

© Copyright 2026

Pavan Sai Guntha

FPGA-Based System for Radiation Energy Histogram Computation using Time-Over-Threshold

Pavan Sai Guntha

A thesis
submitted in partial fulfillment of the
requirements for the degree of

Master of Science

University of Washington

2026

Committee:

Scott Hauck, Chair

Robert Miyaoka

Program Authorized to Offer Degree:
Electrical and Computer Engineering

University of Washington

Abstract

FPGA-Based System for Radiation Energy Histogram Computation using
Time-Over-Threshold

Pavan Sai Guntha

Chair of the Supervisory Committee:
Scott Hauck
Department of Electrical and Computer Engineering

Personalized dosimetry is essential for optimizing radiopharmaceutical therapy, yet current clinical practice relies on serial hospital imaging that is expensive and limits the number of time-point samples. This thesis presents an FPGA-based system used in the Portable Organ Dosimetry Device (PODD), a compact wireless gamma spectroscopy system for monitoring ^{177}Lu therapy.

This thesis focuses on the design and implementation of an FPGA-based real-time energy histogram computation system. The FPGA processes time-over-threshold encoded signals from 16 parallel detector channels, achieving 2.5ns time binning for pulse width measurement and maintaining 512-bin histograms per channel. The digital processing core integrates with GAGG:Ce/SiPM scintillation detectors and communicates wirelessly via Bluetooth to an Android application.

Validation with ^{177}Lu and ^{22}Na radioactive sources demonstrated clearly resolved photopeaks at 113 keV, 208 keV, and 511 keV. The PODD's compact form factor and wireless operation enable accessible point-of-care dosimetry for personalized radiopharmaceutical therapy.

TABLE OF CONTENTS

	Page
List of Figures	iv
Glossary	vi
Chapter 1: Introduction	1
1.1 Radiotheranostics Primer	1
1.2 Project Background	3
1.2.1 Portable Organ Dosimetry Device	3
1.3 Signal Processing Requirements and Approach	5
1.3.1 Processing Requirements and Energy Histograms	5
1.3.2 Time-Over-Threshold Energy Measurement	7
1.4 Thesis Organization	9
1.5 Individual Contributions	9
Chapter 2: System Description	11
2.1 Three-Board Sandwich Architecture	11
2.2 Radiation Detectors	13
2.2.1 GAGG:Ce Scintillator Crystals	14
2.2.2 Silicon Photomultipliers	15
2.2.3 Detector Probe PCB Design	16
2.3 Analog Front-End PCB	17
2.3.1 Channel Architecture	19
2.3.2 Power Supply and Temperature Monitoring	21
2.4 FPGA Processing System	22
2.4.1 DE0-Nano Development Board	22
2.4.2 Multi-Clock Domain Architecture	23

2.4.3	Double-Edge Sampling for Enhanced Resolution	25
2.4.4	Pulse Width Measurement	25
2.4.5	Clock Domain Crossing	26
2.4.6	Histogram Computation and Storage	28
2.4.7	UART Communication Protocol	29
2.4.8	FPGA Resource Utilization	31
2.5	Microcontroller Subsystem	31
2.5.1	PSoC 6 Hardware Platform	32
2.5.2	Firmware Architecture	33
2.6	Mobile Application Interface	36
2.6.1	BLE Communication	36
2.6.2	Operational Modes	37
Chapter 3:	System Testing & Results	39
3.1	Component-Level Testing	39
3.1.1	Analog Front-End PCB Functionality Tests	39
3.1.2	FPGA Logic Functionality Tests	42
3.1.3	Full System Test	47
Chapter 4:	Conclusions	54
Chapter 5:	Future Work	55
Bibliography	56
Appendix A:	Analog Circuit Design Details	60
A.1	Microcontroller Connector Circuitry	60
A.2	Power Supply Circuitry	62
A.3	Thermistor and Detector Connector Circuitry	63
Appendix B:	UART Handshake Protocol Specifics	66
B.1	Command Codes	66
B.2	Status Codes	66
B.3	Protocol Operation	66

Appendix C: Mobile Application Specifics	69
C.1 Histogram Visualization	69
C.2 Temperature Profile Visualization	69
C.3 Application Settings	70
C.4 Peak Detection Algorithm	71
C.4.1 Algorithm Overview	71
C.4.2 Peak Information Data Structure	74
Appendix D: PODD Scan Startup and Operational Sequence	76
D.1 Stage 1: Hardware Power-On	76
D.2 Stage 2: Bluetooth Connection via the Mobile Application	77
D.3 Stage 3: Scan Initiation	79
D.4 Stage 4: Scan Completion and Data Transfer	81
D.5 Summary of Patient-Observable Indicators	82

LIST OF FIGURES

Figure Number	Page
1.1 Radiotheranostics concept	2
1.2 PODD system architecture	4
1.3 Lu-177 energy histogram	6
1.4 Time-over-threshold principle	8
2.1 Complete PODD system	12
2.2 Three-board sandwich assembly	12
2.3 PODD radiation detector components	14
2.4 Silicon photomultiplier structure	15
2.5 Detector probe PCB schematic	17
2.6 Detector probe PCB layouts	18
2.7 Analog front-end PCB	19
2.8 Single channel schematic	20
2.9 Effect of comparator hysteresis on noisy signals	21
2.10 DE0-Nano FPGA development board	23
2.11 FPGA histogram compute architecture	24
2.12 UART command-response timing	30
2.13 Clicker 2 for PSoC 6 microcontroller board	32
2.14 PSoC 6 peripheral configuration	34
2.15 Mobile application interface	37
3.1 Analog front-end PCB test setup	40
3.2 Analog channel signal chain waveforms	41
3.3 ModelSim BRAM controller simulation	43
3.4 FPGA-in-the-Loop test setup	44
3.5 FIL histogram verification	45
3.6 Energy histograms from radioactive sources	48

3.7	Temperature-controlled test setup	50
3.8	^{177}Lu histograms at multiple temperatures	51
A.1	Microcontroller connector circuitry	61
A.2	Power supply circuitry	62
A.3	Thermistor and detector connector circuitry	64
C.1	Mobile application visualization and settings screens	70

GLOSSARY

ADC: Analog-to-Digital Converter. An electronic circuit that converts continuous analog signals into discrete digital values for processing by digital systems.

BLE: Bluetooth Low Energy. A wireless communication protocol designed for short-range data transfer with minimal power consumption, used for communication between the PODD and mobile application.

CCCD: Client Characteristic Configuration Descriptor. A standard BLE descriptor attached to each notifiable GATT characteristic. The client writes to the CCCD to enable or disable notifications from the server, controlling whether the server pushes value updates over the connection.

COMPTON SCATTERING: An interaction between a gamma photon and an electron in which the photon transfers part of its energy to the electron and scatters at a reduced energy. Results in partial energy deposition in radiation detectors.

ENERGY HISTOGRAM: A representation of detected radiation events organized by deposited energy. Each bin accumulates counts of events within a specific energy range, revealing characteristic spectral features.

FREERTOS: Free Real-Time Operating System. An open-source real-time operating system that allows microcontrollers to run multiple concurrent tasks with deterministic

scheduling, used in the PODD microcontroller to manage BLE communication, UART data transfer, sensor reading, and timer functions as independent tasks.

GATT: Generic Attribute Profile. A BLE protocol that defines a client-server architecture for exchanging data over established connections. The server exposes a hierarchical database of services and characteristics that the client can read, write, or subscribe to for notifications.

MTU: Maximum Transmission Unit. The maximum size (in bytes) of an Attribute Protocol (ATT) data unit exchanged between a BLE client and server. The default ATT_MTU is 23 bytes, allowing a usable notification payload of 20 bytes after the 3-byte ATT header. The PODD application requests an MTU of 512 bytes to accommodate the 256-byte histogram notification packets used during data transfer.

NET: Neuroendocrine Tumor. A type of cancer arising from neuroendocrine cells, often expressing somatostatin receptors targeted by DOTATATE-based radiotheranostics.

OAR: Organ At Risk. Healthy organs that may receive radiation dose during therapy and must be monitored to prevent toxicity. Kidneys and bone marrow are primary OARs for ^{177}Lu therapy.

PET: Positron Emission Tomography. A nuclear medicine imaging technique using positron-emitting radionuclides to produce three-dimensional images of metabolic processes.

PLL: Phase-Locked Loop. An electronic circuit that generates an output signal whose phase is synchronized to an input reference signal, used to multiply clock frequencies

in FPGAs.

