exhibit reliable operation within the expected performance tolerances [36; 150; 151; 152]. A set of SOA low-SWaP mixed radiation instruments are described next, exploring their SWaP and performance metrics. These systems settle a reference framework for novel developments targeting similar operation conditions, including the platform presented in Chapter 5.

Huang [153] propose a portable neutron-gamma detector based on CLYC, enabled by a pile-up rejection stage followed by a partial charge-to-peak ratio algorithm. Their SWaP features were 4 L, 4.9 kg, and average consumption of 8.8 W. The maximum measured neutron count-rate was about 1 kcps.

Zhao et al. [18] document a low-power and low-size γ/n discrimination embedded system targeting environmental radiation monitoring using a CLYC detector coupled to a PMT. Their design is compliant with the IP66 standard in a 3.34 L volume, while operating with an average power draw of 3.5 W. Although no pile-up recovery was implemented, a maximum count-rate of 270 cps was certified for events with energy higher than 1.25 MeV. Figure 1.17 shows a 3D render of the embedded system cross-section, featuring a complete autonomous instrument for γ/n surveying.

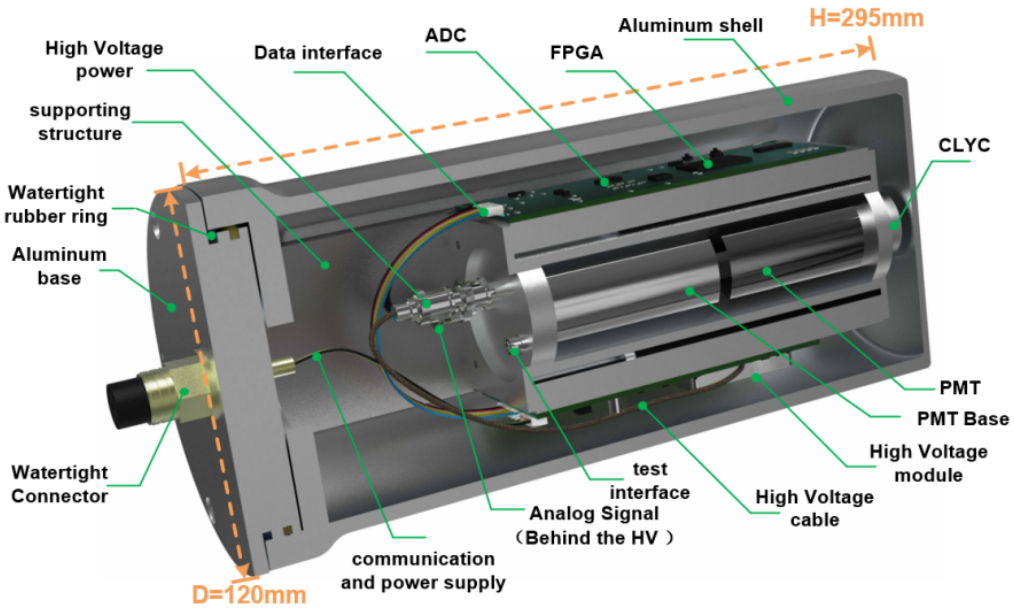

Figure 1.17: Low-power embedded system for autonomous radiation monitoring. Cross-section 3D render of the components, adapted from [18].

Hu et al. [19] exhibit excellent low-size and low-weight features using a CLYC SiPM detector: 0.24 L and 0.3 kg. However, neither power consumption nor γ/n discrimination performance measurements were reported and no pile-up recovery or rejection strategies were implemented, suggesting low count-rate capabilities. Still, a complete embedded system prototype was devised, looking towards a low-SWaP solution, including onboard detector, power supply, and casing. A block diagram of

this implementation is shown in Figure 1.18.

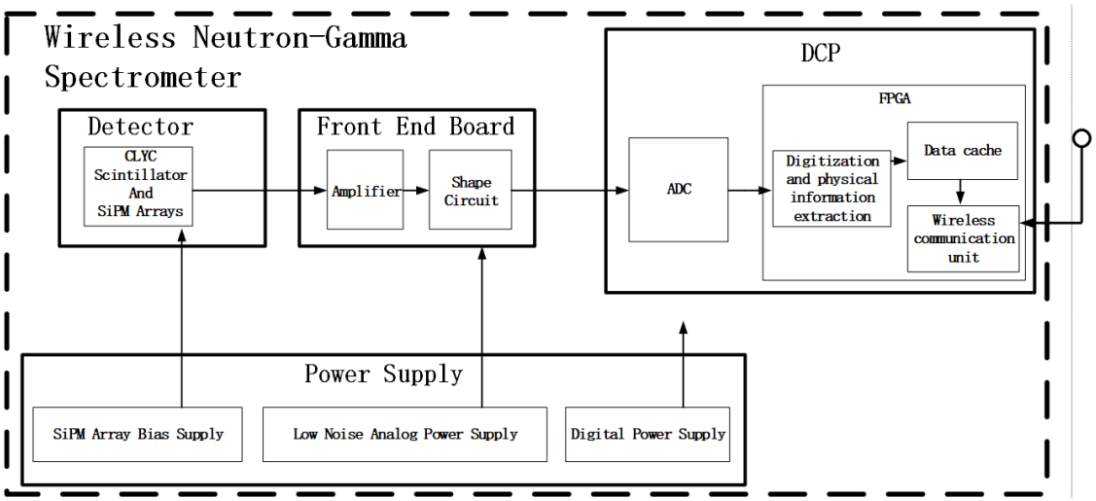

Figure 1.18: Block diagram of a low-size and low-weight embedded system for γ/n instrumentation, including detector, FPGA for DPP, wireless communication, and power supplies. Figure adapted from [19].

The latest commercial γ/n dose-rate meter from ThermoFischer Scientific (RadEye GN+) [105] is based on a CLYC scintillator. It exhibits outstanding low-SWaP features: 0.18 L and 0.16 kg, while offering a battery lifetime of at least 300 hours. However, the maximum neutron count-rate is 1 kcps with maximum gamma energy range of 1.3 MeV, providing reduced detection limits for typical mixed-radiation field applications. This device is an example of how SWaP parameters represent a biased trade-off in disfavor of performance, exhibiting the imposed limitations that challenge developers have to deal with when optimizing both aspects simultaneously.

Furthermore, it is important to stress on the fact that CLYC was chosen by many of the latest low-SWaP research and commercial systems, which along the advantages described in Section 1.1.1, support the detector selection for this thesis.

1.6 Summary

This chapter synthesizes the most relevant concepts and the SOA that support the developments documented in this thesis. It also describes the main components utilized in this research made by third parties, such as the commercial CLYC detector in Section 1.1.3 and the low-power NSIL hardware board from Section 1.2.3. The following chapters describe a sequence of developments that contributed to create an efficient embedded system for mixed radiation applications, supported by SoC/FPGA technologies and real-time machine learning.

Chapter 2

Remote diagnostics platform for DAQ in nuclear and HEP instrumentation

The development of novel data acquisition instruments for nuclear and HEP applications usually involves multiple subsystems coexisting in an embedded system platform, featuring real-time signal processing blocks along with communications and analog signal conditioning. Such increased complexity demands a diagnosing system to assess the functionality of the subsystems as a whole, both in controlled environments and on the final deployment conditions.

This chapter presents an open-source diagnostics platform for instrumentation hardware prototypes, targeting nuclear and HEP applications. The goal is to enable remote access to instrumentation electronics located in radiation-controlled areas by leveraging existent network infrastructures (research objective I).

The advances of this work were presented at the International Conference on Applications in Electronics Pervading Industry, Environment and Society (ApplePies) 2023 in Genoa, Italy. Derived from this presentation, the paper "Open Source Remote Diagnostics Platform for Custom Instrumentation in Nuclear Applications" was published in the conference proceedings on 13 January, 2024 [1]. This chapter contains an extended version of the original conference paper, with broaden details on the user

interface functionality and the evaluated setup. Specific target applications for the remote diagnostics platform in direct link with this research are also discussed.

2.1 Introduction

Testing custom hardware prototypes for DAQ in experiments related to nuclear or HEP instrumentation usually requires scheduling time slots for measurements during a beam run [154]. Therefore, having a graphical user interface (GUI) to assess the behavior of the detectors in real time is crucial to quickly verify the system response under controlled stimuli.

Since no physical access to radiation-controlled areas is allowed during the beam operation, personnel are forced to be located in a control room away from the detectors. In this regard, commercial solutions exist for both remote oscilloscope and multichannel analyzer (MCA) applications [155; 156], which are the most commonly used tools for this testing stage [157]; however, the compatibility is limited to the vendor-specific products and no open alternatives are available to create wrappers for custom hardware solutions. Moreover, many of them require a dedicated computer to be placed in the radiation-controlled area to act as a bridge with the instrument prototype under test to enable the connectivity via remote desktop.

The developed platform addresses both features in a single solution: a GUI for real-time diagnostics and a hardware bridge (HB) to enable remote access for instrument prototypes without network connectivity. The open-source GUI was created using the GNURadio software development kit, taking advantage of its building blocks, which are already optimized for fluid visualization and real-time multi-core processing, regardless of the operating system [158]. Hence, the intensive computation tasks are carried out in the computer where the GUI is being executed, making possible to use low-end and low-power devices as the remote HB.

Furthermore, no third-party libraries are required, enabling a flawless deployment

to easily start testing DAQ prototypes. To ensure the reliability of the transmitted data, the communication between the DAQ and the GUI is carried out using the transport control protocol (TCP).

In terms of alternative communication protocols, some commercial solutions rely on dedicated high-level solutions built atop TCP, such as LAN eXtensions for Instrumentation (LXI) [159], which consumes significant computational resources at the DAQ side and increases the development time for in-situ diagnostics purposes. Other protocol alternatives exist to address real-time remote data transmission: Padded Jittering Operative Network (PJON) designed to exploit noisy or long transmission distances on wireless interfaces [160], or Lab Streaming Layer (LSL), designed for multiple channel brain signals [161]; however, they require third-party or custom libraries to provide GNURadio compatibility and may hinder the deployment process with missing dependency or version issues [162].

This diagnostics platform overcomes such issues by using TCP transactions managed by GNURadio built-in blocks, leveraging existing cabled or wireless network infrastructures. Moreover, the GUI is compatible with multiple operating systems and the HB can be deployed in many different models of inexpensive off-the-shelf low-power single-board computers (SBC). These features provide flexibility to use the available hardware at the laboratories without compromising the fluid oscilloscope and MCA visualization.

This chapter is organized as follows. Section 2.2 gives a system overview of the diagnostics platform, detailing its components and how they interact among each other. In Section 2.3 the reliability and flexibility of the GUI elements are validated, along with the benchmarks for the remote data transmission. Details of the available tools in the GUI are also documented (Subsection 2.3.1), stressing the versatility of the included oscilloscope and MCA for real-time visualization in nuclear and HEP experiments. The hardware bridge performance results are also reported (Subsection

2.3.2), mentioning the provided datasets from multiple detectors and radiation sources. In Section 2.4 the final remarks are discussed, highlighting the contributions of the remote diagnostics platform and the utility for new hardware developments. Lastly, Section 2.5 summarizes this chapter.

2.2 System overview

The diagnostics platform is composed by two main elements: the GUI and the HB. The GUI displays an oscilloscope and MCA in a computer located in a control room (low-radiation area). It also handles the communication with the remote instrument under test (in the radiation-controlled area) through an existing network infrastructure, such as Ethernet or WiFi.

The instrument under inspection may either establish direct connection with the GUI or use the provided HB to generate a transparent TCP link via an universal asynchronous receiver-transmitter (UART) interface. Figure 2.1 depicts a general overview of the system. The HB is an optional element that may be disregarded if the instrument being diagnosed features TCP/IP communication capabilities.

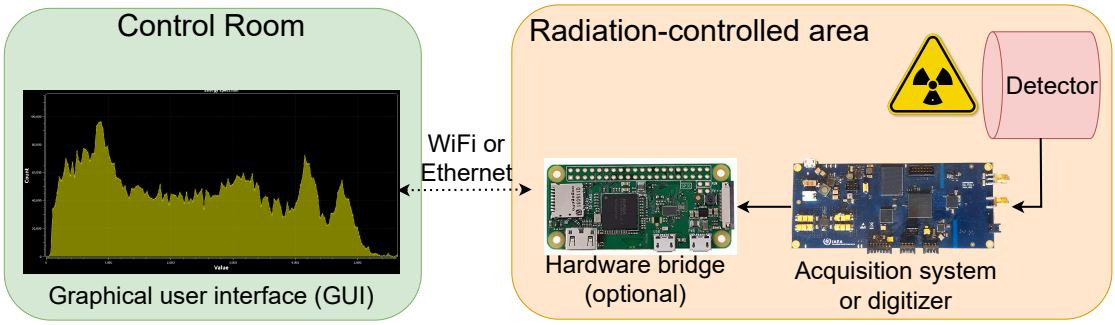

Figure 2.1: Block diagram of the system components distribution, including the optional hardware bridge between the remote data acquisition system and the network infrastructure.

2.2.1 Graphical user interface

The GUI was developed in GNURadio, making it compatible with multiple operating systems (Linux, MacOS, Windows). Moreover, a GNURadio port for Android (currently under development) may also enable its execution on a smartphone or tablet [163]. Such flexibility provides access to the DAQ or detector under diagnostics from mostly any available device in the control room of the laboratory. The GUI features some of the most common tools required for debugging and diagnostics of nuclear and HEP detectors and digitizers, i.e. oscilloscope (real-time trace visualization) and MCA (energy spectrum) [164].

2.2.2 Hardware Bridge

Albeit the HB can be deployed using several Linux-compatible SBC models, a low-end and low-power Raspberry PI Zero W (version 1.1) was selected to demonstrate the few computational requirements needed for its operation and the flexibility to power it from any available USB port or phone charger. This SBC features only 512 MB of DDR2 RAM and a single-core 1 GHz ARMv6KZ CPU: humble specifications compared to other models available nowadays.

2.3 Platform validation

The custom hardware board developed by NSIL presentned in Section 1.2.3 was chosen as the DAQ to be diagnosed for validation. It was connected to the HB via universal serial bus (USB) emulating a UART interface. Figure 2.2 shows the the physical setup in the radiation-controlled area, featuring the DAQ under test with a γ detector, plugged to the HB through a USB-UART converter.

The GUI was executed on a desktop computer (located in the control room) with 16 GB of DDR4 RAM and a Core i7-8700 processor running GNURadio version 3.10.5.1

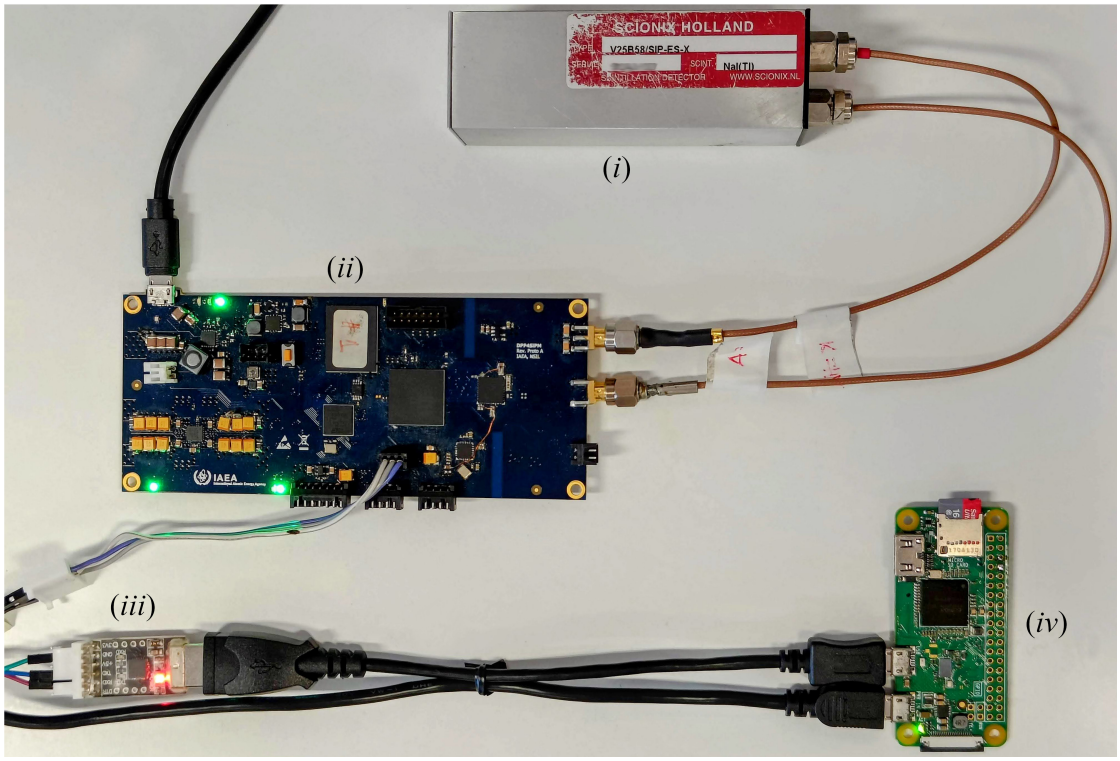

Figure 2.2: Setup of remote deployment in radiation-controlled area: (*i*) SiPM NaI(Tl) γ detector, (*ii*) NSIL's hardware board used for diagnostics, (*iii*) USB to UART converter, (*iv*) hardware bridge based on Raspberry PI Zero-W.

on Linux Ubuntu 20.04.6. The DAQ and detectors were situated about ten meters away, in the radiation-controlled area. The GUI stability was tested during long-term measurements of 24-hours continuous sessions with no issues reported or detected.

The oscilloscope and MCA were validated with the remote DAQ data using multiple detectors (one at a time) stimulated with a Cobalt-60 (Co-60) γ source: a sodium-iodide (NaI(Tl)) crystal with a silicon photomultiplier (SiPM), a NaI(Tl) crystal with a photomultiplier tube (PMT), and a CLYC (Cs2LiYCl6:Ce) crystal coupled to a SiPM. The source was located close to the detectors and served as a reference for pulse shape visualization and energy spectrum computation.

The HB demonstrated stable transmission through the wireless interface between

the DAQ (at the radiation-controlled area) and the computer executing the GUI (in the control room). A sustained throughput close to 2.1 Mbps was measured, while the system drew an average of 934 mW from the power supply. Alternatively, rates close to 27 Mbps were reached with a cabled Ethernet connection.

2.3.1 Graphical user interface

The following settings are required to be defined prior to the first execution of the application: remote IP address and port number (either from the HB or the remote DAQ), sampling rate of the DAQ (in MHz), and trace length (in sample units). After establishing the correct values the GUI can be launched from GNURadio, leading to a window with two tabs (oscilloscope and MCA), described next.

Oscilloscope

Based on GNURadio oscilloscope block, it provides the user the basic tools required for the real-time diagnostics of detectors and the DAQ prototypes attached to them, such as baseline, pulse shape and pulse height distribution. Figure 2.3 shows the oscilloscope tab featuring the pulse shape of a NaI(Tl) scintillator coupled to a SiPM and digitized by the remote DAQ under test. During run-time, this tool provides the following features and settings: cross-level trigger threshold with pulse polarity selection, multiple trigger sources (signal, peak detector), baseline restorer (BLR), peak detector and pulse height analysis (PHA), and recording of individual traces in binary format. An update rate close to 15 traces per second was measured with short trace lengths (128 samples) in cabled Ethernet mode.

Multichannel analyzer

Based on GNURadio histogram block, it displays the estimated deposited energy spectrum in the detector. The individual energy of each event is computed in real-

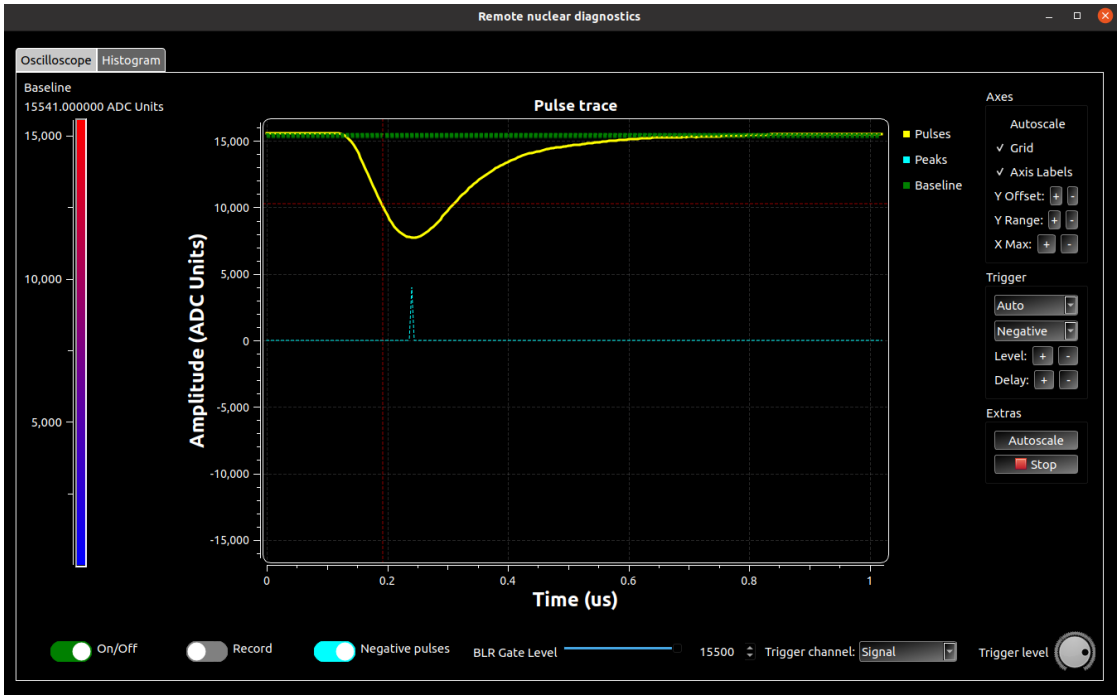

Figure 2.3: Oscilloscope tab: the yellow trace is the raw data received from the acquisition system. The horizontal dashed green line represents the computed baseline of each trace and the cyan pulse shows location of the maximum pulse amplitude. The vertical bar on the left allows the user to visualize the computed baseline in real-time as an aid to set the BLR gate level. The cyan slider sets whether the system expects negative or positive pulse polarity, while the record selector starts the binary recording of the visualized traces. Zooming operation is possible with a click and drag mouse operation as well as with the autoscale button.

time as the peak amplitude (using the PHA) of the pulse minus its baseline (from the BLR value). Flexibility to the user is provided by tuning the following settings during execution: number of histogram bins, auto-scale and window zooming, limited number of events or accumulative spectrum. The spectrum computation can be toggled off to reduce the overall CPU utilization in the control room computer, if required.

Figure 2.4 shows the MCA tab with experimental data built upon a Co-60 γ source located close to the NaI(Tl) detector with SiPM, expressed in ADC channels using 256 bins resolution. The two characteristic photopeaks at 1.17 MeV (above channel 4000) and 1.33 MeV (below channel 5000) as well as the Compton back-scatter peak (210 keV)

and the Compton edge (936 keV) evidence the expected behavior of the system.

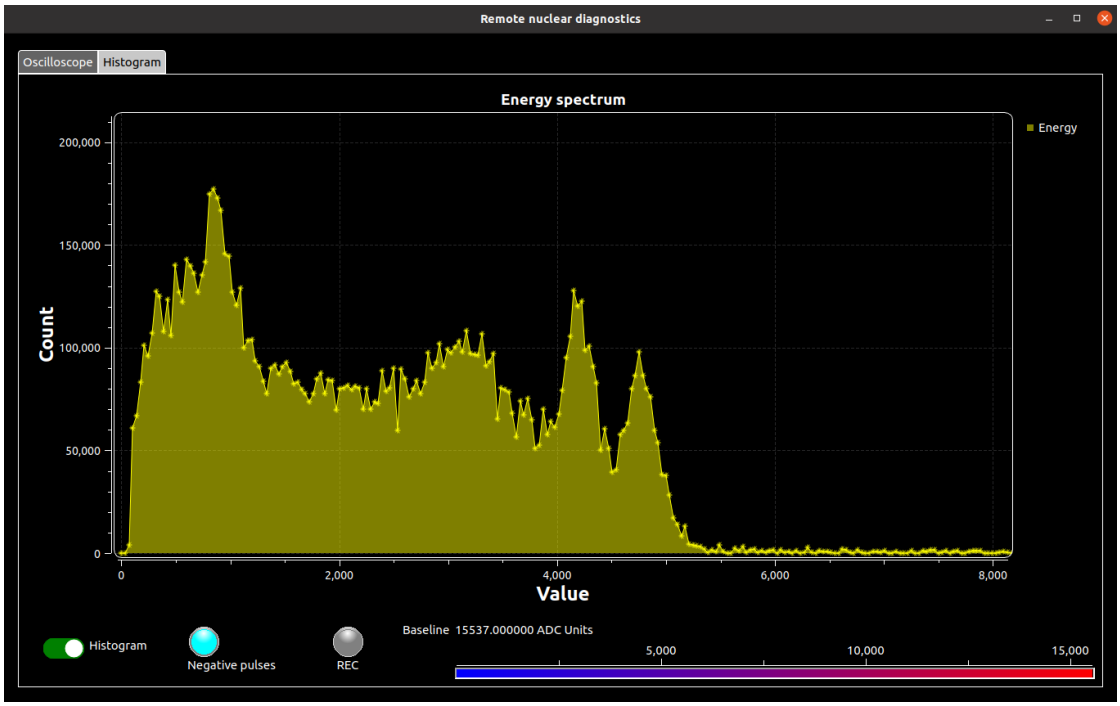

Figure 2.4: MCA tab: the histogram represents the accumulated energy spectrum of measured pulses from the NaI(Tl) crystal with SiPM stimulated with a Co-60 γ source. The spectrum matches the expected detector response.

Besides, a higher-resolution (1024 bins) energy spectrum from a γ/n mixed field measurement performed with the integrated CLYC detector from Section 1.1.3 is shown in Figure 2.5. It features events of moderated neutrons from a deuterium-deuterium generator plus a Cs-137 γ calibration source. The thermal neutron capture reaction peak (3160 keV) [3] is observable close to ADC channel 3000, while the Cs-137 gamma photopeak (661.7 keV) reveals a notable maximum just above 600 ADC units. The backscatter peak (180 keV) is also present and might have been enhanced by the lead shielding surrounding the detector. A prominent lead X-ray fluorescence line (75 keV) between channels 50 and 100 is clearly visible as well. Placing the cursor on the region of interest reveals the precise ADC channel, as shown by the green coordinates close to

the Cs-137 photopeak. Owing to the MCA flexibility, users may decide the trade-off between spectrum build-up time and energy resolution.

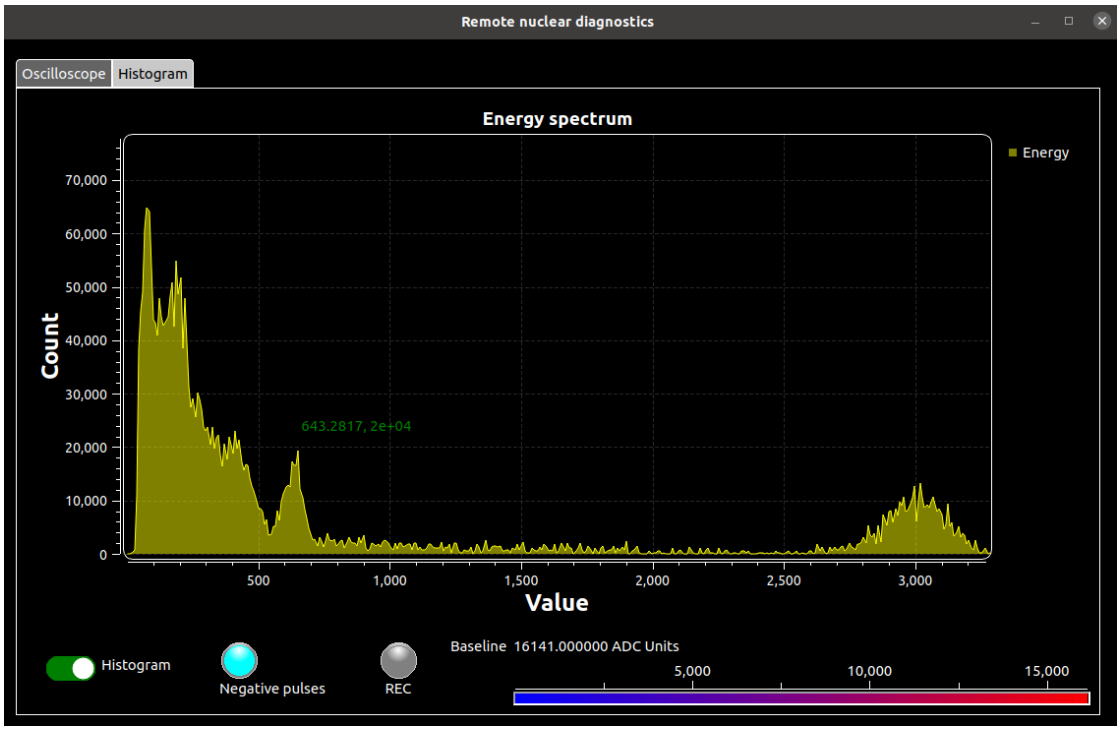

Figure 2.5: MCA tab: the histogram represents the accumulated energy spectrum of events from a CLYC crystal coupled to a SiPM in a mixed-radiation field. Moderated neutrons from a deuterium-deuterium generator and a Cs-137 γ source create distinctive features in the spectrum, matching the expected detector behavior.

2.3.2 Hardware bridge with SBC

The tested HB demonstrated consistent remote data transmission using the embedded WiFi 802.11n peripheral. Configuring the SBC simply required installing the default operating system (Raspberry PI OS, release 20230503) into an 8 GB micro-SD card and setting up the WiFi access-point configuration. Taking the hardware bridge script up and running only consisted in downloading the source code from the provided public repository and executing the Python script.

A set of pre-recorded experimental datasets are provided in the project repository as well. Such data can be used to emulate a remote DAQ plugged into the HB and test the network connectivity using a known detector response. To switch between acquiring data from an external DAQ or emulating a signal output, a simple flag in the Python script is required to be set. These datasets are the result of a sequence of triggered traces from different detectors and radiation sources:

1. NaI(Tl) with SiPM: Co-60 (γ source)
2. NaI(Tl) with SiPM: Cs-137 (γ source)
3. NaI(Tl) with SiPM: Co-60 + Cs-137 (γ sources)
4. NaI(Tl) with SiPM: Background radiation
5. CLYC with SiPM: Cs-137 (γ source)
6. CLYC with SiPM: Cs-137 + moderated DD (γ and thermal neutrons)

2.4 Discussion

An open-source platform for remote DAQ diagnostics has been developed and tested with diverse detectors using common radiation sources. The system comprises an easy setup to debug and control DAQ prototypes with a GUI featuring an oscilloscope and a multichannel analyzer from remote locations. Such remote connectivity is provided by a reliable TCP socket, taking advantage of existent wireless or wired network infrastructures.

The GUI is optimized for multi-core processing to enable a fluid visualization of the data transmitted from the DAQ under test without requiring third-party libraries. The oscilloscope GUI tab allows real-time pulse shape and peak detection visualization, baseline restorer configuration and individual event recording. The multichannel

analyzer tab computes online the deposited energy of the events from the remote DAQ, enabling a versatile tool with configurable resolution for fast diagnostics, aiding in isotope identification and fault detection tasks. An optional hardware bridge is also available to provide connectivity for remote instruments in early development stages that lack of network capabilities.

Flexibility is foreseen, by supporting low-end sub-watt single-board computers as the hardware bridge element, compatible with multiple off-the-shelf power source types. The system is complemented with experimental datasets recorded as raw traces from diverse detectors, allowing to diagnose the remote connection and the hardware bridge with known energy spectra and pulse shapes. A future release of the platform may include energy spectrum calibration, custom command transmission, digital input/output interaction, among others.

The GUI source files, the script for the HB, and the datasets are openly available in the online repository under the MIT license [165].

2.5 Summary

The platform presented in this chapter serves as a drop-in alternative to vendor-locked products for diagnosing custom data acquisition systems in nuclear and HEP experiments, thereby accomplishing research objective I. Experimental embedded system developments for mixed radiation environments might especially benefit from this platform, on which multiple sources may stimulate a single detector. In this regard, the incorporated real-time energy spectrum and the oscilloscope tools might aid in the diagnostics of gain, event discrimination, pulse processing performance, among others. Thus, the system is intended for debugging and controlling custom embedded systems and detector probes by providing an easy visualization and recording tool from remote radiation-controlled areas, such as (but not limited to) the developments introduced in the next chapters:

- Advanced real-time event detection: presented in Chapter 3, focused on research objective II.
- Event discrimination in mixed radiation fields: further presented in Chapter 4, targeting research objective III.
- Event discrimination in high event rate scenarios with pile-up distortion: detailed in Chapter 5, linked to research objective V.

Testing embedded instrumentation platforms for nuclear and HEP applications in development stage can be a rough duty. However, this task can be alleviated with the right tools that safely allow debugging and controlling prototypes located in radiation-controlled areas.

Chapter 3

Fast real-time pulse shape recognition

Detectors in nuclear and HEP experiments produce continuous streams of digitized data that may feature high dynamic ranges or low signal-to-noise ratios, requiring advanced techniques to accurately identify the arrival of relevant events. Thus, a simplified correlation index is presented to be used in real-time pulse shape recognition systems. This index is similar to the classic Pearson's correlation coefficient, but it can be efficiently implemented in SoC/FPGA devices with few logic resources and excellent performance. Numerical simulations with synthetic data and comparisons with the Pearson's correlation show the suitability of the index in applications such as the discrimination and counting of pulses with a predefined shape. Superior performance is evident in signal-to-noise ratio scenarios close to unity. FPGA implementation of Person's method and the optimized correlation index have been successfully tested and the main results are summarized.

Aiming at efficiently detecting individual pulses in continuous streams through pulse shape recognition (research Objective II), the manuscript "A Simplified Correlation Index for Fast Real-Time Pulse Shape Recognition" [2] was published in the *Sensors* journal on 11 October, 2022. In this chapter, an enhanced version of the original publication is reported, extending the validation section of the algorithm implementation in the SoC/FPGA platform presented in Section 1.2.3.

3.1 Introduction

In the context of real-time signal processing and pulse shape discrimination, amplitude cross-level triggering is one of the most commonly used methods for event signaling. Such a technique is widely exploited in fast event detection applications like time-to-digital converters [35], multi-feature discriminators [166], and image processing [167; 168].

Indeed, pulse shape recognition may be carried out with different methodologies and is used in many contexts where a known signal is acquired for further feature extraction [141]. In addition, these methodologies have been categorized according to the metrics or algorithms used, such as pulse shape parameters, template comparison, amplification of pulse shape variations, and statistical models [169]. Specifically, this study relies on cross-correlation between a signal of interest and a known static pattern or template. This pattern represents, to some extent, the intended pulse shape to be recognized, resulting in an optimized matching filter.

It is demonstrated how the traditional cross-level trigger method can be improved by introducing a digital preprocessing correlation stage to the signal under study [170]. Subsequently, a simplified correlation algorithm that targets real-time applications is presented. An in-depth analysis is carried out to exemplify how the signal recognition capabilities are preserved using the simplified correlation method. The main advantage of this algorithm is its reduced computational complexity, which leads to faster execution and lower hardware resource utilization when implemented in real-time event-recognition scenarios. A simulation framework was developed in Python 3.8 and the NumPy library to test the recognition capabilities of a traditional cross-correlation preprocessing algorithm based on Pearson's correlation. Then, a detailed comparison is carried out to also quantify the recognition capabilities of the simplified correlation algorithm. Both correlation methods undergo extensive tests under several peak signal-to-noise ratios (PSNR) and detection threshold levels. Well-known recognition metrics,

such as Precision-Recall (PR) and Critical Success Index (CSI), are used to reliably estimate and summarize the accuracy of the simulations under the testing scenarios.

Relying on a digital processing algorithm enables repeatability of the results while maintaining the advantage of portability among different platforms. Hence, to prove the real-time potential of the proposal and test the reproducibility of both methods, the correlation algorithms were ported to a platform based on a SoC/FPGA ZedBoard development board [122]. High-level synthesis (HLS) based on C++ was chosen as the firmware development tool [171], allowing to deploy complex algebraic operations in the FPGA design. Making use of the math library included in HLS not only permitted the implementation of the nonlinear operations required by correlation algorithms, but also enabled an unbiased comparison between both methods. Tests were executed with synthetic input signals (generated by the simulator), and the outcomes were compared with the expected results. The resource utilization of the SoC/FPGA, latency, and estimated power consumption are summarized, providing valuable information about the two kinds of optimizations that were evaluated for each method.

This chapter is organized as follows. Section 3.2 briefly demonstrates how preprocessing a signal with a correlation pattern can improve the event recognition chances in a continuous stream. In Section 3.3, a simplified correlation index is introduced, based on Pearson's correlation, targeting higher performance for real-time applications. The simulation framework used to compare the optimized index with Pearson's correlation is detailed in Section 3.3.3. The experimental setup to test the algorithms on the SoC/FPGA platform is described in Section 3.3.4. Further, in Section 3.4, the results of running the simulation over multiple noise and threshold levels are given, including the recognition performance comparison between the original correlation index and the simplified one. A demonstration of how the correlation preprocessing provides noise immunity at some extent is also shown using the simulation data. Quantitative evaluations and comparison with the Pearson correlation index regarding resource

utilization and execution performance on the hardware implementation are summarized in Section 3.4.2. Section 3.5 documents the implementation of the Pearson's and simplified correlation indices in hardware, detailing the pseudocode used in high-level synthesis. In Section 3.6, the results are discussed. Finally, in Section 3.7 the chapter is summarized.

3.2 Pulse shape recognition through pattern correlation

The comparison between a traditional cross-level triggering system and a two-stage correlation-based preprocessing algorithm is presented in this section. Similar to what was carried out in [172], pulse shape recognition is performed first, and then a trigger over the detected signal is executed.

Let x be a discrete-time signal equal to the summation of a noiseless sequence w and noise n. From now on, x will be called as stimulus or input signal, where its samples x_i are defined by

$$x_i = w_i + n_i \tag{3.1}$$

The sequence w is composed by individual pattern signals of fixed length. Moreover, if the pattern signal is represented by an analytical model, it can be evaluated at regular intervals to obtain a set of samples, as shown in Figure 3.1 for the case of a double exponential pulse [35].