PODD: Portable Organ Dosimetry Device. The compact radiation detection system developed in this thesis for at-home monitoring of ^{177}Lu therapy patients.

PRRT: Peptide Receptor Radionuclide Therapy. A targeted cancer treatment using radiolabeled peptides that bind to receptors overexpressed on tumor cells.

SCINTILLATOR: A material that emits visible light when excited by ionizing radiation. The light output is proportional to deposited energy, enabling spectroscopic measurements.

SIPM: Silicon Photomultiplier. A solid-state photodetector consisting of an array of single-photon avalanche diodes (SPADs), providing high gain and fast response for scintillation detection.

TOT: Time-Over-Threshold. An analog-to-digital conversion technique that encodes signal amplitude as the duration a pulse remains above a fixed voltage threshold, enabling energy measurement without high-speed ADCs.

ACKNOWLEDGMENTS

I would like to express my sincere gratitude to my advisor, Professor Scott Hauck, for providing me with the opportunity to work on this impactful project. His guidance and mentorship have deepened my enthusiasm for FPGA-based design and its application to meaningful real-world problems. Working in his lab has been an invaluable experience that has helped me grow as an engineer.

I am grateful to Dr. Robert Miyaoka, a distinguished member of my thesis committee, for his patience in explaining the clinical requirements and for providing crucial insights from a radiology perspective. His expertise helped bridge the gap between the engineering implementation and its medical application.

I extend my appreciation to William C. J. Hunter of the Imaging Research Lab for his tremendous support with PCB assembly, testing, and soldering. His hands-on assistance was instrumental in bringing the hardware to a functional state.

Special thanks to Ethan Ingalls for his contributions to the software components of this project. His collaborative spirit and thoughtful input during team meetings were a driving force that kept the project moving forward.

I would also like to thank Geoff Jones for sharing his extensive experience in FPGA design and providing valuable feedback that improved the quality of this work.

Finally, I thank God for the strength and perseverance to complete this journey. I am deeply grateful to my parents and friends for their unwavering support and encouragement throughout my graduate studies.

Chapter 1

INTRODUCTION

This thesis presents the design and implementation of an FPGA-based system for real-time radiation energy histogram computation, forming the digital signal processing core of PODD. The system enables quantitative monitoring of gamma-ray emissions from patients undergoing radiotheranostic treatment, addressing a critical need for patient-friendly dosimetry solutions in nuclear medicine. This chapter establishes the clinical motivation, introduces the technical approach, and outlines the scope of subsequent chapters.

1.1 *Radiotheranostics Primer*

Radiotheranostics is a nuclear medicine paradigm that integrates diagnostic imaging with targeted radionuclide therapy using molecularly identical or similar radiopharmaceuticals [26]. The term combines “radiotherapy” and “diagnostics,” reflecting the dual capability of radiolabeled compounds to visualize disease through emitted gamma radiation and deliver cytotoxic doses to targeted tissues through beta or alpha particle emission [11].

As illustrated in Figure 1.1, the diagnostic phase employs a gamma-emitting radionuclide attached to a tumor-targeting molecule, enabling visualization of disease sites through single-photon emission computed tomography (SPECT) or positron emission tomography (PET). The same targeting molecule, when labeled with a therapeutic radionuclide, delivers localized radiation to identified tumor sites while sparing healthy tissue [3].

A prominent clinical example is peptide receptor radionuclide therapy (PRRT) for neuroendocrine tumors (NETs). Gallium-68 (^{68}Ga) labeled DOTATATE enables PET

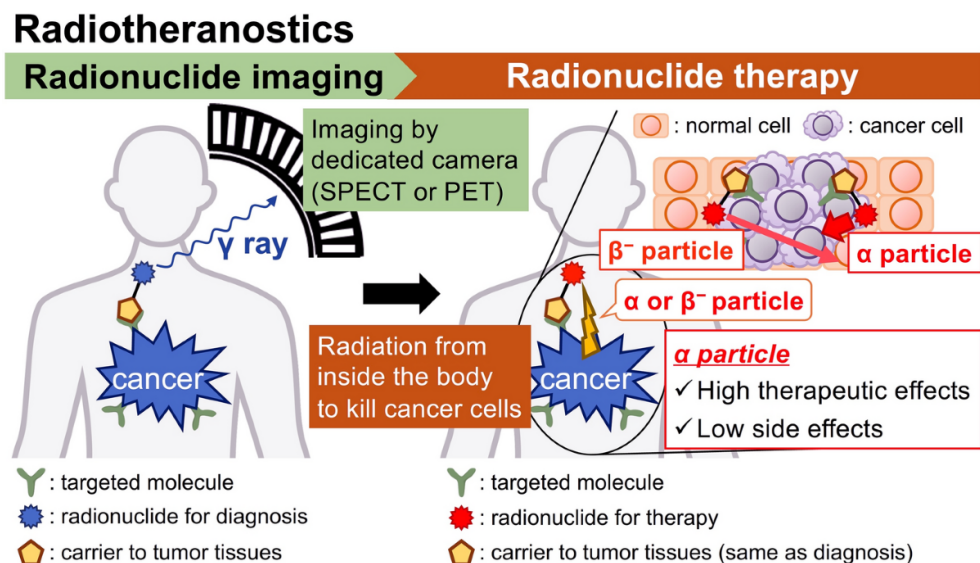


Figure 1.1: The radiotheranostics paradigm combines diagnostic imaging (left) with targeted radionuclide therapy (right). A tumor-targeting molecule carries either a diagnostic radionuclide for gamma-ray imaging via SPECT or PET, or a therapeutic radionuclide emitting alpha or beta particles for treatment. Figure adapted from Ogawa [20].

imaging to identify tumors expressing somatostatin receptors, while lutetium-177 (^{177}Lu) labeled DOTATATE delivers therapeutic beta radiation to those same targets [29]. The NETTER-1 phase III clinical trial demonstrated significant improvements in progression-free survival for NET patients, leading to FDA approval of ^{177}Lu -DOTATATE (marketed as Lutathera) in 2018 [30].

The isotope ^{177}Lu emits beta particles with a maximum energy of 497 keV and mean tissue penetration of approximately 0.67 mm, delivering localized radiation to tumor cells [10]. Importantly for this work, ^{177}Lu also emits gamma photons at 113 keV (6.2% abundance) and 208 keV (10.4% abundance), enabling post-treatment imaging for biodistribution assessment and dosimetric calculations [17].

1.2 *Project Background*

A critical challenge in radiotheranostic treatments is accurate dosimetry i.e., determining the radiation dose absorbed by tumors and healthy organs. Unlike external beam radiation therapy where dose delivery is controlled by the treatment machine, internal radiotherapy dose depends on pharmacokinetics, organ uptake, and biological clearance rates that vary substantially between patients [7]. The kidneys and bone marrow are typically the dose-limiting organs, and real-time dosimetric monitoring can inform dose adjustments and prevent toxicity [25].

Current clinical practice relies on serial SPECT/CT imaging sessions, typically at 24 hours, 4 days, and 7 days post-administration [28]. While quantitative SPECT/CT achieves errors within 5% for organ-sized structures [23], this approach requires multiple hospital visits, consumes significant scanner time, and limits temporal resolution to discrete imaging timepoints. The European Association of Nuclear Medicine (EANM) has mandated patient-specific dosimetry for ^{177}Lu therapies in the European Union, yet the standard of care in the United States prescribes four fixed cycles of 7.4 GBq regardless of individual patient characteristics [29]. Studies demonstrate that patients receiving additional cycles until reaching organ dose limits experience significantly improved survival [25].

This gap between the clinical value of personalized dosimetry and practical implementation barriers motivates development of portable monitoring approaches that complement single-timepoint SPECT/CT imaging with continuous at-home measurements.

1.2.1 *Portable Organ Dosimetry Device*

The PODD enables quantitative gamma-ray measurements outside the hospital setting. PODD measures energy spectra from patients undergoing ^{177}Lu therapy in their homes over a 14-day monitoring period following each treatment cycle.

Figure 1.2 presents the system architecture comprising four subsystems. The radiation detectors, consisting of scintillator crystals coupled to silicon photomultipliers (SiPMs), convert incident gamma photons into electrical signals. The analog front-end board conditions these signals and generates digital pulses whose widths encode the deposited gamma-ray energy using the time-over-threshold technique. The FPGA processing board, the focus of this thesis, performs real-time pulse width measurement and histogram accumulation across all detector channels. A microcontroller provides system coordination and Bluetooth connectivity to a mobile application.