The individual patterns are randomly placed, in such a way that the time interval t between successive pulses follows an exponential distribution with parameter λ, as shown in Equation (3.2).

$$f(t, \lambda)dt = \lambda e^{-\lambda t}dt, t > 0 \tag{3.2}$$

The expected value of the exponential distribution is denoted by β, which is equivalent to $1/\lambda$. Using this probability distribution, a simulator capable of emulating events found in natural sources that follow a Poisson process [35] was designed. As expected,

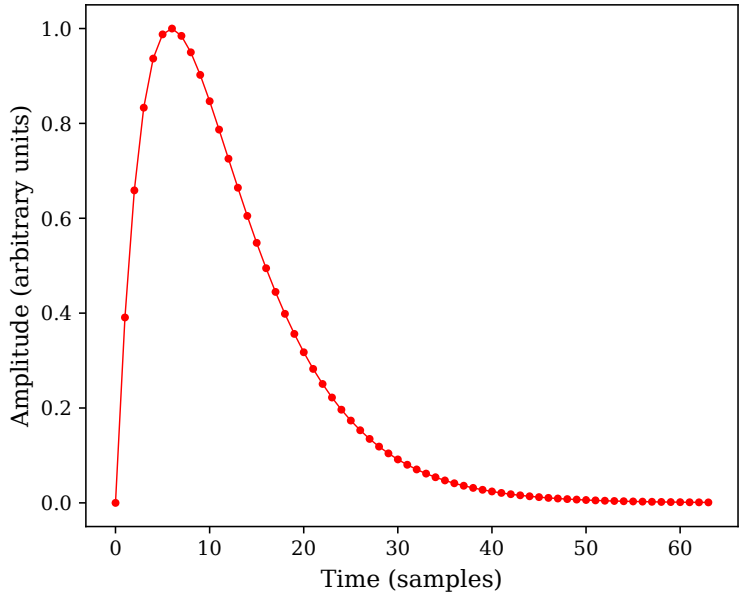

Figure 3.1: Example of a pattern signal used as template, comprised by 64 successive samples.

the pile-up phenomenon is present, increasing its rate at lower β values [173].

An example of a sequence w is shown in Figure 3.2, where a constant amplitude was set for all the pattern signals that generate the noiseless trace.

3.2.1 Simple cross-level trigger

According to Equation (3.1), the stimulus sequence x is the addition of w and noise. If x is shifted into a cross-level trigger (CLT) system, any pair of successive samples x_{i-1} and x_i may indicate the presence of an event of interest. A hypothetical case of this situation is shown in Figure 3.3, where an arbitrary stimulus trace x (based on the same sequence w from Figure 3.2) is evaluated over a fixed threshold level. It is known that an amplitude discrimination system may trigger too many events if the threshold level is set too low; conversely, if the threshold is gradually raised, the number of detected events will decrease accordingly until events are no longer detected. Consequently,

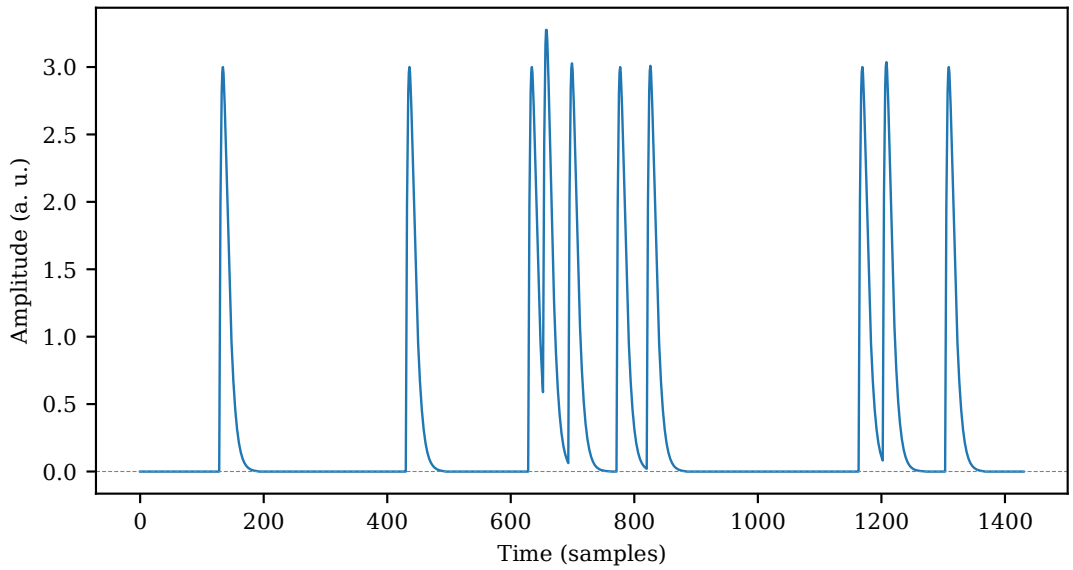

Figure 3.2: Noiseless trace w composed of ten individual patterns. Although individual templates were generated with the same amplitude, pile-up may occasionally cause higher peaks due to superposition of pulses.

setting a constant threshold value to accurately discriminate only the expected events over a signal may become tricky in noisy situations.

3.2.2 Two-stage triggering

A more elaborate method for distinguishing patterns within a signal trace x involves comparing the signal with a static template that reliably represents the target pulse shape to be recognized [174]. If this computation is carried out in the time domain using a correlation index such as Pearson's correlation index (PCI), a normalized measurement of likelihood is obtained for each new discrete sample, independent of the input signal amplitude and offset. The Equation (3.3) defines the PCI (ρ_{xy}) between two segments of N consecutive samples $\mathbf{x}$ and $\mathbf{y}$ of two time-discrete signals (windowed x and y), using z-scores $\mathbf{z_x}$ and $\mathbf{z_y}$, respectively.

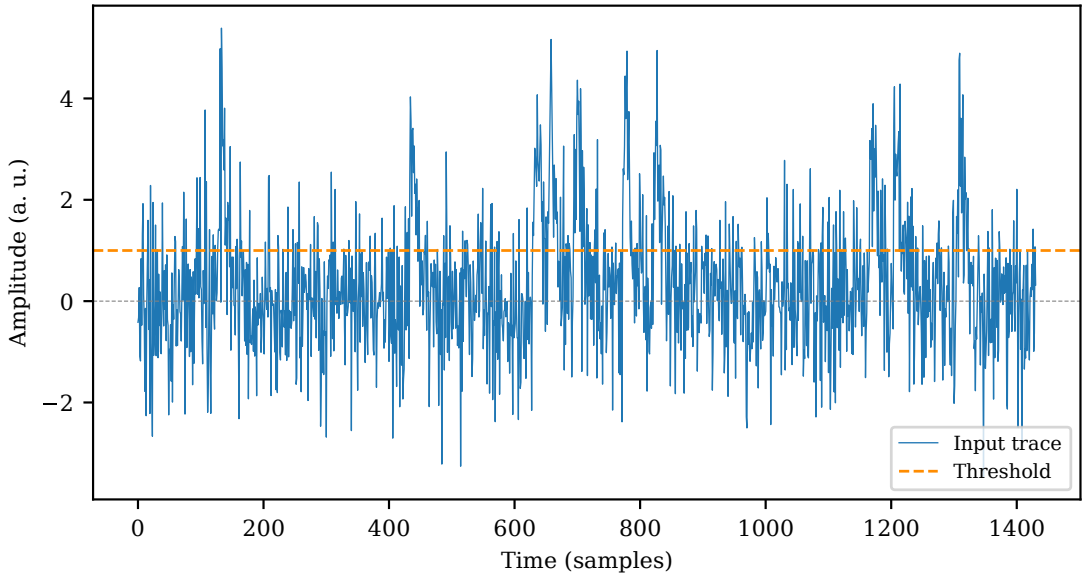

Figure 3.3: Input signal **x** passed through a simple cross-level trigger system. A constant threshold value is set.

$$\rho_{xy} = \mathbf{z_x} \cdot \mathbf{z_y} = \frac{1}{N} \sum_{i=0}^{N-1} \left(\frac{x_i - \bar{x}}{\sigma_x} \right) \left(\frac{y_i - \bar{y}}{\sigma_y} \right) \tag{3.3}$$

By setting a threshold that triggers over the computed PCI, it is more likely to find an event related to an expected pattern within a signal, even in case of lower peak signal-to-noise ratios (PSNR) [126]. Accurately detecting pulses based solely on CLT in noisy signals is less efficient than preprocessing the data using pattern correlation, as demonstrated by Faisal et al [175]. The dynamic detection range can also be improved using correlation, since pulse recognition can be achieved regardless of the peak amplitude. An example of this scenario is shown in Figure 3.4, where the static pattern sequence **c** in Figure 3.1 is substituted into Equation (3.3) to replace the **y** signal. A more detailed explanation is provided in Section 3.3.

As described in Section 3.2.1, the threshold level is set to one third of the expected

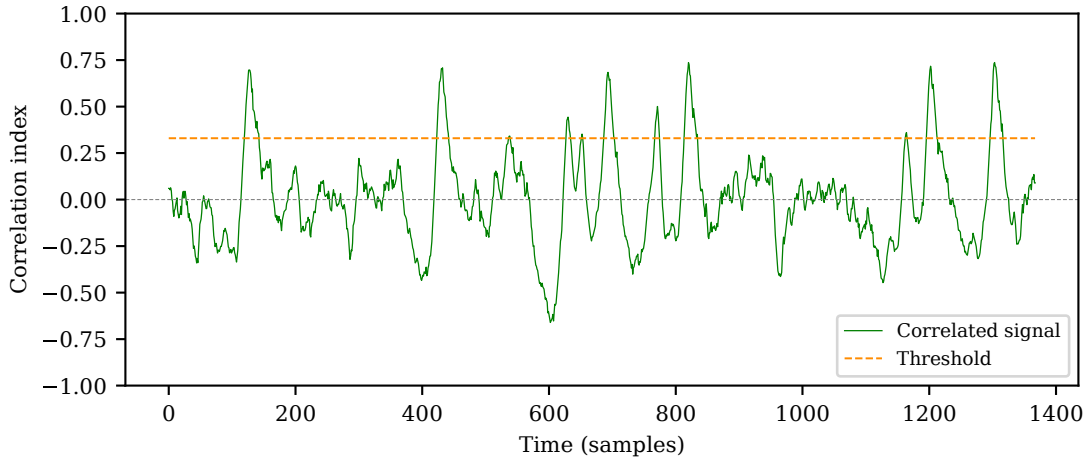

Figure 3.4: Threshold over pre-processed (correlated) signal trace (PCI).

peak amplitude. The correlation index ranges from -1 to $+1$, resulting in $\rho = 0$ if no correlation exists, and $\rho = 1$ if the maximum likeliness between the input signal and the pattern is achieved [176]. The detected events in Figure 3.4 more accurately represent the expected pulses from the original noiseless sequence w, as shown in Figure 3.2.

Pulse-count scenario

In addition to the qualitative analysis shown in Figures 3.3 and 3.4, a numerical simulation was executed (detailed in Section 3.3.3) to quantify the differences between CLT and PCI in a pulse-count scenario, typical in nuclear and HEP instrumentation for event-rate estimation [35]. A trace x with one thousand pulses was simulated starting from a noiseless sequence w. The detected events on x for each method (CLT and PCI) were classified as follows:

- True positive (TP) events: since the original noiseless trace w is known, the simulator is capable of tagging the expected pulse positions and look for triggered events in the current threshold level.

- False negative (FN) events: following the same reasoning than with TPs, but looking for missing expected triggers.
- False positives (FP) events: after seeking the TPs, the triggered events list for the current threshold level is analyzed again, but excluding every TP index. The remaining triggers in the list belong to the unexpected count set. This class corresponds to events that were detected but were not meant to be there.

Counting accuracy was computed using a metric based on precision and recall. Precision, as shown in Equation (3.4), measures the capability of a classifier to discern between the expected (TP) and unexpected (FP) events. Recall (also known as *sensitivity*) represents how well a classification system can detect the absence of an expected event (see Equation (3.5)) by penalizing real events (TP) with missing counts (FN).

$$Precision = \frac{TP}{TP + FP} \tag{3.4}$$

$$Recall = \frac{TP}{TP + FN} \tag{3.5}$$

The precision-recall (PR) curve is a well-studied relationship used to estimate the discrimination performance for imbalanced datasets [177; 178; 179]. The PR curve shown in Figure 3.5 compares the detection capabilities over the full range of threshold values for CLT and PCI. Note the higher precision and recall combination achieved by the two-stage (PCI) method compared with simple cross-level trigger (CLT).

3.3 Simplified correlation index

In this section, an alternative correlation index to the original PCI is introduced, aiming to provide similar event recognition capabilities, but with lower computational complexity for online operation at high data rates. Targeting real-time data processing, the Pearson correlation index is adapted to a continuous data stream over a window in

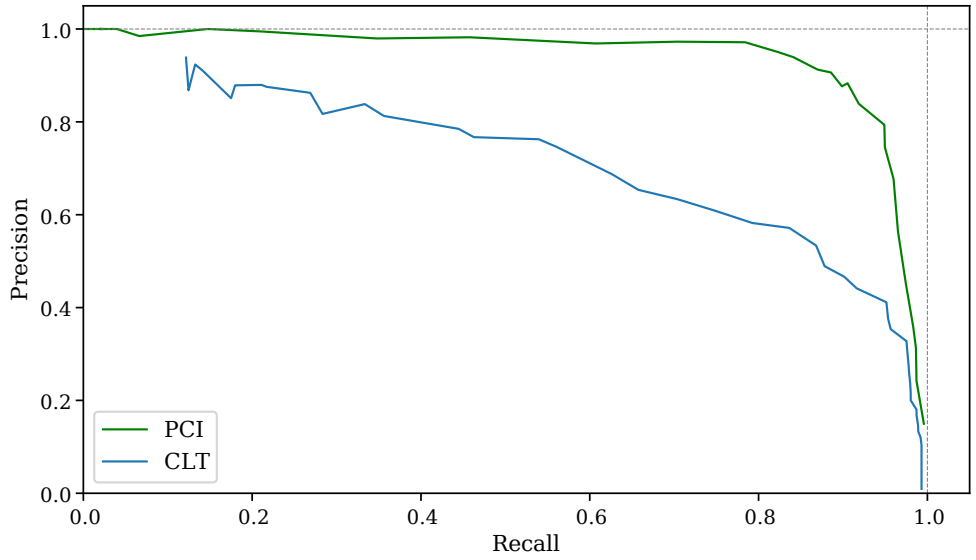

Figure 3.5: Precision-recall curve of event recognition counts for simple cross-level trigger (CLT) and two-stage triggering (PCI). Peak signal-to-noise ratio was set to 3 units for this test.

Section 3.3.1. In Section 3.3.2, a derivation of the simplified correlation index (based on PCI) is provided as an expression that can be implemented in hardware. A simulation developed to measure the recognition performance of the optimized correlation index is described in Section 3.3.3. The capabilities of the method are quantified using metrics that permit a fair comparison with the traditional correlation index in a pulse detection scenario. Moreover, the design of a hardware deployment is detailed in Section 3.3.4, aiming towards an experimental use case in a SoC/FPGA deployment for real-time pulse shape recognition.

3.3.1 Pearson correlation for a fixed-length sliding window

The input data is a continuous stream in typical real-time signal processing applications. Thus, a fixed-length sliding window containing N samples of the input signal $\mathbf{x}$ is fed into the system in a first-in-first-out manner for each discrete period, as conducted by [180]. Using Equation (3.3), the windowed $\mathbf{x}$ is steadily correlated with the pattern $\mathbf{y}$, which shares the same length N.

Note that $\mathbf{y}$ is the representation of the ideal signal used as a reference for correlation, and shall be defined as a constant length vector of coefficients, called pattern or template from now on. These values must be carefully determined either by numerically evaluating an analytical model or by averaging several experimental samples of the signal of interest [53; 175].

Moreover, the z-score vector $\mathbf{z_y}$ from Equation (3.3) is divided by its norm $||\mathbf{z_y}||$, thereby replacing the $1/N$ factor. Consequently, the correlation computation is accelerated, benefiting real-time applications [127]. A re-normalized version of the template vector $\mathbf{c}$ is obtained, as shown in Equation (3.6). This operation ensures that the PCI output range varies only between -1 and +1, independent of the chosen window length and stimulus signal amplitude [181].

$$\mathbf{c} = \frac{\mathbf{z_y}}{\mathbf{z_y} \cdot \mathbf{z_y}} \tag{3.6}$$

Thus, if the PCI of a segment $\mathbf{x}$ is computed against a constant normalized vector $\mathbf{c}$, the expression in (3.3) is reduced to ρ_{xc} (or simply ρ) and can be expressed as:

$$\rho = \sum_{i=0}^{N-1} \left(\frac{x_i - \bar{x}}{\sigma_x} \right) c_i \tag{3.7}$$

3.3.2 Simplified correlation index

An alternative to PCI is documented, aiming to obtain similar correlation results, with the advantage of lower computational complexity. This simplified Pearson-like correlation index (SCPI) uses the absolute mean deviation D as the dispersion metric rather than the standard deviation σ.

$$D_x = \frac{1}{N} \sum_{i=0}^{N-1} |x_i - \bar{x}| \tag{3.8}$$

$$D_y = \frac{1}{N} \sum_{i=0}^{N-1} |y_i - \bar{y}| \tag{3.9}$$

By removing the square root and squaring steps within the summation (required by the standard deviation), the computational resources are reduced [182], providing a substantial advantage in a real-time SoC/FPGA implementation. Consequently, an alternative version of the standard score $\mathbf{z'_y}$ of $\mathbf{y}$ is determined using Equation (3.9), as follows.

$$\mathbf{z'_y} = \left\{ \frac{y_i - \bar{y}}{D_y} \right\} \tag{3.10}$$

Similarly, a new vector of coefficients $\mathbf{c'}$ is determined based on the absolute mean deviation. The same normalization approach from Equation (3.6) is used to obtain this new pattern for SPCI:

$$\mathbf{c'} = \frac{\mathbf{z_{y'}}}{\mathbf{z_{y'}} \cdot \mathbf{z_{y'}}} \tag{3.11}$$

The SPCI is represented by $\rho_{xc'}$ (or simply ρ') and can be obtained using the Equations (3.8) and (3.11), leading to the following expression:

$$\rho' = \sum_{i=0}^{N-1} \left(\frac{x_i - \bar{x}}{D_x} \right) c'_i \tag{3.12}$$

3.3.3 Simulation

To test the SPCI feasibility compared to the PCI recognition performance, a simulation software was developed using the NumPy numeric library version 1.20.1 (NumPy Developers, https://numpy.org/) on Python version 3.8 (Python Software Foundation, https://www.python.org/). Parallel processing was achieved using the built-in multiprocessing library to significantly reduce the simulation time. The software emulates a continuous data stream (stimulus signal x) passing through a sliding window of a fixed length N. Thereafter, the capabilities and differences in detecting an event or pattern are quantified in diverse scenarios:

- Using a simple cross-level trigger (set to a static threshold value) over the original data stream x.
- Computing in a continuous fashion ρ between $\mathbf{x}$ (the windowed portion of x) and the pattern $\mathbf{c}$, and subsequently, triggering over the obtained PCI trace.
- Continuously computing sample-by-sample ρ' to obtain the SPCI and triggering in the same way than with the original PCI algorithm.

Some configurable settings were chosen as global parameters to provide flexibility in the simulation. The noise was also synthesized in the code to verify the behavior of the system at diverse PSNR values.

The stimulus signal x is created by appending multiple patterns $\mathbf{c}$, which are separated from each other by a random number of null samples. Subsequently, x is scaled in amplitude to match the required PSNR value for each simulation run. The sequence is complemented by additive white Gaussian noise, which is defined by an unbiased random Gaussian distribution with a unitary standard deviation.

The peak signal-to-noise ratio (PSNR) is defined as the noiseless peak amplitude of the stimulus signal ($s = \max\{w_i\}$) divided by the noise standard deviation σ_n as follows:

$$PSNR = \frac{s}{\sigma_n} \tag{3.13}$$

Since the added Gaussian noise is set to a unitary standard deviation $\sigma_n = 1$, the Equation (3.13) can be reduced to

$$PSNR = s \tag{3.14}$$

Simulation validation

Tests were conducted to ensure the functionality of the simulation, including:

- Self-correlation of a pattern signal with itself (validation of perfect correlation);

- Correlation of a static-length stimulus signal with a pattern;
- Correlation of a streaming signal in a sliding window with the pattern.

The last item emulates a continuous stream of data x, which is the starting point of the remaining tests for the pulse-count simulation.

Simulation parameters

The settings described in this subsection determine how the simulation is executed in terms of the pulse model template $\mathbf{c}$ and how $\mathbf{c}$ will be replicated to synthesize the continuous streams x. Multiple simulation runs are required to estimate the performance of the correlation indices in diverse noise and sensitivity scenarios. Thus, a new stimulus signal x is synthesized for every PSNR value, implying sequential executions to sweep through the expected ranges. The amplitudes of the pattern $\mathbf{c}$, stimulus signal x as well as noise standard deviation are all expressed using the same arbitrary units. The numerical precision of the simulation results can be controlled by regulating the granularity of the parameters, and so is the time required to execute all the runs. The simulation can be configured by means of the variables described next, set to the values presented in Section 3.4:

- Pattern type: double-exponential pulse model [35], triangular, rectangular, and Kronecker delta.
- Pattern length N: defines the number of discrete samples of the template.
- Asymmetry factor p: affects the asymmetry of the pattern pulse shape, related to leading and trailing edges.
- Number of pulses per trace k: sets how many times the pattern c is replicated to synthesize the stimulus signal x.

- PSNR range: each simulation run comprises the performance grading of both correlation algorithms (ρ and ρ') over diverse noise levels. The PSNR range sets the lower and upper PSNR limits for which the stimulus signal x is synthesized on each run.

- PSNR step: the step size sets the granularity of the expected results. The smaller the step is set, the larger the number of simulation runs are executed. Multiple stimulus signals x are synthesized and evaluated with diverse PSNR values within the imposed range.

- Threshold level range: the algorithms' performance evaluation depends on how well they detect real events, and their ability to reject spurious ones. Thus, multiple runs are executed to sweep over diverse threshold values at each PSNR step. A cross-level trigger algorithm is run over each correlated output (ρ and ρ') as a means of two-stage discrimination. Since the trigger is meant to be executed over a correlated index, real values between 0 and 1 are expected.

- Threshold level step: similarly to the PSNR step, the threshold level step sets the granularity of the threshold level sweep within the corresponding range.

- The exponential parameter β sets the mean interval time between successive templates to build the stimulus signal x. The larger this constant, the lower the probability of pulse overlapping (pile-up) [35]. This constant is expressed in units of pattern length N. As a special case, if $\beta = 0$ pulse overlapping never occurs.

Amplitude discrimination using threshold level

As a means of discrimination, the simulator first synthesizes a random stimulus trace x, based on the given parameters and initial $PSNR$ value. Then, the correlation over the sliding window is computed using each of the algorithms (PCI and SPCI) over the sequence $\mathbf{x}$, as explained in Section 3.2.2. Once the correlated traces are obtained (ρ

and ρ'), a threshold sweep test is performed with the parameters detailed in Section 3.4. A flow diagram representing the simulation steps is detailed in Figure 3.6.

Detection performance estimation

As mentioned in Sections 3.2.2 and 3.3.3, CSI and PR curves are used to assess the recognition performance of the correlation indices. The simulation has been prepared to detect the known pulse patterns within a continuous data trace $\mathbf{x}$. Each TP count corresponds only to the first sample that exceeds the threshold within a predefined window of time (the detection dead time). However, the so-called true negative (TN) events correspond to the absence of such pulses. Thus, the low proportion of TP relative to the TN count leads to an imbalanced distribution of classes. For instance, in Figure 3.2, only ten pulses are expected to trigger a TP outcome; however, the trace contains thousands of samples with "absence of events" (TN).

The CSI provides a measurement of the detection accuracy for events that matter in the triggered system. That is, CSI is concerned only with the expected pulses (TP) and how well the missing pulses (FN) and false alarms (FP) are rejected, whereas the absence of events (TN) is not important [183]. Moreover, the CSI detection performance does not change as a function of event frequency [184], making it suitable for diverse count-rate scenarios. Although CSI is widely used to forecast weather events, it has been applied in other disciplines, when discrimination of rare events is required [185].

The Equation (3.15) shows how the TPs are penalized by missed events and unexpected triggers. CSI ranges from 0 to 1, where the unit value is the perfect event detection metric.

$$\mathrm{CSI} = \frac{TP}{TP + FN + FP} \tag{3.15}$$

Moreover, the area under the curve (AUC) of the PR serves as a normalized indicator, capable of quantifying the recognition performance through all the threshold levels [186; 187; 188].

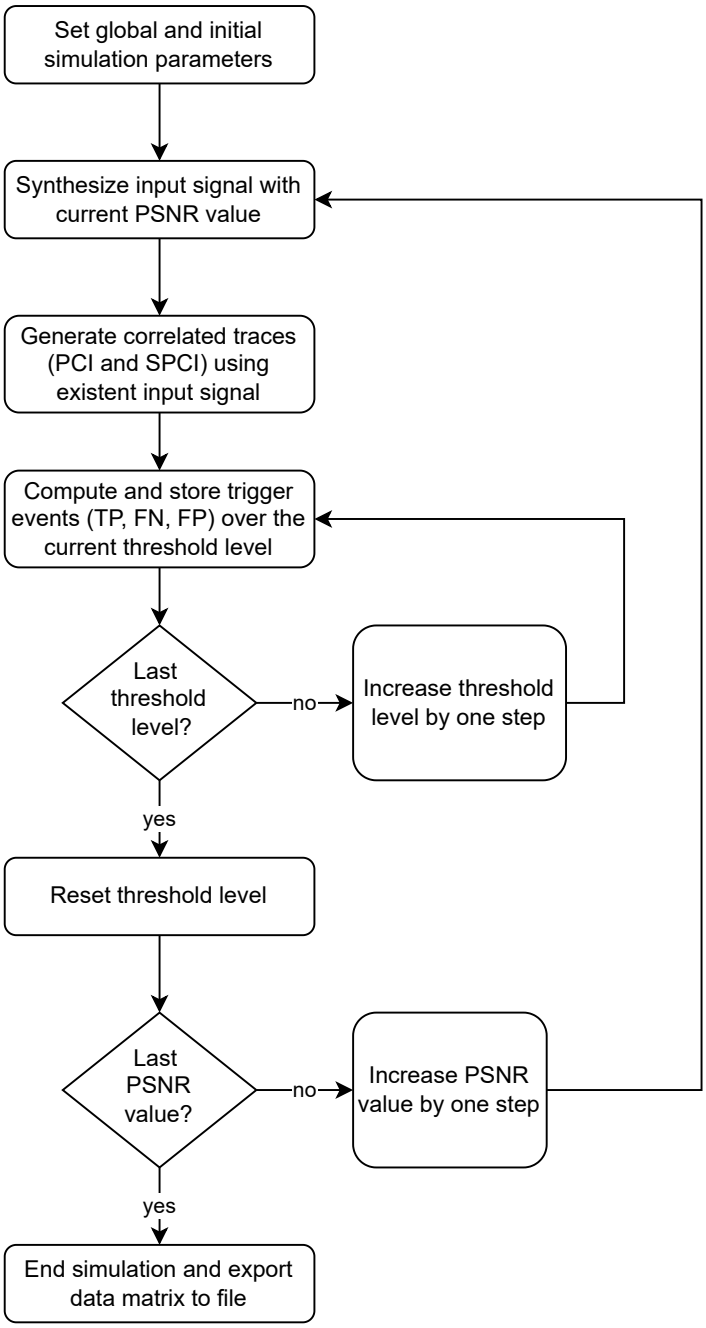

Figure 3.6: Summarized simulation flow diagram.

The CSI and the PR curve (explained in Section 3.2.2) use the same input parameters to estimate the system performance; however, their application to demonstrate the recognition capabilities are used in different contexts. The CSI is used to show the existence of an optimal threshold setting using both correlation methods (PCI and SPCI) and the improvement of signal-to-noise ratio in low PSNR scenarios. Meanwhile, the PR area-under-curve (PR-AUC), computed using the trapezoidal rule [189], summarizes in a single plot the pattern recognition performance of both correlation methods under every simulated condition.

3.3.4 Hardware implementation

In order to verify the capabilities of the simplified correlation index (SPCI) tested in the simulations, a comparison with the classic PCI was deployed in a real-time processing environment using a Xilinx Zynq-7000 SoC Zedboard development board, introduced in Section 1.2.3. Two individual processing blocks (IP cores or simply IPs) were designed using C++ high-level synthesis (HLS), capable of executing each of the algorithms within the programmable logic (PL) section of the SoC. Both HLS blocks share most of the source code, except for the arithmetic expressions that differentiate the algorithms from each other. The standard deviation (SD) was used for PCI computation, as in Equation (3.7), whereas SPCI featured the mean average deviation (MAD) from Equation (3.12). The implemented arithmetic operations were written using Xilinx's integrated Vitis HLS math library to code the SD and MAD, which also permitted an unbiased comparison between correlation algorithms.

Two versions of the IPs were tested for each correlation algorithm by enabling or disabling one directive in the HLS code. The first version was left without any optimization directive, resulting in implementations that required few hardware resources, which will be referred to as area-optimized. The other version featured a *pipeline* directive inside the main processing loop, explicitly forcing the compiler to optimize that section of code for performance (throughput increase).

HLS tools have automatic rules when a specific directive is applied [171]. In a nested loop, if the outer loop is pipelined, the inner loop is unrolled if static bounds are defined [190]. If the top-level function is pipelined, all loops inside the functions are unrolled. In the performance version of the IP cores, the outer loop that buffers the data has the directive *PIPELINE* applied so that each operation can run in parallel on different input data. Moreover, due to the presence of a nested loop in the source code, the insertion of this directive in the outer loop leads to an automatic unroll of the inner loops [191].

A common architecture design was devised to serve as a shared validation platform for each processing unit under test (the IP cores), allowing them to be easily swapped without affecting the test parameters. Such implementation included an instantiation of the processing system (PS7) based on an Arm Cortex-A9 dual-core processor embedded in the SoC, as well as a configurable interface block (ComBlock) to manage the communication between the custom design and PS7 [192], as conducted in [193]. Moreover, to achieve the maximum throughput allowed by the PL, an online data exchange is carried out using the AXI4-Stream protocol, which supports single-cycle transmission between the IP under test and the testbed. The behavior and synchronization of the custom design elements in the PL are also controlled via the ComBlock registers from the PS7.

A dataset was generated by the simulator explained in Section 3.3.3, used to stimulate the correlation blocks under test. A pattern pulse $\mathbf{c}$, a synthetic noisy trace input signal $\mathbf{x}$, and a pair of PCI (ρ) and SPCI (ρ') outputs were generated to validate the IP cores in a real hardware implementation. Prior to feeding the synthetic trace $\mathbf{x}$, type conversion was performed to match the IP's fixed-point representation. Then, using the integrated logic analyzer (ILA) from the Xilinx Vivado Tool, the input and output of the actual IP cores were captured in real time and exported to a text file for offline verification.

Each IP core was tested on the aforementioned platform using a Xilinx Vivado 2020.2 block-level design environment as the development tool. A sketch of the system design is depicted in Figure 3.7, which shows the exchangeable IPs as *Correlation HLS Block**.

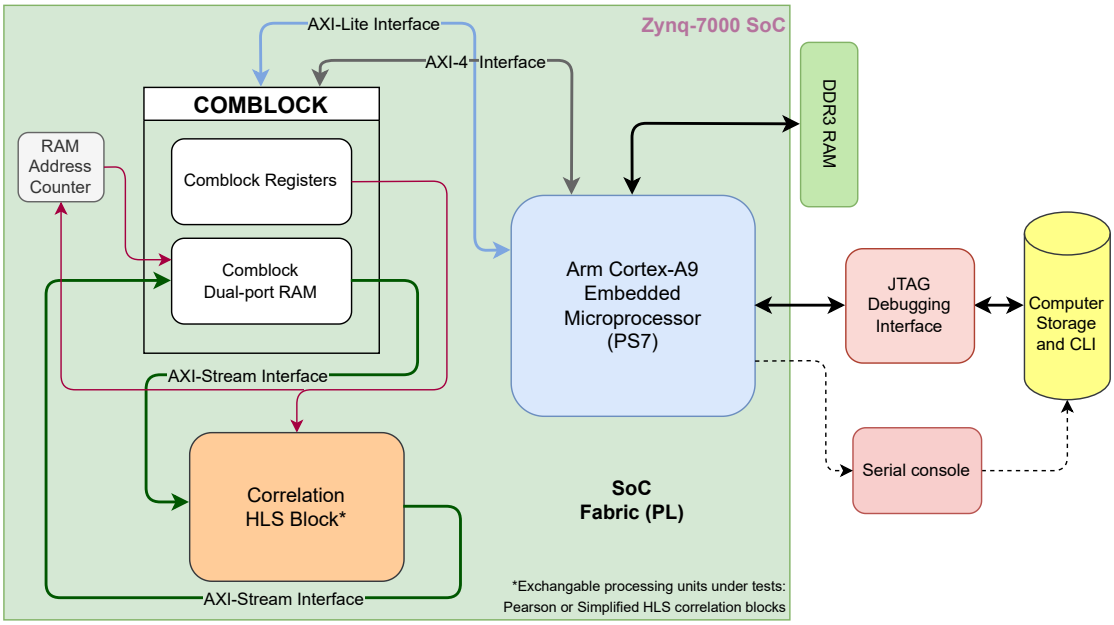

Figure 3.7: Hardware design diagram featuring IP processing blocks. The element named *Correlation HLS block** represents the algorithms under test (PCI or SPCI), implemented as exchangeable IP cores developed using HLS. Each correlation index IP core was individually tested under the same conditions.

Constraints were set to allow IP cores implementation along with the surrounding stimulus and control blocks. This action was necessary due to the limitation of the available resources in the PL of the target SoC/FPGA device. This approach avoids a significant accuracy loss, as in [194]:

- An important compression ratio was achieved by quantizing the data to a 14-bit fixed-point representation (as done by [195; 196]), rather than using the double-precision floating-point numerical resolution of the original Python simulation. Such optimization methods have been proven to reduce the required hardware

resources in PSD and machine learning applications without significantly affecting the accuracy [7; 53; 17]. The 14-bit representation was chosen to match the ADC resolution of the hardware board from Section 1.2.3.

- The stimulus signal **x** was fed into the IPs from a circular buffer in a triggered fashion. This allowed to easily align the processed output data and compare them with the expected (simulated) results.

Accordingly, the pattern signal model and its length were maintained the same as those in the Python simulation. In addition, correlation computations (including the averaging and deviation calculations) are executed on every clock cycle (10 ns), demonstrating the real-time operational capabilities of the processing blocks. Thus, a fair comparison of the performance of the algorithms in a hardware platform was achieved.

3.4 Results

Both, numerical simulations and hardware tests are important for demonstrating the behavior of the correlation algorithms. Therefore, their specific results are divided into Sections 3.4.1 and 3.4.2. To keep the simulation and hardware experiments as homogeneous as possible, the working parameters were set equal in both cases.

- Pattern signal vector size: $N = 64$ samples;
- Pattern signal type: double exponential pulse;
- Asymmetry factor: $p = 0.45$;
- Number of pulses per trace: $k = 1000$;
- Exponential distribution constant: $\beta = 5 \times N$;
- Variable PSNR between 1.0 and 8.0 with 0.25 step size;

- Variable threshold level between 0.1 and 1.75 with 0.025 step size.

The double exponential model is used to synthesize the patterns (**c** and **c'**) and the traces **w**. The model is defined by the following discrete-time equation [35]:

$$f[n] = A(e^{\frac{-n}{(1-p)\tau}} - e^{\frac{-n}{p\tau}}) \tag{3.16}$$

where the constant A represents a scaling factor applied to normalize the model according to the required PSNR value for each simulation run, computed using Equation (3.14). Similarly, the time constant τ in Equation (3.16) depends on the template length ($\tau = N/5$). As the pattern is defined with length $N = 64$, the time constant value is $\tau = 64/5$. The parameter p establishes the pulse rise time and decay time relationship, which is constant in all simulation runs $p = 0.45$. Figure 3.8 shows a sample plot of the synthetic pattern **c**, highlighting the individual coefficients (samples) inherent to the discrete-time definition with small dots.

The stimuli traces **x** are synthesized by appending multiple pattern signals **c** consecutively and adding unitary Gaussian noise. A random separation between each pattern signal is applied, as explained in Section 3.2. An extract of how such traces may look is shown in Figure 3.9, where the pulses were configured to create a trace with PSNR equal to 3, for this particular example. The signal in the plot shows only 10 of the original 1000 pulses to improve visual interpretation of the sequence.

The length of the correlated signal is the difference between the length of the input signal trace x and the length of the pattern N. A representative plot of the single-run results (ρ and ρ') is shown in Figure 3.10a: for both cases, an output range of $[-1, +1]$ is expected [176]. The high qualitative likeliness of both algorithms is evident along the output traces, as shown in detail in Figure 3.10b, where the residuals of PCI and SCPI are plotted. A numerical test suite is detailed in Section 3.4.1 to demonstrate the quantitative similarity between the two methods.

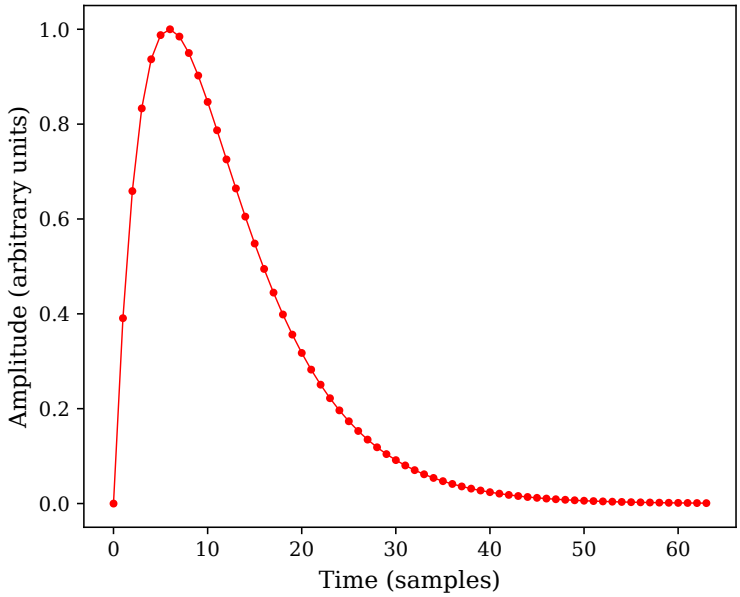

Figure 3.8: Pattern signal (c) with parameters $N = 64$, $\tau = N/5$, and $p = 45/100$, as it was used in the simulation runs.

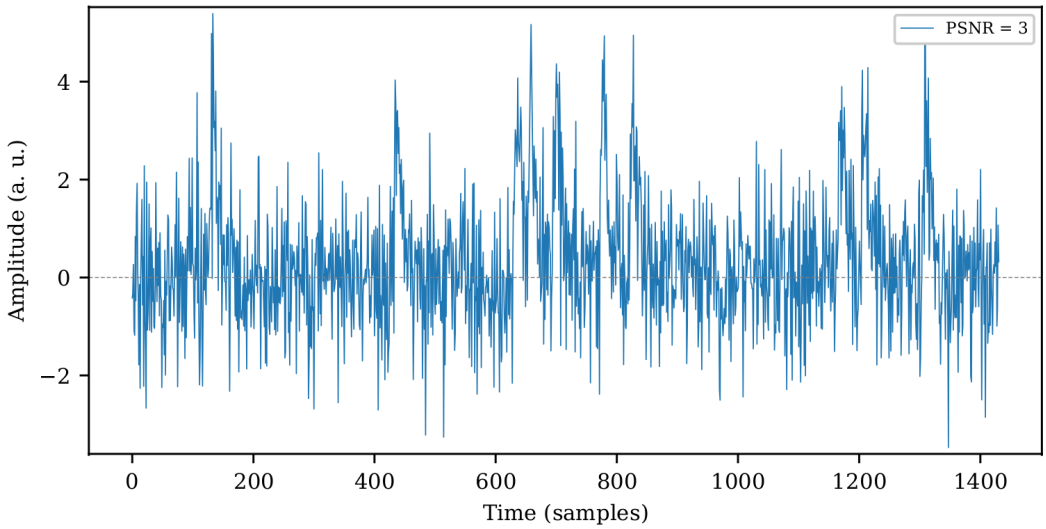

Figure 3.9: Synthetic stimulus signal composed by 10 individual pulses with additive white Gaussian noise emulating a $PSNR$ equal to 3.

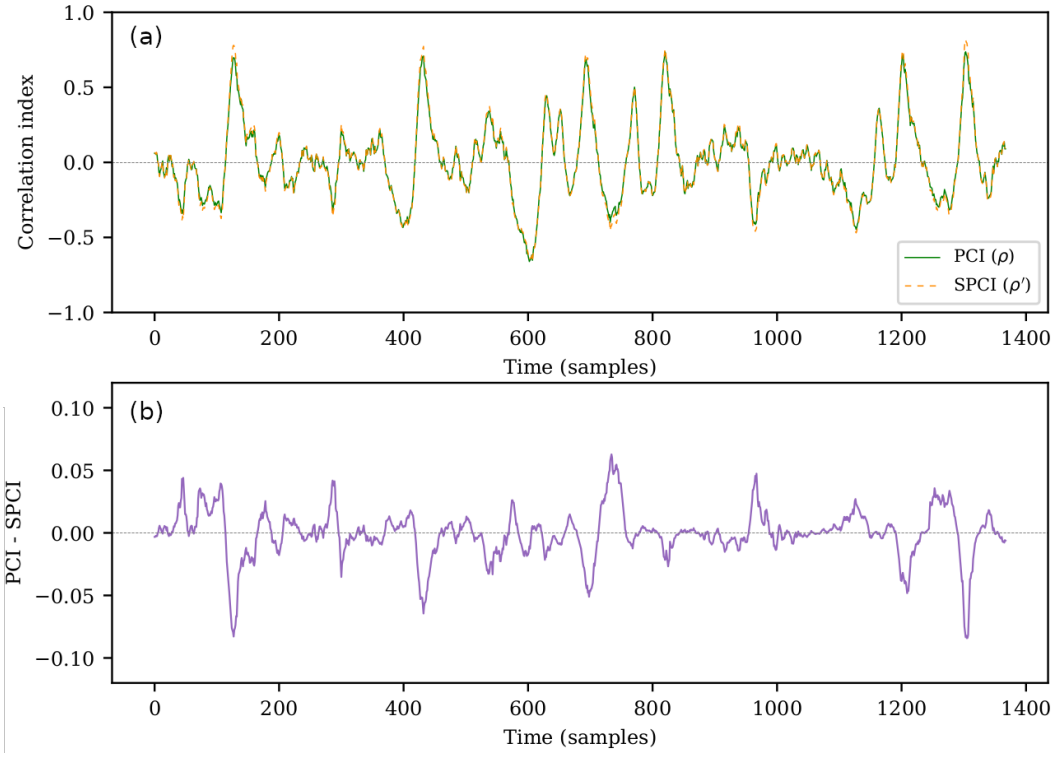

Figure 3.10: (a) Pearson (ρ) and simplified (ρ') correlation results of stimulus (x) with $PSNR = 3$ over a sliding window, based on a double exponential pattern (c) of size $N = 64$. (b) Residuals of correlation indices ρ and ρ'.

3.4.1 Simulation

A set of scripts was built atop of the core simulation tests from Section 3.3.3. This implementation simulates multiple parameter variations and measures the performance of the algorithms. A summary of the results of both indices is then exported to a comma-separated file. Such information is further analyzed to quantitatively compare both correlation indices and main results are explained in the following subsections.

Noise immunity

A family of curves is shown in Figure 3.11 to verify the recognition performance of the algorithms (PCI and SPCI) at diverse thresholds, where the CSI is referenced

to the trigger level. In addition, the PSNR corresponding to each simulation run is shown in the corresponding curve. Naturally, the pulse recognition capabilities in noisy environments (such as PSNR values close to unity) imply low CSI values across all the threshold levels. However, at PSNR $>$ 2, the CSI improves dramatically at threshold levels close to 0.5. This characteristic allows the user to fix a default trigger value regardless of the noise level, thereby providing immunity to unexpected noise variations. Besides, both algorithms perform similarly within the entire threshold range for all PSNR values. Thus, the SPCI behaves on-pair with its more computationally complex counterpart. A depiction of the consistent discrimination performance through multiple noise levels with a fixed threshold level is shown in Figure 3.12.

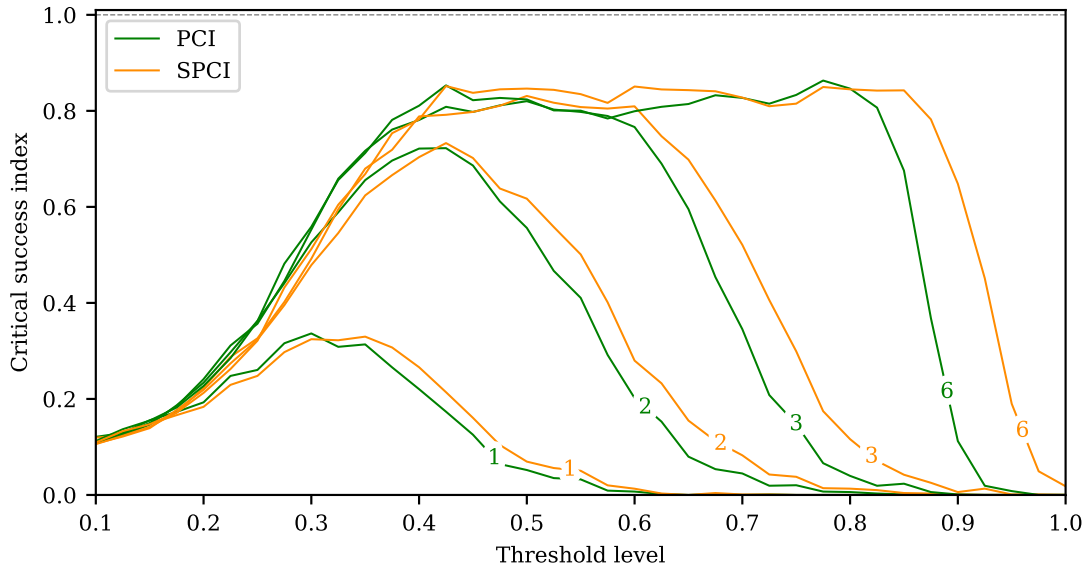

Figure 3.11: Critical Success Index estimation of PCI and SPCI versus threshold level. A family of curves represents the different PSNR values (1, 2, 3, 6) evaluated along the threshold values.

Recognition performance

A PR curve is used to quantitatively assess the performance likeliness of SPCI and PCI, which also evidences the improved recognition capabilities compared to CLT.

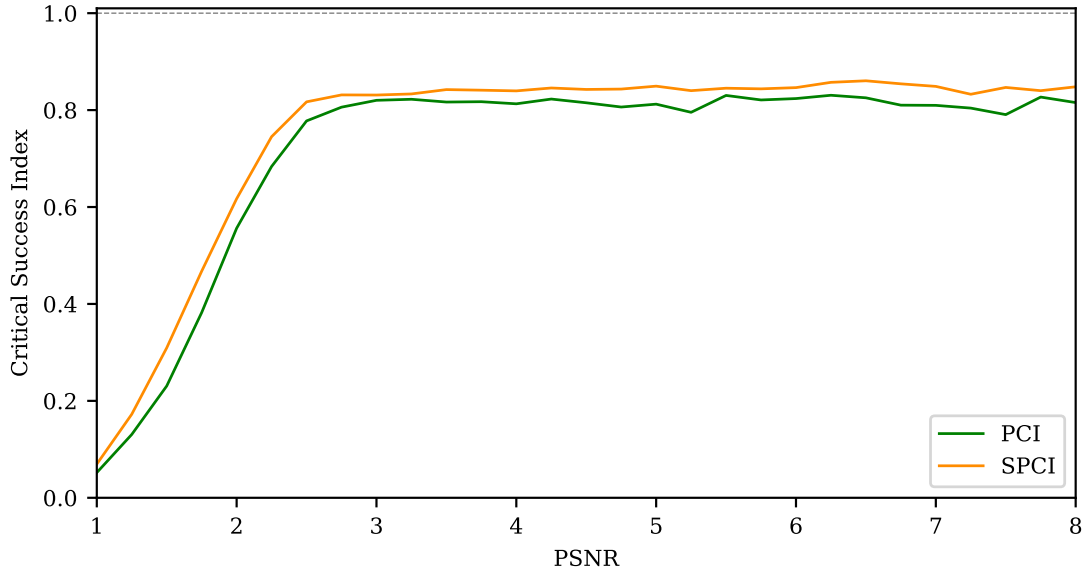

Figure 3.12: Critical Success Index estimation of both correlation indices (PCI and SPCI) versus PSNR. The recognition performance is shown to be similar and close to 80% in all cases, even at PSNR values as low as 3 and pile-up caused by the parameter $\beta = 5$, while remaining practically unchanged up to the maximum evaluated limit. A convenient threshold value (0.5) is set to reinforce the discrimination robustness under diverse noise scenarios.

The PR curve in Figure 3.13 is set at a fixed noise level $PSNR = 3$. Note that even at such PSNR values, both implementations perform quite well [186], similar to the results obtained in [133]. The area under the curve (AUC) is also computed for each PR curve through all PSNR conditions evaluated in the simulation, providing a general indication of the recognition capabilities. This comparison is shown in Figure 3.14, as the AUC for both correlation algorithms and for CLT.

A dissimilarity index (Δ) has been defined to quantify the variation in recognition performance between PCI and SPCI, based on the AUC of PR curves (PR-AUC) from Figure 3.14. The Δ is expressed as the absolute difference of PR-AUC values (y-axis) between the PCI and SPCI curves, relative to their average at every PSNR value (x-axis). By naming the y-axis variables $y_P(x)$ for PCI and $y_S(x)$ for SPCI, the Δ is expressed

as follows:

$$\Delta(x) = \frac{|y_S(x) - y_P(x)|}{\frac{1}{2}[y_S(x) + y_P(x)]} \tag{3.17}$$

The largest Δ obtained by evaluating Equation (3.17) through all the PSNR values was less than 2%. This result shows the high similarity between both correlation indices for recognition performance under all the simulated scenarios.

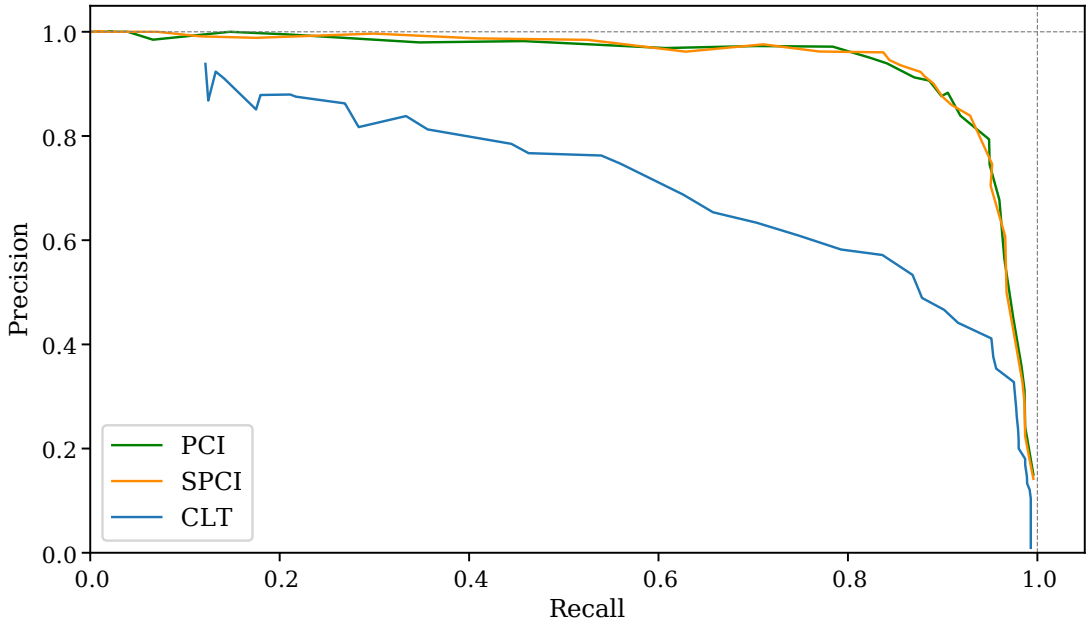

Figure 3.13: PR curves of both correlation indices (PCI and SPCI) and simple cross-level trigger (CLT) at PSNR = 3. The plot axes were set using the threshold ranges specified in Section 3.4.

Simulation execution benchmark

Several simulation trials were run to compare the required execution time for both correlation indices. A workstation equipped with an Intel Xeon E5-268 v2 processor, 64 GB of RAM and Ubuntu 20.04.5 was used for this experiment. Execution time of the SPCI simulation in single-core mode was, on average, 24% faster than PCI after one hundred trials.

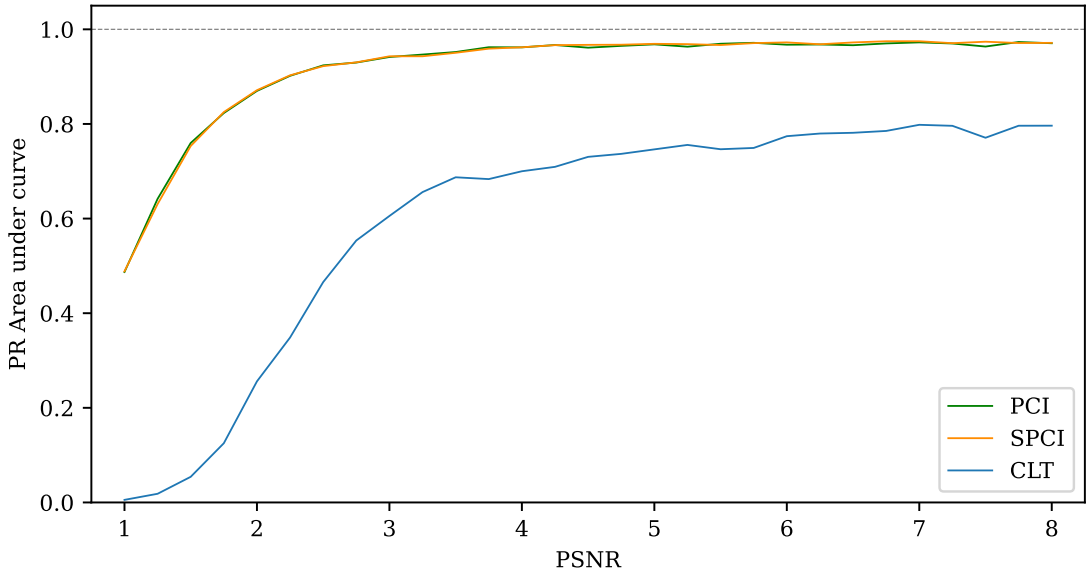

Figure 3.14: Area-under-curve (AUC) for PCI, SPCI and CLT PR curves. The abscissa axis represents the PSNR. Unit AUC value indicates perfect pattern recognition capabilities under the tested conditions.

3.4.2 Hardware implementation

In total, four IP cores were implemented using high-level synthesis to verify the behavior in hardware. For each correlation algorithm, two versions were tested: area optimization and performance optimization, which yielded excellent results for different types of applications. Starting with the area-optimized algorithm comparison, where the HLS and synthesis tools were left with default settings, both correlation blocks occupied similar hardware resources.

The SPCI implementation stood out in the *DSP Block* utilization, saving such resources by a factor greater than 60 times compared to the PCI. Being DSP blocks one of the most limiting resources in several FPGA applications [146; 197], the SPCI has a significant advantage for constrained area applications, or projects that require parallel implementations of the same algorithm. However, some tasks demand fast processing and response [196; 198; 199], constraining the latency to a ceiling value and requiring

higher throughput: which is the case where a performance boost may be better suited. The performance optimization consisted of an explicit pipeline directive in the main processing loop inside the code, causing the HLS compiler to further infer loop unroll parameters inside the pipelined stage. Further details of the HLS implementation can be found in Section 3.5.

Consequently, a significant latency reduction was achieved in both algorithms in performance mode, at the cost of an enormous increase in resource utilization. A throughput improvement was also evident in the performance optimization (according to the post-implementation reports) compared to the area-optimized versions. The *Performance Explore Post-Route* strategy was chosen in the place and route settings, aiming to optimize the available resources in the PL of the SoC. Nevertheless, all the IP core versions are capable to keep constant throughput greater than 100 MHz. Besides, SPCI is expected to use less power, while reaching higher throughput than its counterpart, in addition to consuming less critical resources in the PL, due to the lower algorithmic complexity.

Table 3.1 summarizes the aforementioned results, where the resources utilization of the correlation IP core implementations are detailed with absolute units quoted in parentheses. Power consumption was computed from the estimations reported in the post-place and route tool.

The accuracy of both indices was also measured in the 14-bit hardware implementation, and compared respect to the double-precision floating point simulation results. Figure 3.15a and 3.15b depict the correlation results from the simulation and the hardware implementations for PCI and SPCI, respectively. In Figure 3.15c the sample-by-sample residuals between the simulations and hardware deployments are shown, evidencing remarkable similarity in both cases.

The mean absolute error (MAE) [200] was used to quantify this resemblance, and further normalized to obtain a relative error value represented in proportion

	Area optimization		Performance optimization	
	PCI	**SPCI**	**PCI**	**SPCI**
Resources utilization				
LUT (53,200)	11.21% (5962)	15.15% (8058)	40.13% (21,349)	42.70% (22,718)
Registers (106,400)	4.71% (5016)	4.74% (5040)	20.23% (21,524)	22.06% (23,468)
Block RAM (140)	0.00% (0)	0.00% (0)	0.00% (0)	0.00% (0)
DSP Blocks (220)	29.55% (65)	0.45% (1)	54.55% (120)	24.09% (53)
Timing results				
Max. frequency (MHz)	119.3	122.4	137.8	143.4
Latency (clock cycles)	2.23×10^6	2.23×10^6	1.1×10^3	1.1×10^3
Estimated power consumption @ 100 MHz				
Average power (mW)	190	118	796	705

Table 3.1: Comparison of correlation methods implemented in SoC/FPGA target. Resource utilization and timing characteristics are summarized according to post-implementation reports. The values in this table correspond only to the IP correlation indices. Two types of optimizations are summarized for each correlation IP core: area and performance. The area optimization resulted in fewer resources utilization, whereas the performance optimization provided higher throughput and reduced latency. A great advantage in computational resources is evident for SPCI compared to PCI, particularly considering the reduced DSP blocks utilization and power consumption.

to the envisioned result for each correlation index. This normalized mean absolute error resulted lower than 2% in both cases (PCI and SPCI), proving the high accuracy expected from the hardware implementation, despite the applied quantization.

Although the area-optimized SPCI IP core barely fitted in the FPGA of the custom NSIL hardware board presented in Section 1.2.3, heavy congestion in the implementation did not allow to compare the remaining designs in the same system.

3.5 Pseudocode of HLS IP cores

Two IP cores were developed in HLS to evenly compare PCI and SPCI. The only code section that differs from each other is the correlation standardization. Standard

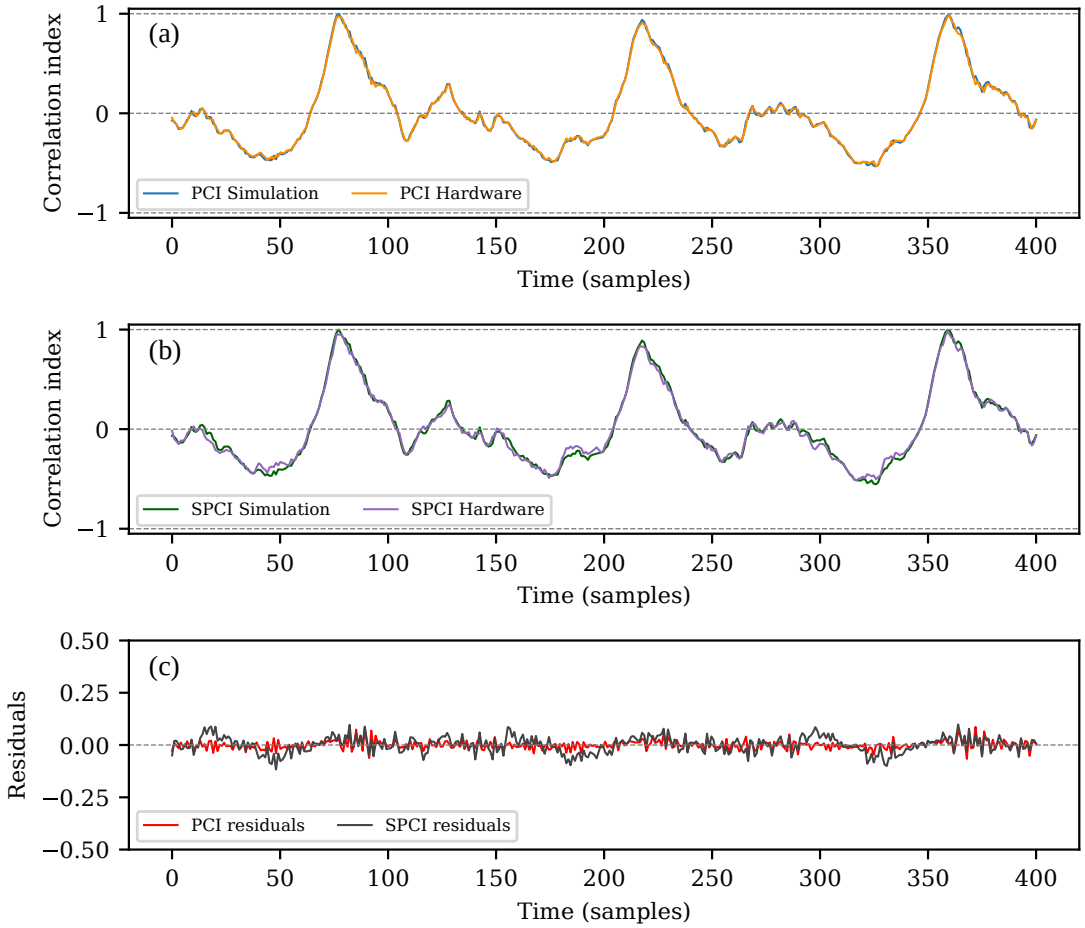

Figure 3.15: (a) Comparison of the original Pearson's correlation index with a common input trace for both the simulation (double-precision floating point) and SoC/FPGA (14-bit fixed point) implementations. (b) Assessment of the simplified correlation index in simulation and hardware deployment. (c) Residuals of simulation and hardware deployments for PCI and SPCI.

deviation (SD) is used for PCI, whereas mean average deviation (MAD) is computed for SPCI. In both cases, the arithmetic operations involving square root (for SD) and absolute value (for MAD) were implemented using the HLS Math library. In Algorithm 1 the pseudocode illustrates the definition of both correlation methods. More specifically, the code in line 19 (*[Deviation code snippet]*) is meant to be replaced by the

standardization pseudocode snippets that distinguish each correlation index, being Algorithm 2 for PCI and Algorithm 3 for SPCI.

Moreover, two versions of the IP cores were tested for each correlation algorithm, without requiring source code modifications: area and performance optimizations. The absence or presence of the *PIPELINE* directive in line 2 from Algorithm 1 was the only change required to define the area or performance variants, accordingly. In the performance version, the presence of the *PIPELINE* directive infers a complete array partitioning for the 64 slots of the data buffer *fifo*. Unrolling of *fifoLoop*, *avgLoop*, and *deviationLoop* loops are inferred completely with a factor of 64 as well. Meanwhile, in the area-optimized version the *fifoLoop* is not unrolled and the other inferred directives remain the same.

Algorithm 1 Pseudocode of HLS IP Cores

Input: $BUFFER\_SIZE = 1024$
Input: $N\_COEFFICIENTS = 64$
Input: $fifo[N\_COEFFICIENTS]$
1: **for** i = 0, i < BUFFER_SIZE **do** ▷ *samplingLoop*
2: #pragma PIPELINE
3: **for** j = N_COEFFICIENTS -1, $j \geq 0$ **do** ▷ *fifoLoop*
4: **if** j = 0 **then**
5: $fifo.read(inStream)$
6: **end if**
7: **for** k = 0, k < N_COEFFICIENTS **do** ▷ *avgLoop*
8: $meanVal \leftarrow fifo[k]$
9: **end for**
10: $meanVal \leftarrow meanVal / N\_COEFFICIENTS$
11: $thisR \leftarrow PATTERN\_COEFF\_LIST[j] * (fifo[j] - meanVal)$
12: **if** j = N_COEFFICIENTS - 1 **then**
13: $correlation \leftarrow thisR$
14: **else**
15: $correlation \leftarrow correlation + thisR$
16: **end if**
17: **if** j = 0 **then**
18: $deviation \leftarrow 0$
19: [Deviation code snippet] ▷ *deviationLoop*
20: $correlation \leftarrow correlation / deviation$
21: $outStream.write(correlation)$
22: **end if**
23: **end for**
24: **end for**

Algorithm 2 Pseudocode snippet of standardization (SD) for Pearson's correlation index (PCI)

1: **for** k = 0 to N_COEFFICIENTS **do** ▷ Replace in *deviationLoop*
2: $deviation \leftarrow deviation + (fifo[k] - meanVal)^2$
3: **end for**
4: $deviation \leftarrow \sqrt{deviation}$

Algorithm 3 Pseudocode snippet of standardization (MAD) for simplified correlation index (SPCI).

```
1: for k = 0 to N_COEFFICIENTS do                ▷ Replace in deviationLoop
2:     deviation ← deviation + abs(fifo[k] − meanVal)
3: end for
```

3.6 Discussion

A simplified version of the Pearson correlation index (SPCI) was presented, targeting real-time pulse recognition systems implemented in nuclear and HEP instrumentation embedded systems. As a starting point, the improvement in discrimination capabilities has been evidenced with the classic Pearson correlation-based (PCI) triggers, compared to cross-level triggering over a raw signal trace (CLT). Once the discerning performance is shown to increase with the Pearson correlation as a preprocessing step, the rest of the chapter demonstrates its similarity with a version of the algorithm based on the simplified correlation index. A simulation framework was implemented to support the hypothesis, and by emulating multiple noise and threshold level scenarios, well-known statistical tools were used to measure the similarity in terms of noise immunity and recognition ability.

Moreover, both correlation index-based algorithms were deployed in hardware (as IP cores) using high-level synthesis (HLS) in the ZedBoard SoC/FPGA target board introduced in Section 1.2.3. Quantization was applied to reduce the hardware complexity and improve the throughput, at the cost of negligible change in recognition performance. In addition, two types of optimization arose for each algorithm: one aimed at low resource utilization, while the other dramatically reduced latency. In total, the four IP cores were individually deployed in a common unit test surrounded by control and test logic, which allowed to stimulate them and obtain the output results using a logic real-time debugger (ILA).

Besides, both SPCI implemented optimizations (area and performance) outper-

formed the PCI counterpart, at very low cost in terms of accuracy. The concordance of the hardware outputs was measured using an input/output validation dataset generated by the aforementioned simulator. Thus, not only was the superior recognition ability of correlation-based trigger systems demonstrated, but the more efficient algorithm resulted in an excellent alternative to the classic methods in terms of hardware resource usage and performance. In addition, one has the flexibility to choose between high-performance or low-power implementations, according to the target design requirements.

3.7 Summary

In this chapter an advanced pulse recognition algorithm for SoC/FPGA platforms was presented. Reliable event detection in nuclear and HEP applications with low signal-to-noise ratio can benefit from this work, owing to its superior recognition performance and optimizations for real-time embedded deployments. These features support the fulfillment of research objective II. Moreover, the presented IP cores targeting high performance or low area utilization demonstrate the possibility to choose an adequate trade-off to meet specific design parameters, such as low-latency constraints or systems demanding multiple detectors, accordingly.

Furthermore, the FPGA firmware framework used in this development serves as the starting point for more complex designs presented in Chapters 4 and 5. Besides, studies demanding a faster method for pulse-shape recognition may also take advantage of the presented method, as detailed in Chapter 5.

Chapter 4

Embedded real-time γ/n discrimination

This chapter presents a method for γ/n discrimination in mixed radiation environments, based on frequency-domain analysis. In contrast to the traditional pulse-shape discrimination, which requires precise pulse alignment, this mechanism does not need any preprocessing of the digitized data, apart from removing saturated traces and sporadic piled-up events. It also features the identification of neutron events spanning through the full energy range using a single device: from thermal neutrons to fast neutrons, including low-energy pulses. Superior discrimination figure-of-merit is achieved in all the evaluated scenarios.

The analysis in the frequency domain consists of computing the fast Fourier transform of a triggered trace and integrating it through a simplified version of the transform magnitude components that distinguish the neutron features from those of the gamma rays. Owing to this simplification, the discrimination method was envisioned for and tested in real-time embedded systems based on SoC/FPGA (research objective III).

The experimental validation is conducted with an off-the-shelf detector based on a low-SWaP CLYC detector described in Chapter 1, aiming at lightweight embedded mixed radiation monitors and dosimeter applications. A dataset of tagged γ/n raw events resulting from this outcome is also documented and publicly released (research objective IV). Both the labeled data and the methodology data might be of interest to

further develop supervised machine learning applications, dealing with challenging scenarios (research objective V), as detailed in Chapter 5.

The potential of the introduced method is further explored for fast organic scintillators, by evaluating a simulation of a commercial probe with an EJ276 plastic scintillator. The classification performance with pulse-shape discrimination is compared down to the lowest energy ranges, underlining the γ contamination rejection in neutron detection applications.

This chapter features two scientific contributions in direct relationship with the research scope of the thesis: the journal paper [3], published in *Nuclear Engineering and Technology* on February, 2024 as well as the public γ/n tagged dataset [4], available in the *Zenodo* platform since 14 June, 2023. The present development resulted from the collaboration between the Multidisciplinary Laboratory (MLab), ICTP and the Nuclear Science and Instrumentation Laboratory (NSIL), IAEA.

4.1 Introduction

In this chapter, a method to exploit the γ/n discrimination features of a small CLYC crystal is detailed, relying on a silicon photomultiplier (SiPM) sensor array with an integrated preamplifier and bias power supply.

Using a simplified computation of the energy spectral density magnitude of each captured pulse, the studied mechanism exhibits better γ/n separation than the traditional charge-comparison method (CCM) for PSD. It also outperforms other systems based on frequency-domain analysis in terms of discrimination performance for both CLYC and liquid (organic) scintillators. These advantages, combined with the ability to use the full CLYC detector energy range, turn this mechanism into a good alternative to PSD. Moreover, the documented frequency-based method has been conceived for efficient implementation in real-time embedded applications based on SoC/FPGA.

A tagged dataset of γ and n events resulted as a sub-product of the data analysis

performed to demonstrate the features of the system. This dataset includes raw pulse traces resampled to 100 MHz, suitable for online SoC/FPGA embedded systems. The calibrated energy in gamma-equivalent units along with a γ/n classification index are also provided for each pulse.

Leveraging the proven discrimination performance of the method, a study of viability for fast organic scintillators is also conducted. The event discerning metrics are evaluated and compared with PSD, evidencing remarkable γ/n discrimination in all the energy ranges. The superior false-event rejection capabilities are noticeable and might be relevant for detecting neutron sources of concern [148; 201].

The remainder of this chapter is organized as follows: the detailed contributions of this chapter are enumerated in Section 4.2. In Section 4.3 the experimental design, data analysis methodology and the γ/n discrimination used as reference are documented. Section 4.4 details the frequency-based methodology for γ/n tagging. In Section 4.5 the results of applying the method to experimental data are documented, including the comparison with the traditional discrimination based on PSD. Quantitative and qualitative performance metrics are discussed as well. The FPGA firmware design used to verify the method in a real-time embedded system is also reported. In section 4.6, a study to validate the frequency-based method for γ/n discrimination in organic scintillators is documented, based on a simulation of a commercial detector probe plugged into an embedded data acquisition platform. Final thoughts are discussed in Section 4.7. Section 4.8 presents a wrap-up of the chapter.

4.2 Contributions

Based on existing developments and methods for γ/n discrimination in the frequency domain, and the current requirements in the field with SiPM detectors, the contributions of this chapter are described as follows:

- An improved discrimination FoM compared to PSD with for SiPM-based CLYC detectors, featuring a low-SWaP setup with superior immunity to strong magnetic fields.
- A method for γ/n discrimination in the full detector energy range: from the thermal neutron region and tested up to 4.5 MeV fast neutrons without time-of-flight (ToF) calibration.
- A γ/n discrimination method with a straightforward deployment capability to a real-time embedded platform using SoC/FPGA, taking advantage of the proven city-block-like simplification for approximate spectral density magnitude computation.
- The creation of an open dataset of tagged γ/n mixed radiation events using an integrated SiPM-based CLYC detector, targeting real-time embedded systems.

4.3 Materials and Methods

The development of the γ/n discrimination system was experimentally validated using the integrated CLYC detector described in Section 1.1.3. The experimental setup and data analysis used for validation are described next.

4.3.1 Experimental setup

The experiment to record the data used in this work was conducted at the Neutron Science Facility (NSF) of the Nuclear Science and Instrumentation Laboratory (NSIL), International Atomic Energy Agency (IAEA), in Seibersdorf, Austria [202].

Detector settings

Taking advantage of the integrated preamplifier in the detector, no external signal conditioning was required at its output, which allowed to directly match the digitizer's 50 Ω input impedance. The bias voltage of the SiPM array was set to 28.7 V, as reported in the manufacturer's calibration sheet. Since shaping and gain parameters are not configurable, the detector is expected to behave as detailed in the calibration report.

Mixed radiation sources

One gamma-only dataset and two mixed neutron and gamma datasets were recorded. The gamma source used is a Caesium-137 (Cs-137) with a registered activity of 0.25 µCi in June 2015. Deuterium-tritium (DT) and deuterium–deuterium (DD) were used as neutron generators.

The DD emits a pulsed beam of monochromatic neutrons of 2.45 MeV with an estimated 4π neutron flux in the range of 10^6 neutrons/s, whereas the DT generates pulses of monochromatic neutrons of 14.1 MeV at an average rate set to 1.12×10^8 neutrons/s. The scattered neutrons in the neighborhood of the generators reduce their kinetic energy, causing a continuum in the neutron spectrum up to the peak energy. Moreover, neutron activation of the surrounding matter and shielding lead to gamma radiation in both cases, generating a mixed-radiation field [35].

The events were digitized from the CLYC detector using a CAEN DT5761 digitizer, sampling at 4 Gsps (giga samples per second) with 10 bits of amplitude resolution. A cross-level trigger was used to capture individual traces containing single events with a threshold set at four standard deviations of the baseline noise level. The CLYC detector described in Section 1.1.3 was used throughout the experiment, subject to a stable temperature of 21 °C.

Data curation

The data were subsampled from the original 4 Gsps sampling rate down to 100 Msps (mega samples per second), aiming at using a reasonable execution frequency for embedded real-time deployments [53; 111; 203; 204].

Three datasets were originally used, each featuring specific conditions, as described below.

- Gamma: 20,000 events recorded from a Caesium-137 (Cs-137) gamma source.
- DD neutrons: 20,000 events recorded in a mixed γ/n radiation field with neutrons up to 2.45 MeV.
- DT neutrons: 20,000 events recorded in a mixed γ/n radiation field with neutrons up to 14.1 MeV.

The data were processed using the Pandas Framework version 1.5.2 and the NumPy numerical library version 1.23.5, running over Python version 3.10. Pile-up (PU) events and saturated traces were further removed as described in [169]. This process yielded three datasets containing the most representative pulse shapes under diverse experimental conditions.

Subsequently, after filtering out the undesired events, a joint dataset was created by adding the recorded traces from the gamma events to the mixed γ/n data (DD and DT). Hereafter, this dataset is referred to as DDTg, or simply the combined dataset, which is the basis of the study. Since the maximum energy range of the detector plugged to the used digitizer is about 4.5 MeVee, only fast neutrons and gamma pulses up to that limit are included. The energy for each event was estimated as the peak amplitude of the trace, referenced to the baseline average over 20 samples. Figure 4.1 shows a subset containing some captured pulse shapes.

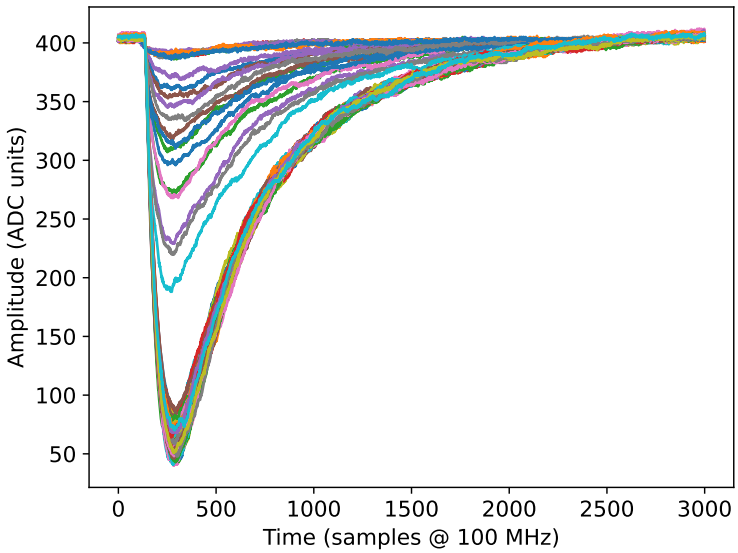

Figure 4.1: Raw pulse traces after PU and saturated event removal, captured with a CLYC detector and resampled to 100 Msps.

Energy calibration

Evidence exists on the linearity of detectors with a CLYC crystal coupled to a SiPM (based on SensL J-60035 arrays) when sensing events of up to 8 MeVee [34; 96]. This permitted to conduct a precise calibration procedure by leveraging only three points (including baseline), as shown in Table 4.1. Figure 4.2 shows the obtained linear regression curve, yielding a determination coefficient $r^2 = 0.999998$ and the equation defined as: $E = (9.024 \pm 0.015)x + (-2.54 \pm 3.29)$ [keVee], where x is the pulse peak amplitude in ADC units. The 3160 keVee neutron capture reaction peak was set according to the manufacturer's calibration report for the detector used in the experiment.

The Cs-137 gamma and ^{6}Li neutron capture calibration peaks (in ADC units) were obtained using Gaussian distribution regressions of the energy histogram. The full-width at half-maximum (FWHM) energy resolutions were 8% and 7%, respectively.