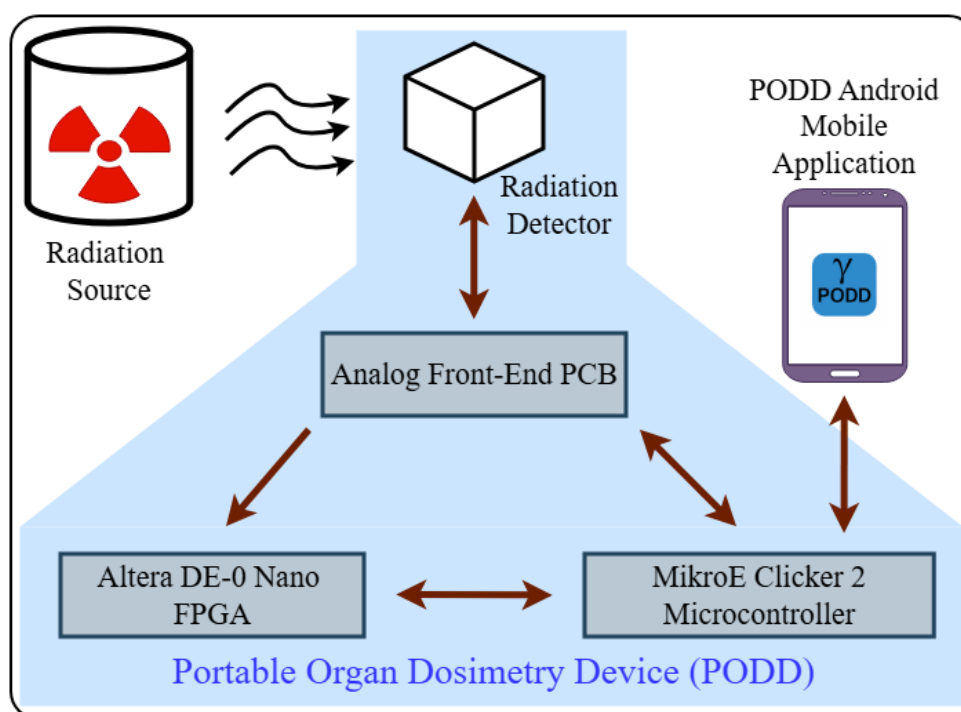


Figure 1.2: System architecture of the PODD. Gamma radiation from the patient interacts with the detectors, producing signals processed by the analog front-end PCB. The FPGA performs real-time histogram computation, communicating with the microcontroller which provides Bluetooth connectivity to the mobile application.

Unlike conventional gamma cameras that produce spatial images, the PODD focuses on

energy histogram, i.e., capturing the energy distribution of detected events. The patient lies in a custom-fitted frame with detector placement optimized based on pre-treatment PET/CT imaging. This registration between detector positions and internal anatomy enables the system to distinguish contributions from different organs. Following ^{177}Lu -DOTATATE administration and a standard 24-hour SPECT/CT scan, patients take the PODD home for daily measurements. The accumulated time-activity data, combined with the calibrated SPECT/CT measurement, enables calculation of organ-specific absorbed doses with temporal resolution exceeding serial imaging alone.

1.3 Signal Processing Requirements and Approach

The digital signal processing architecture must satisfy requirements derived from the clinical application and physical characteristics of radiation detection. This section describes the processing requirements, introduces energy histograms as the primary data representation, and explains the time-over-threshold technique for energy measurement.

1.3.1 Processing Requirements and Energy Histograms

The PODD system must fulfill four key processing requirements. First, the system must capture radiation events across 16 independent detector channels simultaneously, each potentially experiencing event rates of several thousand counts per second during peak activity [19, 18]. Second, energy measurement resolution must distinguish the characteristic photopeaks of ^{177}Lu at 113 keV and 208 keV from scattered radiation and background noise. Third, the system must accumulate data over acquisition periods of minutes at a time over the course of multiple days (e.g., 10–14 daily measurements) while maintaining data integrity. Fourth, the processing chain must operate within power and space constraints of a portable, battery-operated device.

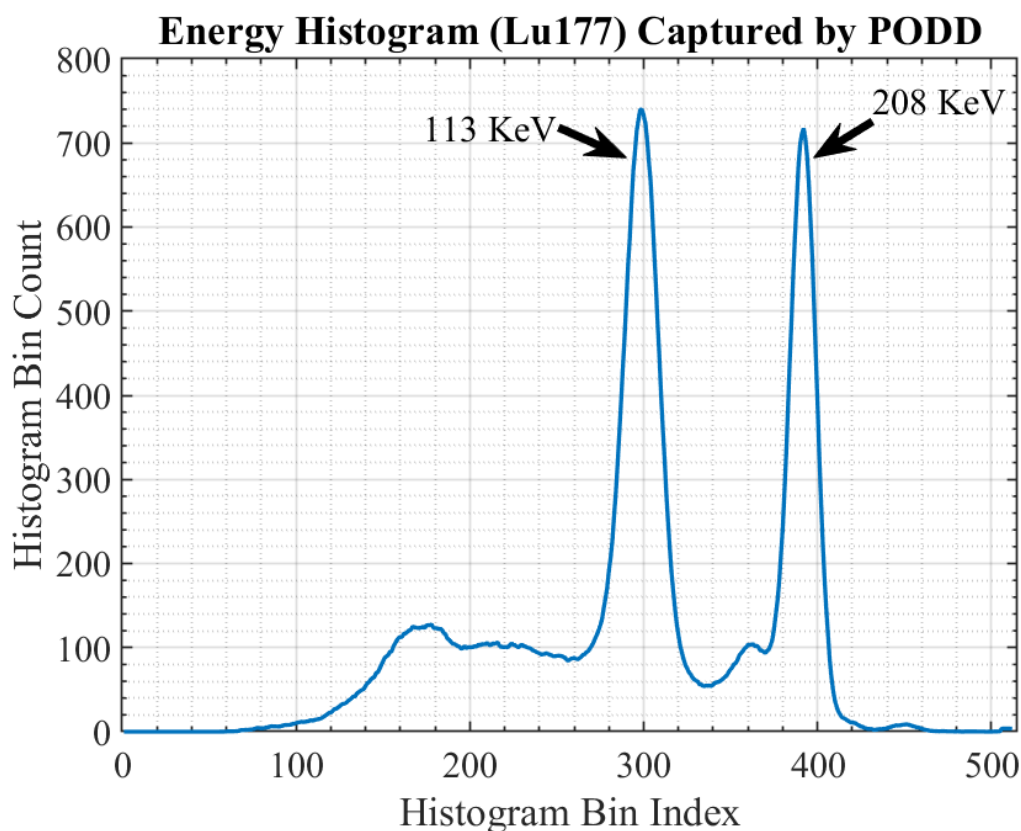


Figure 1.3: Energy histogram of ^{177}Lu captured by the PODD system showing the characteristic photopeaks at 113 keV and 208 keV. The x-axis represents histogram bin index (proportional to pulse width and deposited energy), while the y-axis shows accumulated counts per bin.

The fundamental data structure produced is an *energy histogram*, a statistical representation of gamma-ray energies detected over an acquisition period [16]. Each detection event deposits energy in the scintillator crystal, producing a signal proportional to that energy. The histogram accumulates events by incrementing a counter (bin) corresponding to each measured energy value, building a spectrum that reveals characteristic energy signatures of the radiation source.

Figure 1.3 shows an example ^{177}Lu energy histogram captured by the PODD system.

The two prominent peaks correspond to the 113 keV and 208 keV gamma emissions. Photopeak events occur when a gamma photon deposits its full energy through photoelectric absorption. The continuum at lower energies results from Compton-scattered events where photons scatter within the patient or detector before escaping, depositing only partial energy. By analyzing the histogram shape and quantifying counts in photopeak regions, the system determines detected activity while rejecting scattered contributions.

The FPGA implements a 512-bin histogram for each of the 16 detector channels. Each bin accumulates 16-bit counts, supporting up to 65,535 events per bin, adequate for expected clinical count rates. The histogram bin index maps directly to measured pulse width values, with the relationship between bin number and deposited energy established through calibration.

1.3.2 Time-Over-Threshold Energy Measurement

Time-over-threshold (ToT) is an analog-to-digital conversion technique that encodes signal amplitude as pulse duration [14]. Figure 1.4 illustrates the principle. When a gamma photon interacts with the scintillator, the resulting optical flash produces a current pulse in the SiPM. After amplification, this pulse is compared against a fixed voltage threshold by a high-speed comparator. The comparator output is high while the pulse exceeds the threshold; the duration of this digital pulse is proportional to the original analog amplitude and thus to deposited energy.