Source	Peak amplitude	Expected energy peak
Baseline	0.0 ADC bins	0.00 keVee
Cs-137 photopeak	74 ADC bins	662 keV
^{6}Li neutron capture	350 ADC bins	3160 keVee

Table 4.1: Energy calibration points

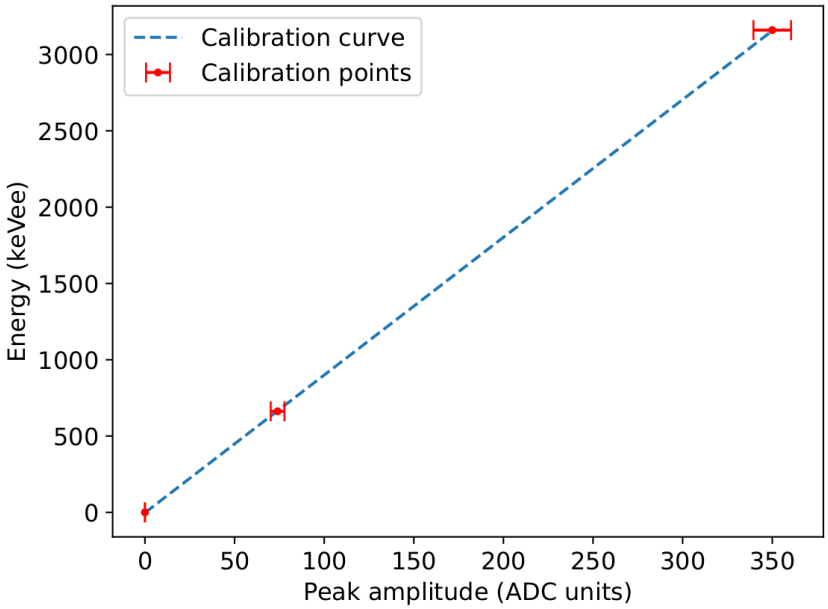

Figure 4.2: Energy calibration curve of the CLYC detector with three reference points: baseline (0 keVee), Cs-137 photopeak (662 keV), and thermal neutron peak (3160 keVee).

4.3.2 Pulse-shape discrimination

A pulse-shape discrimination index was computed with the CCM, similar to [205; 206]. It was used only as a reference for the traditional techniques applied to the detection system. The baseline was removed for each trace and the pulses were inverted on the x-axis to make them positive. Horizontal alignment was also performed using constant fraction discrimination (CFD) to ensure precise charge estimation for the CCM [207].

The intrinsic capacitance of the SiPM acts as an integrator, deforming the original pulse shape of the crystal response. By comparing the pulse decay times of a

standard CLYC crystal coupled to a fast PMT [12] with the detector output in Figure 4.3a (about 5 µs), it is apparent that some shaping is carried out in integrated detectors case [34], presumably by the contribution of the SiPM capacitance and of the built-in preamplifier. To deal with the pulse shape deformation without requiring external hardware or pulse shape pre-processing during discrimination, a single gamma trace and neutron event were selected to further analyze the key points in their discrete derivatives. As usual [87; 95; 208], two discrete integration window ranges were defined to compute the PSD index: short or prompt (W_s) and long or delayed (W_l). The short and long areas of each pulse in these ranges are computed as $A_s = \sum_{W_s} x_n$ and $A_l = \sum_{W_l} x_n$, being x_n each discrete value of the pulse under analysis. The PSD index is computed as a ratio of the difference of the areas, as shown in Equation 4.1:

$$PSD = \frac{A_l - A_s}{A_l} \tag{4.1}$$

The limits of such integration windows, expressed in sampling periods, were chosen according to the derivative plots shown in Figure 4.3b, featuring the inflection points and slope values of both types of pulses. The window ranges were initially set as $[55 - 100]$ for W_s and $[55 - 150]$ for W_l. The highest FoM was achieved using the window ranges $W_s = [50 - 105]$ and $W_l = [50 - 170]$, after iteratively tuning these limits. Details of how the FoM is computed are provided in Subsection 4.5.2.

The derivatives were used as a reference to set the window ranges, with the pulse-shape discrimination analysis directly applied to the aligned data without amplitude normalization, as shown in Figure 4.3a. It is also important to notice in advance that the derivative plots of the normalized events indicate higher abundance of low-frequency components in the neutron pulse shape (smaller average slope), giving a hint to how frequency analysis may be beneficial for discrimination.

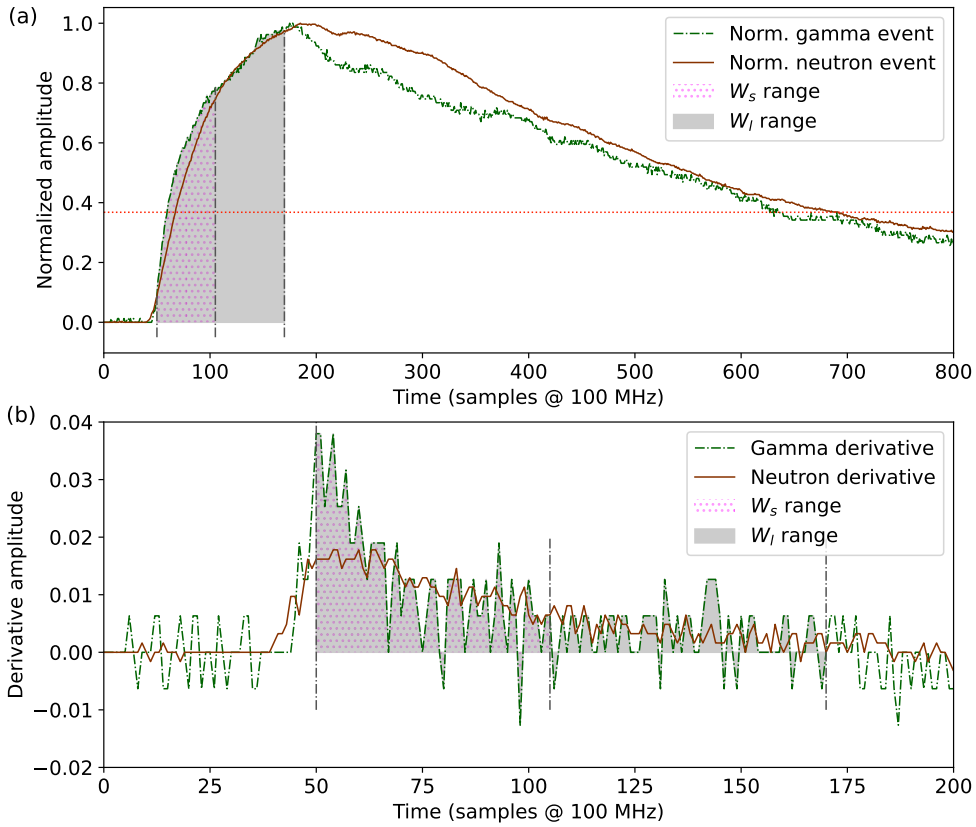

Figure 4.3: (a) Normalized γ and n pulse traces with short (W_s) and long (W_l) integration windows for PSD. The horizontal line indicates the $(1/e)$-decay amplitude. (b) Trace derivatives of γ and n events with short and long integration windows. The integration limits for the PSD are $W_s = [50 - 105]$ and $W_l = [50 - 170]$.

4.4 Fourier-based classification

Contrary to the method based on discrete Fourier transform (DFT) presented by Safari et al. [140], the zero-frequency component is disregarded in this work, leading to a γ/n classification index with substantial immunity to baseline offset variations. In addition, no lower-energy neutron data are discarded, allowing the inclusion of all

the recorded (non-distorted) events regardless of their energy. Moreover, this method does not require any further time-of-flight (ToF) calibration (as carried out in [209]) to achieve γ/n separation, even at the lowest energy range of fast neutrons. A more computationally efficient method is used to obtain the DFT magnitude based on the FFT, as detailed in Subsection 4.4.1. The limitation of using the FFT is the requirement of a power-of-two sample count[210]: in this development the first 2^{11} (2048) samples were used for each triggered event.

Similar to the goals exposed in the work of Dutta et al. [16], a CLYC crystal is used to perform γ/n event discrimination based on the features extracted from the frequency-domain representation, resulting into a binary classification. However, the method presented in this thesis provide the advantages of not requiring pulse alignment and analyzing individual events at a time. Thus, a Frequency Classification Index (FCI) based on the partial spectral area (PSA) of the computed FFT magnitude for each detected pulse is defined. The data were elaborated using the frequency transformation functions in the SciPy library version 1.10.1.

A rectangular window was used prior to the frequency domain transformation. Other shapes, including Hanning, Kaiser, Blackman, and Hamming, provoked severe pulse distortion and hindered the discrimination capabilities of the method.

4.4.1 Partial spectral area

Partial spectral area (PSA) has been defined as the discrete sum of the energy spectral density magnitude (ESD) components of the DFT from a finite-length trace, within a defined frequency range. Let a single pulse trace be described as a discrete signal x_n, which is further transformed into the frequency domain using FFT for the DFT computation. This operation results in a complex expression X_k containing the frequency and phase components of the signal [211], as described in Equation 4.2:

$$X_k = \sum_{n=0}^{N-1} x_n e^{-\frac{2\pi i}{N} kn}, \quad k = 0, \ldots, N-1 \tag{4.2}$$

Differences in the rise and decay times of both event types can be observed in Figure 4.3a, where the neutron events contribute to a higher ESD in the lower frequency components. Consequently, the same differences in the frequency domain plot from Figure 4.4 are expected, where the ESD of the first 2048 samples of the raw pulses (as in Figure 4.1) is shown. It is straightforward to infer that events with higher low-frequency ESD components are more likely to be related to neutron events. It is also worth noting that the differences between the two types of events are apparent within the first 100 discrete frequency components, providing a hint of the integration limits for PSA computation.

Aiming at reducing the computational complexity in an embedded target, the approximate ESD (AESD) is calculated as the sum of the absolute values of the real and imaginary parts of the DFT, also known as the city-block (CB) approach [133]. Instead of using the original definition involving the computation of a square root and two products, the CB alternative provides similar magnitude results with fewer computational requirements [2]. This approximation demonstrated to work exceptionally in Chapter 3 as well. The AESD is used from now on to compute the area of the PSA.

4.4.2 Frequency classification index

In this chapter, the FCI is defined to determine the likelihood of an event being a gamma or neutron based on its PSA features. However, FCI might be applied in other mixed radiation fields for diverse particle discrimination requirements as well. Following a similar approach to those explained in Subsections 1.3.2 and 4.3.2, the FCI is defined as the ratio of the integrated short and long windows of the AESD. In this case, the integration comprises the PSA within two ranges: low-frequency components PSA_l and wide-frequency components PSA_w. The integration ranges

were optimized as described in Subsection 4.3.2 to achieve the maximum FoM, leading to $[1-25]$ for the low-frequency integration window and $[1-90]$ for the wide range in the discrete frequency axis, as highlighted in Figure 4.4. It is worth mentioning that the zero-frequency component is not used to compute the FCI, providing the method with considerable immunity to baseline variations. Equation 4.3 defines the FCI based on the described integration windows:

$$FCI = \frac{PSA_w - PSA_l}{PSA_w} \tag{4.3}$$

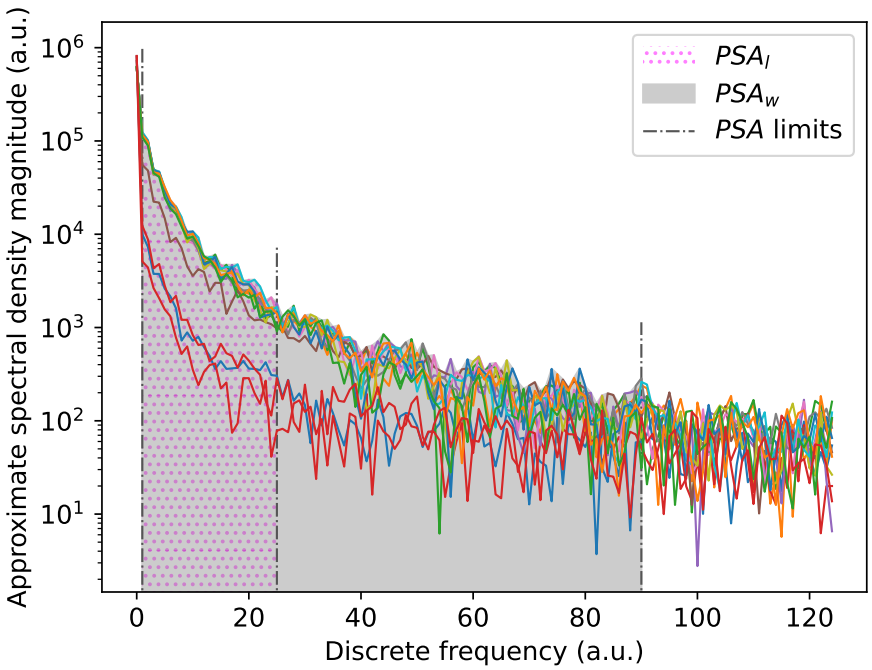

Figure 4.4: Approximate spectral density magnitude (computed using FFT) of a set of pulses from the mixed γ/n dataset. The integration windows used to compute the FCI are delimited in the vertical dotted-dashed lines and marked as PSA_l and PSA_w, respectively. Events with higher AESD in the PSA_l range are more likely to be related to neutrons.

4.5 Results

The FCI γ/n discrimination method has been evaluated using two quantitative criteria:

- The energy range on which the γ/n classification can be achieved
- The discrimination performance compared to other methods based on frequency-domain analysis

The energy range and performance metrics are detailed in Subsections 4.5.1 and 4.5.2, respectively.

In terms of the qualitative results, a two-class dataset with more than 38,000 gamma and neutron events was created and is publicly available at the Zenodo data sharing platform [4]. Each class contains not only the individual pulse traces resampled to 100 Msps, but also the calibrated energy and the FCI label. Moreover, the possibility of making a precise implementation of the method targeting embedded hardware deployments renders this solution an attractive alternative for applications that require real-time mixed-radiation event discrimination.

4.5.1 Neutron energy range

Because of CLYC's ^{35}Cl reaction with fast neutrons, the inherent Q-value and quenching losses create a scintillation energy equivalent to 517 keVee for the lowest-energy n_f [82]. By applying this base neutron energy-equivalent limit to the resolution of the detector (8%) at this range, the detectable neutron bottom is established to be about 476 keVee. It is important to notice this mandatory energy cut mandatory in PSD with CCM to achieve an acceptable separation, leading to FoM = (1.11 ± 0.02). However, the method presented in this chapter, based on frequency analysis, does not require any energy cut, and γ/n classification is achieved automatically by a straight separation limit in the FCI. A detailed comparison of both methods is shown

in Figure 4.5, where the lower neutron energy limit depicts how well FCI distinguishes γ from n events, as expected from the physical behavior of the detector. Neutron interactions were grouped above the FCI classification line, whereas gamma events remained below this limit.

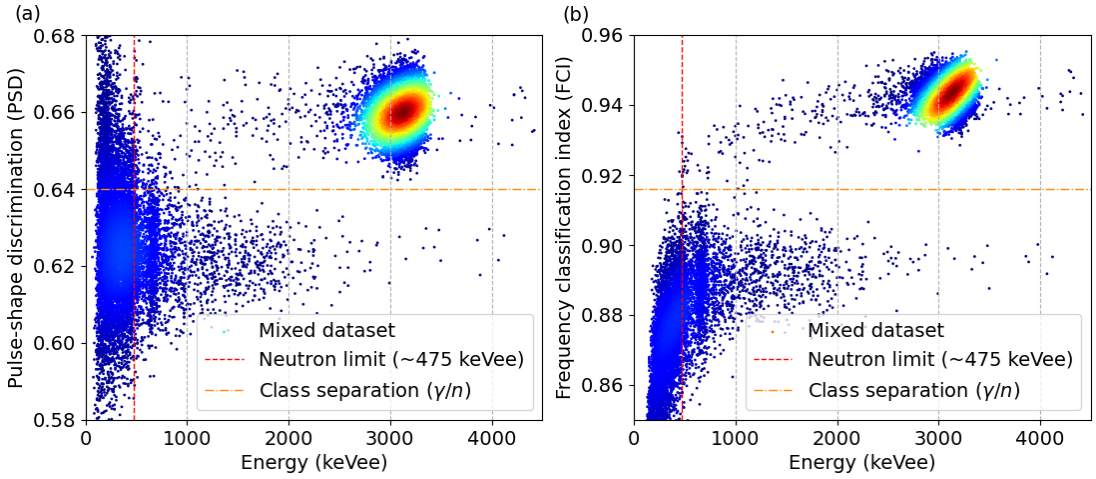

Figure 4.5: Comparison of γ/n separation with PSD (a) and FCI (b). The neutron limit represents the lowest energy-equivalent value at which neutron events can be physically detected using a CLYC-based setup. The classification limit defines the discrimination between γ (lower cluster) and neutron (upper cluster) events. The FCI (b) qualitatively shows superior performance over the entire energy range. More evident is the comparison below the neutron detection energy limit, on which multiple γ events are misclassified using PSD, but an almost perfect classification is obtained with FCI.

As expected, thermal neutrons exhibit a dense cluster close to 3160 keVee because of the ^{6}Li neutron capture reaction in the crystal, detailed in Section 1.1.2. Most of the remaining neutron events are caused by ^{35}Cl neutron scattering, and provide not only counting capabilities, but also valuable spectroscopy information [89].

Figure 4.6 reveals the actual classification power of the FCI method (Figure 4.6b) compared to PSD with CCM (Figure 4.6a), showing a clear view of how none of the γ events in the Cs-137 dataset have been misclassified as neutrons, even at the lowest measured energy.

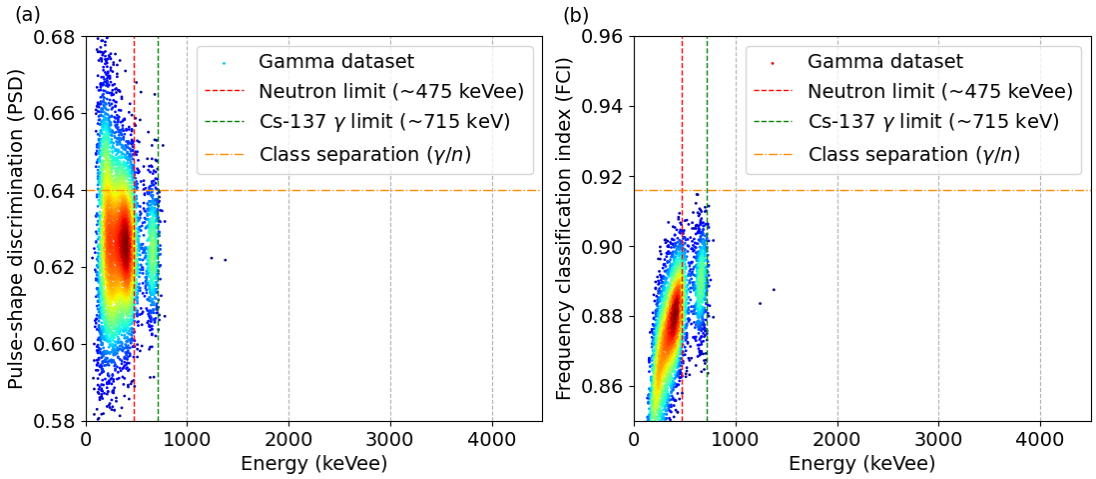

Figure 4.6: The same axes as in Figure 4.5 but with a Cs-137 γ-only dataset. Neutron-related events are meant to be clustered over the classification limit in both plots (a) and (b). The FCI method (b) shows a perfect classification for the events in the entire energy range, whereas the PSD with CCM (a) requires a lower energy cut to successfully separate both types of pulses. The maximum expected energy of the photons is marked as the Cs-137 limit in both plots, based on the known photopeak energy plus the detector FWHM resolution.

4.5.2 Discrimination performance

The same FoM defined in [16; 141; 140; 142] is used as the discrimination performance metric, as shown in Equation 4.4. Variable S represents the distance between the gamma and neutron mean values of the distributions, whereas $FWHM_\gamma$ and $FWHM_n$ are the FWHM of each Gaussian curve for the γ and n events, respectively.

$$FoM = \frac{S}{FWHM_\gamma + FWHM_n} \tag{4.4}$$

To achieve equal comparison conditions between PSD and FCI, an energy cut was applied to the dataset for FoM estimation: events with energies below the neutron limit (approximately 475 keVee) were disregarded. This action does not affect the neutron detection capabilities with FCI though, since no neutron events are present below this energy limit within the gamma cluster, as shown in Figures 4.5b and 4.6b.

The FoM was computed by integrating along the energy axis of both methods (PSD

and FCI) down to the neutron energy limit, resulting in a histogram. The estimation of such a distribution is obtained from the sum of the two Gaussian functions. Figure 4.7 shows the histograms and corresponding regressions performed to compute the FoM for both discrimination results.

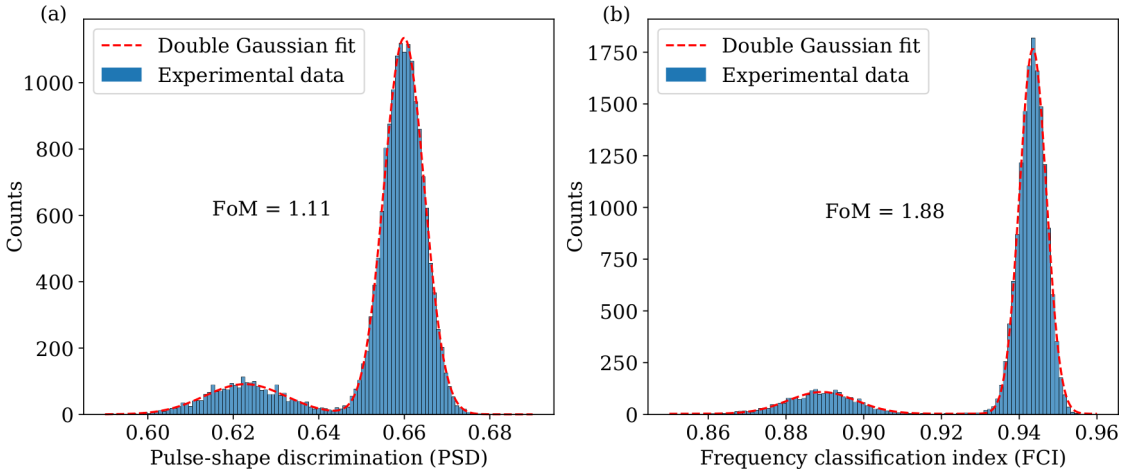

Figure 4.7: FoM comparison of PSD (a) and FCI (b) for the same mixed γ/n dataset

Superior discrimination is possible using the FCI (Figure 4.7b) compared to the PSD with CCM (Figure 4.7a) after evaluating the experimental data discrimination metrics. Moreover, the best FoM is achieved compared to other methods that use frequency domain analyses, accounting for the discrimination capabilities in the complete detector energy range. As a reference with other implementations, a further test was carried out to verify the performance of FCI with an energy cut at 1 MeVee (the limit discussed in [142]), naturally leading to a higher FoM. Table 4.2 summarizes the results of FCI and previous studies related to γ/n discrimination based on frequency analysis.

4.5.3 Hardware deployment

A hardware deployment was conducted to validate the method using the embedded hardware platform from NSIL presented in Section 1.2.3, based on an AMD Artix-7 XC7A35T FPGA. The discrimination algorithm was ported using high-level synthesis

Method	Best FoM	Lowest energy	Detection system
[16]	1.35	1930 ADC bins	CLYC + PMT
[140]	1.48	200 keVee	BC501A + PMT
[141]	1.51	200 keVee	BC501A + PMT
This work	1.88	No lower limit	CLYC + SiPM
This work	2.13	1 MeVee	CLYC + SiPM

Table 4.2: The lowest energy range refers to the lowest neutron energy-equivalent limit on which discrimination from gamma events is possible in each study. ADC bins units are used in [16] due to the absence of energy calibration in the referenced work. The method presented in this chapter was tested under two conditions: without any neutron being disregarded (no lower limit) and with an energy cut at 1 MeVee (the limit discussed by Polack et al.)

(HLS), resulting in an IP-core operating at 100 MHz. This core has been devised with 10-bit input amplitude resolution and a 2048-samples FFT window length, matching the requirements presented in Sections 4.3.1 and 4.4.

Aiming at real-time operation, the FPGA firmware has been designed in a pipelined streaming scheme based on the AXI4-Stream protocol, by instantiating the optimized AMD (Xilinx) FFT LogiCore IP [212], and further computing the integration windows to obtain the FCI. The highest resource utilization in the FPGA is in the DSP blocks with a total of 37%, while the LUT and BRAM usage remain below 17% and 12%, respectively. Both, the timing closure of the design (as reported by the development tool) and the resource utilization evidence the viability of implementing an embedded hardware solution to execute this γ/n classification method using a commercial SoC/FPGA device.

Figure 4.8 shows a block diagram of the validation framework in the Artix-7 FPGA, featuring the γ/n classification IP core based on FCI. The Communications Block (ComBlock) is used to enable seamless user interaction through a serial interface [157]. Curated traces resulting from the data processing in Section 4.3.1 are streamed from a computer to the FCI-based IP core using AXI4-Stream via the ComBlock output FIFO.

The computed FCI is then yielded to the ComBlock input FIFO, which can be easily transformed into a classification result (γ/n) by limiting the classes with a defined

threshold. AXI4-Lite interfaces from the MicroBlaze soft-core are used to configure the classification IP core and to retrieve the results from the ComBlock. These values are sent back to the computer through the UART interface. This validation framework is also compatible with the remote diagnostics platform presented in Chapter 2, useful to diagnose the detector plugged to the system through the onboard AFE.

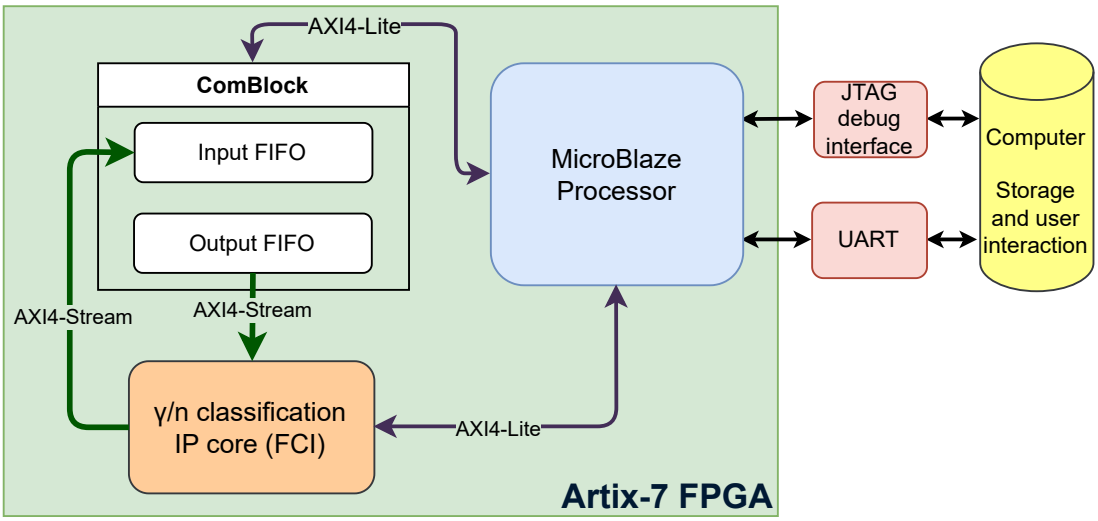

Figure 4.8: Block diagram of hardware validation framework, deployed in an Artix-7 FPGA.

The γ/n classification IP core was also tested in the ZedBoard SoC development board mentioned in Section 1.2.3, reporting an overall resource utilization lower than 8.1%. Such statistics suggest the potential to exploit the parallelization capabilities in multi-detector systems.

4.6 Discrimination in organic scintillators

A simulation was developed to explore the potential of the FCI method for γ/n discrimination in fast organic scintillators. This approach consisted in synthesizing unique γ and n pulses as if they were captured by a cross-level trigger mechanism. Then, discrimination was performed using both the reference PSD and the FCI-based meth-

ods. The same FoM from Section 4.5.2 was used to compare their performances, swept through a set of low-level discriminator (LLD) values (energy cuts). Since the ground-truth of each event type (label γ/n) is known from the pulse synthesis, a confusion matrix was also computed for each discrimination method. This allowed to determine the overall accuracy, without disregarding low-energy events.

4.6.1 Pulse synthesis

The pulse models for γ and n events belong to an EJ276 plastic scintillator coupled to a PMT [213], defined as a quadruple exponential curve with the decay times and amplitudes detailed in [214]. The pulse length was set to 2.048 µs: an heuristic trade-off to maximize the Fourier transform resolution with a reasonable capture time window. Individual pulse amplitudes were obtained by uniformly sampling independent γ and n energy spectra from a dataset recorded at NSIL (IAEA): these traces were acquired with the same digitizer from Section 4.3.1, capturing events from a PMT-based EJ276 probe under the stimulus of the DT generator and a Cs-137 γ calibration source. White additive Gaussian noise was summed to each trace equating the standard deviation of the measured baseline noise from the experimental dataset: $\sigma_n = 91$ keVee. Since the recorded data were captured with a cross-level trigger set at four times σ_n, the event energies range from the threshold at 364 keVee up to 6591 keVee, where ADC saturation is reached.

Energy calibration was conducted using the 2.2 MeV prompt γ events derived from the hydrogen activation, and the 661.7 keV pulses from Cs-137 source. Since there is no photoelectric absorption in the plastic scintillator, the Compton edges for each known energy were used to compute the calibration curve in gamma-equivalent units [215]. The intercept was also used, set as zero keV and equivalent to the baseline level.

The synthesized pulses were emulated with a sampling rate of 250 Msps and 14 bits of amplitude resolution, matching with the technical specifications of NSIL's DAQ

board from Section 1.2.3. In total, one hundred thousand events were created for each class (γ/n), randomly picking the energy distributions from the experimental dataset with the EJ276 PMT probe.

4.6.2 Event discrimination

Adhering to the same approaches from Sections 4.3.2 and 4.4.2, γ/n discrimination was performed with the simulated traces. The optimal integration windows to achieve the best PSD performance were obtained from [215]: 80 and 320 ns for the short and long gates, respectively.

Besides, Nakhostin demonstrated that most of the discrimination features in organic scintillators can be extracted from the lower frequencies (below 18 MHz) [141]. Hence, a constrained (within Nakhostin's limit) two-dimensional grid search was executed using the synthesized dataset to determine the optimal narrowband (PSA_l) and wideband (PSA_w) integration windows to compute FCI, as in Equation 4.3. The AESD featured a resolution of 488 kHz, derived from the 2.048 µs traces sampled at 250 MHz (512 samples FFT). The best discrimination was achieved by including the zero-frequency component, resulting in $PSA_l = [0 - 0.98]$ MHz and $PSA_w = [0 - 15.1]$ MHz. Similarly to Figure 4.4 with CLYC, Figure 4.9 shows the approximate energy spectral density magnitude, exhibiting the discrimination features for single γ and n events from the synthesized EJ276 dataset. Baseline restoring was mandatory in this case, due to the inclusion of the first discrete frequency component.

4.6.3 Discrimination performance

Figure 4.10 shows the γ/n discrimination matrices computed with PSD and FCI. Contrary to CLYC, fast neutron events in organic scintillators evidence no minimum energy caused by the Q-value of the reaction. This hinders the discrimination performance in lower energy ranges. Still, relevant differences between PSD and FCI are

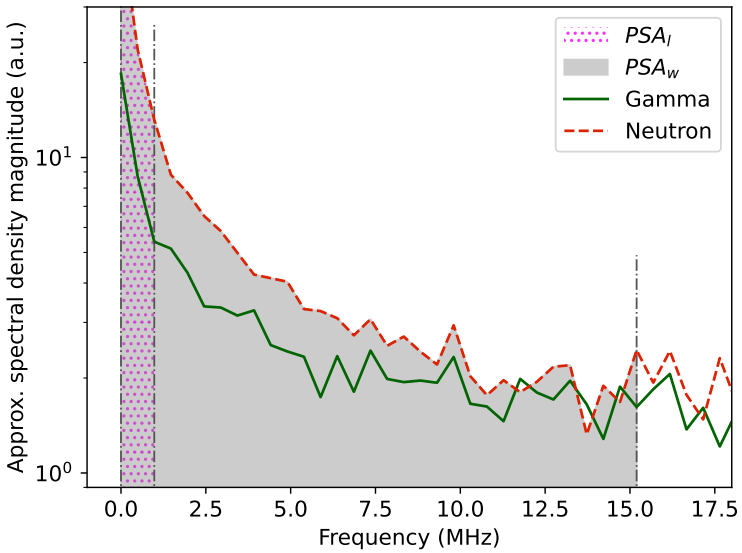

Figure 4.9: Approximate energy spectral density magnitude of simulated EJ276 γ and n traces.

clearly visible in events with low energies ($<$ 1000 keVee), where PSD usually fails to yield reliable results in this type of detectors [142]. Although a polynomial or step-wise function might significantly increase the performance of FCI, a constant classification line was applied to both results, preserving a fair comparison between both methods. The class separation limit value was optimized using a grid search with 1×10^{-4} a. u. step resolution to achieve the best overall accuracy for each case.

The FoM (from Equation 4.4) was computed to quantify the γ/n discrimination performance. Table 4.3 summarizes the FoM values for PSD and FCI after applying different energy cuts (LLDs), emphasizing the improvement for the frequency-based discrimination method in all the evaluated cases.

Outlining the worst-case discrimination scenario in the simulation (no LLD applied), Figure 4.11 shows the confusion matrices from the classification results for PSD and FCI. The latter resulted in superior performance, especially in reliable neutron identification: false positive counts lower than 1% (compared to $\sim$ 8% of PSD). This outcome may

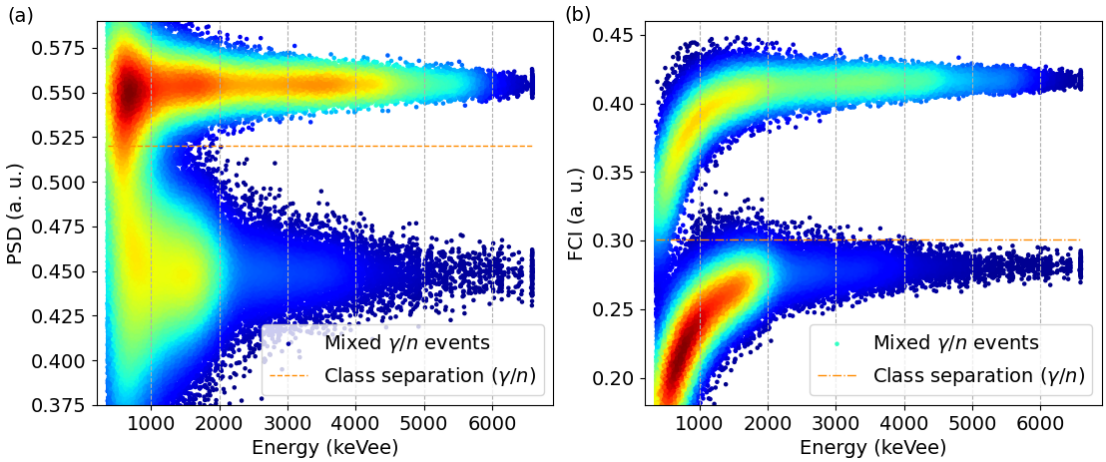

Figure 4.10: γ/n discrimination matrices from EJ276 simulation using a) PSD and b) FCI. In both cases γ events are clustered in the lower end and n are represented by the upper lobes.

LLD (keVee)	PSD FoM (a. u.)	FCI FoM (a. u.)	Improvement (%)
363	0.90	1.45	61.5
750	1.31	2.15	63.4
1000	1.51	2.56	69.2
1500	1.87	3.16	68.7
2000	2.31	3.94	70.7

Table 4.3: PSD and FCI γ/n FoM from EJ276 simulation runs over diverse LLDs.

be of particular interest in applications demanding dependable detection of neutron sources, where reduced γ contamination counts are critical [148].

4.7 Discussion

The FCI-based method achieved better discrimination performance (FoM) for gamma and neutron events when compared to other solutions based on frequency-domain analysis. It also demonstrated a superior FoM compared to PSD when tested with the same experimental data. Moreover, reliable confirmation of the observed γ/n distributions with FCI is supported by the expected response of the CLYC crystal, even at the lowest n_f energy-equivalent range of the detector. No ToF calibration is

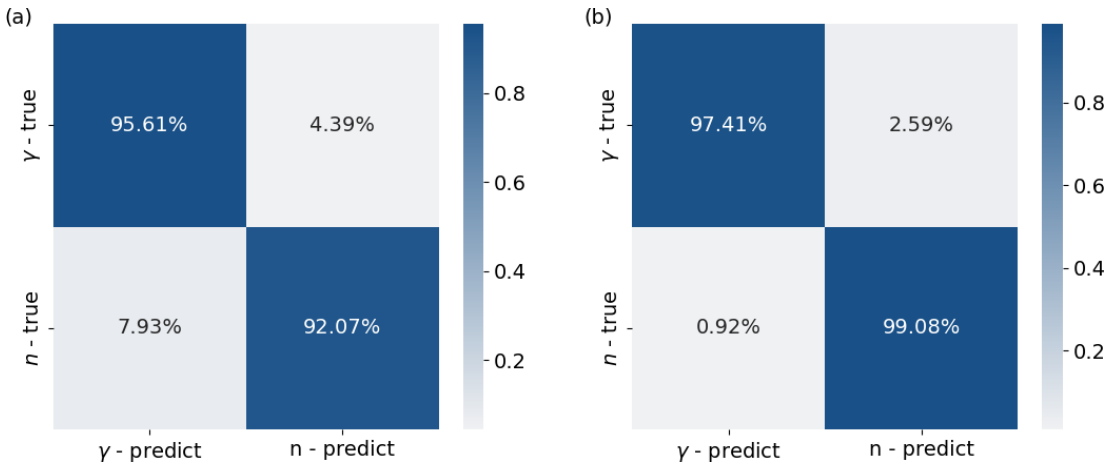

Figure 4.11: Confusion matrices of γ/n classification from EJ276 detector simulation without LLD using a) PSD with overall accuracy 93.8% and b) FCI with overall accuracy 98.2%.

required to achieve such results, demonstrating the CLYC capabilities as a standalone n_t detection and $n_f + \gamma$ spectroscopy system when coupled with SiPM sensors.

In contrast to other methods based on frequency-domain analysis for γ/n discrimination [141; 16], FCI was conceived for its deployment in embedded applications. The city-block simplification used in the approximate energy spectral density magnitude (AESD) computation along with the 100 Msps sampling rate render it an attractive solution for either SoC/FPGA or DSP real-time implementations.

An open dataset of tagged gamma and neutrons was obtained after the data curation process. Such dataset is valuable due to the scarce availability of classified γ/n traces sampled using a small integrated detector based on CLYC coupled to a SiPM array. It is worth mentioning that all analyzed traces were labeled, including pulses related to low-energy events. The dataset includes two types of events (classes): gamma photons and neutrons. The raw pulse traces and their corresponding calibrated energy are stored for each detected event up to 4.5 MeVee.

The simulation results using an EJ276 probe model suggest the applicability of the FCI method to organic scintillators, rendering useful for applications demanding

fast neutron spectroscopy. Since this analysis was conceived with real-time embedded deployment in mind, the sampling rate and amplitude resolution were limited to match the capabilities of the NISL embedded hardware platform from Chapter 1.

4.8 Summary

This chapter details the development of a discrimination method for γ/n events in mixed radiation fields based on frequency-domain analysis. Similarly to pulse-shape discrimination, it was conceived to be implemented in SoC/FPGA platforms, yet offering superior performance spanning through all the event energy ranges. Testing these features in simulation and in a hardware deployment support the accomplishment of research objective III. Besides, the results after evaluating a fast organic scintillator suggest the potential to exploit the discrimination method using diverse detectors.

Moreover, the resulting open dataset from this work with tagged gamma and neutron events fulfills research objective IV. The dataset along with the FCI method are very likely to benefit further algorithm developments that demand labeled raw data [143], such as supervised machine learning applications for real-time γ/n discrimination in challenging scenarios (research objective V). Further details of this specific contribution are explained in Chapter 5.

Chapter 5

High event-rate online discrimination with machine learning on mixed radiation fields

This chapter focuses on the development of a flexible instrument for event identification and quantification under challenging conditions. The limitations of the system from Chapter 4 are impressively overcome, by summing up several contributions described in previous chapters.

An embedded instrumentation platform for particle discrimination in mixed radiation fields is introduced, experimentally demonstrating its potential with the commercial γ/n SiPM-CLYC detector presented in Chapter 1. This embedded system executes a real-time ML classification model based on a multilayer perceptron running on the NSIL hardware board from Section 1.2.3.

The supervised ML model was trained with the reliable labeling method from Chapter 4. This training dataset was carefully curated using the correlation filter from Chapter 3, to avoid including events or noise unrelated to γ and n events in the training stage. Besides, the instrument is designed with upgradability in mind, equipped with an open serial port for debugging, compatible with the diagnostics platform presented in Chapter 2.

The embedded system provides single-event online γ/n identification with 98.2%

overall accuracy at rates higher than 200 kilocounts per second. Thermal and fast neutrons up to 5 MeV can be detected and discriminated from γ events, even under severe pile-up scenarios with a dead-time lower than 2.5 µs.

Moreover, the system exhibits excellent size, weight, and power consumption (SWaP) characteristics, packed in a volume smaller than 0.6 liters and weighing less than 0.5 kilograms, while ensuring continuous operation with only 1.5 watts. Such features render this solution suitable for embedded applications where low SWaP is critical and radiation levels manifest large count rate variability, such as space exploration, portable dosimeters, radiation surveillance on uncrewed aerial vehicles (UAV), and soil moisture monitoring.

The flexibility provided by the system allows to easily adopt several types of detectors for mixed radiation fields, enabling real-time event discrimination without requiring significant hardware updates. Diverse types of detectors can leverage the system capabilities, by simply tuning the FPGA firmware and the optimized ML model to tailor the target detector specifications.

This chapter is written upon the publication "Gamma/neutron online discrimination based on machine learning with CLYC detectors" in the *IEEE Transactions on Nuclear Science* journal on 14 November, 2024 [5]. The developed system is a result from the collaboration between the Multidisciplinary Laboratory (ICTP) and the Nuclear Sciences and Instrumentation Laboratory (IAEA).

5.1 Introduction

The use of CLYC scintillators in mixed-radiation detectors has recently become widespread due to their remarkable γ/n discrimination capability, enabling gamma and fast neutron spectroscopy as well as thermal neutron sensing [88; 216]. These features may be exploited in several applications where SWaP justify a design based on a single multi-purpose detector, such as space and planetary exploration [217; 218; 219;

220; 221; 106], portable γ/n dosimeters [153; 105], remote radiation surveillance on uncrewed aerial vehicles (UAV) [100; 222], and soil moisture monitoring [223; 224; 225]. Although traditional photomultiplier tubes (PMT) coupled to CLYC crystals offer excellent pulse-shape discrimination and notorious gamma spectroscopy resolution, silicon photomultipliers (SiPM) exhibit better immunity to external magnetic fields, withstand higher mechanical stress, and require a less complex bias power supply. Moreover, SiPMs evince much lower SWaP characteristics than PMT, making them excellent candidates for applications demanding portability and mechanical robustness [82]. However, the parasitic capacitance of SiPM arrays not only severely penalizes the inherent pulse-shape discrimination feature of scintillators, but also increases the pulse decay times [95; 44]. Hence, higher probability of pile-up and reduced γ/n separation harm the overall performance compared to PMT-based detectors, encouraging the development of advanced techniques for event discrimination and pile-up recovery in large pulse count-rate scenarios [138].

Machine learning (ML) algorithms have shown promising results in γ/n discrimination for CLYC [226; 227; 228] and organic scintillators [229; 230; 231]. Offline ML methods have also demonstrated to recover information in high event rates where pile-up usually distorts individual pulse shapes [232; 233; 234]. Consequently, applications requiring γ/n discrimination at high count rates are very likely to need similar advanced processing techniques. If real-time processing is added to this already-complex endeavour, simultaneously achieving these goals turns into an ambitious task targeting high-performance digital signal processing [20].

By exploiting the potential of integrated mixed-radiation detectors under these challenging conditions, an efficient solution for real-time event discrimination is presented, enabling high count-rates and low dead-time, featuring pile-up recovery in a low-SWaP embedded system (ES). These outstanding characteristics were achieved by tackling the discrimination challenge with a FPGA for real-time signal processing,

executing a pipeline with a fast feature extraction system, followed by an optimized ML classification model.

The remainder of this chapter is organized as follows. In Section 5.2 the specific contributions of this work are described. In Section 5.3 the embedded real-time system is summarized, detailing the FPGA firmware design for the target hardware board. Section 5.4 delves into the experimental data recording, offline data analysis and tagging for the ML training datasets, highlighting the attributes that enhance the efficiency of the system compared to other ML-based implementations. The real-time feature extraction system and its implementation are detailed in Section 5.5. Section 5.6 describes the ML model and the performed optimizations to provide outstanding γ/n discrimination performance in mixed radiation fields combined with low latency under continuous operation. Section 5.7 explains the tests conducted to evaluate the performance and SWaP aspects, including comparisons with SOA. Final thoughts and motivation for future work are presented in Section 5.8. The summary of the chapter can be consulted in Section 5.9.

5.2 Contributions

By identifying the current trends in low-SWaP applications that demand high count-rate event discrimination in mixed radiation fields, an embedded instrumentation platform has been developed featuring the following contributions:

- Event discrimination with overall accuracy higher than 98% at count-rates greater than 200 kcps.
- Continuous recovery under severe pile-up conditions with dead-time as low as 2.5 µs.
- Low SWaP using an off-the-shelf integrated CLYC detector and a low-end FPGA on a custom hardware board.

- Discrimination performed by a compressed machine learning model trained with a ground-truth dataset, endorsed by the expected scintillator response under mixed γ/n radiation fields.

5.3 Embedded system design

The embedded system (ES) is a self-contained autonomous device composed by the off-the-shelf CLYC detector and the custom NSIL hardware board from Sections 1.1.3 and 1.2.3, enabled by custom firmware developments on the microcontroller (MCU), FPGA fabric, and FPGA soft processor. The ES is housed in a 3D-printed enclosure printed using acrylonitrile butadiene styrene and optimized to minimize the overall volume, while providing a robust lightweight structure to safely pack the whole system for diverse use-case scenarios.

Figure 5.1a shows the block diagram of the most relevant ES components, including the detector and the hardware board. Figure 5.1b shows the physical appearance of the embedded system as used in the experiments.

In this development, the integrated CLYC detector demonstrated to saturate at 1.5 V_{pk-pk} with event energies of 6.1 MeV, yielding a calibrated dynamic range of 67.8 dB and energy resolution of 496 eV/LSB (4.07 keV/mV), when matched to the AFE of the NSIL DAQ board.

5.3.1 FPGA firmware

The firmware for the FPGA was designed from ground-up, based on ComBlock for slow-control, as presented in Chapter 4. Flexibility is granted in the design, enabling updates to be built atop of this implementation, looking towards diverse discrimination applications involving mixed radiation fields.

The FPGA performs the data acquisition and real-time digital signal processing in the ES. An overall view of the firmware design integration is shown in Figure 5.2, on

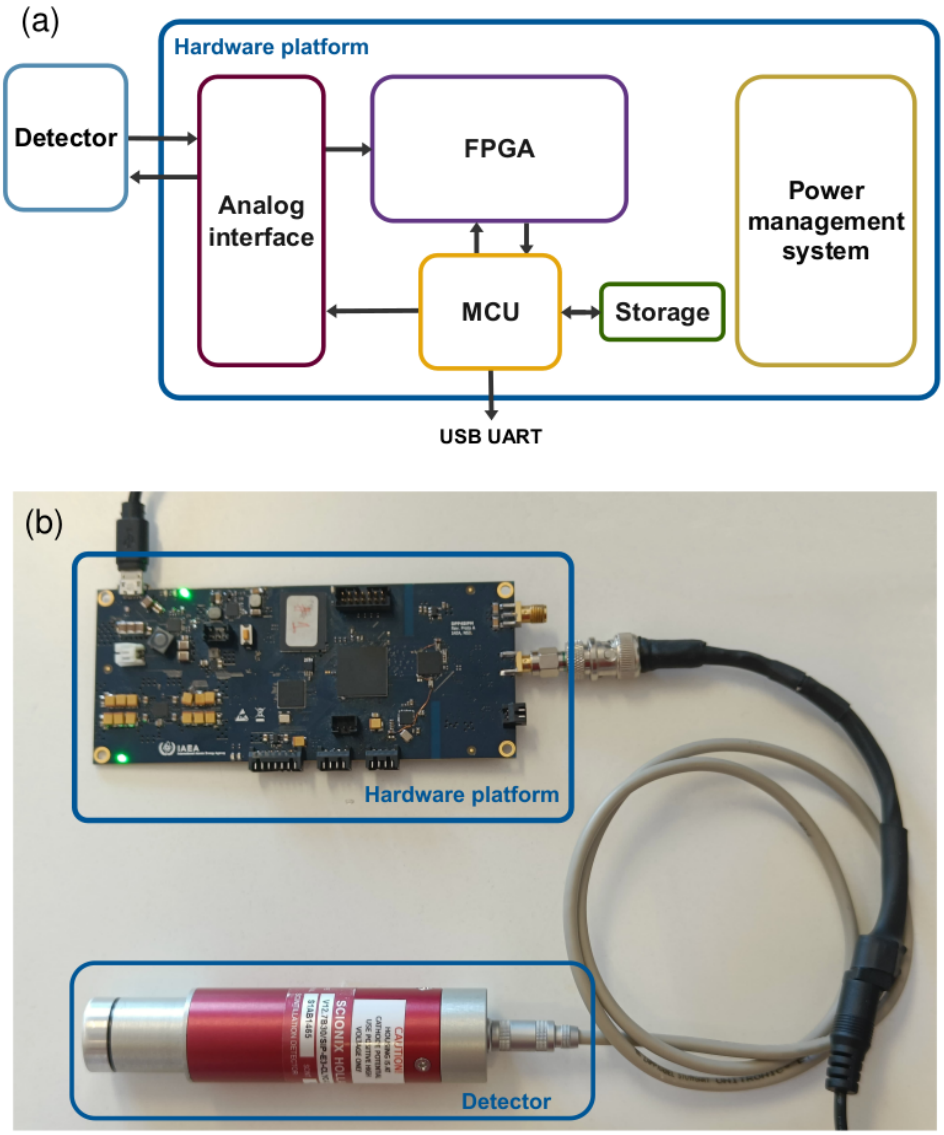

Figure 5.1: Embedded system design, featuring the commercial SiPM-based CLYC detector and NSIL custom low-SWaP hardware board. (a) A block diagram highlighting the most relevant elements of the platform. (b) A photograph of the system components as used in the experiments.

which the most relevant elements are presented. This implementation was developed using custom VHDL modules for the real-time processing stages. The main dataflow is managed by standard AXI4-Stream interfaces [235], enabling automatic handshaking and flow control through all the online processing chain.

Following this methodology, the output from the detector is continuously digitized by the ADC at 100 Msps and streamed into the feature extraction block, where the most relevant data are captured and transferred out. Subsequently, a clock-domain-crossing (CDC) module is used to translate these samples from the low-power 100 MHz clock domain to a faster 200 MHz high-performance processing section.

The ML-based classification model computes the event type (γ/n) in less than 1 µs and feeds the performance profiling block with the inferred outcome. At this stage, γ total count, n total count, γ event rate, and n event rate are continuously computed. Finally, another CDC transfers these measurements to the Communications Block (ComBlock) input registers [157], used to periodically gather the performance metrics and set the working parameters of the feature extraction, ML-based classification system, and profiling blocks.

An AMD MicroBlaze soft-processor running at 100 MHz is used to connect the ComBlock with an external serial interface (UART), handled by a USB-UART bridge coded in the onboard STM32 MCU. This port is used to output the performance metrics once a second. The UART also worked as the debugging port in the development stage, conducted with the diagnostics platform from Chapter 2. The CDC modules are implemented with AXI4-Stream FIFO blocks, simultaneously serving as a pipeline buffer for real-time processing and as a safe clock domain translator.

Power consumption optimization in the FPGA firmware design consists in accelerating only the ML inference and performance profiling blocks at 200 MHz, while keeping the remaining components at 100 MHz. The ADC sampling frequency was also set to 100 Msps instead of the maximum 250 Msps available; this does not affect the

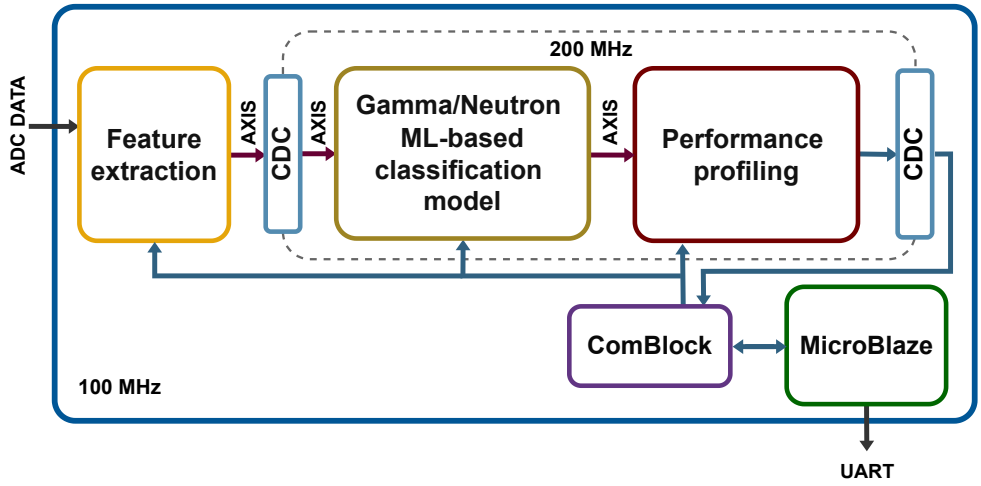

Figure 5.2: FPGA firmware design integration. Feature extraction block captures a set of samples from the leading edges of events digitized by the system's ADC. The ML-based classification model performs the γ/n discrimination and feeds the performance profiling block for real-time operation monitoring. ComBlock and MicroBlaze are used for parameter setup and results visualization via a serial (UART) interface. AXI Stream (AXIS) protocol was used through the critical real-time processing chain: from the feature extraction block output to the performance profiling input.

γ/n discrimination capabilities and contributes to dynamic power reduction in both the ADC and FPGA. Indeed, at 50 Msps the feature extraction system was not able to accurately determine the time of arrival of the pulses and the ML model complexity was not reduced proportionally, whereas at 250 Msps the ML model complexity increased significantly without providing relevant accuracy improvement, while increasing dead-time and power consumption. Thus, the feature extraction and the ML model benefited from the 100 Msps sampling rate.

5.4 Data preparation

The experimental data for this research were recorded at the Neutron Science Facility (NSF), IAEA Laboratories, in Seibersdorf, Austria. Two neutron sources were used: an

Americium-Beryllium (AmBe) passive source and a pulsed Deuterium-Deuterium (DD) active neutron generator [202]. A Cs-137 γ source was placed beside the CLYC detector during most of the measurements to provide valuable energy calibration data. Strong thermalization of the fast neutrons was caused by the DD operation housing module (high-density polyethylene shield) and the surrounding room objects. Figure 5.3 shows a close-up view of the CLYC detector placed beside the beam port of the DD generator and the Cs-137 source, as used during the data recording session at the NSF. A lead block was positioned between the DD output and the detector to maximize the γ counts from the calibration source. Similarly, fast neutrons from the AmBe source suffered some scattering due to the surrounding matter in the room, even if no polyethylene shield was present during the experiment.

Figure 5.3: Integrated CLYC detector placed at the output window of the DD generator, separated by a lead block to maximize the Cs-137 γ counts, used for energy calibration.

The detector's output signal port was coupled through a 50 Ohm RG-316 cable to a CAEN DT5761 digitizer to collect a dataset sampled at 4 Gsps with 10 bit amplitude

resolution and virtually no dead-time [49]. More than 100,000 raw pulses were recorded for offline processing.

5.4.1 Data wrangling

Data wrangling was performed using the Pandas library to filter out undesirable events in the recorded dataset, such as saturation or noise. Piled-up pulses were also removed, leaving only single clean traces to further synthesize piled-up events under controlled conditions. This cleanup process consisted of consecutive stages to deal with the different spurious pulses:

1. The dataset was downsampled from 4 Gsps to 100 Msps, matching the sampling rate of the low-SWaP ES.

2. Individual events from multiple data recording sessions were mixed and indexed in a single data frame.

3. Baseline restoring was applied with the average of the initial 20 samples of each pulse. Amplitude polarity was flipped to obtain a set of positive traces with their first samples located in an average of zero amplitude units.

4. High-energy amplitude-saturated events were removed.

5. Pile-up rejection was performed using the matching filter to distinguish single events from distorted traces detailed in Chapter 3, dropping events with correlation lower than 98%. Because gamma and neutron pulses exhibit similar shape features in the SiPM-based detector, the same template was used in both cases.

6. A second cleaning stage was performed using residuals analysis, where the normalized average pulse shapes (γ/n) were subtracted from each amplitude-normalized trace. Once the sum of the residuals was computed, the outliers (distorted pulses) were discarded. Less than 100 traces ($< 0.1\%$) were affected by

this cleanup, promoted by low-energy events piled-up over high-energy pulses that passed undetected in the previous stage.

7. Energy calibration was performed over the clean dataset, computed as the amplitude of the pulse, referenced to the well-known Cs-137 gamma 661.6 keV photopeak and the characteristic CLYC thermal neutron cluster close to 3.2 MeVee [3].

Then, individual event tagging (γ/n) was carried out using the frequency-based classification algorithm from Chapter 4, leveraging the ground-truth labeling method supported by the interactions of the CLYC crystal under mixed radiation fields [3]. Figure 5.4 depicts the γ/n discrimination matrix of the labeled events from both the DD generator and the AmBe source, shown as a 2D histogram of discrimination index vs energy. The thermal neutron cluster is merged in the fast neutron spectrum, close to 3.2 MeVee due to the Q-value of the ^{6}Li reaction channel [236]. Thus, fast and thermal neutrons have been seamlessly discriminated from gamma events, following the same methodology utilized to generate the open dataset released in [4].

The data wrangling process resulted in a clean dataset with γ/n tags containing original (raw) pulse traces, calibrated in keVee units. These events were used to synthesize piled-up pulses under controlled conditions: individual event type (γ/n), energy in gamma-equivalent units, and time of arrival. As from the nuclear reactions involving fast and thermal neutron interaction in CLYC crystals, it is well known that neutron events cannot be measured as pulses with energies less than 500 keV in the gamma energy scale (keVee) [89]. Thus, events with energy lower than this limit are disregarded in the pile-up dataset synthesis, as no γ/n discrimination is required and gamma-only pile-up recovery may be processed using an existing real-time FPGA-based method [237], prior to amplitude discrimination.

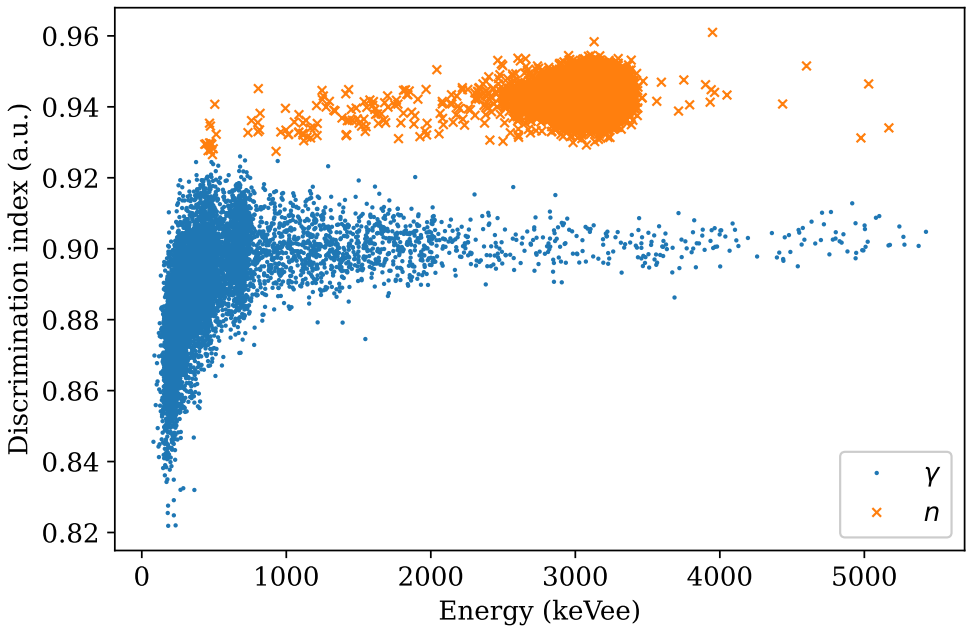

Figure 5.4: Labeled γ/n discrimination matrix of the experimental data, including events from both DD generator and AmBe source. The fast neutron spectrum in gamma-equivalent energy units is merged with the thermal neutron events around 3.2 MeVee.

5.4.2 Pile-up synthesis

Although some piled-up traces were observed directly in the recorded data from the neutron sources, no ground-truth of the event type (γ/n) or time of arrival can be certainly defined. Therefore, a dataset of pile-up events was created using individual (tagged) events from the clean dataset obtained through the process presented in Section 5.4.1, following the approach used in [143; 229]. Several types of event combinations resulted of this process: $\gamma+\gamma$, $\gamma+n$, and $n+n$, with traces containing up to 8 consecutive pulses.

The distribution of the time of arrival of piled-up events follows a Poisson process, modeled as an exponential distribution with rate parameter λ. Based on this, several

combinations of traces at different average event rates were created: 10, 25, 100, and 200 kHz.

It was also observed that a minimum separation between the time of arrival of each piled-up pulse was required to avoid losing a low-energy event in the trailing edge of a high-energy trace, caused by the long decay time of the CLYC detector. A value of 2.5 µs was set as the minimum inter-pulse distance (MIPD), aiming at reducing the probability of missing an event with the aforementioned characteristics in a trace. MIPD defines the overall dead-time of the system, leaving the possibility to tune it according to individual implementation requirements. Also, a minimum of 5% overlap between pulses was forced in the synthesis to accurately quantify the pile-up recovery of the real-time feature extraction stage (detailed in Section 5.5) in a worst-case multiple pile-up scenario.

After applying these hard constraints, the Poisson behavior was assessed by fitting the time distance between the synthesized pulses with an exponential distribution. Figure 5.5 shows the results of analyzing 60,000 piled-up traces, evidencing a perfect coefficient of determination ($R^2 = 1$) and a measured pulse rate $\lambda_{fit} = (199.8 \pm 0.7)$ kHz that matches the value set for the synthesis $\lambda_{set} = 200.0$ kHz. Therefore, the established hard constraints did not influence the average rate of events nor the arrival times distribution.

Individual events were randomly picked from the clean dataset (detailed in Section 5.4.1) and further piled-up as shown in Figure 5.6, where an interleaved $n + \gamma$ synthetic trace with MIPD = 2.5 µs and $\lambda = 200$ kHz is depicted.

5.4.3 Pulse shape features

Traditional methods usually analyze the trailing edge of the detector traces to compute a γ/n discrimination factor [16]. On the contrary, the work presented in this chapter analyzes the leading edge features of gamma and neutron events to determine

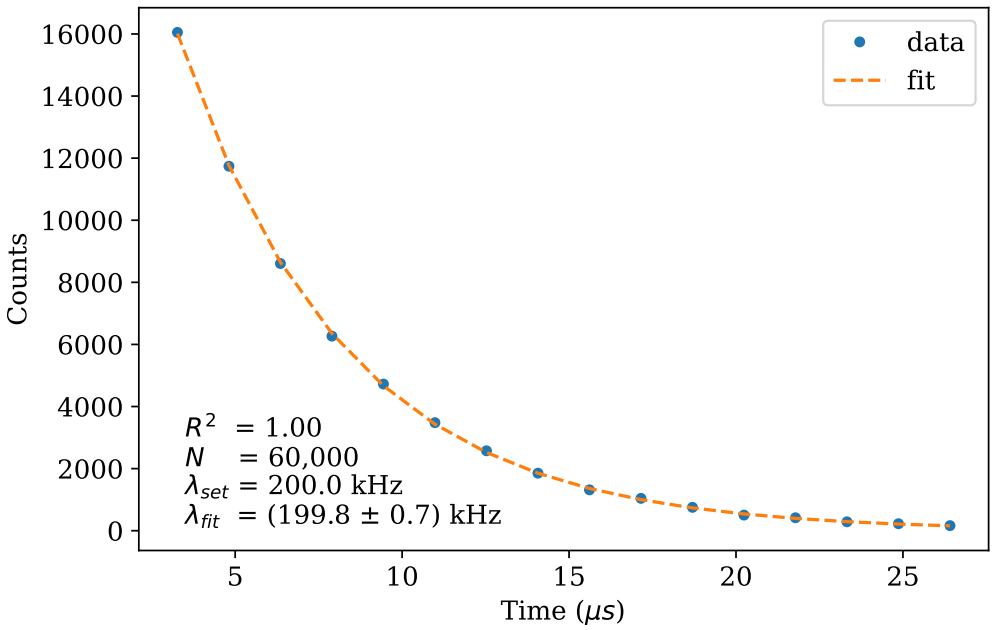

Figure 5.5: Exponential fit for time distance between 60,000 events with an excellent coefficient of determination $R^2 = 1$. The expected rate of the synthesized piled-up events $\lambda_{set} = 200.0$ kHz remained within the fitted uncertainty limits $\lambda_{fit} = (199.8 \pm 0.7)$ kHz.

the most significant differences of both types, aiming at developing a reliable system under high event rate scenarios. The developed approach extracts just the right amount of features to enable a quick γ/n identification, simultaneously decreasing the overall dead-time and significantly reducing the complexity of the machine learning model used for discrimination.

Offline analysis with the clean dataset from Section 5.4.1 was carried out, obtaining the amplitude-normalized average pulse representation of gamma and neutron events: the leading edges of these shapes are shown in Figure 5.7a. The residuals between the average neutron and gamma traces of Figure 5.7a were computed sample by sample to determine where the neutron and gamma traces show the highest variability.

The derivative of the residuals trace further highlights in a region of interest (ROI)

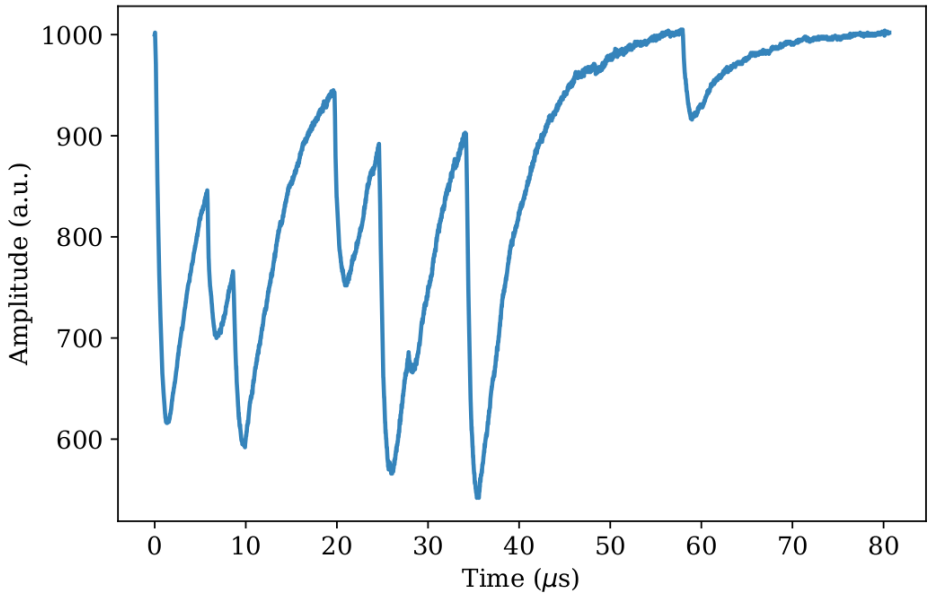

Figure 5.6: Synthetic piled-up trace with eight $n + \gamma$ interleaved events from experimental data. Minimum time between pulses set to 2.5 µs and average event rate $\lambda = 200$ kHz.

the most significant differences in the build-up process of an event type, as shown in Figure 5.7b. The optimal ROI length was minimized to 350 ns, equivalent to the first 35 samples of an event digitized at 100 Msps. This value is the number of samples used to further feed the ML γ/n classification model, as detailed in Sections 5.5 and 5.6.1.

Although the amount of samples could be slightly reduced, a subsequent precise pulse alignment mechanism would be required, significantly impacting the overall dead-time by introducing a stochastic delay [20]. Minimizing the number of samples in the leading edge ROI was crucial to notably improve the efficiency of the system compared to other offline and real-time ML-based solutions, which typically analyze the complete pulse shape or even multiple successive events [20; 44; 227; 228; 229; 230].

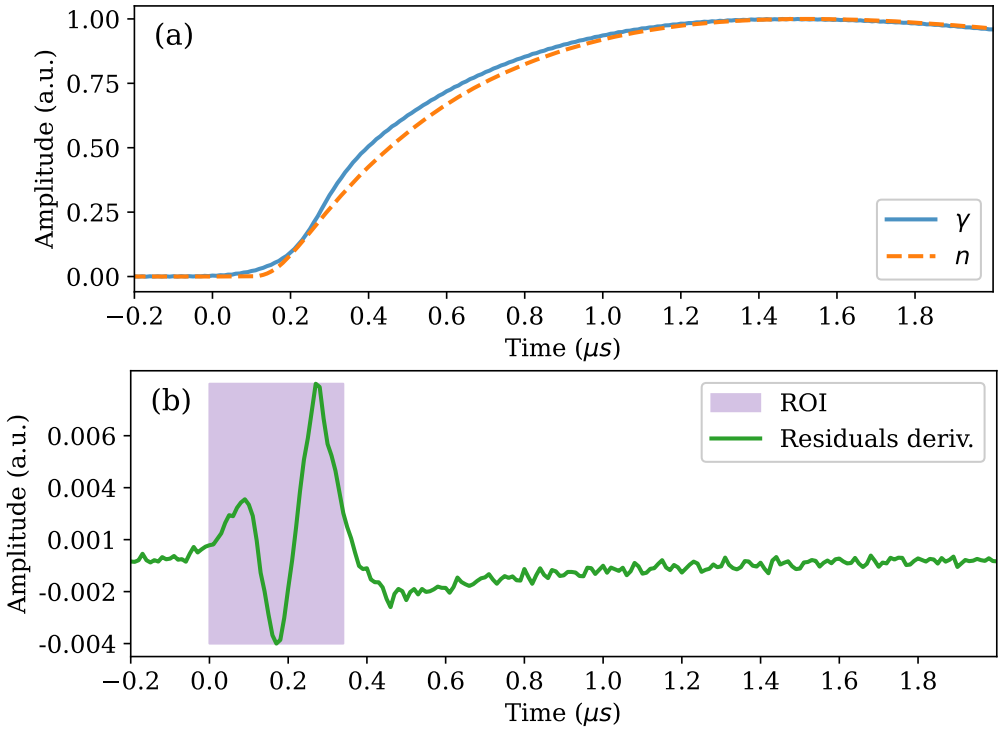

Figure 5.7: (a) Gamma and neutron leading edges of normalized average pulse shapes. (b) Gamma and neutron pulse shape residuals derivative and ROI used to define the most relevant samples for the ML-based discrimination (350 ns).

5.5 Real-time feature extraction

A real-time feature extraction (FE) method was devised aiming at sourcing the ML model with the right amount of data to accurately and efficiently perform the pulse discrimination for each incoming event from the detector.

First, the input signal's discrete second derivative is used to determine the time of arrival of an event, which is then passed through a cross-level trigger (CLT) with a constant threshold θ. Similarly to [138; 234], a smoothed variation of the discrete derivative is used to enable an efficient event detection despite of the pile-up scenario or variation in baseline. Their band-pass filtering was also carried out by adding

consecutive derivative and smoothing stages, obtaining a smoothed second-order derivative (SSD).

Since the processing is carried out in the digital domain, discrete implementations are mandatory to model the real-time computations. The discrete derivative $w(n)$ was approximated by the central-difference operator detailed in Equation 5.1:

$$w(n) = \frac{1}{2}\left[x(n) - x(n-2)\right] \tag{5.1}$$

Its band-pass frequency response, recognizable from the transfer function (TF) in $\mathcal{Z}$-domain shown in Equation 5.2, avoids the amplification of undesired high frequency components.

$$H_d(z) = \frac{1}{2}\left(1 - z^{-2}\right) \tag{5.2}$$

Smoothing is further performed by an N-point filter based on a moving-average implementation. The amplitude normalization $1/N$ was ommited to keep integer representation throughout the entire feature extraction process. This smoothing operator was chosen to avoid fine-tuning the filter for specific pulse shapes of different detectors, resulting in a reliable generic FE solution for multiple leading edge timing parameters. Moreover, the discrete-time filter $y(n)$ is recursively defined in the computationally-efficient implementation shown in Equation 5.3 [238], whereas its $\mathcal{Z}$-domain TF is represented in Equation 5.4.

$$y(n) = y(n-1) + w(n) - w(n-N) \tag{5.3}$$

$$H_s(z) = \frac{1 - z^{-N}}{1 - z^{-1}} \tag{5.4}$$

The N averaging parameter was set to 64 after successive tuning iterations, providing a good noise level reduction while preserving the required features for leading

edge detection in pile-up scenarios [239; 240].

Being a linear time-invariant system, the TFs used to compute the SSD (two successive smoothed discrete derivatives) can be combined into a single $\mathcal{Z}$-domain expression to define the feature extraction filter, shown in Equation 5.5.

$$H(z) = H_d^2(z)H_s^2(z) = \frac{1}{4}\left[\frac{(1-z^{-2})(1-z^{-64})}{(1-z^{-1})}\right]^2 \tag{5.5}$$

This FE filter evidences a stable and causal behavior for an infinite input trace. However, since the initial values of the delay lines z^{-2} and z^{-64} are reset at startup in the FPGA deployment, a slight warm-up time close to 1.3 µs is required. After this short period, the SSD computation runs uninterruptedly with the data stream provided from the ADC.

Whenever the CLT is activated over the SSD signal, a time-over-threshold (ToT) counter of length M is immediately started. The ToT is used to ensure that the CLT was certainly triggered by an event instead of noise: only if the total number of samples surpassing the threshold is more than $M/2$, the event is registered. Also, the CLT functionality is disabled after the end of the ToT M period, getting re-enabled only after a defined time-after-ToT (ϕ), used to filter out events or noise that may incorrectly detect an event too close one from each other. This safeguard introduces a deterministic dead-time that matches the MIPD from Section 5.4.2.

After the arrival of an event has been successfully identified by the FE system, the first R samples of its leading edge are captured and sent to the ML-based classification model. It was experimentally observed that the SSD pulses related to expected events in the original trace kept a ToT higher than 31 samples, with a threshold equal to 4 times the average SSD baseline ($\theta = 16$). Thus, the ToT parameter was set to $M = 64$ in the final implementation. A piled-up input trace along the computed SSD are shown in Figure 5.8, depicting the actual threshold value where events may be detected only after the ToT requirement is met.

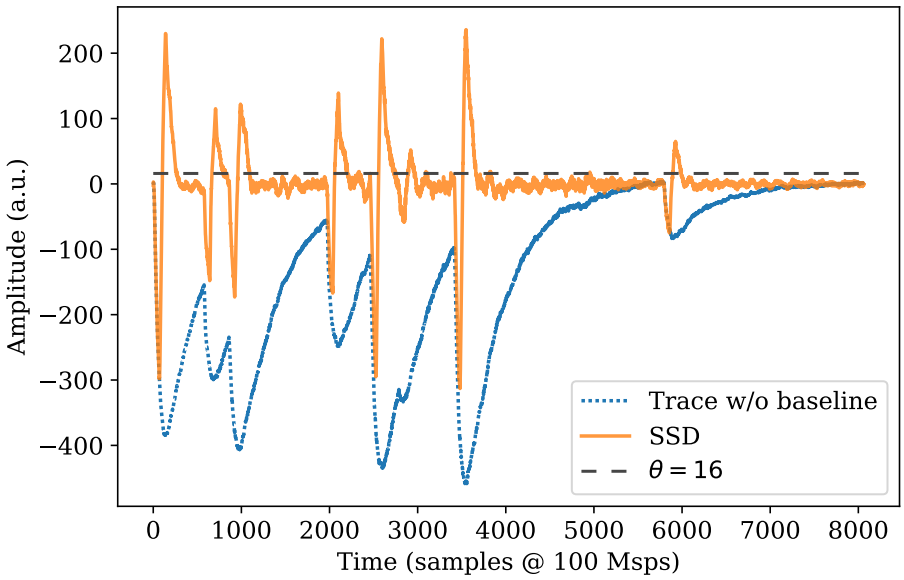

Figure 5.8: Piled-up trace with eight successive events and smoothed second derivative (SSD) filter output. Baseline was removed on purpose from the original trace to facilitate the visual comparison of the plots. The SSD computation inherently removes any baseline component. Threshold for detection on SSD set to $\theta = 16$ ADC channels.

As for the residuals analysis described in Section 5.4.3, the chosen number of samples fed to the ML model is $R = 35$. In accordance, the inter-pulse safeguard parameter is set to $\phi = 151$, which summed with M and R equals to the expected overall dead-time: 250 samples at 100 MHz, or 2.5 µs. Although shorter dead-times are easily achievable by reducing ϕ, the detection efficiency would not increase owing to the integrated CLYC long decay.