Scintillation pulses exhibit fast rise times followed by exponential decay. For pulses of consistent shape but varying amplitude, larger pulses remain above threshold longer. This monotonic relationship enables energy measurement through calibration [21].

The ToT approach offers critical advantages for portable applications. It eliminates high-speed analog-to-digital converters (ADCs) that would require tens to hundreds of




Figure 1.4: Time-over-threshold (ToT) energy measurement. Two scintillation pulses of different amplitudes cross a fixed threshold. The larger pulse (green) remains above threshold longer (ToT_2) than the smaller pulse (red, ToT_1), encoding energy information as pulse width.

megahertz sampling rates to capture fast scintillation pulses [5]. High-speed ADCs consume substantial power and require complex circuitry unsuitable for battery-operated devices. The comparator output interfaces directly with FPGA logic, reducing component count and noise coupling. Pulse width measurement achieves high precision using the FPGA's internal clock.

The PODD system achieves 2.5 ns time binning through double-edge sampling at 200 MHz, implemented within the FPGA fabric. This resolution provides sufficient granularity for ^{177}Lu spectroscopy. Measured pulse widths serve directly as histogram bin addresses, with each event incrementing the corresponding bin count. Over the acquisition period, the histogram reveals the characteristic spectral shape shown in Figure 1.3, enabling quantitative activity estimation.

1.4 *Thesis Organization*

This thesis documents the design, implementation, and validation of the FPGA-based histogram computation system.

Chapter 2 presents the complete system description, detailing the three-board hardware architecture, scintillation detector physics, analog front-end circuitry, FPGA processing logic including double-edge sampling and clock domain crossing, microcontroller coordination, and mobile application interface.

Chapter 3 describes testing and validation methodology, including circuit simulations, digital testbenches, and experimental validation with radiation sources.

Chapter 4 presents conclusions and evaluates system performance.

Chapter 5 discusses opportunities for enhancement and future scope of the project.

1.5 *Individual Contributions*

The PODD system is a collaborative project. My primary contribution to this project is the complete FPGA development and testing, including the pulse width measurement architecture, histogram computation logic, clock domain crossing mechanisms, and the UART communication protocol implementation. Ethan developed the initial microcontroller firmware, the first version of the Android mobile application, and the majority of the analog front-end PCB layout.

Beyond the FPGA work, I made several additional contributions across the system:

Android Mobile Application: I revamped the mobile application to include features not present in the initial version:

- Implementation of distinct clinical and developer operation modes
- Real-time visualization of the temperature profile during data acquisition

- Histogram visualization and analysis capabilities post data acquisition
- User interface for configuring analog front-end board settings (threshold, hysteresis, bias voltage)
- Integration of audible heartbeat feedback for scan progress, adhering to medical device sound standards

Microcontroller Firmware: I integrated the thermistor sensing capability into the firmware, enabling temperature monitoring of the 16 SiPM channels.

PCB Hardware: I tested the analog front-end PCB and identified and fixed issues in the layout. Additionally, I designed the detector probe PCB layouts and validated their functionality.

Chapter 2

SYSTEM DESCRIPTION

This chapter presents the hardware architecture of the PODD, describing each subsystem’s design and interconnection. The system employs a modular three-board architecture that separates digital processing, analog front-end, and system control functions onto dedicated printed circuit boards. This approach facilitates independent development and testing of each subsystem while enabling compact integration for portable deployment.

Figure 2.1 presents the complete PODD system, showing the hardware stack connected to a subset of the radiation detectors via ribbon cable. The system operates from a compact lithium-polymer battery, enabling fully portable operation during patient measurements.

2.1 *Three-Board Sandwich Architecture*

The PODD electronics are organized as a vertically stacked assembly of three interconnected PCBs, referred to as the “three-board sandwich.” This architecture physically separates distinct functional domains while maintaining a compact form factor through board-to-board header connections. Figure 2.2 shows the assembled stack with each layer labeled.

The **bottom layer** houses the Altera DE0-Nano FPGA development board, which serves as the computational core responsible for real-time histogram computation on the incoming time-over-threshold signals from all 16 detector channels. The FPGA interfaces with the middle layer through GPIO headers that carry the digital ToT input signals and control lines.

The **middle layer** is the custom analog front-end PCB that converts analog scintillation pulses from the detectors into digital signals. Each of the 16 channels includes amplification




Figure 2.1: The complete PODD system showing the three-board hardware stack, battery, ribbon cable interface, and a subset of the radiation detectors. The USB connector provides power input and serial debugging capability during development.




Figure 2.2: The three-board sandwich assembly showing the vertically stacked architecture. The Altera DE0-Nano FPGA board (bottom) connects to the custom analog front-end PCB (middle) via GPIO headers, while the MikroE Clicker 2 microcontroller (top) provides system coordination and wireless connectivity.

and comparison stages that produce time-over-threshold pulses encoding the deposited gamma energy. This board also generates the high-voltage bias for the SiPMs and provides temperature monitoring through an array of thermistors. The analog board connects to the external detectors via 50-pin ribbon cables.

The **top layer** contains the Mikro-e Clicker 2 for PSoC 6 microcontroller board, which serves as the system coordinator. The microcontroller communicates with the FPGA via UART to retrieve histogram data, manages system configuration parameters, and provides Bluetooth Low Energy connectivity to the mobile application. The microcontroller board supplies power to the entire stack through its USB or battery input.

The scintillation detectors, comprising 16 GAGG:Ce crystals coupled to SiPMs, is external to the three-board sandwich and connects to the middle analog board via ribbon cable. This separation allows the detector assembly to be positioned against the patient while the electronics stack remains accessible.

2.2 *Radiation Detectors*

The radiation detection subsystem converts incident gamma photons into electrical signals through a two-stage process: scintillation and photodetection. When a high-energy gamma photon enters the scintillator crystal, it interacts with the material through photoelectric absorption or Compton scattering, depositing energy that excites electrons in the crystal lattice. The excited electrons produce visible scintillation photons proportional in number to the deposited energy [16].

These scintillation photons are detected by a silicon photomultiplier (SiPM), a solid-state photodetector composed of an array of independently operating microcells connected in parallel on a common silicon substrate. Figure 2.4 illustrates the SiPM structure: each microcell consists of a single-photon avalanche diode (SPAD) biased above its breakdown

voltage and an integrated quenching resistor. When a scintillation photon triggers a microcell, a standardized current pulse is produced. The total SiPM output is the sum of all simultaneously firing microcells, providing an analog signal amplitude proportional to the number of detected photons and thus to the deposited gamma energy [24].

The PODD system uses SiPMs for light detection. SiPMs offer several advantages for portable applications: low operating voltage (typically 53–56 V), compact solid-state form factor compatible with wearable detector modules, insensitivity to magnetic fields, and mechanical ruggedness [1]. Figure 2.3 shows the detector components used in the PODD, including the GAGG:Ce scintillator crystal, SiPM, and tungsten collimators that define the detector field of view.




Figure 2.3: Components of the PODD radiation detector module: 14 mm³ tungsten collimators with apertures for directional sensitivity, a 6×6 mm² SiPM (Hamamatsu S13360-6050VE), and a 6×6×3 mm³ GAGG:Ce scintillator crystal. A US dime provides scale reference.

2.2.1 GAGG:Ce Scintillator Crystals

The PODD employs cerium-doped gadolinium aluminum gallium garnet (GAGG:Ce) scintillator crystals with dimensions of 6×6×3 mm³. GAGG:Ce was selected for its combination

of high light yield (approximately 50,000–60,000 photons/MeV), fast decay time (88 ns primary component), and non-hygroscopic nature that simplifies packaging requirements [15]. The material’s effective atomic number ($Z_{\text{eff}} \approx 54$) provides good photoelectric absorption efficiency for the 113 keV and 208 keV gamma emissions characteristic of ^{177}Lu .

When a gamma photon interacts with the scintillator via photoelectric absorption or Compton scattering, the deposited energy excites electrons in the crystal lattice. These electrons transfer energy to cerium activator ions, which subsequently de-excite by emitting visible photons. The number of scintillation photons produced is proportional to the energy deposited, forming the physical basis for spectroscopic measurement [16].

2.2.2 Silicon Photomultipliers




Figure 2.4: Structure of a silicon photomultiplier (SiPM). Left: cross-section of three microcells showing the avalanche region, quenching resistors (R_Q), and common substrate. Center: top view of the microcell array layout. Right: equivalent electrical circuit with each microcell consisting of an avalanche photodiode and quenching resistor connected in parallel. Figure from Piatek, Hamamatsu Photonics [24].

Each scintillator crystal is optically coupled to a Hamamatsu S13360-6050VE silicon photomultiplier (SiPM). The S13360 series SiPMs feature a $6 \times 6 \text{ mm}^2$ active area matched to the scintillator face, with $50 \mu\text{m}$ microcell pitch providing approximately 14,400 microcells per device [9]. As illustrated in Figure 2.4, each microcell operates as an independent

avalanche photodiode biased above its breakdown voltage, producing a standardized current pulse when triggered by a photon. The total SiPM output is proportional to the number of detected photons and thus to the deposited gamma energy.

The SiPM bias voltage, typically 53–56 V for these devices, is supplied by the programmable high-voltage module on the analog board. Operating voltage must be carefully controlled because SiPM gain exhibits strong dependence on overvoltage (the difference between bias and breakdown voltage) [1]. The breakdown voltage itself varies with temperature, necessitating either temperature compensation or monitoring to maintain consistent energy calibration. The thermistor array on the analog board provides the temperature data required for offline gain correction.

2.2.3 Detector Probe PCB Design

To integrate the SiPM, thermistor, and connector into a compact detector module, custom detector probe PCBs were designed. Figure 2.5 shows the circuit schematic common to both PCB variants. The design includes the SiPM (M1), a decoupling capacitor (C1) on the cathode, a current-limiting resistor (R1), an NTC thermistor (RT1) for temperature monitoring, and a miniature board-to-board connector (J1) for interfacing with the main analog front-end PCB via ribbon cable.

Two PCB variants were designed to accommodate different SiPM sizes for application flexibility. The smaller variant measures 6×8 mm and accommodates the Hamamatsu S13360-3050VE SiPM with a 3.4×3.4 mm² active area. The larger variant measures 7×8 mm and accommodates the S13360-6050VE SiPM with a 6.4×6.4 mm² active area used in the current PODD prototype. Figure 2.6 presents the PCB layouts for both variants.

Both PCB designs use a four-layer construction with components placed on the top layer and the SiPM mounted on the bottom layer facing the scintillator crystal. The two

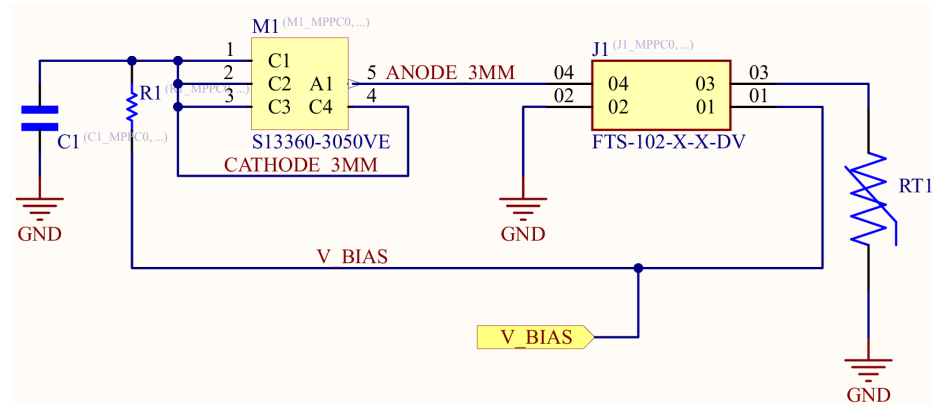



Figure 2.5: Schematic of the detector probe PCB showing the SiPM (S13360-3050VE), decoupling capacitor (C1), current-limiting resistor (R1), NTC thermistor (RT1) for temperature compensation, and miniature connector for ribbon cable interface (J1).

inner layers provide dedicated ground and power planes, improving signal integrity and noise immunity in the compact form factor. Each detector probe is housed within its own individual tungsten collimator, providing radiation shielding and directional sensitivity on a per-channel basis. This modular collimator design enables flexible probe placement while maintaining channel isolation.


2.3 Analog Front-End PCB

The analog front-end board implements 16 parallel signal conditioning channels, power supply generation, and temperature monitoring circuitry on a custom 4-layer PCB measuring 5.3×2.9 inches, with all components mounted on a single side. Figure 2.7 shows the assembled board layout, with the 16 signal conditioning channels arranged symmetrically on either side of the central power supply section.


Each channel converts the analog scintillation pulse from its corresponding SiPM into a digital time-over-threshold signal suitable for FPGA processing. The board interfaces with the detectors via 50-pin ribbon connectors and communicates with the digital processing




(a) 3.4 mm SiPM variant – top



(b) 3.4 mm SiPM variant – bottom



(c) 6.4 mm SiPM variant – top



(d) 6.4 mm SiPM variant – bottom

Figure 2.6: Detector probe PCB layouts for two SiPM size variants: (a, b) 6×8 mm PCB for the 3.4×3.4 mm² S13360-3050VE SiPM; (c, d) 7×8 mm PCB for the 6.4×6.4 mm² S13360-6050VE SiPM. Top views show component placement including passive components and connector; bottom views show the SiPM landing pattern and routing.




Figure 2.7: The analog front-end PCB (5.3×2.9 inches) with 16 parallel ToT channels arranged symmetrically around the central power supply section. Ribbon cable connectors (J1, J2) interface with the detector array, while board-to-board headers connect to the FPGA below and microcontroller above.

board through dedicated GPIO headers.

2.3.1 Channel Architecture

Each of the 16 analog channels implements a two-stage signal conditioning topology: a high-speed transimpedance amplifier followed by a fast comparator with programmable threshold. Figure 2.8 presents the schematic for a single channel.

The input signal from the SiPM passes through an attenuator network that scales large scintillation pulses within the amplifier’s linear operating range. The amplification stage employs an LT1818 operational amplifier, selected for its 400 MHz gain-bandwidth product and 2500 V/ μ s slew rate, which are essential for preserving sub-microsecond pulse timing. The comparator stage uses an ADCMP601 with 3.5 ns propagation delay; when the amplified

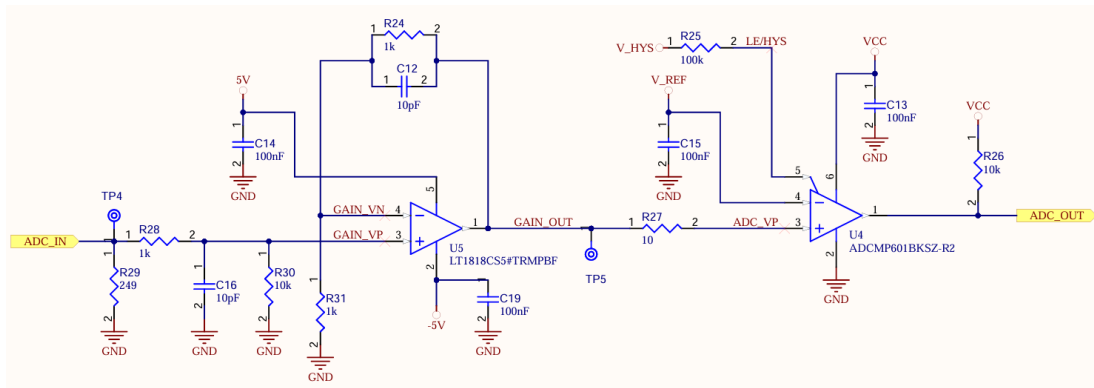


Figure 2.8: Schematic of a single analog channel showing the amplification stage (LT1818) and comparator stage (ADCMP601). The input signal (ADC_IN) from the SiPM is amplified and compared against the programmable threshold voltage (V_REF) to produce the digital ToT output (ADC_OUT).

pulse exceeds the threshold reference voltage (V_REF), the comparator output transitions high, and when the pulse falls below threshold, the output returns low. The resulting digital pulse width encodes the original analog amplitude, the fundamental basis of time-over-threshold energy measurement.

The threshold voltage (V_REF) is generated by a digital-to-analog converter on the microcontroller board and distributed to all 16 channels, enabling software-controlled adjustment of the discrimination level. A dedicated hysteresis control input (V_HYS) provides noise immunity against false triggering from baseline fluctuations.

Figure 2.9 illustrates the importance of comparator hysteresis. Without hysteresis, a noisy input signal crossing a single threshold causes multiple spurious output transitions (red waveform). With hysteresis, the comparator uses separate upper and lower thresholds: the output transitions high only when the input exceeds the upper threshold and returns low only when it falls below the lower threshold (green waveform). This separation prevents noise-induced oscillations near the switching point, ensuring clean digital pulses for accurate time-over-threshold measurement [8].