The FE implementation was initially developed as an offline simulation in Python 3.10, and then further translated to a real-time deployment for FPGA using VHDL. Integer representation is used in both the Python simulation and the FPGA design, avoiding unnecessary complexity related to floating-point in terms of resource utilization, while simultaneously improving throughput and latency [7].

The real-time feature extraction in FPGA was verified to match with the Python simulation behavior using a VHDL test-bench and further validated in hardware using the AMD's Integrated Logic Analyzer (ILA), as shown in Figure 5.9.

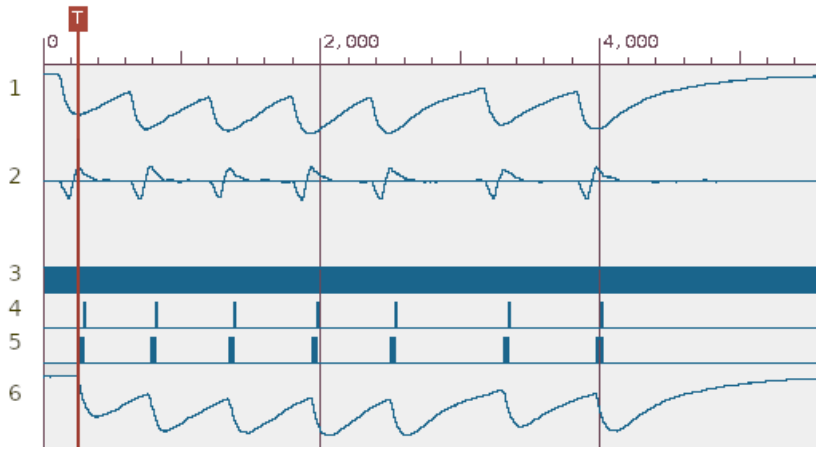

Figure 5.9: Screenshot of real-time feature extraction in FPGA shown in ILA. Signals from upper to lower: (1) Digitized piled-up input trace from ADC. (2) Smoothed second derivative of the input signal. (3) Valid data flag shows continuous single-clock cycle execution. (4) End of capture flag (AXI4-Stream TLAST). (5) Capture signal to fetch the first 35 leading edge samples of each detected event (AXI4-Stream TVALID). (6) Delayed input signal, used to align the capture signal with the deterministic latency of the feature extraction system. Horizontal axis represents time in sample units at 100 MHz.

5.6 γ/n discrimination based on machine learning

From a computational perspective, a low-SWaP solution implies reduction in power consumption during computation, storage, and memory access operations. As observed in SOA [241; 242; 243], ML-based models are computationally costly, impacting the energy efficiency of the overall system. Nevertheless, it has been demonstrated that these models can be implemented on edge devices without compromising the system performance [7; 53; 244; 245]. Aiming at real-time online event discrimination on the edge for mixed radiation fields, an optimized ML model was designed employing the

end-to-end workflow for deploying deep neural networks on SoC/FPGA proposed in a previous work [8] related to this research. A good compromise between effectiveness, memory footprint, and inference time was achieved while improving energy efficiency.

5.6.1 Model training and compression

The workflow combines three techniques to achieve efficient model compression: pruning (P), quantization (Q), and knowledge distillation (KD). P and Q are orthogonal to KD, aiding to achieve better performance and reducing the size of the model with minimum loss in accuracy [246]. P aims to decrease the number of parameters by removing neurons and connections. Q reduces the memory footprint by selecting the number of bits that represent the weights and biases. Finally, KD [247] is devoted to transferring knowledge from a teacher network to a smaller and faster target network (i.e., distilled or student), which can reproduce the teacher's behavior while being less computationally expensive. Moreover, the workflow uses Bayesian optimization to heuristically define hyperparameters.

To implement the workflow, the teacher and student architectures should be defined for the KD method. Furthermore, P and Q strategies are established within the student architecture because the KD learning process is performed through quantization-aware pruning (QAP) [248].

A multilayer perceptron (MLP) was trained as a teacher model to discriminate between neutron and gamma events. The MLP-based model was composed of six hidden layers chosen heuristically with 2,623 parameters, exhibiting an overall accuracy of 99.0%. The input of the model was the first 35 amplitude samples of the leading edge of the pulse.

The training dataset was generated by processing the piled-up traces from Section 5.4.2 using the feature extraction method from Section 5.5. In total, 100,000 leading edges were generated, with a perfectly balanced distribution for each output class:

50,000 gamma and 50,000 neutron events. Events covering all the energy range under study were used to enable the ML model generalization capabilities to accurately infer the expected class (γ/n), no matter the average slope of the leading edges.

After the teacher model had undergone training, the student architecture was defined as an MLP with 8-bit fixed-point precision, target sparsity of 30.00%, and 217 hyperparameters distributed in six layers. Subsequent to student architecture training through KD and QAP, the overall accuracy of the compressed model showed a slight decrease: from 99.0% to 98.2%. Specifically, the accuracy was 96.8% for γ and 99.5% for n.

A reduction in the memory footprint was achieved through model compression while maintaining high accuracy. The compressed ML-based γ/n classification architecture is illustrated in Fig. 5.10.

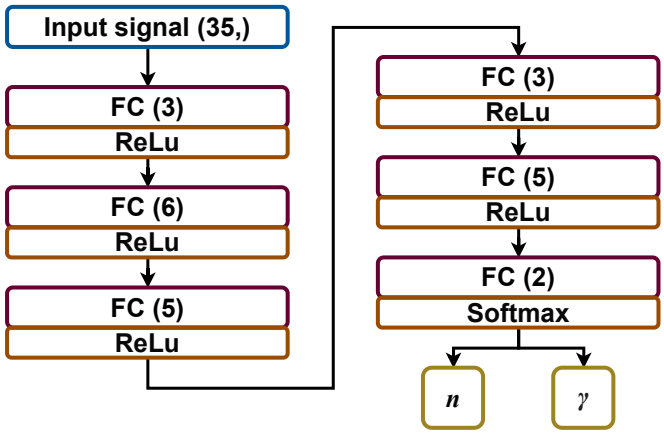

Figure 5.10: Compressed architecture based on MLP for γ/n event classification.

Once the compressed model was created, the hls4ml package [196] was employed to obtain a high-level synthesis (HLS) project to synthesize and export the model to the FPGA register transfer level (RTL). The resulting IP-core allows stream data transmission through dedicated interfaces optimized for low latency through the maximum parallelization of computations. Furthermore, the block generated for the

inference process was fully deployed in the on-chip memory of the FPGA, significantly reducing the overall power consumption compared to implementations based on off-chip memory transactions mandatory for bigger ML models in low-end FPGA parts [249; 250].

5.7 Results

The ES performance assessment and SWaP validation were conducted by quantifying experimental data of the most relevant indicators under diverse conditions. The measured performance metrics were count-rate, dead-time, discrimination accuracy, and pile-up recovery capabilities. The FPGA resource utilization is also reported, highlighting the efficiency of the ML model and the overall firmware design. In terms of SWaP, the evaluated features were weight, volume, and power consumption. The experiment descriptions, their outcomes, and quantitative comparisons with the SOA are presented next.

5.7.1 Performance assessment

Individual performance parameters are quantified and discussed, stressing the real-time capabilities of the system for high event rate scenarios. These parameters were experimentally measured with the performance profiling block in the FPGA, updating the results at 1 Hz through the UART port. Two tests have been devised with the CLYC detector aiming at providing a broad verification method as follows: evaluating maximum performance figures with synthesized piled-up events and using experimental raw data to validate the system behavior in real conditions. Besides, a complementary assessment was carried out with two different detector pulse shapes to showcase the flexibility of the system to work with other types of scintillators.

Verification with synthetic traces

Aiming at pushing the performance of the system to the design limits in terms of count-rate and dead-time, a new set of piled-up traces was synthesized following the methodology described in Section 5.4.2. This validation dataset was never used in the ML training or testing stages. Three types of pulse combinations were created on each trace, with up to 16 piled-up sequential events: all-gamma events (γ), all-neutron events (n), and alternated $\gamma + n$. These traces were stored and sequentially reproduced in three 24-hour individual sessions using an arbitrary waveform generator (AWG) to stimulate the input of the system's AFE.

Variability was ensured by summing 1% of white Gaussian noise atop of the traces in the AWG, reducing the probability to generate equal events during the long period of experimental validation in hardware. The following parameters were assessed at the intermediate stages and verified at the end of the individual tests: total gamma counts, total neutron counts, gamma count-rate, and neutron count-rate. These unique validation traces provided a solid reference for the total expected gamma and neutron counts and count-rates.

The measurements revealed a worst-case scenario for the feature extraction false-negative rate: 74 counts out of 1 million expected events were missed (**74 ppm**), equivalent to a sensitivity higher than 99.9%. Moreover, the ML experimental γ/n discrimination overall accuracy matched the expected model performance with three significant digits: **98.2%**. Since the dataset used for the validation was generated with an average event rate of **200 kHz** and a dead-time of **2.5 µs**, the 74 ppm proportion of missed pulses is negligible, demonstrating the actual performance of the ES as a whole with peak count-rates up to 400 kcps or $(2.5\ \mu s)^{-1}$.

Verification with experimental data

A verification stage with the experimental data was performed using raw traces recorded at the NSF with the CLYC detector. Three continuous traces were streamed sequentially into the ES analog front-end using the AWG at 100 Msps. Since no pulse cleanup was carried out, pile-up, saturation, and other perturbations were present in the traces. The data streams were selected as follows: (*i*) a γ-only Cs-137 source placed beside the detector, (*ii*) the AmBe neutron source along with the Cs-137 calibration source, (*iii*) the DD generator at 100% duty cycle with a lead shield in front of the detector, exhibiting higher neutron activity due to the thermalization caused by the surrounding high-density polyethylene.

Following the quantification approach from [251], the AmBe and DD γ/n relative abundances showed agreement in (*ii*) and (*iii*), respectively. Regarding (*i*) the false-positive rate for neutrons remained below 2.54%, matching the expected gamma accuracy of the ML model. Similarly to [20; 153], Figure 5.11 shows the real-time count-rates during this verification stage, computed by the performance profiling block and transmitted via UART to the SerialPlot application [252].

Validation for other detectors

Besides the performance measurements with the integrated CLYC, additional profiling was conducted with a SiPM-based NaI(Tl) detector. This device features similar leading edge characteristics to those of the CLYC gamma pulses, but much faster decay times: 1.2 µs compared to 5.0 µs. The NaI(Tl) was directly plugged to the system's AFE input and a Co-60 gamma source was placed beside. The measured discrimination performance was even better than with the CLYC detector, with 99.1% of the events classified as gamma at an average 35.1 kcps count-rate (limited by the 183 kBq source activity). Figure 5.12 shows the experiment results while displacing the source position nearby the detector.

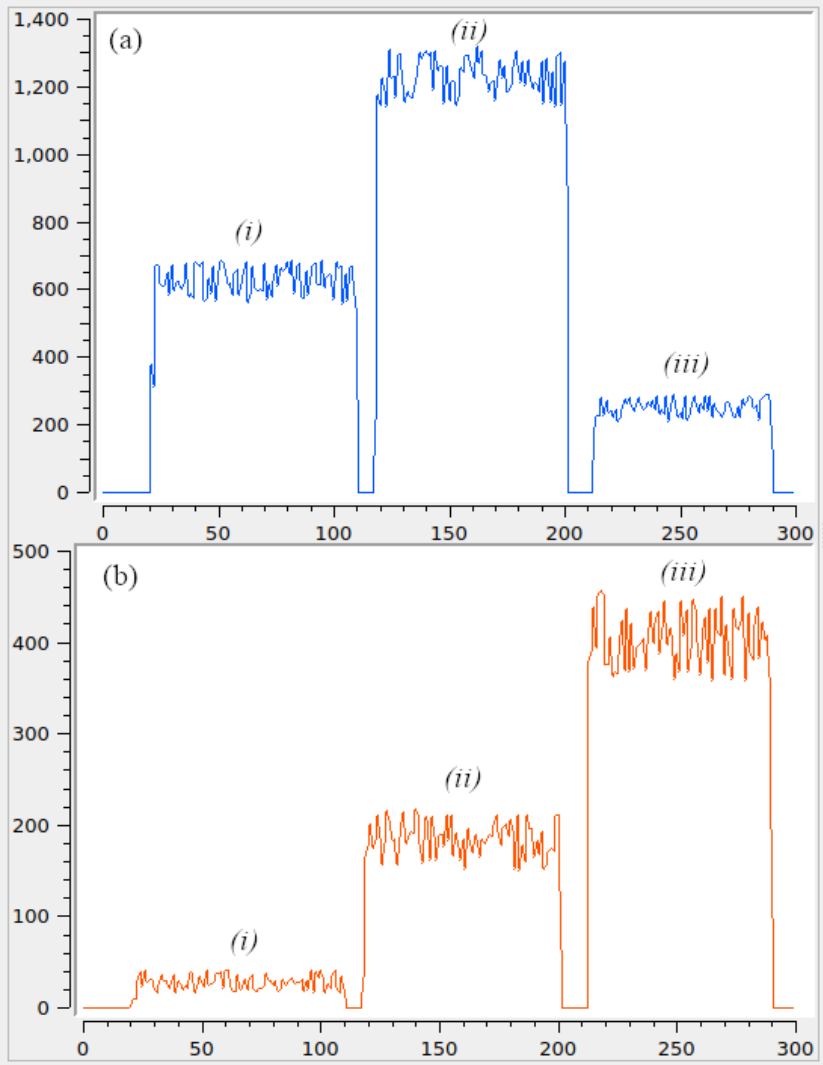

Figure 5.11: Screenshot from SerialPlot shows the count-rates computed online by the performance profiling block. Event quantification rates are shown for (a) gamma and (b) neutrons, as discriminated in real-time by the developed system. In both cases horizontal axes represent time in seconds, whereas vertical axes express the count-rates for each class (γ/n) in counts/second. Three stimuli were tested using experimental traces recorded at NSF with the CLYC detector: (*i*) γ-only Cs-137 source, (*ii*) AmBe source plus Cs-137 calibration source, (*iii*) DD generator plus lead shield. Stimuli were reproduced as a continuous stream at 100 Msps with an AWG plugged directly into the ES analog front-end input.

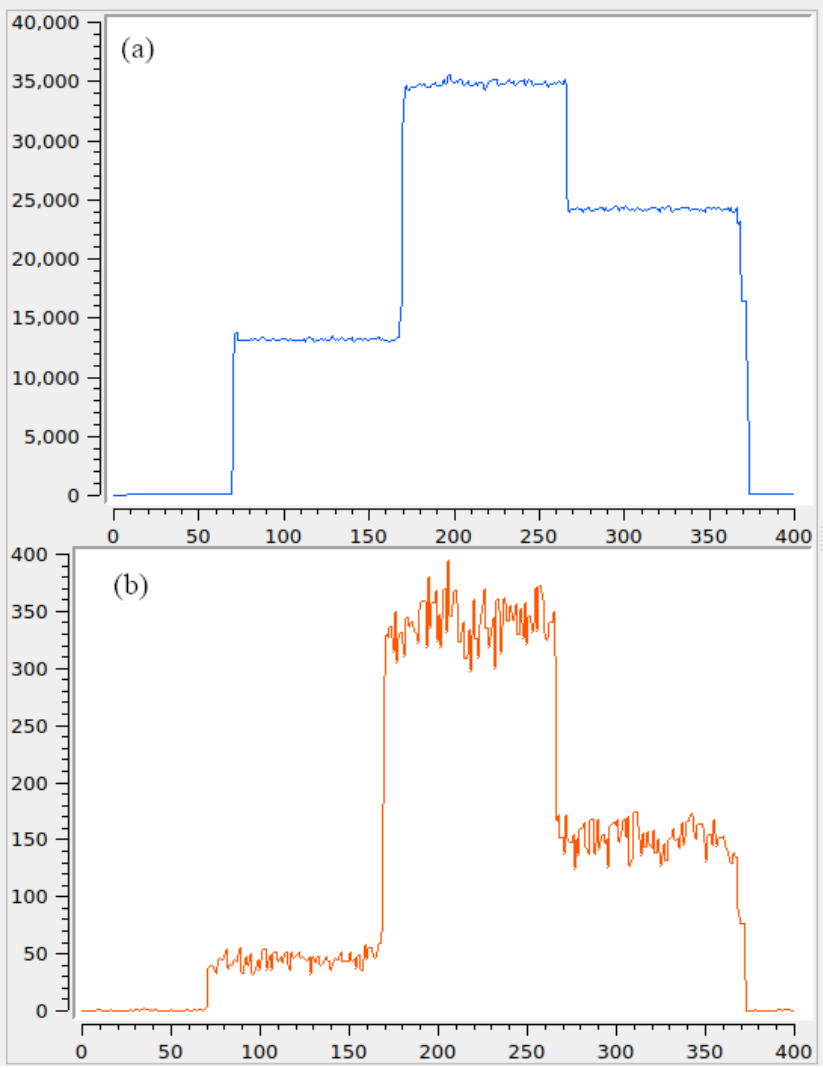

Figure 5.12: Screenshot from SerialPlot application shows the count-rates computed online by the performance profiling block. A Co-60 gamma source was used to stimulate the SiPM-based NaI(Tl) detector, while displaced at three fixed positions during the measurements. Event quantification rates are shown for (a) gamma and (b) neutrons, as discriminated in real-time by the developed system. In both cases horizontal axes represent time in seconds, whereas vertical axes express the count-rates for each class (γ/n) in counts/second. The highest average count-rate was 35.1 kcps, limited by the source activity. Neutron false positive rate remained below 1% during the experiment.

Additionally, aiming at measuring the maximum count-rate of the digital signal processing system and verifying the absolute minimum dead-time, the pulse shape of a γ event from a fast organic scintillator was synthesized using the double-exponential model with a single 50 ns decay time [253]. These pulses were continuously reproduced with the AWG at diverse rates adding 10% of white Gaussian noise, plugged to the AFE input of the hardware board.

The ϕ safeguard parameter (from Section 5.5) was removed in the FPGA configuration for this test, unlocking the maximum event rate of the system for fast detectors. The highest measured figures yielded 1.01 Mcps with an overall dead-time lower than 995 ns and an average 129 ppm error in counting sensitivity: on pair with the event rate limit reported by [20] and the lowest dead-time among the reviewed SOA. Values higher than 1.01 Mcps resulted in a drastic amount of missed events, setting the absolute maximum rate for this system clock frequency without further optimization.

These impressive metrics evidence the advantages of the FPGA reconfigurability and the actual capabilities of the flexible ES design when paired with faster scintillators. Naturally, the ML model requires proper training for different pulse shapes to accurately discriminate γ/n on each type of detector.

5.7.2 FPGA resource utilization

Although the ES hardware board is based on a low-end FPGA part, the firmware design (from Section 5.3.1) demanded less than one third of the total resources available. Table 5.1 summarizes the AMD Vivado post-implementation report, where each relevant design element is described, including the overall occupation of the firmware. Even though some components such as the UART interface controller are not representative and were not detailed in the table, they were accounted for the overall utilization.

As in Chapters 3 and 4, the total occupation suggests possibility to develop multi-

Design element	LUTs	Registers	BRAMs	DSPs
Available resources	20800	41600	50	90
Overall utilization	30.4%	17.4%	24.0%	17.8%
Feature Ext. + CDC	8.6%	3.9%	8.0%	0.0%
MicroBlaze + memory	6.4%	2.5%	16.0%	3.3%
Performance profiling	0.9%	1.5%	0.0%	0.0%
ML classification	11.2%	7.0%	0.0%	14.4%
ComBlock	0.8%	0.7%	0.0%	0.0%

Table 5.1: FPGA resource utilization per firmware design element

detector systems on SoC/FPGA platforms, useful in multichannel designs like nuclear fusion diagnostics.

5.7.3 SWaP validation

The SWaP parameters were measured using common laboratory instruments with the reported digit significance, leading to the following results: the external case volume is **0.56 liters**, while the overall weight of the system is **0.41 kilograms**. Power consumption was measured using an off-the-shelf USB power meter during a processing stress-test session, evidencing an average draw of **1.49 Watt** when performing real-time discrimination on a stream of piled-up events at the maximum count-rate.

5.7.4 Comparison with SOA

A broad comparison of the ES SWaP and performance metrics with SOA is carried out, providing overall quantitative scores for each indicator.

In terms of low-SWaP, Table 5.2 summarizes the parameters of recent CLYC-based γ/n discrimination implementations optimized for the same SWaP reduction goal. Size is measured in liters, weight in kilograms, and power consumption in Watts.

An overall SWaP score could not be computed due to the lack of information in one field for most of the reference works. In this regard, size and weight score (SaW) as well as weight and power draw score (WaP) were defined individually: SaW is the inverse of the product of weight and size ($1/SW$), while WaP is computed as the inverse product of weight and power ($1/WP$). In both cases the highest value serves as the reference for normalization and represents the best score in the range [0, 1].

Work	S (L)	W (kg)	P (W)	SaW	WaP
This work	0.56	0.41	1.5	0.13	1.00
Huang [153]	4.08	4.9	8.8	0.00	0.01
Zhao [18]	3.34	-	3.5	-	-
[1]Mesick [221]	-	7.0	14	-	0.01
[1]Hardgrove [106]	-	3.4	9.6	-	0.02
Thermofischer [105]	0.18	0.16	-	1.00	-

Table 5.2: SWaP comparison of recent CLYC-based γ/n discrimination systems, including SaW and WaP scores

The ES presented in this work exhibits the highest WaP score and the lowest power consumption among the compared developments. Although [105] evidences the smallest and lightest housing, the lack on performance may be a deciding factor in demanding applications: their maximum neutron count-rate is 200 times lower than this development, while the maximum detectable gamma energy is about one-fourth.

Regarding performance metrics, this system was contrasted with recent works targeting real-time γ/n discrimination in high event-rate scenarios. A numerical score for pile-up rejection or recovery (PuP R/R) features is suggested due to their relevance in high count-rate applications, as detailed in the following categories. Score 1.00 is assigned to continuous R/R, where a system is able to continuously recover or reject

[1]These systems were optimized for reliability in outer-space environments, commonly requiring rugged (heavier and bigger) housings.

events, no matter the total amount of successive piled-up pulses; score 0.50 relates to implementations on which up to four piled-up pulses can be recovered or rejected; score 0.00 is assigned to systems without R/R capabilities. Equal grading is designated to rejection and recovery on each category, since the impact of rejection will be reflected in the count-rate performance of the system.

Table 5.3 and Figure 5.13 quantitatively summarize the performance features of each work: minimum sustained count-rate (CR), inverse of maximum dead-time (1/DT), discrimination accuracy (Acc), and pile-up rejection or recovery capabilities (PuP R/R) rating.

An overall performance score (OP) is also defined, computed as the product of the evaluated parameters, further normalized in the range [0,1] respect to the highest score. This metric serves as an integral performance grade involving all the measured aspects. Values in Figure 5.13 are normalized to facilitate the visualization in a unit circle representation.

Work	CR (kcps)	1/DT (s^{-1})	Acc (%)	PuP R/R	OP
This work	200	400×10^3	98.2	1.00	1.00
[2]Michels [20]	1100	130×10^3	98.2	0.50	0.89
[3]Wen [15]	25	375×10^3	100	1.00	0.12
[4]Cruz [21]	2000	83.3	100	0.50	0.01
[5]Astrain [22]	79	20×10^3	98.0	0.50	0.01

Table 5.3: Performance indicators and overall performance score

[2]Michels et al. were able to discriminate two piled-up events only. Triple or quadruple pile-ups could not be individually distinguished.

[3]Wen et al. reported several count-rate and accuracy values: the highest accuracy/count-rate pair was chosen to maximize their OPS score. Their accuracy metric was defined as the system's sensitivity to neutrons, not the overall accuracy. Continuous pile-up rejection is used in this work.

[4]For Cruz et al. discrimination accuracy is assumed as 100%, since no value was reported besides mentioning they obtained the expected γ/n spectra. Dead-time includes pulse processing and emissivity profile reconstruction.

[5]Astrain et al. demonstrated pile-up recovery of up to two successive events.

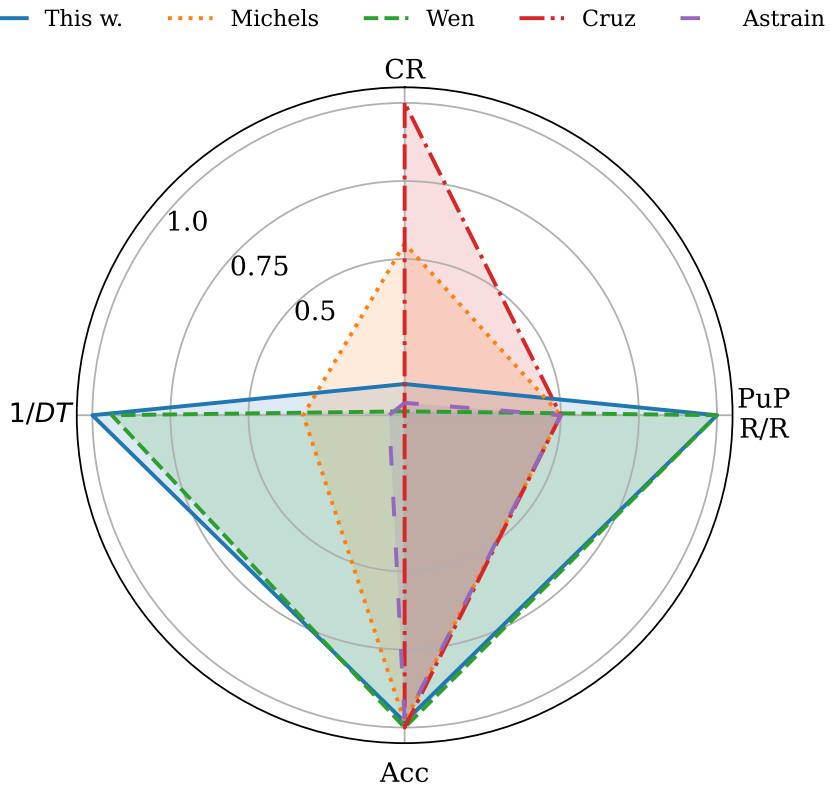

Figure 5.13: Performance comparison of real-time γ/n discrimination systems, based on measurements respect to the highest value per category. Parameters normalized within the [0,1] range, being the highest value the top performance indicator. *CR* defines the minimum γ/n discrimination sustained rate at which the reported overall accuracy is reached. $1/DT$ is the inverse of the maximum dead-time of each implementation. *Acc* represents the overall accuracy reported at the evaluated count-rate. *PuP R/R* is a score based on the implemented pile-up rejection or pile-up recovery method. The system developed for this thesis is compared against *Michels et al.* [20], *Wen et al.* [15], *Cruz et al.* [21], and *Astrain et al.* [22].

Notwithstanding the impressive low SWaP metrics achieved, the embedded instrumentation platform for mixed radiation fields presented in this thesis also exhibited the best overall performance (OP) score and the lowest dead-time among the evaluated works, evidencing the high efficiency of the system.

5.8 Discussion

The developed system demonstrated to be a low-SWaP embedded solution suitable for event discrimination in mixed radiation scenarios, featuring 98.2% overall constant accuracy at rates higher than 200 kcps in γ/n fields, accomplishing research objective V. The feature extraction mechanism paired to an efficient ML model in a pipeline guaranteed real-time discrimination, achieved uninterrupted pile-up recovery, and ensured a total dead-time lower than 2.5 µs with detection sensitivity higher than 99.9%.

These characteristics position this development in an overall high-performance range similar to that of the benchtop solution from [20], with the added portability optimizations enabled by the low-SWaP parameters described in [153]. As from the latest revised works, the ES presented in this thesis is the first in combining low-SWaP and real-time features into an device capable of continuous pile-up recovery based on a CLYC detector. It also outperforms in terms of count-rate and maximum energy range the commercial portable active γ/n dosimeter from [105]. SWaP reduction was achieved with a small and light yet robust enclosure, housing an ES designed with power consumption optimizations at multiple layers: hardware board, FPGA firmware, and ML model.

This implementation is complementary to existing γ/n discrimination methods, relying on precise data tagging for the ML model training and validation, as presented in Chapter 4. Employing a solid labeling method in combination with an optimized ML model was crucial to ensure an efficient real-time solution with high discrimination accuracy [143]. The flexibility of the embedded instrumentation platform was evidenced by adopting diverse types of detectors, where keeping reduced SWaP features not necessarily compels a performance compromise for pulses with faster decay times. Thus, other scintillators with discrimination capabilities such as TLYC, CLLB, NaI(Tl+Li), or CLLBC would potentially benefit from this work.

The FPGA firmware has also been experimentally verified with standard third-party

digital pulse processing blocks for nuclear instrumentation, including shaper, peak detector, histogram builder, and baseline restorer. Thus, using the design presented in this chapter as a starting point, independent γ/n spectroscopy under high count-rate scenarios might be performed by employing any typical technique for γ and n methods [237; 86], further appending the γ/n tag computed by the embedded system. Moreover, proper gamma and neutron total ambient dose equivalent rates can be conducted using standard calibration methods [254; 255].

5.9 Summary

The development presented in this chapter exploited contributions from previous chapters to create an embedded instrument devoted to event discrimination in mixed radiation fields under high event rates (research objective V), experimentally proven through γ/n identification. The flexibility provided by the SoC/FPGA architectures allowed to expand the system capabilities by building upon an existing hardware board, enabling considerable improvements without compelling redesigns in this regard.

Compatibility with diverse type of detectors was tested, demanding only firmware updates to match the ML model with the expected pulse shapes to perform event discrimination. Thus, nuclear and high-energy physics applications can benefit from this development [6; 53; 59; 60; 125], by leveraging the pulse discrimination capabilities under severe pile-up scenarios.

Reliable identification at high count-rates even on low-energy events was possible by collectively utilizing features reported in previous chapters, carefully optimized to achieve real-time operation in a low-power platform. The specific goals that aided in the realization of this ultimate research objective are summarized as follows:

- The versatile diagnostics conducted with the remote platform from Chapter 2 (research objective I), valuable in the experimental setup preparation and prelimi-

nary data recording stages.

- The efficient pulse shape event recognition from Chapter 3 (research objective II) for the data wrangling process, tuned with a filter matching the detector features.
- The reliable γ/n classification using experimental data (research objective III) for the ML model training, performed with the development reported in Chapter 4 .

Hence, this chapter consolidates the developments envisioned and validated in this thesis to produce an embedded instrumentation platform for mixed radiation fields based on FPGA. Potential applications of the built system include (but are not limited to) fusion reactor diagnostics, radiological contamination surveying, nuclear threat detection, space exploration, HEP experiments, and portable active dosimetry.

Chapter 6

Conclusions

In this thesis an embedded instrumentation platform for event discrimination in mixed radiation fields is presented. The requirements of the implementation demanding high-throughput and low-latency signal processing were fulfilled with SoC/FPGA technologies, enabling efficient execution of custom algorithms optimized for real-time hardware deployments. The embedded system was tested for γ/n event discrimination relying on an efficient ML real-time model running on a FPGA. Nevertheless, such development can be adapted to other mixed radiation applications with minimal hardware modifications, mostly depending on fine tuning the firmware components according to the specific pulse shape parameters to process.

A series of contributions were carried out aiming at systematically conceiving the embedded instrument elements. Besides making part of the overall system, these developments exhibit important advancements by their own to the state-of-the-art in nuclear and HEP instrumentation. First, an open-source tool to diagnose SoC/FPGA-based data acquisition systems was created to enable faster development cycles in experimental setups involving remote radiation measurements [1], accomplishing research objective I and aiding in the fulfillment of research objectives II, III, and V.

Next, a fast event detection system for nuclear and HEP experiments was envisioned and deployed in a SoC/FPGA platform [2]. An efficient alternative to a correlation

(matching) filter was tested, demonstrating superior efficiency on resource utilization with on-pair detection performance and reduced latency compared to conventional correlation methods. Reliable event identification in low SNR scenarios was also evidenced, enabling single event detection of lower energies. This system satisfies the research objective II and contributes in the accomplishment of objective V.

Then, enhanced event discrimination in mixed radiation fields was tackled by developing an algorithm optimized for SoC/FPGA implementations. This method analyzes features in the frequency domain of different pulse shapes in γ/n commercial detectors [3]. Better performance was obtained compared to PSD, significantly improving event separation, being even more relevant at lower energies. These results demonstrated to complete the requirements to attain research objective III. Moreover, aiming at accomplishing research objective IV, a public dataset of tagged γ/n raw pulse shapes was released [4], opening the possibility for other researchers to explore potential improvements in such challenging endeavour. These developments demonstrated a reliable methodology further employed to generate a labeled experimental dataset, crucial in the training stage of the ML model utilized to achieve the last research goal.

Finally, the embedded instrumentation platform for mixed radiation fields grasps the reliable event tagging from objective III and overcomes its inherent limitations under intense radiation activity [5], accomplishing research objective V. This system exploits the potential of an optimized real-time machine learning model running on a low-power FPGA-based custom board, generated through an in-house workflow to efficiently compress and deploy supervised ML architectures on SoC/FPGA [8]. Both achievements from research objectives I and II assisted in the fulfilment of research objective V as well, providing a fast diagnostics system during data recording sessions and accurately identifying low-energy events in the data curation process prior to the ML model training, respectively. The documented embedded platform demonstrated remarkable discrimination performance in a γ/n application, efficient

operation through the measured SWaP parameters, and flexibility to adopt diverse type of detectors for event classification in other mixed radiation scenarios in nuclear and HEP applications.

6.1 Future work

In Chapter 5, fast neutron and gamma spectroscopy were discussed as a possibility by adapting known spectrum reconstruction techniques. However, scintillators enriched with ^{6}Li (such as CLYC) evidence much higher detection efficiency to thermal neutrons than to fast neutrons, due to the high cross section of the neutron capture with the ^{6}Li nuclei [93]. This reaction is exothermic, releasing a pair of charged particles, yielding a noticeable event cluster close to 3.2 MeVee [89], as detailed in Chapter 1, severely intervening in the fast neutron spectrum deconvolution process.

This issue is usually overcome by filtering out information related to the thermal neutrons: either by physically suppressing their arrival to the detector with boron or cadmium shielding [34], using crystals enriched with ^{7}Li (CLYC-7) [256; 257; 258], or by applying energy cuts around the 3.2 MeVee cluster [259]. However, if an experiment demands the information retrieval from thermal and fast neutron events, two detectors must be used simultaneously to run over the limitations [260].

Nonetheless, enabling CLYC-6 detectors for simultaneous thermal neutron counting and spectroscopy of fast neutrons and gamma rays may be possible [134]. This innovation would significantly reduce the detector count or SWaP requirements in constrained applications, such as space exploration and portable instruments. Some studies have been recently performed in this regard without successful results [86; 15; 236; 256; 261; 262], whereas some of them suggest machine learning as a viable alternative to deal with this challenge [263]. By taking advantage of the flexibility provided in the presented embedded instrumentation platform, such application might be realized with a fine-tuned pulse leading edge extraction [84; 134], followed by a

classification system based on a ML model, similarly to what has been carried out in Chapter 5.

Another possible application for the flexible embedded instrumentation platform would be adapting a water Cherenkov detector for cosmic ray studies. Since diverse particle types leave distinguishable leading edges, a drop-in replacement of the ML model along with a parameter tuning of the feature extraction system presented in Chapter 5 would be enough to accurately discriminate muon and electron events [6].

In terms of γ/n discrimination with frequency-domain analysis, the simulation in Chapter 4 showed a promising possibility to improve the performance in lower energy ranges for plastic detectors. Conducting an experiment under stimuli of a mixed radiation field would validate these results, broadening the method to fast organic scintillators with the support of experimental data.

Bibliography

[1] I. R. Morales, M. L. Crespo, S. Carrato, Open Source Remote Diagnostics Platform for Custom Instrumentation in Nuclear Applications, in: Applications in Electronics Pervading Industry, Environment and Society, Springer Nature Switzerland, Cham, 2024, pp. 424–430. doi:10.1007/978-3-031-48121-5_61.

[2] A. Cicuttin, I. R. Morales, M. L. Crespo, S. Carrato, L. G. García, R. S. Molina, B. Valinoti, J. Folla Kamdem, A Simplified Correlation Index for Fast Real-Time Pulse Shape Recognition, Sensors 22 (20) (2022). doi:10.3390/s22207697.

[3] I. R. Morales, M. L. Crespo, M. Bogovac, A. Cicuttin, K. Kanaki, S. Carrato, Gamma/neutron classification with SiPM CLYC detectors using frequency-domain analysis for embedded real-time applications, Nuclear Engineering and Technology 56 (2024) 745–752. doi:10.1016/j.net.2023.11.013.

[4] I. R. Morales Argueta. Gamma and neutron tagged dataset from CLYC SiPM detector [online] (jun 2023). Accessed: 2024-08-12. doi:10.5281/zenodo.8037058.

[5] I. R. Morales, R. S. Molina, M. Bogovac, N. Jovalekic, M. L. Crespo, K. Kanaki, G. Ramponi, S. Carrato, Gamma/neutron online discrimination based on machine learning with CLYC detectors, IEEE Transactions on Nuclear Science (2024). doi:10.1109/TNS.2024.3498321.

[6] L. G. García Ordóñez, R. S. Molina, I. R. Morales Argueta, M. L. Crespo, A. Cicuttin, S. Carrato, G. Ramponi, H. E. Pérez Figueroa, M. G. Ballina Escobar, Pulse

Shape Discrimination for Online Data Acquisition in Water Cherenkov Detectors Based on FPGA/SoC, in: Proceedings of 37th International Cosmic Ray Conference — PoS(ICRC2021), Sissa Medialab, 2021, p. 274. doi:10.22323/1.395.0274.

[7] R. S. Molina, L. G. Garcia, I. R. Morales, M. L. Crespo, G. Ramponi, S. Carrato, A. Cicuttin, H. Perez, Compression of NN-based pulse-shape discriminators in front-end electronics for particle detection, in: International Conference on Applications in Electronics Pervading Industry, Environment and Society, Springer, 2021, pp. 93–99. doi:10.1007/978-3-030-95498-7_13.

[8] R. S. Molina, I. R. Morales, M. L. Crespo, V. G. Costa, S. Carrato, G. Ramponi, An End-to-End Workflow to Efficiently Compress and Deploy DNN Classifiers on SoC/FPGA, IEEE Embedded Systems Letters 16 (3) (2024) 255–258. doi:10.1109/LES.2023.3343030.

[9] D. Prelipcean, G. Lerner, I. Slipukhin, D. Lucsanyi, H. Sandberg, J. Storey, P. Martin-Holgado, A. Romero-Maestre, Y. Morilla García, R. García Alía, Towards a timepix3 radiation monitor for the accelerator mixed radiation field: Characterisation with protons and alphas from 0.6 mev to 5.6 mev, Applied Sciences 14 (2) (2024) 624. doi:10.3390/app14020624.

[10] P. Homola, D. Beznosko, G. Bhatta, Ł. Bibrzycki, M. Borczyńska, Ł. Bratek, N. Budnev, D. Burakowski, D. E. Alvarez-Castillo, K. Almeida Cheminant, et al., Cosmic-ray extremely distributed observatory, Symmetry 12 (11) (2020) 1835. doi:10.3390/sym12111835.