Figure 2.9: Effect of comparator hysteresis on noisy input signals. Without hysteresis (red output), noise causes multiple false transitions. With hysteresis (green output), separate upper and lower thresholds prevent spurious switching. Note that this inverting comparator produces an active-low output: the output transitions low when the input rises above the upper threshold and returns high when the input falls below the lower threshold [8].

Detailed circuit analysis and component specifications are provided in Appendix A.

2.3.2 Power Supply and Temperature Monitoring

The power supply section generates two critical voltage rails: a programmable high-voltage bias (50–65 V) for the SiPM array using a compact DC-DC converter module, and a negative supply rail (–5 V) for the operational amplifiers using a switched-capacitor voltage inverter. This topology eliminates inductors, reducing electromagnetic interference in the sensitive analog section.

Temperature monitoring is essential for SiPM-based systems because photomultiplier gain exhibits significant temperature dependence (typically 1–3% per °C, depending on overvoltage) [22]. Experimental characterization of the PODD system, presented in Section 3.1.3.2, confirmed a gain temperature coefficient of $-0.90\%/^{\circ}\text{C}$ at the operating over-

voltage of 6 V (Figure 3.8). The board incorporates 16 NTC thermistor channels distributed across the detector interface. To minimize microcontroller ADC requirements, the thermistor signals are routed through two 8:1 analog multiplexers, enabling sequential sampling of all 16 temperatures using only two ADC inputs. The real-time temperature data supports on the fly gain correction during histogram acquisition on the FPGA. The specifics of this process are out of the scope of the current project and will be developed as part of future work.

2.4 *FPGA Processing System*

The FPGA serves as the computational core of the PODD, executing real-time pulse width measurement and histogram accumulation for all 16 detector channels simultaneously. Unlike microcontroller-based approaches that process events sequentially, the FPGA's parallel hardware architecture enables true concurrent processing; each channel operates independently with dedicated logic resources, ensuring no events are missed during high count-rate periods. The design is implemented in Verilog HDL and synthesized using Intel Quartus Prime for the Cyclone IV E device family [13].

2.4.1 *DE0-Nano Development Board*

The PODD utilizes the Altera DE0-Nano development board, shown in Figure 2.10, a compact FPGA platform well-suited for embedded signal processing applications. The board features an Intel Cyclone IV EP4CE22F17C6N FPGA with 22,320 logic elements, 594 Kbits of embedded memory, and 66 multipliers [31].

The design leverages several key DE0-Nano resources. The onboard 50 MHz crystal oscillator serves as the input reference clock, which is multiplied to 200 MHz using an internal phase-locked loop (PLL) to achieve the required time binning for pulse width measurement.




Figure 2.10: The Altera DE0-Nano development board featuring the Intel Cyclone IV E FPGA. The dual 40-pin GPIO headers visible at top and bottom provide connectivity to the analog board.

The dual 40-pin GPIO headers provide connectivity to the analog board, carrying the 16 digital ToT input signals from the comparators along with UART lines for microcontroller communication. Histogram data for all channels is stored in the FPGA's internal M9K embedded memory blocks, configured as dual-port RAM to allow simultaneous histogram updates and data readout. The board's compact form factor and low power consumption make it suitable for integration into the portable device stack.

2.4.2 Multi-Clock Domain Architecture

The histogram computation system employs a multi-clock domain architecture to balance time binning requirements with power efficiency. Figure 2.11 presents the complete signal processing datapath.

The **high-frequency domain** operates at 200 MHz and performs pulse acquisition with 2.5 ns time binning. This domain contains the input synchronizers, edge detection logic,

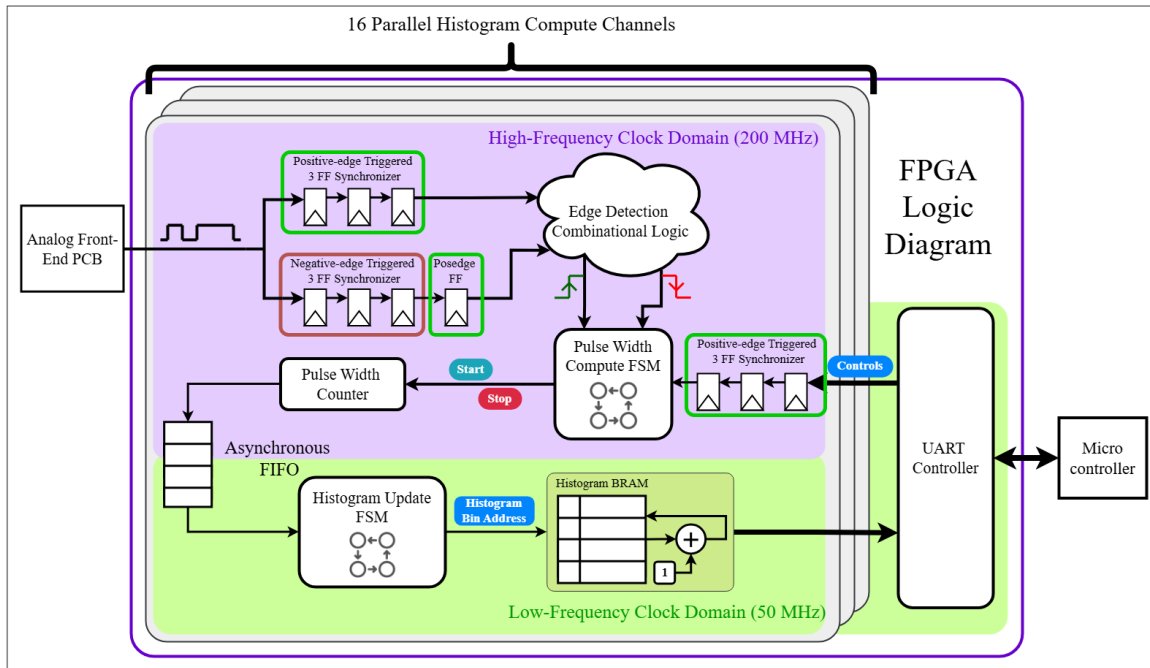


Figure 2.11: Block diagram of the FPGA-based histogram computation system. The high-frequency domain (200 MHz, shown in purple) performs pulse width measurement using double-edge sampling, while the low-frequency domain (50 MHz, shown in green) handles histogram updates and UART communication. Asynchronous FIFOs bridge the clock domains.

and pulse width counters for all 16 channels. The 200 MHz clock is generated from the 50 MHz board oscillator using the FPGA's internal PLL.

The **low-frequency domain** operates at 50 MHz and handles histogram computation, memory management, and UART communication. This domain consumes less power per logic element than the high-speed acquisition logic and provides adequate throughput for the histogram update rate and serial data transfer requirements.

Reliable data transfer between domains is achieved through dual-clock FIFOs (DCFI-FOs) that safely pass pulse width measurements from the acquisition logic to the histogram computation module. Control signals crossing domain boundaries pass through multi-stage

synchronizer chains to prevent metastability.

2.4.3 Double-Edge Sampling for Enhanced Resolution

The pulse acquisition subsystem employs double-edge sampling to achieve 2.5 ns time binning from the 200 MHz clock. The input pulse signal is sampled independently on both the positive and negative edges of the clock, providing an effective 400 MHz sampling rate. This technique enables detection of pulse transitions occurring between consecutive positive clock edges without requiring a faster system clock.

The architecture implements parallel synchronizer chains: a 3-stage positive-edge triggered chain and a complementary 3-stage negative-edge triggered chain. Both chains mitigate metastability from the asynchronous comparator inputs. The negative-edge sampled value is then synchronized to the positive-edge domain through an additional register stage, enabling unified edge detection logic.

Edge detection compares current and previous samples from both synchronizer outputs to identify rising and falling transitions. The logic handles four transition scenarios: edges detected on positive clock edges, edges detected on negative clock edges, and the special cases of very short pulses (less than one clock period) where both transitions occur within a single sampling window.

2.4.4 Pulse Width Measurement

Upon detecting a rising edge, a dedicated counter begins incrementing by 2 on each 5 ns clock cycle, representing the pulse width in units of 2.5 ns. The counter continues until a falling edge is detected, at which point the accumulated count represents the uncorrected pulse width.