[11] S. Min, B. Seo, C. Roh, S. Hong, J. Cheong, Phoswich detectors in sensing applications, Sensors 21 (12) (2021) 4047. doi:10.3390/s21124047.

[12] A. Giaz, L. Pellegri, F. Camera, N. Blasi, S. Brambilla, S. Ceruti, B. Million, S. Riboldi, C. Cazzaniga, G. Gorini, M. Nocente, A. Pietropaolo, M. Pillon, M. Rebai,

M. Tardocchi, The CLYC-6 and CLYC-7 response to γ-rays, fast and thermal neutrons, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 810 (2016) 132–139. doi:10.1016/j.nima.2015.11.119.

[13] D. Nitz, New Electronics for the Surface Detectors of the Pierre Auger Observatory, in: Proceedings of 36th International Cosmic Ray Conference — PoS(ICRC2019), Vol. 358, 2019, p. 370. doi:10.22323/1.358.0370.

[14] B. Valinoti, W. F. Samayoa, L. Garcia, M. Crespo, A. Cicuttin, S. Levorato, F. Tessarotto, A. Bressan, I. Konorov, G. Pastuszak, et al., A SoC-FPGA based readout platform for the free-running AMBER data acquisition system, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment (2024) 169546doi:10.1016/j.nima.2024.169546.

[15] X. Wen, A. Enqvist, Pulse shape discrimination of Cs2LiYCl6:Ce3+ detectors at high count rate based on triangular and trapezoidal filters, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 866 (2017) 129–133. doi:10.1016/j.nima.2017.06.007.

[16] A. Dutta, P. Chandhran, K. E. Holbert, E. B. Johnson, Using decay time to discriminate neutron and gamma ray pulses from a CLYC detector, in: 2015 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC), IEEE, 2015, pp. 1–7. doi:10.1109/NSSMIC.2015.7581902.

[17] F. Fahim, B. Hawks, C. Herwig, J. Hirschauer, S. Jindariani, N. Tran, L. P. Carloni, G. D. Guglielmo, P. Harris, J. Krupa, D. Rankin, M. B. Valentin, J. Hester, Y. Luo, J. Mamish, S. Orgrenci-Memik, T. Aarrestad, H. Javed, V. Loncar, M. Pierini, A. A. Pol, S. Summers, J. Duarte, S. Hauck, S.-C. Hsu, J. Ngadiuba, M. Liu, D. Hoang,

E. Kreinar, Z. Wu, hls4ml: An Open-Source Codesign Workflow to Empower Scientific Low-Power Machine Learning Devices (2021). arXiv:2103.05579, doi:10.48550/arXiv.2103.05579.

[18] K. Zhao, Y. Huang, C. Shan, C. Feng, Z. Wang, K. Chang, W. Xie, F. Liu, S. Liu, A low-power neutron and gamma-ray detector for environmental radiation monitoring using CLYC scintillator, Journal of Instrumentation 18 (2023) P09043. doi:10.1088/1748-0221/18/09/P09043.

[19] Q. Hu, H. Ye, G. Jin, A SiPM-Based Wireless Neutron-Gamma Spectrometer for Remote Monitoring, in: 2023 6th International Conference on Electronics Technology (ICET), IEEE, 2023, pp. 955–959. doi:10.1109/ICET58434.2023.10212084.

[20] N. M. Michels, A. J. Jinia, S. D. Clarke, H.-S. Kim, S. A. Pozzi, D. D. Wentzloff, Real-Time Classification of Radiation Pulses With Piled-Up Recovery Using an FPGA-Based Artificial Neural Network, IEEE Access 11 (2023) 78074–78083. doi:10.1109/ACCESS.2023.3298208.

[21] N. Cruz, B. Santos, A. Fernandes, P. F. Carvalho, J. Sousa, B. Gonçalves, M. Riva, C. Centioli, D. Marocco, B. Esposito, C. M. B. Correia, R. C. Pereira, The Design and Performance of the Real-Time Software Architecture for the ITER Radial Neutron Camera, IEEE Transactions on Nuclear Science 66 (7) (2019) 1310–1317. doi:10.1109/TNS.2019.2907056.

[22] M. Astrain, M. Ruiz, A. V. Stephen, R. Sarwar, A. Carpeño, S. Esquembri, A. Murari, F. Belli, M. Riva, Real-Time Implementation of the Neutron/Gamma Discrimination in an FPGA-Based DAQ MTCA Platform Using a Convolutional Neural Network, IEEE Transactions on Nuclear Science 68 (8) (2021) 2173–2178. doi:10.1109/TNS.2021.3090670.

[23] A. Tootell, K. Szczepura, P. Hogg, An overview of measuring and modelling dose and risk from ionising radiation for medical exposures, Radiography 20 (4) (2014) 323–332. doi:10.1016/j.radi.2014.05.002.

[24] M. Hossain, Output trends, characteristics, and measurements of three megavoltage radiotherapy linear accelerators, Journal of Applied Clinical Medical Physics 15 (4) (2014) 137–151. doi:10.1120/jacmp.v15i4.4783.

[25] J. Marques, Evolution of nuclear fission reactors: Third generation and beyond, Energy Conversion and Management 51 (9) (2010) 1774–1780. doi:10.1016/j.enconman.2009.12.043.

[26] R. P. Taleyarkhan, Monitoring neutron radiation in extreme gamma/X-ray radiation fields, Sensors 20 (3) (2020) 640. doi:10.3390/s20030640.

[27] H. Pant, A. Kundu, K. Nigam, Radiotracer applications in chemical process industry, Reviews in chemical engineering 17 (3) (2001) 165–252. doi:10.1515/REVCE.2001.17.3.165.

[28] N. Othman, S. Kamarudin, Radiotracer technology in mixing processes for industrial applications, The Scientific World Journal 2014 (1) (2014) 768604. doi:10.1155/2014/768604.

[29] S. M. Alam, R. Ansari, M. A. Khan, Application of Radioisotopes and Radiation in the Field of Agriculture: Review, Online J. Biol. Sci 1 (2001) 82–86. doi:10.3923/jbs.2001.82.86.

[30] A. Bakri, K. Mehta, D. R. Lance, Sterilizing Insects with Ionizing Radiation, Springer Netherlands, Dordrecht, 2005, pp. 233–268. doi:10.1007/1-4020-4051-2_9.

[31] M. Claessens, ITER: The giant fusion reactor, Springer, 2020. doi:10.1007/978-3-031-37762-4.

[32] O. Dendene, V. V. Afanasiev, S. V. Lushin, A. C. Chergui, A. A. Stifutkin, L. Boukerdja, Study of silicon photomultipliers (SiPM) with organic scintillator for neutron diagnostics of thermonuclear plasma, Fusion Engineering and Design 178 (2022) 113113. doi:10.1016/j.fusengdes.2022.113113.

[33] M. Smith, Y. Akatov, H. Andrews, V. Arkhangelsky, I. Chernykh, H. Ing, N. Khoshooniy, B. Lewis, R. Machrafi, I. Nikolaev, et al., Measurements of the neutron dose and energy spectrum on the International Space Station during expeditions ISS-16 to ISS-21, Radiation protection dosimetry 153 (4) (2013) 509–533. doi:10.1093/rpd/ncs129.

[34] W. Zhou, T. Cui, Z. Zhang, Y. Yang, H. Yi, D. Hou, Measurement of wide energy range neutrons with a CLYC(Ce) scintillator, Journal of Instrumentation 18 (2023) P02014. doi:10.1088/1748-0221/18/02/P02014.

[35] G. F. Knoll, Radiation detection and measurement, 4th Edition, John Wiley & Sons, 2010.

[36] G. V. Riley, J. E. Barney, Low Size, Weight, Power and Cost (SWaP-C) Neutron Detection for Space Missions, Tech. rep., Los Alamos National Laboratory (LANL), Los Alamos, NM (United States) (4 2023). doi:10.2172/1973791.

[37] M. Laubenstein, I. Lawson, Low background radiation detection techniques and mitigation of radioactive backgrounds, Frontiers in Physics 8 (2020) 577734. doi:10.3389/fphy.2020.577734.

[38] N. Rachkova, I. Shuktomova, A. Taskaev, The state of natural radionuclides of uranium, radium, and thorium in soils, Eurasian Soil Science 43 (2010) 651–658. doi:10.1134/S1064229310060050.

[39] I. Hajdas, P. Ascough, M. H. Garnett, S. J. Fallon, C. L. Pearson, G. Quarta,

K. L. Spalding, H. Yamaguchi, M. Yoneda, Radiocarbon dating, Nature Reviews Methods Primers 1 (1) (2021) 62. doi:10.1038/s43586-021-00058-7.

[40] S. Miyamoto, S. Nagahara, Tomographic Imaging of Volcano Structures with Cosmic-Ray Muons, Muography: Exploring Earth's Subsurface with Elementary Particles (2022) 19–32doi:10.1002/9781119722748.ch2.

[41] K. Bernloehr. Cosmic Ray / Gamma Ray / neutrino and similar experiments [online] (Feb 2020). Accessed: 2024-09-10.

[42] O. Brüning, L. Rossi, The High Luminosity Large Hadron Collider–HL-LHC, in: THE HIGH LUMINOSITY LARGE HADRON COLLIDER: New Machine for Illuminating the Mysteries of the Universe, World Scientific, 2024, pp. 1–53. doi:10.1142/9789811278952_0001.

[43] E. Karantzoulis, A. Fabris, S. Krecic, The Elettra 2.0 Project, in: Proc. IPAC'22, no. 13 in International Particle Accelerator Conference, JACoW Publishing, Geneva, Switzerland, 2022, pp. 1459–1462. doi:10.18429/JACoW-IPAC2022-TUPOMS023.

[44] J. Lu, X. Tuo, H. Yang, Y. Luo, H. Liu, C. Deng, Q. Wang, Pulse-Shape Discrimination of SiPM Array-Coupled CLYC Detector Using Convolutional Neural Network, Applied Sciences 12 (2022) 2400. doi:10.3390/app12052400.

[45] B. W. Blackburn, J. T. Johnson, S. M. Watson, D. L. Chichester, J. L. Jones, F. H. Ruddy, J. G. Seidel, R. W. Flammang, Fast digitization and discrimination of prompt neutron and photon signals using a novel silicon carbide detector, in: T. T. Saito, D. Lehrfeld, M. J. DeWeert (Eds.), Optics and Photonics in Global Homeland Security III, Vol. 6540, International Society for Optics and Photonics, SPIE, 2007, p. 65401J. doi:10.1117/12.722921.

[46] X. Hong, H. Wang, J. Zhou, X. Yang, M. Wang, Y. Ma, W. Zhou, Y. Liu, X. Zhu, Peak tailing cancellation techniques for digital CR-(RC)n filter, Applied Radiation and Isotopes 167 (2021) 109471. doi:10.1016/j.apradiso.2020.109471.

[47] CAEN, DPP-PSD, Accessed: 2024-09-20 (Nov 2023).
URL https://www.caen.it/products/dpp-psd/

[48] Z. Liu, J. Peters, J.-I. Kim, S. Das, K. M. McCall, B. W. Wessels, Y. He, W. Lin, M. G. Kanatzidis, Noise sources and their limitations on the performance of compound semiconductor hard radiation detectors, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 916 (2019) 133–140. doi:10.1016/j.nima.2018.11.013.

[49] CAEN, CAEN DT5761 Digitizer, https://www.caen.it/products/dt5761/, Accessed: 2024-03-12 (2012).

[50] O. McCormack, L. Giacomelli, G. Croci, A. Muraro, G. Gorini, G. Grosso, R. Pasqualotto, E. P. Cippo, M. Rebai, D. Rigamonti, M. Tardocchi, Characterization and operational stability of EJ276 plastic scintillator-based detector for neutron spectroscopy, Journal of Instrumentation 16 (2021) P10002. doi:10.1088/1748-0221/16/10/P10002.

[51] R. Winyard, J. Lutkin, G. McBeth, Pulse shape discrimination in inorganic and organic scintillators. I, Nuclear Instruments and Methods 95 (1) (1971) 141–153. doi:10.1016/0029-554X(71)90054-1.

[52] V. Galkin, A. Anokhina, E. Konishi, A. Misaki, A Discrimination Procedure between Muon and Electron in Superkamiokande Experiment Based on the Angular Distribution Function Method, arXiv preprint hep-ex/0501058 (2005). doi:10.48550/arXiv.hep-ex/0501058.

[53] L. G. Garcia, R. S. Molina, M. L. Crespo, S. Carrato, G. Ramponi, A. Cicuttin, I. R. Morales, H. Perez, Muon–Electron Pulse Shape Discrimination for Water Cherenkov Detectors Based on FPGA/SoC, Electronics 10 (2021) 224. doi:10.3390/electronics10030224.

[54] E. Bodewits, D. Cester, M. Lunardon, S. Moretto, P. Schotanus, L. Stevanato, G. Viesti, Characterization of a large area ZnS (Ag) detector for gross alpha and beta activity measurements in tap water plants, IEEE Transactions on Nuclear Science 63 (3) (2016) 1565–1569. doi:10.1109/TNS.2015.2495207.

[55] M.-Y. Cui, F. Alemanno, C. Altomare, Q. An, P. Azzarello, F. C. T. Barbato, P. Bernardini, X. J. Bi, I. Cagnoli, M. S. Cai, E. Casilli, E. Catanzani, J. Chang, D. Y. Chen, J. L. Chen, Z. F. Chen, Z. X. Chen, P. Coppin, M. Y. Cui, T. S. Cui, Y. X. Cui, I. De Mitri, F. de Palma, A. Di Giovanni, M. Di Santo, Q. Ding, T. K. Dong, Z. X. Dong, G. DONVITO, D. Droz, J. L. Duan, K. K. Duan, R. R. Fan, Y. Z. Fan, F. Fang, K. Fang, C. Q. Feng, L. Feng, M. Fernandez Alonso, J. M. Frieden, P. Fusco, M. Gao, F. Gargano, E. Ghose, K. Gong, Y. Z. Gong, D. Y. Guo, J. H. Guo, S. X. Han, Y. M. Hu, G. S. Huang, X. Y. Huang, Y. Y. Huang, M. Ionica, L. Y. Jiang, W. Jiang, Y. Z. Jiang, J. Kong, A. Kotenko, D. Kyratzis, S. J. Lei, W. L. Li, W. H. Li, X. Li, X. Q. Li, Y. M. Liang, C. M. Liu, H. Liu, J. Liu, S. B. Liu, Y. Liu, F. Loparco, C. N. Luo, M. Ma, P. X. Ma, T. Ma, X. Y. Ma, G. Marsella, M. N. Mazziotta, D. Mo, X. Y. Niu, X. Pan, A. Parenti, W. X. Peng, X. Y. Peng, C. Perrina, E. Putti-Garcia, R. Qiao, J. N. Rao, A. Ruina, Z. Shangguan, W. H. Shen, Z. Q. Shen, Z. T. Shen, L. Silveri, J. X. Song, M. Stolpovskiy, H. Su, M. Su, H. R. Sun, Z. Y. Sun, A. Surdo, X. J. Teng, A. Tykhonov, J. Z. Wang, L. G. Wang, S. Wang, X. L. Wang, Y. F. Wang, Y. Wang, Y. Z. Wang, D. M. Wei, J. J. Wei, Y. F. Wei, D. Wu, J. Wu, L. B. Wu, S. S. Wu, X. Wu, Z. Q. Xia, E. H. Xu, H. T. Xu, J. Xu, Z. H. Xu, Z. Z. Xu, Z. L. Xu, G. F. Xue, H. B. Yang, P. Yang, Y. Q. Yang, H. J. Yao, Y. H. Yu, G. W. Yuan, Q. Yuan, C. Yue, J. J. Zang, S. X. Zhang, W. Z. Zhang, Y. Zhang, Y. P.

Zhang, Y. Zhang, Y. J. Zhang, Y. Q. Zhang, Y. L. Zhang, Z. Zhang, Z. Y. Zhang, C. Zhao, H. Y. Zhao, X. F. Zhao, C. Y. Zhou, Y. Zhu, A PCA-based Method for Electron-Proton Discrimination of the DAMPE Experiment, in: Proceedings of 38th International Cosmic Ray Conference — PoS(ICRC2023), Vol. 444, 2023, p. 131. doi:10.22323/1.444.0131.

[56] J. Köhler, B. Ehresmann, C. Zeitlin, R. Wimmer-Schweingruber, D. Hassler, G. Reitz, D. Brinza, J. Appel, S. Böttcher, E. Böhm, S. Burmeister, J. Guo, H. Lohf, C. Martin, A. Posner, S. Rafkin, Measurements of the neutron spectrum in transit to Mars on the Mars Science Laboratory, Life Sciences in Space Research 5 (2015) 6–12. doi:10.1016/j.lssr.2015.03.001.

[57] D. Wilkinson, The phoswich—a multiple phosphor, Review of Scientific Instruments 23 (8) (1952) 414–417. doi:10.1063/1.1746324.

[58] M. Tyagi, T. Patel, P. Q. Vuong, P. Sarkar, H. Kim, et al., Discrimination of Fast and Thermal Neutrons Using a Novel Phoswich Detector of LaCl 3 and LiI: Eu Single Crystal Scintillators, IEEE Transactions on Nuclear Science 70 (7) (2023) 1325–1330. doi:10.1109/TNS.2023.3271460.

[59] C. Frangville, G. H. Bertrand, F. Carrel, C. Lynde, V. Buridon, M. Hamel, Omniscinti™: Simultaneous detection of alpha/beta/gamma/fast and thermal neutron with a triple-layered phoswich detector, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 1068 (2024) 169790. doi:10.1016/j.nima.2024.169790.

[60] M. Grodzicka- Kobylka, T. Szczesniak, L. Swiderski, K. Brylew, M. Moszyński, J. Valiente-Dobón, P. Schotanus, K. Grodzicki, P. Mazerewicz, J. Szymanowski, Z. Mianowska, Comparison of detectors with pulse shape discrimination capability for simultaneous detection of gamma-rays, slow and fast neutrons,

Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 1019 (2021) 165858. doi:10.1016/j.nima.2021.165858.

[61] Q. Wei, S. Wang, T. Ma, L. Lu, T. Dai, Y. Liu, Design of a high resolution phoswich PET detector, in: 2011 IEEE Nuclear Science Symposium Conference Record, IEEE, 2011, pp. 2874–2876. doi:10.1109/NSSMIC.2011.6152508.

[62] I. Giomataris, R. De Oliveira, S. Andriamonje, S. Aune, G. Charpak, P. Colas, G. Fanourakis, E. Ferrer, A. Giganon, P. Rebourgeard, et al., Micromegas in a bulk, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 560 (2) (2006) 405–408. doi:10.1016/j.nima.2005.12.222.

[63] Z. Fang, Z. Zhang, B. Shi, W. Jiang, X. Liu, S. He, J. Chen, P. Cao, J. Liu, Y. Zhou, M. Shao, B. Qu, S. Zhang, Q. Wang, A large area, high counting rate micromegas-based neutron detector for BNCT, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 1053 (2023) 168361. doi:10.1016/j.nima.2023.168361.

[64] F. H. Ruddy, L. Ottaviani, A. Lyoussi, C. Destouches, O. Palais, C. Reynard-Carette, Performance and applications of silicon carbide neutron detectors in harsh nuclear environments, in: EPJ Web of Conferences, Vol. 253, EDP Sciences, 2021, p. 11003. doi:10.1051/epjconf/202125311003.

[65] M. H. KUSHORO, Development and Characterization of Silicon Carbide Neutron Detectors for Fusion Reactors, Ph.D. thesis, Università degli Studi di Milano-Bicocca (2024). doi:10281/457121.

[66] B. Valinoti, Real-time Features Extraction for Trigger-less Data Acquisition Sys-

tems in Particle PHysics Experiments, Ph.D. thesis, Università degli Studi di Trieste (2023). doi:11368/3059202.

[67] R. T. Kouzes, A. T. Lintereur, E. R. Siciliano, Progress in alternative neutron detection to address the helium-3 shortage, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 784 (2015) 172–175. doi:10.1016/j.nima.2014.10.046.

[68] J. Glodo, Y. Wang, R. Shawgo, C. Brecher, R. H. Hawrami, J. Tower, K. S. Shah, New Developments in Scintillators for Security Applications, Physics Procedia 90 (2017) 285–290. doi:10.1016/j.phpro.2017.09.012.

[69] R. Fronk, S. Bellinger, L. Henson, D. Huddleston, T. Ochs, T. Sobering, D. McGregor, High-efficiency microstructured semiconductor neutron detectors for direct 3He replacement, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 779 (2015) 25–32. doi:10.1016/j.nima.2015.01.041.

[70] D. Rezaei-Ochbelagh, Comparison of 3He and BF3 neutron detectors used to detect hydrogenous material buried in soil, Radiation Physics and Chemistry 81 (4) (2012) 379–382. doi:10.1016/j.radphyschem.2011.12.031.

[71] E. van Loef, G. Markosyan, U. Shirwadkar, M. McClish, K. Shah, Gamma-ray spectroscopy and pulse shape discrimination with a plastic scintillator, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 788 (2015) 71–72. doi:10.1016/j.nima.2015.03.077.

[72] D. Bazin, A. Mueller, W.-D. Schmidt-Ott, A NE213 liquid scintillator, neutron detector designed for lifetime measurements of very neutron-rich nuclei, Nuclear Instruments and Methods in Physics Research Section A: Accelerators,

Spectrometers, Detectors and Associated Equipment 281 (1) (1989) 117–127. doi:10.1016/0168-9002(89)91222-9.

[73] K. D. Ngo, C. Cazzaniga, M. Paoletti, D. Rigamonti, M. Kastriotou, C. Frost, M. Tardocchi, J. Sykora, S. Mann, B. Lutz, et al., Fast neutron response characterization of an EJ-276 plastic scintillator for use as a neutron monitor, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 1051 (2023) 168216. doi:10.1016/j.nima.2023.168216.

[74] K. Yang, P. R. Menge, V. Ouspenski, Li co-doped NaI: Tl (NaIL)—A large volume neutron-gamma scintillator with exceptional pulse shape discrimination, IEEE Transactions on Nuclear Science 64 (8) (2017) 2406–2413. doi:10.1109/TNS.2017.2721398.

[75] P. Menge, K. Yang, V. Ouspenski, Large Format Li Co-Doped Nai: Tl (Nail™) Scintillation Detector for Gamma-Ray and Neutron Dual Detection, in: Proceedings of the 12th Pacific Rim Conference on Ceramic and Glass Technology: Ceramic Transactions, Vol. 264, Wiley Online Library, 2018, pp. 201–208. doi:10.1002/9781119494096.ch20.

[76] M. Tao, Z. Wang, Q. Chen, F. Li, J. Qi, P. Qi, T. Gao, Q. Zhao, Z. Zhang, B. Zhu, et al., Design and performance of a NaIL detector for neutron/gamma discrimination, Journal of Instrumentation 16 (08) (2021) P08067. doi:10.1088/1748-0221/16/08/P08067.

[77] F. Liang, H. Brands, L. Hoy, J. Preston, J. Smith, Lithium-loaded scintillators coupled to a custom-designed silicon photomultiplier array for neutron and gamma-ray detection, IEEE Transactions on Nuclear Science 65 (8) (2018) 2162–2168. doi:10.1109/TNS.2018.2820658.

[78] T. Jin, S. Hao, Y. Shang, Z. Lei, C. Yang, Recent Trends in Elpasolite Single Crystal Scintillators for Radiation Detection, Crystals 12 (7) (2022). doi:10.3390/cryst12070887.

[79] R. Hawrami, E. Ariesanti, L. Soundara-Pandian, J. Glodo, K. S. Shah, Tl2LiYCl6:Ce: A New Elpasolite Scintillator, IEEE Transactions on Nuclear Science 63 (6) (2016) 2838–2841. doi:10.1109/TNS.2016.2627523.

[80] J. Henderson, R. Casperson, R. Hughes, C. Wu, Neutron response of the TLYC scintillator, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 993 (2021) 165050. doi:10.1016/j.nima.2021.165050.

[81] M. Smith, M. McClish, T. Achtzehn, H. Andrews, M. Baginski, D. Best, B. Budden, E. Clifford, N. Dallmann, C. Dathy, J. Frank, S. Graham, H. Ing, L. Stonehill, Assessment of photon detectors for a handheld gamma-ray and neutron spectrometer using Cs2LiYCl6:Ce (CLYC) scintillator, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 715 (2013) 92–97. doi:10.1016/j.nima.2013.03.023.

[82] M. C. Recker, E. J. Cazalas, J. W. McClory, J. E. Bevins, Comparison of SiPM and PMT Performance Using a $Cs_2LiYCl_6 : Ce^{3+}$ (CLYC) Scintillator With Two Optical Windows, IEEE Transactions on Nuclear Science 66 (8) (2019) 1959–1965. doi:10.1109/TNS.2019.2926246.

[83] A. M. Okowita, Characterization of Lithium-6 as a Commercial Helium-3 Alternative for Nuclear Safeguards and Security, Master's thesis, University of Tennessee - Knoxville (12 2014).
URL https://trace.tennessee.edu/utk_gradthes/3166

[84] N. D'Olympia, P. Chowdhury, E. Jackson, C. Lister, Fast neutron response of

6Li-depleted CLYC detectors up to 20MeV, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 763 (2014) 433–441. doi:10.1016/j.nima.2014.06.074.

[85] C. M. Whitney, L. Soundara-Pandian, E. B. Johnson, S. Vogel, B. Vinci, M. Squillante, J. Glodo, J. F. Christian, Gamma–neutron imaging system utilizing pulse shape discrimination with CLYC, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 784 (2015) 346–351. doi:10.1016/j.nima.2014.09.022.

[86] F. Ferrulli, M. Labalme, M. Silari, Investigation of CLYC-6 for thermal neutron detection and CLYC-7 for fast neutron spectrometry, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 1029 (2022) 166460. doi:10.1016/j.nima.2022.166460.

[87] B. Budden, L. Stonehill, N. Dallmann, M. Baginski, D. Best, M. Smith, S. Graham, C. Dathy, J. Frank, M. McClish, A Cs2LiYCl6:Ce-based advanced radiation monitoring device, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 784 (2015) 97–104. doi:10.1016/j.nima.2014.11.051.

[88] M. Bourne, C. Mussi, E. Miller, S. Clarke, S. Pozzi, A. Gueorguiev, Characterization of the CLYC detector for neutron and photon detection, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 736 (2014) 124–127. doi:10.1016/j.nima.2013.10.030.

[89] M. C. Recker, Enabling Mobile Neutron Detection Systems with CLYC, Ph.D. thesis, Air Force Institute of Technology (12 2019).

[90] P. A. Rodnyi, Core–valence luminescence in scintillators, Radiation Measurements 38 (2004) 343–352. doi:10.1016/j.radmeas.2003.11.003.

[91] M. Smith, T. Achtzehn, H. Andrews, E. Clifford, P. Forget, J. Glodo, R. Hawrami, H. Ing, P. O'Dougherty, K. Shah, U. Shirwadkar, L. Soundara-Pandian, J. Tower, Fast neutron measurements using Cs2LiYCl6:Ce (CLYC) scintillator, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 784 (2015) 162–167. doi:10.1016/j.nima.2014.09.021.

[92] J. Glodo, U. Shirwadkar, R. Hawrami, T. Achtzehn, H. R. Andrews, E. T. H. Clifford, H. Ing, V. D. Kovaltchouk, M. B. Smith, K. S. Shah, Fast Neutron Detection With Cs2LiYCl6, IEEE Transactions on Nuclear Science 60 (2013) 864–870. doi:10.1109/TNS.2012.2227499.

[93] A. Mentana, F. Camera, A. Giaz, N. Blasi, S. Brambilla, S. Ceruti, L. Gini, F. Groppi, S. Manenti, B. Million, et al., Measurement of fast neutron detection efficiency with 6Li and 7Li enriched CLYC scintillators, in: Journal of Physics: Conference Series, Vol. 763, IOP Publishing, 2016, p. 012006. doi:10.1088/1742-6596/763/1/012006.

[94] S. H. B.V. Thermal neutron detector V12.7B30/SIP-E3-CLYC-x datasheet - Scionix [online] (2017). Accessed: 2023-05-19.

[95] N. Dinar, D. Celeste, M. Silari, V. Varoli, A. Fazzi, Pulse shape discrimination of CLYC scintillator coupled with a large SiPM array, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 935 (2019) 35–39. doi:10.1016/j.nima.2019.04.099.

[96] S. West, D. Beckman, D. Coupland, N. Dallmann, C. Hardgrove, K. Mesick,

L. Stonehill, Compact readout of large CLYC scintillators with silicon photomultipler arrays, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 951 (2020) 162928. doi:10.1016/j.nima.2019.162928.

[97] T. Huang, Z. Zhang, Characterization of 1-inch CLYC scintillator coupled with 8 × 8 SiPM array, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 999 (2021) 165225. doi:10.1016/j.nima.2021.165225.

[98] T. Huang, Q. Fu, S. Lin, B. Wang, NaI(Tl) scintillator read out with SiPM array for gamma spectrometer, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 851 (2017) 118–124. doi:10.1016/j.nima.2017.01.068.

[99] D. D. Vita, L. Buonanno, F. Canclini, G. Ticchi, F. Camera, M. Carminati, C. Fiorini, A 144-SiPM 3" LaBr3 readout module for PMTs replacement in Gamma spectroscopy, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 1040 (2022) 167179. doi:10.1016/j.nima.2022.167179.

[100] J. Hartman, A. Barzilov, I. Novikov, Remote sensing of neutron and gamma radiation using aerial unmanned autonomous system, in: 2015 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC), IEEE, 2015, pp. 1–4. doi:10.1109/NSSMIC.2015.7581763.

[101] M. Altayeb, M. Zennaro, E. Pietrosemoli, TinyML Gamma Radiation Classifier, Nuclear Engineering and Technology 55 (2023) 443–451. doi:10.1016/j.net.2022.09.032.

[102] B. S. Budden, A. J. Couture, L. C. Stonehill, A. V. Klimenko, J. R. Terry, J. O.

Perry, Analysis of Cs2LiYCl6:Ce3+ (CLYC) waveforms as read out by solid state photomultipliers, in: 2012 IEEE Nuclear Science Symposium and Medical Imaging Conference Record (NSS/MIC), IEEE, 2012, pp. 347–350. doi:10.1109/NSSMIC.2012.6551123.

[103] R. Camposano, J. Wilberg, Embedded system design, Design Automation for Embedded Systems 1 (1996) 5–50. doi:10.1007/BF00134682.

[104] M. L. Crespo, A. Cicuttin, J. D. D. Gazzano, F. R. Calle, Reconfigurable virtual instrumentation based on fpga for science and high-education, in: Field-Programmable Gate Array (FPGA) Technologies for High Performance Instrumentation, IGI Global, 2016, pp. 99–123. doi:10.4018/978-1-5225-0299-9.ch005.

[105] Thermofischer, RadEyeTM GN+ Gamma Neutron pager, Accessed: 2024-04-13. URL https://www.thermofisher.com/order/catalog/product/4250631

[106] C. Hardgrove, R. Starr, I. Lazbin, A. Babuscia, B. Roebuck, J. DuBois, N. Struebel, A. Colaprete, D. Drake, E. Johnson, J. Christian, L. Heffern, S. Stem, S. Parlapiano, M. Wiens, A. Genova, D. Dunham, D. Nelson, B. Williams, J. Bauman, P. Hailey, T. OBrien, K. Marwah, L. Vlieger, J. Bell, T. Prettyman, T. Crain, E. Cisneros, N. Cluff, G. Stoddard, M. Kaffine, The Lunar Polar Hydrogen Mapper CubeSat Mission, IEEE Aerospace and Electronic Systems Magazine 35 (2020) 54–69. doi:10.1109/MAES.2019.2950747.

[107] I. G. Mitrofanov, M. L. Litvak, A. B. Varenikov, Y. N. Barmakov, A. Behar, Y. I. Bobrovnitsky, E. P. Bogolubov, W. V. Boynton, K. Harshman, E. Kan, A. S. Kozyrev, R. O. Kuzmin, A. V. Malakhov, M. I. Mokrousov, S. N. Ponomareva, V. I. Ryzhkov, A. B. Sanin, G. A. Smirnov, V. N. Shvetsov, G. N. Timoshenko, T. M. Tomilina, V. I. Tret'yakov, A. A. Vostrukhin, Dynamic Albedo of Neutrons (DAN) Experiment

Onboard NASA's Mars Science Laboratory, Space Science Reviews 170 (2012) 559–582. doi:10.1007/s11214-012-9924-y.

[108] P. McNelles, L. Lu, A review of the current state of FPGA systems in nuclear instrumentation and control, in: International Conference on Nuclear Engineering, Vol. 55836, American Society of Mechanical Engineers, 2013, p. V006T16A057. doi:10.1115/ICONE21-16819.

[109] K. S. Mannatunga, L. G. Ordóñez, M. B. Amador, M. L. Crespo, A. Cicuttin, S. Levorato, R. Melo, B. Valinoti, Design for Portability of Reconfigurable Virtual Instrumentation, in: 2019 X Southern Conference on Programmable Logic (SPL), 2019, pp. 45–52. doi:10.1109/SPL.2019.8714446.

[110] L. Buonanno, D. Di Vita, G. Ticchi, F. Canclini, F. Camera, M. Carminati, C. Fiorini, Real-time Imaging with Thick LaBr3: FPGA-Embedded Machine and Deep Learning for Nuclear Physics, in: 2021 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC), 2021, pp. 1–3. doi:10.1109/NSS/MIC44867.2021.9875479.

[111] L. H. Arnaldi, D. Cazar, M. Audelo, I. Sidelnik, The new data acquisition system of the LAGO Collaboration based on the Redpitaya board, in: 2020 Argentine Conference on Electronics (CAE), IEEE, 2020, pp. 87–92. doi:10.1109/CAE48787.2020.9046374.

[112] L. G. García Ordóñez, M. L. Crespo, S. Carrato, A. Cicuttin, W. O. Florian Samayoa, D. D'Ago, S. Levorato, Multichannel Time Synchronization Based on PTP through a High Voltage Isolation Buffer Network Interface for Thick-GEM Detectors, Instruments 6 (1) (2022) 11. doi:10.3390/instruments6010011.

[113] W. O. Florian Samayoa, HyperFPGA: SoC-FPGA Cluster Architecture for Su-

percomputing and Scientific applications, Ph.D. thesis, Università degli Studi di Trieste (2024). doi:20.500.14242/63181.

[114] L. H. Crockett, R. A. Elliot, M. A. Enderwitz, R. W. Stewart, The Zynq book: embedded processing with the ARM Cortex-A9 on the Xilinx Zynq-7000 all programmable SoC, Strathclyde Academic Media, 2014.

[115] F. Giordano. Novel acquisition systems and SMART FPGA programmability [online] (Nov 2020). Accessed: 2024-10-16.

[116] A. Taylor, D. Binnun, S. Srivastava, A Hands-On Guide to Designing Embedded Systems, Artech House, 2021.

[117] ADS4149 - 14-Bit, 250-MSPS Analog-to-Digital Converter (ADC) [online] (2016). Accessed: 2024-10-21.

[118] L. G. García Ordoñez, High Voltage Power Supply System and Front-End DAQ on FPGA/SoC for High Energy Particle Detectors, Ph.D. thesis, Università degli Studi di Trieste (2022). doi:11368/3030919.

[119] M. Romer, S. Murray, J. Schmitz, S. Balkır, M. Hoffman, A low-power analog front-end amplifier for SiPM based radiation detectors, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 1048 (2023) 167897. doi:10.1016/j.nima.2022.167897.

[120] Nuclear Science and Instrumentation Newsletter No. 1, February 2021, INTERNATIONAL ATOMIC ENERGY AGENCY, 2021.

[121] A. Bressan, S. Carrato, C. Chatterjee, A. Cicuttin, M. Crespo, D. D'Ago, S. Dalla Torre, S. Dasgupta, L. G. Ordóñez, M. Gregori, et al., The high voltage system of the novel MPGD-based photon detectors of COMPASS RICH-1 and its development towards a scalable High Voltage Power Supply System

for MPGDs, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 1056 (2023) 168558. doi:10.1016/j.nima.2023.168558.

[122] Digilent. ZedBoard Hardware User's Guide [online] (9 2014). Accessed: 2024-09-11.

[123] V. T. Jordanov, G. F. Knoll, Digital synthesis of pulse shapes in real time for high resolution radiation spectroscopy, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 345 (2) (1994) 337–345. doi:10.1016/0168-9002(94)91011-1.

[124] A. Karathanasis, B. Hajek, Elemental analysis by X-ray fluorescence spectroscopy, Methods of Soil Analysis: Part 3 Chemical Methods 5 (1996) 161–223. doi:10.2136/sssabookser5.3.c7.

[125] H. Salazar, L. Villaseñor, Separation of cosmic-ray components in a single water Cherenkov detector, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 553 (1) (2005) 295–298, proceedings of the fifth International Workshop on Ring Imaging Detectors. doi:10.1016/j.nima.2005.08.013.

[126] Y. Huang, H. Bao, X. Qi, Seismic Random Noise Attenuation Method Based on Variational Mode Decomposition and Correlation Coefficients, Electronics 7 (2018) 280. doi:10.3390/electronics7110280.

[127] D. Pani, G. Barabino, L. Citi, P. Meloni, S. Raspopovic, S. Micera, L. Raffo, Real-Time Neural Signals Decoding onto Off-the-Shelf DSP Processors for Neuroprosthetic Applications, IEEE Transactions on Neural Systems and Rehabilitation Engineering 24 (2016) 993–1002. doi:10.1109/TNSRE.2016.2527696.

[128] F. J. Iniguez-Lomeli, Y. Bornat, S. Renaud, J. H. Barron-Zambrano, H. Rostro-Gonzalez, A real-time FPGA-based implementation for detection and sorting of bio-signals, Neural Computing and Applications 33 (2021) 12121–12140. doi:10.1007/s00521-021-05853-7.

[129] L. Wang, S.-H. Pun, P. U. Mak, A. Klug, B.-J. Zhang, M. I. Vai, T. C. Lei, A real-time correlational combination algorithm to improve SNR for multi-channel neural recordings, in: 2021 IEEE Asia Pacific Conference on Circuit and Systems (APCCAS), 2021, pp. 213–216. doi:10.1109/APCCAS51387.2021.9687737.

[130] F. Pollastrone, M. Riva, D. Marocco, F. Belli, C. Centioli, Automatic pattern recognition on electrical signals applied to neutron gamma discrimination, Fusion Engineering and Design 123 (2017) 969–974. doi:10.1016/j.fusengdes.2017.03.009.

[131] C. H. Moore, W. Lin, FPGA Correlator for Applications in Embedded Smart Devices, Biosensors 12 (2022) 236. doi:10.3390/bios12040236.

[132] A. Glenn, Q. Cheng, A. D. Kaplan, R. Wurtz, Pulse pileup rejection methods using a two-component Gaussian Mixture Model for fast neutron detection with pulse shape discriminating scintillator, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 988 (2021) 164905. doi:10.1016/j.nima.2020.164905.

[133] T. Alharbi, Distance metrics for digital pulse-shape discrimination of scintillator detectors, Radiation Physics and Chemistry 156 (2019) 205–209. doi:10.1016/j.radphyschem.2018.11.014.