A correction algorithm compensates for the phase relationship between pulse transitions

and clock edges. When the rising edge is detected on the negative clock phase and the falling edge on the positive phase (or vice versa), the measured count is adjusted by ± 1 count to maintain measurement accuracy. The corrected pulse width therefore spans a range from 1 (corresponding to the minimum measurable width of 2.5 ns) to a maximum of 512 (corresponding to 1280 ns), with overflow values clamped to the maximum bin address to prevent wraparound errors.

2.4.5 Clock Domain Crossing

Pulse width data from the 200 MHz acquisition domain to the 50 MHz histogram domain via a dual-clock FIFO instantiated from Intel's DCFIFO megafunction. Each channel's FIFO is configured with 32 entries of 16-bit width, with 5-stage synchronization delay pipes on both read and write interfaces to ensure robust metastability protection [6].

The FIFO provides independent full and empty flags for each clock domain, properly synchronized to prevent metastability. The write-side full flag prevents data loss during high pulse rates by stalling the acquisition state machine, while the read-side empty flag ensures the histogram module only processes valid data. The histogram computation module accounts for the FIFO's read latency by implementing a 3-cycle wait state after asserting the read request.

The FIFO depth of 32 entries is sufficient for the expected event rates in radiation dosimetry, but understanding the overflow boundary is useful for characterizing system limits. On the producer side, although the dual-edge sampling technique enables pulse width resolution down to 2.5 ns, the throughput is limited by the state machine which operates on the positive clock edge and requires a minimum of 2 clock cycles at 200 MHz per event (one cycle for rising edge detection, one for falling edge detection). This holds regardless of pulse width: a 2.5 ns pulse still requires a 7.5 ns gap before the next event can

be detected, yielding a maximum write rate of one entry per 10 ns. On the consumer side, the histogram state machine requires 6 clock cycles at 50 MHz to perform a complete read-modify-write update: 1 cycle to assert the FIFO read request, 3 cycles for FIFO read data synchronization (matching the dual-clock FIFO's internal synchronizer pipeline), 1 cycle for the BRAM read-to-output latency, and 1 cycle to write the incremented count back. This gives a consumer drain rate of one entry per 120 ns. At the maximum sustained producer rate, the net accumulation is approximately 11 entries per 120 ns, and the 32-entry FIFO would overflow in roughly 350 ns, corresponding to a burst of approximately 35 back-to-back minimum-width pulses. The FIFO can therefore sustain an event rate up to the consumer drain rate of one event per 120 ns, or approximately 8.3 million counts per second (Mcps) per channel, before overflow occurs. In practice, per-channel count rates during ^{177}Lu dosimetry are on the order of several thousand counts per second during peak activity (Section 1.3.1), approximately three orders of magnitude below this threshold. Furthermore, radioactive decay is a memoryless Poisson process in which successive photon detections are statistically independent [16], so correlated bursts of events cannot physically occur. At a conservatively high peak rate of $\lambda = 10$ kcps per channel, the mean inter-arrival time between detections is $1/\lambda = 100 \mu\text{s}$, four orders of magnitude longer than the 10 ns minimum producer event spacing. Under Poisson statistics, the probability of two successive events arriving within $\Delta t = 10$ ns is $p = 1 - e^{-\lambda\Delta t} \approx 10^{-4}$, and the probability of the ~ 35 successive such arrivals required to overflow the FIFO is $p^{34} \approx 10^{-136}$. The 32-entry FIFO depth is therefore more than adequate for this application.

Control signals from the UART controller (operating in the 50 MHz domain) are synchronized to the 200 MHz acquisition domain using 3-stage synchronizer chains. Critical signals include the start command (initiating pulse capture), stop command (halting acquisition), and reset completion indicator. Status signals from the acquisition domain are

synchronized back to the control domain for system monitoring through identical synchronizer structures.

2.4.6 *Histogram Computation and Storage*

The histogram computation subsystem aggregates statistical distributions by incrementing bin counts corresponding to measured pulse width values. Each of the 16 channels has a dedicated histogram memory instance, enabling independent and simultaneous histogram generation.

Histogram storage uses dual-port on-chip block RAM (BRAM) configured from the Cyclone IV's M9K memory blocks:

- 512 bins per channel (9-bit address width)
- 16 bits per bin (supporting counts up to 65,535)
- Port A: histogram updates via read-modify-write operations
- Port B: non-blocking read access for the microcontroller interface

The histogram update mechanism implements a state machine that initializes all 512 bins to zero on reset by sequentially writing to each address, requiring 513 clock cycles at 50 MHz (approximately 10.3 μ s) to complete. After initialization, the state machine monitors the FIFO for available pulse width data and performs read-modify-write operations to increment bin counts. The pulse width value (lower 9 bits) directly maps to the bin address; the current count is read from BRAM, incremented by one, and written back to the same address. The dual-port architecture enables concurrent access: Port A handles histogram updates from the acquisition pipeline while Port B provides independent read access for data retrieval without disrupting ongoing measurements.

2.4.7 *UART Communication Protocol*

The UART interface provides bidirectional communication between the FPGA and microcontroller at 115200 baud. A custom command-response protocol ensures reliable data exchange between the asynchronous processing units operating at different clock rates and execution models.

The protocol implements a command-response architecture where each command from the microcontroller must be followed by an `END_COMMAND` delimiter byte (0xFF), and the FPGA responds with an acknowledgment sequence terminating with its own `END_COMMAND` byte. Supported commands include:

- `START_HISTOGRAM`: Begin pulse acquisition and histogram accumulation
- `STOP_HISTOGRAM`: Halt acquisition while preserving histogram data
- `CLEAR_RESULTS`: Reset all histogram bins to zero
- `SET_BIN_ADDRESS`: Configure the starting bin address for histogram data transfer (default: 0)
- `SET_NUM_BINS`: Configure the number of bins to transfer starting from the configured bin address (default: 512)
- `SET_CHANNEL`: Select which channel's histogram data to transfer; values 0–15 select a specific channel, while value 16 (default) transfers data from all channels
- `START_UPLOAD`: Initiate histogram data transfer using the previously configured parameters (bin address, number of bins, and channel selection)

Commands requiring parameter values use multi-byte sequences. For example, the histogram start address can range from 0 to 511, requiring 9 bits to represent all possible values. Since UART transfers 8-bit bytes, two bytes are needed to transmit the address. The command format is:

<SET_BIN_ADDRESS> | <ADDR_LSB[7:0]> | <ADDR_MSB[8]> | <END_COMMAND>

Figure 2.12 illustrates this protocol with an example SET_BIN_ADDRESS transaction. The microcontroller transmits the command byte followed by the address value 283 encoded as two bytes (LSB = 0x1B = 27, MSB = 0x01 = 1), terminated by END_COMMAND. The FPGA updates its internal start_bin_address register and acknowledges with its own END_COMMAND response.




Figure 2.12: Example UART transaction for the SET_BIN_ADDRESS command. The microcontroller transmits the command byte, two address bytes (LSB first), and the delimiter. The FPGA updates its start address register from 0 to 283 and responds with END_COMMAND to acknowledge successful reception.

During histogram upload, the FPGA transmits bin data for the requested channels sequentially. Each 16-bit bin count is transmitted as two bytes, with the upper nibble of the second byte containing the channel identifier to enable the microcontroller to verify data integrity. This embedded channel tagging provides error detection without requiring additional protocol overhead. Complete timing diagrams for all supported UART commands are provided in Appendix B.

2.4.8 FPGA Resource Utilization

Table 2.1 summarizes the FPGA resource utilization after synthesis and place-and-route using Intel Quartus Prime 18.1.1 Lite Edition. The design targets the Cyclone IV EP4CE22F17C6 device on the DE0-Nano board.

Table 2.1: FPGA resource utilization for the histogram computation system on the Cyclone IV EP4CE22F17C6 device.

Resource	Used / Available	Utilization
Logic elements	11,331 / 22,320	51%
Registers	8,030	—
Pins	30 / 154	19%
Memory bits	98,304 / 608,256	16%
Embedded multipliers (9-bit)	0 / 132	0%
PLLs	1 / 4	25%

The design consumes approximately half of the available logic elements, primarily due to the 16 parallel channel instantiations each containing pulse acquisition logic, synchronizers, and histogram computation state machines. The memory utilization of 16% reflects the 16 independent histogram memories, each storing 512 bins of 16-bit counts (8,192 bits per channel, totaling 131,072 bits for histogram storage plus FIFO buffers). The single PLL generates the 200 MHz high-frequency clock from the 50 MHz board oscillator. No embedded multipliers are required since the histogram computation relies solely on increment operations.