[134] N. D'Olympia, P. Chowdhury, C. Lister, J. Glodo, R. Hawrami, K. Shah, U. Shirwadkar, Pulse-shape analysis of CLYC for thermal neutrons, fast neutrons, and gamma-rays, Nuclear Instruments and Methods in Physics Research Section A:

Accelerators, Spectrometers, Detectors and Associated Equipment 714 (2013) 121–127. doi:10.1016/j.nima.2013.02.043.

[135] J. Flores, I. Martel, R. Jiménez, J. Galán, P. Salmerón, Application of neural networks to digital pulse shape analysis for an array of silicon strip detectors, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 830 (2016) 287–293. doi:10.1016/j.nima.2016.05.107.

[136] J. Zhou, A. Abdulaziz, Y. Altmann, A. Di Fulvio, Generalized method for the optimization of pulse shape discrimination parameters, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 1050 (2023) 168184. doi:10.1016/j.nima.2023.168184.

[137] S. Yoon, C. Lee, H. Seo, H.-D. Kim, Improved fast neutron detection using CNN-based pulse shape discrimination, Nuclear Engineering and Technology 55 (11) (2023) 3925–3934. doi:10.1016/j.net.2023.07.007.

[138] E. B. Johnson, C. Whitney, S. Vogel, J. F. Christian, K. Holbert, P. Chandhran, High event rate, pulse shape discrimination algorithm for CLYC, in: 2015 IEEE International Symposium on Technologies for Homeland Security (HST), IEEE, 2015, pp. 1–7. doi:10.1109/THS.2015.7225282.

[139] B. Blair, C. Chen, A. Glenn, A. Kaplan, J. Ruz, L. Simms, R. Wurtz, Gaussian mixture models as automated particle classifiers for fast neutron detectors, Statistical Analysis and Data Mining: The ASA Data Science Journal 12 (2019) 479–488. doi:10.1002/sam.11432.

[140] M. J. Safari, F. A. Davani, H. Afarideh, S. Jamili, E. Bayat, Discrete Fourier Transform Method for Discrimination of Digital Scintillation Pulses in Mixed

Neutron-Gamma Fields, IEEE Transactions on Nuclear Science 63 (2016) 325–332. doi:10.1109/TNS.2016.2514400.

[141] M. Nakhostin, A General-Purpose Digital Pulse Shape Discrimination Algorithm, IEEE Transactions on Nuclear Science 66 (2019) 838–845. doi:10.1109/TNS.2019.2910153.

[142] J. Polack, M. Flaska, A. Enqvist, C. Sosa, C. Lawrence, S. Pozzi, An algorithm for charge-integration, pulse-shape discrimination and estimation of neutron/photon misclassification in organic scintillators, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 795 (2015) 253–267. doi:10.1016/j.nima.2015.05.048.

[143] A. Abdulaziz, J. Zhou, M. Fang, S. McLaughlin, A. Di Fulvio, Y. Altmann, A variational autoencoder for minimally-supervised pulse shape discrimination, Annals of Nuclear Energy 204 (2024) 110496. doi:10.1016/j.anucene.2024.110496.

[144] J. Peña Rodríguez, R. de'León Barrios, A. Ramírez-Muñóz, D. Villabona-Ardila, M. Suárez-Durán, A. Vásquez-Ramírez, H. Asorey, L. A. Núñez, Muography background sources: simulation, characterization, and machine-learning rejection, in: Proceedings of 37th International Cosmic Ray Conference — PoS(ICRC2021), Vol. 395, 2021, p. 400. doi:10.22323/1.395.0400.

[145] C. Glaser, S. McAleer, S. Stjärnholm, P. Baldi, S. Barwick, Deep-learning-based reconstruction of the neutrino direction and energy for in-ice radio detectors, Astroparticle Physics 145 (2023) 102781. doi:10.1016/j.astropartphys.2022.102781.

[146] T. Aarrestad, V. Loncar, N. Ghielmetti, M. Pierini, S. Summers, J. Ngadiuba,

C. Petersson, H. Linander, Y. Iiyama, G. D. Guglielmo, J. Duarte, P. Harris, D. Rankin, S. Jindariani, K. Pedro, N. Tran, M. Liu, E. Kreinar, Z. Wu, D. Hoang, Fast convolutional neural networks on FPGAs with hls4ml, Machine Learning: Science and Technology 2 (2021) 045015. doi:10.1088/2632-2153/ac0ea1.

[147] J. M. Duarte, Novel machine learning applications at the lhc, arXiv preprint arXiv:2409.20413 (2024). doi:10.48550/arXiv.2409.20413.

[148] L. M. Simms, B. Blair, J. Ruz, R. Wurtz, A. D. Kaplan, A. Glenn, Pulse discrimination with a Gaussian mixture model on an FPGA, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 900 (2018) 1–7. doi:10.1016/j.nima.2018.05.039.

[149] S. Qi, J. Li, S. Wang, Y. Chen, W. Xiao, X. Ai, W. Zhao, K. Zhang, W. Chen, Online neutron/gamma discrimination and source detection using CLYC(Ce) scintillator: A sequential approach, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 1014 (2021) 165733. doi:10.1016/j.nima.2021.165733.

[150] A. Magid, L. Shkedy, N. B. Ari, N. Shiloah, C. G. Jakobson, W. Freiman, E. Armon, E. Almog, S. Srur, A. Katz, N. Ashush, G. Tzvieli, M. Alcheck, J. Lazarus, O. Magril, O. Dicker, B. Milgrom, T. Markovitz, New low SWaP high definition MWIR detector with 5um pitch, in: G. F. Fulop, M. H. MacDougal, D. Z. Ting, M. Kimata (Eds.), Infrared Technology and Applications L, Vol. 13046, International Society for Optics and Photonics, SPIE, 2024, p. 1304612. doi:10.1117/12.3013212.

[151] D. Panitzek, C. Romano, M. Eichhorn, C. Kieleck, Temperature investigation of low SWaP thulium-doped fiber lasers, Optics Express 32 (2) (2024) 1890–1901. doi:10.1364/OE.504623.

[152] J. Hopkins, D. Ellis, Survive the Lunar Night Low-SWaP Optical Communication

System, in: 2024 IEEE Aerospace Conference, IEEE, 2024, pp. 1–8. doi:10.1109/AERO58975.2024.10520943.

[153] Y. Huang, K. Zhao, C. Shan, C. feng, Y. Wang, Z. Wang, K. Chang, F. Liu, W. Xie, S. Liu, Development of a Clyc-Based Wide Dose Rate Range Portable Neutron-Gamma Detector, Social Science Research Network (Feb 2024). doi:10.2139/ssrn.4717223.

[154] R. Diener, J. Dreyling-Eschweiler, H. Ehrlichmann, I. Gregor, U. Kötz, U. Krämer, N. Meyners, N. Potylitsina-Kube, A. Schütz, P. Schütze, M. Stanitzki, The DESY II test beam facility, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 922 (2019) 265–286. doi:10.1016/j.nima.2018.11.133.

[155] CAEN, CAENScope, Accessed: 2023-06-15 (2023).
URL https://www.caen.it/products/caenscope/

[156] AMETEK, MAESTRO Multichannel Analyzer Emulation Software, Accessed: 2023-06-15 (2023).
URL https://www.ortec-online.com/products/application-software/maestro-mca

[157] M. L. Crespo, et al., Remote Laboratory for E-Learning of Systems on Chip and Their Applications to Nuclear and Scientific Instrumentation, Electronics 10 (2021) 2191. doi:10.3390/electronics10182191.

[158] D. Valerio, Open source software-defined radio: A survey on gnuradio and its applications, Forschungszentrum Telekommunikation Wien, Vienna, Technical Report FTW-TR-2008-002 (2008).

[159] R. Mumford, The LXI standard: past, present and future, Microwave Journal 50 (3) (2007) 70.

[160] G. B. Mitolo. PJON (Padded Jittering Operative Network) [online] (12 2022). Accessed: 2023-06-22.

[161] T. Stenner, C. Boulay, M. Grivich, D. Medine, C. Kothe, tobiasherzke, G. Grimm, xloem, A. Biancarelli, B. Mansencal, chausner, J. Frey, kyucrane, S. Powell, P. Clisson, phfix., Lab Streaming Layer project, Accessed: 2023-06-15 (3 2022). doi:10.5281/zenodo.6387090.

[162] USRP. LSL Toolkit for GNURadio/USRP Data [online] (2020). Accessed: 2023-06-22.

[163] B. Bloessl, L. Baumgärtner, M. Hollick, Hardware-Accelerated Real-Time Stream Data Processing on Android with GNU Radio, in: Proceedings of the 14th International Workshop on Wireless Network Testbeds, Experimental evaluation & Characterization, ACM, 2020, pp. 103–109. doi:10.1145/3411276.3412184.

[164] J. Yang, M. Xiong, X. Zhao, Improved low power multichannel pulse amplitude analyzer, Open Access Library Journal 7 (8) (2020) 1–7. doi:10.4236/oalib.1106629.

[165] I. R. Morales. Remote diagnostics platform for custom instrumentation in nuclear applications - code repository [online] (2023). Accessed: 2023-07-10. doi:10.5281/zenodo.8082581.

[166] B. Angelucci, G. Anzivino, C. Avanzini, C. Biino, A. Bizzeti, F. Bucci, A. Cassese, P. Cenci, R. Ciaranfi, G. Collazuol, V. Falaleev, S. Galeotti, S. Giudici, E. Iacopini, G. Lamanna, M. Lenti, G. Magazzù, E. M. Marinova, M. Pepe, R. Piandani, M. Piccini, G. Ruggiero, A. Sergi, M. Sozzi, M. Veltri, Pion–Muon separation with a RICH prototype for the NA62 experiment, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors

and Associated Equipment 621 (2010) 205–211. doi:10.1016/j.nima.2010.05.062.

[167] J. Liu, Y. Zhang, Q. Zhao, Adaptive ViBe Algorithm Based on Pearson Correlation Coefficient, in: 2019 Chinese Automation Congress (CAC), IEEE, 2019, pp. 4885–4889. doi:10.1109/CAC48633.2019.8997209.

[168] C.-C. Zhang, J.-D. Fang, Edge Detection Based on Improved Sobel Operator, in: Proceedings of the 2016 International Conference on Computer Engineering and Information Systems, Atlantis Press, 2016, pp. 129–132. doi:10.2991/ceis-16.2016.25.

[169] R. E. Wurtz, Consistent principles for particle ID from PSD systems, in: A. Burger, R. B. James, S. A. Payne (Eds.), Hard X-Ray, Gamma-Ray, and Neutron Detector Physics XXI, SPIE, 2019, p. 34. doi:10.1117/12.2528898.

[170] P. Chandhran, K. E. Holbert, E. B. Johnson, C. Whitney, S. M. Vogel, Neutron and gamma ray discrimination for CLYC using normalized cross correlation analysis, in: 2014 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC), IEEE, 2014, pp. 1–8. doi:10.1109/NSSMIC.2014.7431169.

[171] J. Cong, B. Liu, S. Neuendorffer, J. Noguera, K. Vissers, Z. Zhang, High-Level Synthesis for FPGAs: From Prototyping to Deployment, IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems 30 (2011) 473–491. doi:10.1109/TCAD.2011.2110592.

[172] X. He, L. Li, Y. Liu, X. Yu, J. Meng, A Two-Stage Biomedical Event Trigger Detection Method Integrating Feature Selection and Word Embeddings, IEEE/ACM Transactions on Computational Biology and Bioinformatics 15 (2018) 1325–1332. doi:10.1109/TCBB.2017.2715016.

[173] D. J. R. Sevilla, Probability distributions for Poisson processes with pile-up (10 2013). arXiv:1310.7566, doi:10.48550/arXiv.1310.7566.

[174] R. Romo, A. Hernández, A. Zainos, E. Salinas, Correlated Neuronal Discharges that Increase Coding Efficiency during Perceptual Discrimination, Neuron 38 (2003) 649–657. doi:10.1016/S0896-6273(03)00287-3.

[175] M. Faisal, R. T. Schiffer, M. Flaska, S. A. Pozzi, D. D. Wentzloff, A correlation-based pulse detection technique for gamma-ray/neutron detectors, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 652 (2011) 479–482. doi:10.1016/j.nima.2010.10.072.

[176] J. L. Rodgers, W. A. Nicewander, Thirteen Ways to Look at the Correlation Coefficient, The American Statistician 42 (1988) 59–66. doi:10.1080/00031305.1988.10475524.

[177] Louis, S. J. G. Mark, Oliphant, Learning Ensembles of First-Order Clauses for Recall-Precision Curves: A Case Study in Biomedical Information Extraction, in: Ross, S. A. C. Rui, King (Eds.), Inductive Logic Programming, Springer Berlin Heidelberg, 2004, pp. 98–115. doi:10.1007/978-3-540-30109-7_11.

[178] R. A. Johnson, N. V. Chawla, J. J. Hellmann, Species distribution modeling and prediction: A class imbalance problem, in: 2012 Conference on Intelligent Data Understanding, IEEE, 2012, pp. 9–16. doi:10.1109/CIDU.2012.6382186.

[179] T. Saito, M. Rehmsmeier, The Precision-Recall Plot Is More Informative than the ROC Plot When Evaluating Binary Classifiers on Imbalanced Datasets, PLOS ONE 10 (2015) e0118432. doi:10.1371/journal.pone.0118432.

[180] K. Rana, R. Singh, K. Sayann, Correlation based novel technique for real time

oscilloscope triggering for complex waveforms, Measurement 43 (2010) 299–311. doi:10.1016/j.measurement.2009.11.002.

[181] A. A. A. Etumi, F. J. Anayi, A. A. Fahmy, E. E. Eldukhri, New algorithm based on auto-correlation and cross-correlation scheme to detect the internal fault in single phase transformer, in: 12th IET International Conference on Developments in Power System Protection (DPSP 2014), 2014, pp. 1–5. doi:10.1049/cp.2014.0151.

[182] S. Gorard, An Absolute Deviation Approach to Assessing Correlation, British Journal of Education, Society & Behavioural Science 5 (2015) 73–81. doi:10.9734/BJESBS/2015/11381.

[183] J. T. Schaefer, The Critical Success Index as an Indicator of Warning Skill, Weather and Forecasting 5 (1990) 570–575. doi:10.1175/1520-0434(1990)005<0570:TCSIAA>2.0.CO;2.

[184] M. E. Baldwin, J. S. Kain, Sensitivity of Several Performance Measures to Displacement Error, Bias, and Event Frequency, Weather and Forecasting 21 (2006) 636–648. doi:10.1175/WAF933.1.

[185] A. Larner, Assessing cognitive screeners with the critical success index, Progress in Neurology and Psychiatry 25 (2021) 33–37. doi:10.1002/pnp.719.

[186] A. Tharwat, Classification assessment methods, Applied Computing and Informatics 17 (2021) 168–192. doi:10.1016/j.aci.2018.08.003.

[187] D. Berrar, On the Noise Resilience of Ranking Measures, in: Seiichi, D. Kenji, I. Kazushi, L. Minho, L. D. H. Akira, Ozawa (Eds.), Neural Information Processing, Springer International Publishing, 2016, pp. 47–55. doi:10.1007/978-3-319-46672-9_6.

[188] H. R. Sofaer, J. A. Hoeting, C. S. Jarnevich, The area under the precision-recall curve as a performance metric for rare binary events, Methods in Ecology and Evolution 10 (2019) 565–577. doi:10.1111/2041-210X.13140.

[189] J. Davis, M. Goadrich, The relationship between Precision-Recall and ROC curves, in: Proceedings of the 23rd international conference on Machine learning - ICML '06, ACM Press, 2006, pp. 233–240. doi:10.1145/1143844.1143874.

[190] Y.-k. Choi, J. Cong, HLS-Based Optimization and Design Space Exploration for Applications with Variable Loop Bounds, in: 2018 IEEE/ACM International Conference on Computer-Aided Design (ICCAD), IEEE Press, 2018, p. 1–8. doi:10.1145/3240765.3240815.

[191] V. Tsoutsouras, K. Koliogeorgi, S. Xydis, D. Soudris, An Exploration Framework for Efficient High-Level Synthesis of Support Vector Machines: Case Study on ECG Arrhythmia Detection for Xilinx Zynq SoC, Journal of Signal Processing Systems 88 (2017) 127–147. doi:10.1007/s11265-017-1230-1.

[192] W. Florian, B. Valinoti, L. G. García, M. Cervetto, E. Marchi, M. L. Crespo, S. Carrato, A. Cicuttin, An Open-Source Hardware/Software Architecture for Remote Control of SoC-FPGA Based Systems, in: S. Saponara, A. De Gloria (Eds.), Applications in Electronics Pervading Industry, Environment and Society, Springer International Publishing, Cham, 2022, pp. 69–75. doi:10.1007/978-3-030-95498-7_10.

[193] L. G. García Ordóñez, I. R. Morales Argueta, M. L. Crespo, S. Carrato, A. Cicuttin, H. D. L. T. Perez, D. Barrientos, S. Levorato, B. Valinoti, W. Florian, K. Mannatunga, M. Ballina, M. Cruz, DAQ platform based on SoC-FPGA for high resolution time stamping in cosmic ray detection, PoS ICRC2019 (2019) 266. doi:10.22323/1.358.0266.

[194] S. Roy, P. Banerjee, An Algorithm for Trading Off Quantization Error with Hardware Resources for MATLAB-Based FPGA Design, IEEE Transactions on Computers 54 (2005) 886–896. doi:10.1109/TC.2005.106.

[195] C. N. Coelho, A. Kuusela, S. Li, H. Zhuang, J. Ngadiuba, T. K. Aarrestad, V. Loncar, M. Pierini, A. A. Pol, S. Summers, Automatic heterogeneous quantization of deep neural networks for low-latency inference on the edge for particle detectors, Nature Machine Intelligence 3 (2021) 675–686. doi:10.1038/s42256-021-00356-5.

[196] J. Duarte, S. Han, P. Harris, S. Jindariani, E. Kreinar, B. Kreis, J. Ngadiuba, M. Pierini, R. Rivera, N. Tran, Z. Wu, Fast inference of deep neural networks in FPGAs for particle physics, Journal of Instrumentation 13 (2018) P07027–P07027. doi:10.1088/1748-0221/13/07/P07027.

[197] Q. Xiao, Y. Liang, L. Lu, S. Yan, Y.-W. Tai, Exploring Heterogeneous Algorithms for Accelerating Deep Convolutional Neural Networks on FPGAs, in: Proceedings of the 54th Annual Design Automation Conference 2017, ACM, 2017, pp. 1–6. doi:10.1145/3061639.3062244.

[198] K. Chen, H. Chen, W. Wu, H. Xu, L. Yao, Optimization on fixed low latency implementation of the GBT core in FPGA, Journal of Instrumentation 12 (2017) P07011–P07011. doi:10.1088/1748-0221/12/07/P07011.

[199] N. Nottbeck, D. C. Schmitt, P. D. V. Büscher, Implementation of high-performance, sub-microsecond deep neural networks on FPGAs for trigger applications, Journal of Instrumentation 14 (2019) P09014–P09014. doi:10.1088/1748-0221/14/09/P09014.

[200] C. Willmott, K. Matsuura, Advantages of the mean absolute error (MAE) over

the root mean square error (RMSE) in assessing average model performance, Climate Research 30 (2005) 79–82. doi:10.3354/cr030079.

[201] R. T. Kouzes, E. R. Siciliano, J. H. Ely, P. E. Keller, R. J. McConn, Passive neutron detection for interdiction of nuclear material at borders, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 584 (2-3) (2008) 383–400. doi:10.1016/j.nima.2007.10.026.

[202] Nuclear Science and Instrumentation Newsletter No. 3, February 2022, no. 3 in Nuclear Science and Instrumentation Newsletter, INTERNATIONAL ATOMIC ENERGY AGENCY, Vienna, 2022.

[203] S. Moretto, F. P. Andrades, J. Delgado, C. Fontana, D. Fabris, G. Nebbia, M. Turcato, D. Brunelli, L. Pancheri, A. Quaranta, UAV prototype for localization and identification of radioactive contamination and emitters, EPJ Web of Conferences 253 (2021) 08001. doi:10.1051/epjconf/202125308001.

[204] P. Dejdar, P. Munster, T. Horvath, High-speed Data Acquisition and Signal Processing Using Cost Effective ARM + FPGA Processors, in: 2019 42nd International Conference on Telecommunications and Signal Processing (TSP), IEEE, 2019, pp. 593–596. doi:10.1109/TSP.2019.8769055.

[205] H. Ye, L. Chen, X. Xu, G. Jin, Fast FPGA algorithm for neutron–gamma discrimination, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 1027 (2022) 166256. doi:10.1016/j.nima.2021.166256.

[206] M. E. Hammad, H. Kasban, R. M. Fikry, M. I. Dessoky, O. Zahran, S. M. Elaraby, F. E. A. El-Samie, Pile-up correction algorithm for high count rate gamma ray

spectroscopy, Applied Radiation and Isotopes 151 (2019) 196–206. doi:10.1016/J.APRADISO.2019.06.003.

[207] H. Kasani, S. Ashrafi, N. Ghal-Eh, H. R. Vega-Carrillo, Gamma-ray spectroscopy with anode pulses of NaI(Tl) detector using a low-cost digitizer system, Applied Radiation and Isotopes 176 (2021) 109854. doi:10.1016/j.apradiso.2021.109854.

[208] S.-X. Liu, W. Zhang, Z.-H. Zhang, S. Lin, H.-R. Cao, C.-X. Song, J.-L. Zhao, G.-Q. Zhong, Performance of real-time neutron/gamma discrimination methods, Nuclear Science and Techniques 34 (2023) 8. doi:10.1007/s41365-022-01160-5.

[209] H. Bai, G. Zhang, Z. Xiong, Y. Li, D. Zhao, M. Su, Z. Mo, X. Wang, F. Gao, H. Zhang, Z. Zhang, J. Wen, A method to calibrate the n-γ discrimination property of scintillators in low energy region, Applied Radiation and Isotopes 167 (2021) 109447. doi:10.1016/j.apradiso.2020.109447.

[210] E. O. Brigham, R. E. Morrow, The fast Fourier transform, IEEE Spectrum 4 (1967) 63–70. doi:10.1109/MSPEC.1967.5217220.

[211] A. V. Oppenheim, A. S. Willsky, S. H. Nawab, Signals & systems, 2nd Edition, Prentice-Hall, Inc., USA, 1996.

[212] AMD. LogiCORE IP - Fast Fourier Transform v9.1 [online] (2024). Accessed: 2024-07-25.

[213] E. Technology. PSD PLASTIC SCINTILATOR: EJ-276D & EJ-276G [online] (2023). Accessed: 2024-08-29.

[214] M. Grodzicka-Kobylka, T. Szczesniak, M. Moszyński, K. Brylew, L. Swiderski, J. Valiente-Dobón, P. Schotanus, K. Grodzicki, H. Trzaskowska, Fast neutron and gamma ray pulse shape discrimination in EJ-276 and EJ-276G plastic scintillators,

Journal of Instrumentation 15 (03) (2020) P03030. doi:10.1088/1748-0221/15/03/P03030.

[215] E. Ryabeva, I. Urupa, E. Lupar, V. Kadilin, A. Skotnikova, Y. Kokorev, R. Ibragimov, Calibration of EJ-276 plastic scintillator for neutron–gamma pulse shape discrimination experiments, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 1010 (2021) 165495. doi:10.1016/j.nima.2021.165495.

[216] T. Quist, E. Cazalas, Design, modeling, and initial testing of a multi-volume neutron spectrometer, Radiation Measurements 166 (2023) 106982. doi:10.1016/j.radmeas.2023.106982.

[217] G. Shen, S. Zhang, X. Zhang, H. Zhang, B. Yuan, D. Hou, C. Wang, Z. Quan, Z. Yang, Y. Sun, Using energy particle detection technology on the Tiangong's Space Station's wentian laboratory cabin module, Aerospace 10 (4) (2023) 373. doi:10.3390/aerospace10040373.

[218] Coupland, Daniel D. S. and Stonehill, Laura C. and Mesick, Katherine E. and Dunn, John P., The SENSER CLYC experiment, in: 2016 IEEE Nuclear Science Symposium, Medical Imaging Conference and Room-Temperature Semiconductor Detector Workshop (NSS/MIC/RTSD), 2016, pp. 1–4. doi:10.1109/NSSMIC.2016.8069879.

[219] J. Köhler, B. Ehresmann, C. Zeitlin, R. Wimmer-Schweingruber, D. Hassler, G. Reitz, D. Brinza, J. Appel, S. Böttcher, E. Böhm, S. Burmeister, J. Guo, H. Lohf, C. Martin, A. Posner, S. Rafkin, Measurements of the neutron spectrum in transit to Mars on the Mars Science Laboratory, Life Sciences in Space Research 5 (2015) 6–12. doi:10.1016/j.lssr.2015.03.001.

[220] D. Ölçek, R. A. Jebali, D. Meier, A. F. Haavik, C. Gheorghe, M. S. Gjennestad, S.-E.

Hamran, A. Kohfeldt, G. Maehlum, Development of a compact scintillator-based gamma-ray neutron detector for terrestrial and space applications, Journal of Instrumentation 19 (01) (2024) C01030. doi:10.1088/1748-0221/19/01/C01030.

[221] K. Mesick, D. Coupland, K. Bartlett, D. Beckman, S. West, L. Stonehill, N. Dallmann, S. Storms, W. Feldman, Development of a Low-Resource Combined Gamma-Ray and Neutron Spectrometer for Planetary Science, in: Proceedings of the Small Satellite Conference, 2020.

[222] A. Barzilov, M. Kazemeini, Unmanned Aerial System Integrated Sensor for Remote Gamma and Neutron Monitoring, Sensors 20 (2020) 5529. doi:10.3390/s20195529.

[223] M. Köhli, M. Schrön, M. Zreda, U. Schmidt, P. Dietrich, S. Zacharias, Footprint characteristics revised for field-scale soil moisture monitoring with cosmic-ray neutrons, Water Resources Research 51 (7) (2015) 5772–5790. doi:10.1002/2015WR017169.

[224] A. Cirillo, M. Caresana, Calibration of a large-size wide-range neutron spectrometer, The European Physical Journal Plus 137 (2022) 774. doi:10.1140/epjp/s13360-022-02945-8.

[225] B. Fersch, T. Francke, M. Heistermann, M. Schrön, V. Döpper, J. Jakobi, G. Baroni, T. Blume, H. Bogena, C. Budach, T. Gränzig, M. Förster, A. Güntner, H.-J. Hendricks Franssen, M. Kasner, M. Köhli, B. Kleinschmit, H. Kunstmann, A. Patil, D. Rasche, L. Scheiffele, U. Schmidt, S. Szulc-Seyfried, J. Weimar, S. Zacharias, M. Zreda, B. Heber, R. Kiese, V. Mares, H. Mollenhauer, I. Völksch, S. Oswald, A dense network of cosmic-ray neutron sensors for soil moisture observation in a highly instrumented pre-Alpine headwater catchment in Germany, Earth System Science Data 12 (2020) 2289–2309. doi:10.5194/essd-12-2289-2020.

[226] Y. Ma, S. Hang, P. Gong, Z. Wang, D. Liang, Z. Hu, X. Tang, C. Zhou, X. Zhu, A method for discriminating neutron and gamma waveforms based on a comparison of differences between Pulse Feature Heights, Journal of Radioanalytical and Nuclear Chemistry 333 (1) (2024) 375–386. doi:10.1007/s10967-023-09280-x.

[227] K. Zhao, C. Feng, S. Wang, Z. Shen, K. Zhang, S. Liu, n/γ discrimination for CLYC detector using a one-dimensional Convolutional Neural Network, Journal of Instrumentation 18 (01) (2023) P01021. doi:10.1088/1748-0221/18/01/P01021.

[228] E. Doucet, T. Brown, P. Chowdhury, C. Lister, C. Morse, P. Bender, A. Rogers, Machine learning n/γ discrimination in CLYC scintillators, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 954 (2020) 161201. doi:10.1016/j.nima.2018.09.036.

[229] C. Fu, A. D. Fulvio, S. Clarke, D. Wentzloff, S. Pozzi, H. Kim, Artificial neural network algorithms for pulse shape discrimination and recovery of piled-up pulses in organic scintillators, Annals of Nuclear Energy 120 (2018) 410–421. doi:10.1016/j.anucene.2018.05.054.

[230] J. Griffiths, S. Kleinegesse, D. Saunders, R. Taylor, A. Vacheret, Pulse shape discrimination and exploration of scintillation signals using convolutional neural networks, Machine Learning: Science and Technology 1 (2020) 045022. doi:10.1088/2632-2153/abb781.

[231] M. Gelfusa, R. Rossi, M. Lungaroni, F. Belli, L. Spolladore, I. Wyss, P. Gaudio, A. Murari, Advanced pulse shape discrimination via machine learning for applications in thermonuclear fusion, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 974 (2020) 164198. doi:10.1016/j.nima.2020.164198.

[232] I. Morad, M. Ghelman, D. Ginzburg, A. Osovizky, N. Shlezinger, Model-Based Deep Learning Algorithm for Detection and Classification at High Event Rates, IEEE Transactions on Nuclear Science (2024) 1–1doi:10.1109/TNS.2024.3371573.

[233] C. Kim, S. Ahn, K. Chae, J. Hooker, G. Rogachev, Restoring original signals from pile-up using Deep Learning, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 1055 (2023) 168492. doi:10.1016/j.nima.2023.168492.

[234] C. Yi, J. Han, R. Song, X. Yan, F. Ren, X. Luo, Z. Han, C. Wen, G. Qu, X. Liu, W. Lin, P. Wang, Y. Fan, S. Qian, Z. Wang, G. Tang, L. Qin, X. Wang, J. Liu, Discrimination of piled-up neutron-gamma pulses using charge comparison method and neural network for CLYC detectors, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 1055 (2023) 168561. doi:10.1016/j.nima.2023.168561.

[235] H. Gao, H. Dai, Y. Zeng, High-speed image processing and data transmission based on vivado hls and axi4-stream interface, in: 2018 IEEE International Conference on Information and Automation (ICIA), 2018, pp. 575–579. doi:10.1109/ICInfA.2018.8812379.

[236] N. Blasi, S. Brambilla, F. Camera, S. Ceruti, A. Giaz, L. Gini, F. Groppi, S. Manenti, A. Mentana, B. Million, S. Riboldi, Fast neutron detection efficiency of 6Li and 7Li enriched CLYC scintillators using an Am-Be source, Journal of Instrumentation 13 (11) (2018) P11010. doi:10.1088/1748-0221/13/11/P11010.

[237] V. H. Cao, T. H. Dinh, T. T. Nguyen, K. H. Nguyen, D. K. Pham, X. H. Nguyen, N. A. Nguyen, T.-A. Nguyen, Flexible and low-cost FPGA-based multichannel analyzer for handheld measurement devices, Nuclear Instruments and Meth-

ods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 1018 (2021) 165808. doi:10.1016/j.nima.2021.165808.

[238] R. G. Lyons, Real-Time Computation of Signal Averages and Variances, 3rd Edition, Pearson Education, Inc., 2011, p. 799.

[239] X. Wang, Y. Chen, Q. Yin, S. Zeng, Q. Xie, Advantages of digitally sampling scintillation pulses in pileup processing in PET, in: 2009 IEEE Nuclear Science Symposium Conference Record (NSS/MIC), 2009, pp. 144–147. doi:10.1109/NSSMIC.2009.5401834.

[240] Z. Gu, D. L. Prout, R. Taschereau, B. Bai, A. F. Chatziioannou, A New Pulse Pileup Rejection Method Based on Position Shift Identification, IEEE Transactions on Nuclear Science 63 (1) (2016) 22–29. doi:10.1109/TNS.2015.2495169.

[241] S. M. Mousavi, W. L. Ellsworth, W. Zhu, L. Y. Chuang, G. C. Beroza, Earthquake transformer—an attentive deep-learning model for simultaneous earthquake detection and phase picking, Nature communications 11 (1) (2020) 3952. doi:10.1038/s41467-020-17591-w.

[242] Z. Peng, J. Yang, T.-H. Chen, L. Ma, A first look at the integration of machine learning models in complex autonomous driving systems: a case study on Apollo, in: Proceedings of the 28th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering, 2020, pp. 1240–1250. doi:10.1145/3368089.3417063.

[243] I. Qureshi, J. Yan, Q. Abbas, K. Shaheed, A. B. Riaz, A. Wahid, M. W. J. Khan, P. Szczuko, Medical image segmentation using deep semantic-based methods: A review of techniques, applications and emerging trends, Information Fusion 90 (2023) 316–352. doi:10.1016/j.inffus.2022.09.031.

[244] N. Ghielmetti, V. Loncar, M. Pierini, M. Roed, S. Summers, T. Aarrestad, C. Petersson, H. Linander, J. Ngadiuba, K. Lin, et al., Real-time semantic segmentation on FPGAs for autonomous vehicles with hls4ml, Machine Learning: Science and Technology 3 (4) (2022) 045011. doi:10.1088/2632-2153/ac9cb5.

[245] R. S. Molina, SoC-based FPGA architecture for image analysis and other highly demanding applications, Ph.D. thesis, Università degli Studi di Trieste (2022). doi:11368/3043618.

[246] T. Choudhary, V. Mishra, A. Goswami, J. Sarangapani, A comprehensive survey on model compression and acceleration, Artif. Intell. Rev. 53 (7) (2020) 5113–5155. doi:10.1007/s10462-020-09816-7.

[247] G. Hinton, O. Vinyals, J. Dean, Distilling the Knowledge in a Neural Network, in: NIPS Deep Learning and Representation Learning Workshop, 2015. doi:10.48550/arXiv.1503.02531.

[248] B. Hawks, J. Duarte, N. J. Fraser, A. Pappalardo, N. Tran, Y. Umuroglu, Ps and qs: Quantization-aware pruning for efficient low latency neural network inference, Frontiers in Artificial Intelligence 4 (2021) 676564. doi:10.3389/frai.2021.676564.

[249] C. Zhang, G. Sun, Z. Fang, P. Zhou, P. Pan, J. Cong, Caffeine: Toward uniformed representation and acceleration for deep convolutional neural networks, IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems 38 (11) (2018) 2072–2085. doi:10.1109/TCAD.2017.2785257.

[250] J. Haris, P. Gibson, J. Cano, N. B. Agostini, D. Kaeli, SECDA: Efficient Hardware/Software Co-Design of FPGA-based DNN Accelerators for Edge Inference, in: 2021 IEEE 33rd International Symposium on Computer Architecture and

High Performance Computing (SBAC-PAD), IEEE, Belo Horizonte, Brazil, 2021, pp. 33–43. doi:10.1109/SBAC-PAD53543.2021.00015.

[251] R. Pereira, A. Fernandes, N. Cruz, J. Sousa, M. Riva, D. Marocco, F. Belli, B. Gonçalves, Neutron/Gamma discrimination code based on trapezoidal filter, Fusion Engineering and Design 134 (2018) 118–122. doi:10.1016/j.fusengdes.2018.07.002.

[252] H. Y. Özderya, SerialPlot, https://github.com/hyOzd/serialplot, Accessed: 2024-07-02 (2023).

[253] S. Marrone, D. Cano-Ott, N. Colonna, C. Domingo, F. Gramegna, E. Gonzalez, F. Gunsing, M. Heil, F. Käppeler, P. Mastinu, P. Milazzo, T. Papaevangelou, P. Pavlopoulos, R. Plag, R. Reifarth, G. Tagliente, J. Tain, K. Wisshak, Pulse shape analysis of liquid scintillators for neutron studies, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 490 (1) (2002) 299–307. doi:10.1016/S0168-9002(02)01063-X.

[254] I. Morales, A. Monterroso, S. Urizar, Design, assembly and calibration of a microcontroller-based Geiger-Muller doserate meter, in: 2016 IEEE Central America and Panama Student Conference (CONESCAPAN), IEEE, 2016, pp. 1–5. doi:10.1109/CONESCAPAN.2016.8075210.

[255] T. K. Nguyen, H.-N. Tran, Q. N. Nguyen, T. M. L. Dang, V. L. Bui, T. Liamsuwan, N.-T. Le, Calibration of a neutron dose rate meter in various neutron standard fields, Nuclear Science and Techniques 31 (5 2020). doi:10.1007/s41365-020-00763-0.

[256] P. Chowdhury, E. Jackson, C. Lister, et al., Fast neutron response of 6Li-depleted CLYC detectors up to 20 MeV, Nuclear Instruments and Methods in Physics

Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 763 (2014) 433–441. doi:10.1016/j.nima.2014.06.074.

[257] D. Rigamonti, G. Zhong, G. Croci, L. Giacomelli, G. Gorini, Z. Hu, A. Muraro, M. Nocente, E. P. Cippo, M. Rebai, et al., First neutron spectroscopy measurements with a compact C7LYC based detector at EAST, Journal of Instrumentation 14 (09) (2019) C09025. doi:10.1088/1748-0221/14/09/C09025.

[258] M. Xu, G. Zhong, Y. Zhang, L. Huang, L. Hu, R. Zhang, L. Yang, Y. Zhang, W. Chen, Y. Li, et al., Investigation of Cs27LiYCl6: Ce scintillator energy response for DD fusion neutron spectrometer, Fusion Engineering and Design 204 (2024) 114490. doi:10.1016/j.fusengdes.2024.114490.

[259] N. D'Olympia, P. Chowdhury, C. Guess, T. Harrington, E. Jackson, S. Lakshmi, C. Lister, J. Glodo, R. Hawrami, K. Shah, et al., Optimizing Cs2LiYCl6 for fast neutron spectroscopy, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 694 (2012) 140–146. doi:10.1016/j.nima.2012.07.021.

[260] L. Pellegri, A. Giaz, F. Camera, N. Blasi, S. Brambilla, S. Ceruti, B. Million, S. Riboldi, C. Cazzaniga, M. Nocente, et al., Thermal and fast neutron detection with two CLYC scintillators, in: 2014 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC), IEEE, 2014, pp. 1–3. doi:10.1109/NSSMIC.2014.7431190.

[261] D. Pérez-Loureiro, O. Kamaev, G. Bentoumi, L. Li, C. Jewett, M. Thompson, Evaluation of CLYC-6 and CLYC-7 scintillators for detection of nuclear materials, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 1012 (2021) 165622. doi:10.1016/j.nima.2021.165622.

[262] P. Prusachenko, T. Bobrovskiy, The measurement of the 6Li (n, t) 4He reaction cross-section in the energy range of 4.25–7.50 MeV, The European Physical Journal A 60 (1) (2024) 12. doi:10.1140/epja/s10050-024-01236-3.

[263] R. Song, X. Yan, J. Han, X. Luo, F. Ren, Y. Zhang, Z. Han, C. Wen, X. Zhang, L. Chen, W. Lin, G. Qu, X. Liu, Q. Leng, J. Zhu, S. Qian, Z. Wang, Simultaneous Detection of Energy and Flux for Neutron and Gamma by CLYC Scintillator, IEEE Transactions on Nuclear Science 70 (8) (2023) 2148–2156. doi:10.1109/TNS.2023.3289945.