2.5 Microcontroller Subsystem

The microcontroller serves as the system coordinator, bridging the FPGA's histogram computation engine with the mobile application's user interface. It handles configuration management, orchestrates data transfers with the FPGA via UART, provides wireless con-

nectivity through Bluetooth Low Energy (BLE), and supplies power to the entire three-board stack.

2.5.1 PSoC 6 Hardware Platform

The PODD employs the Mikro-e Clicker 2 for PSoC 6 development board, shown in Figure 2.13, featuring the Infineon PSoC 6 microcontroller. The PSoC 6 provides a dual-core architecture, with an ARM Cortex-M4 running at 150 MHz for application processing and an ARM Cortex-M0+ dedicated to the BLE radio stack [12]. This architecture enables efficient concurrent handling of wireless communication and system control tasks without real-time scheduling conflicts.




Figure 2.13: The Mikro-e Clicker 2 for PSoC 6 development board featuring dual mikroBUS sockets, USB connectivity, and battery charging circuitry. The board provides BLE wireless capability and serves as the system coordinator in the PODD stack.

The firmware utilizes several PSoC 6 peripherals for PODD operation:

- **BLE module:** Provides wireless connectivity to the Android application via GATT-

based data exchange

- **UART:** Communicates with the FPGA for command transactions and histogram retrieval
- **12-bit DAC:** Generates the threshold reference voltage distributed to all comparator channels
- **PWM output:** Produces the filtered hysteresis control voltage
- **SAR ADC:** Samples multiplexed thermistor outputs for temperature monitoring

The Clicker 2 board provides USB and battery power input options, with onboard voltage regulators distributing power to the FPGA and analog boards through the stacking headers.

2.5.2 *Firmware Architecture*

The PSoC 6 firmware leverages the device's rich peripheral set to interface with the FPGA, control analog circuit parameters, and communicate wirelessly with the mobile application. Figure 2.14 presents the peripheral configuration implemented in PSoC Creator, showing the interconnection of communication interfaces, analog control outputs, and sensor acquisition channels.

2.5.2.1 *Communication Interfaces*

The microcontroller employs two UART interfaces: one dedicated to communication with the DE0-Nano FPGA for histogram data transfer and control commands, and a secondary UART for external debugging and development purposes. The integrated BLE 5.0 module handles wireless communication with the Android mobile application, providing GATT-based services for device configuration, scan control, and histogram data streaming. Status

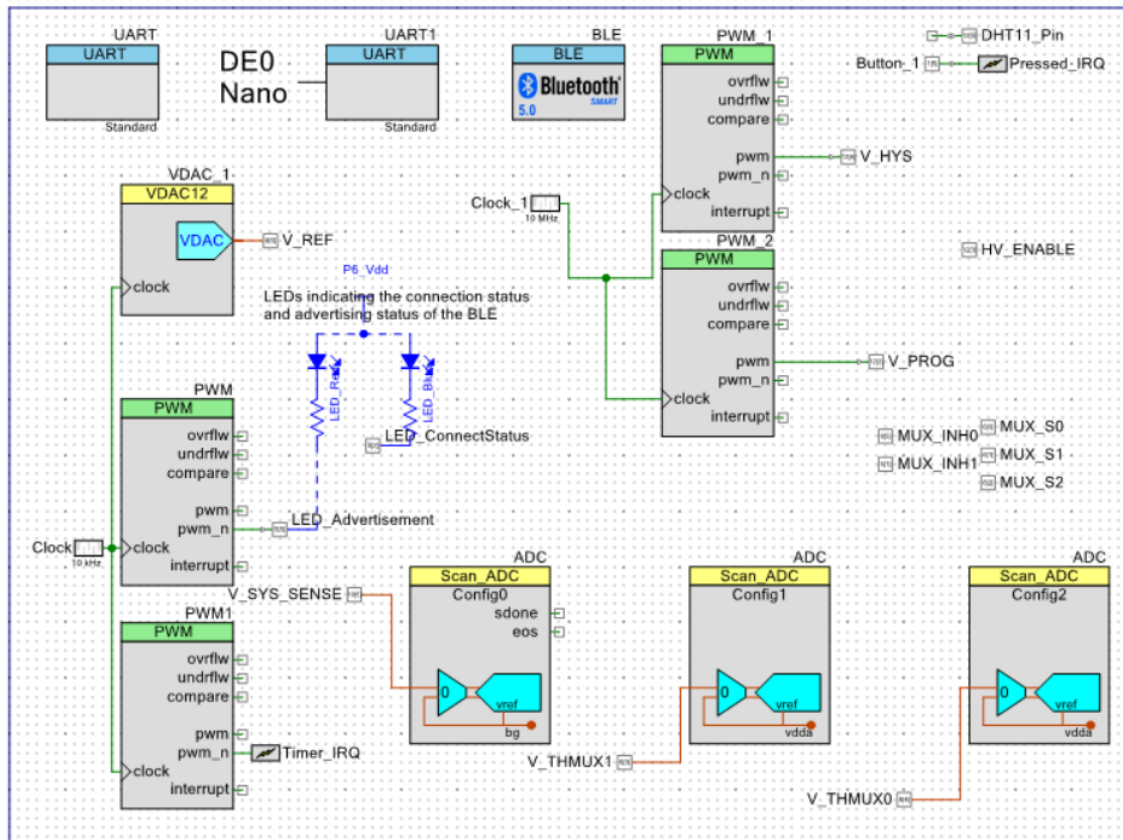


Figure 2.14: PSoC 6 microcontroller block diagram showing peripheral configurations: dual UART interfaces for debugging and FPGA communication, BLE 5.0 module for wireless connectivity, 12-bit VDAC for threshold voltage generation (V_REF), PWM outputs for hysteresis control (V_HYS) and programmable bias voltage (V_PROG), and three SAR ADC configurations for thermistor temperature acquisition through two 8:1 analog multiplexers.

LEDs indicate BLE advertising and connection states, providing visual feedback to the developer.

2.5.2.2 Analog Control Outputs

The 12-bit voltage DAC (VDAC12) generates the programmable threshold voltage (V_REF) that sets the comparator discrimination level on the analog front-end board. Two PWM modules produce filtered analog voltages: V_HYS controls the comparator hysteresis, while V_PROG sets the programmable SiPM bias voltage through the high-voltage DC-DC converter. The HV_ENABLE signal provides software control over the high-voltage supply, allowing the bias to be disabled during configuration changes or fault conditions.

2.5.2.3 Temperature Acquisition

Temperature monitoring of the 16 SiPM channels is accomplished through two 8:1 analog multiplexers (MUX_INH0, MUX_INH1 with shared select lines MUX_S0–S2) that route thermistor voltages to the SAR ADC inputs. The firmware sequentially scans all 16 channels, converting the measured voltages to temperature values using stored calibration coefficients.

2.5.2.4 Real-Time Operating System

The firmware is structured around FreeRTOS with dedicated tasks for BLE communication, FPGA control, and analog parameter management. The BLE task handles connection management, characteristic notifications, and incoming command parsing. The FPGA control task manages the UART protocol state machine, issuing commands and processing histogram data during upload sequences. The analog management task handles two independent functions: periodic sampling of the thermistor array for temperature monitoring, and DAC output adjustment for threshold control. A hardware timer generates periodic

interrupts (Timer_IRQ) for time-critical operations including ADC scan scheduling.

Configuration parameters, including scan duration, bias voltage, threshold, and hysteresis settings, are stored in non-volatile memory and restored at startup. The microcontroller validates all parameter changes against hardware limits before applying them, preventing out-of-range values from damaging the analog circuitry or corrupting measurements.

2.6 *Mobile Application Interface*

The Android mobile application serves as the primary user interface for the PODD, enabling clinicians and researchers to interact with the embedded hardware through an intuitive touchscreen interface. The application handles the complete data acquisition workflow from device discovery and connection through histogram visualization and data export. Figure 2.15 presents the key application screens.

The main screen (Figure 2.15a) provides the clinical interface with device connection status, scan timer display, and primary control buttons. Developer mode (Figures 2.15b and 2.15c) exposes the full parameter configuration interface, allowing adjustment of scan duration, bias voltage, threshold voltage, and hysteresis voltage, with options to save and load configuration presets. Additional visualization screens for histogram and temperature data are documented in Appendix C.

2.6.1 *BLE Communication*

The application communicates with the microcontroller through BLE, establishing a GATT client-server architecture where the microcontroller operates as the peripheral server. Upon startup, the application scans for devices advertising the PODD service, initiates connection with Maximum Transmission Unit (MTU) negotiation (requesting 512 bytes for efficient histogram transfer), and discovers available characteristics.

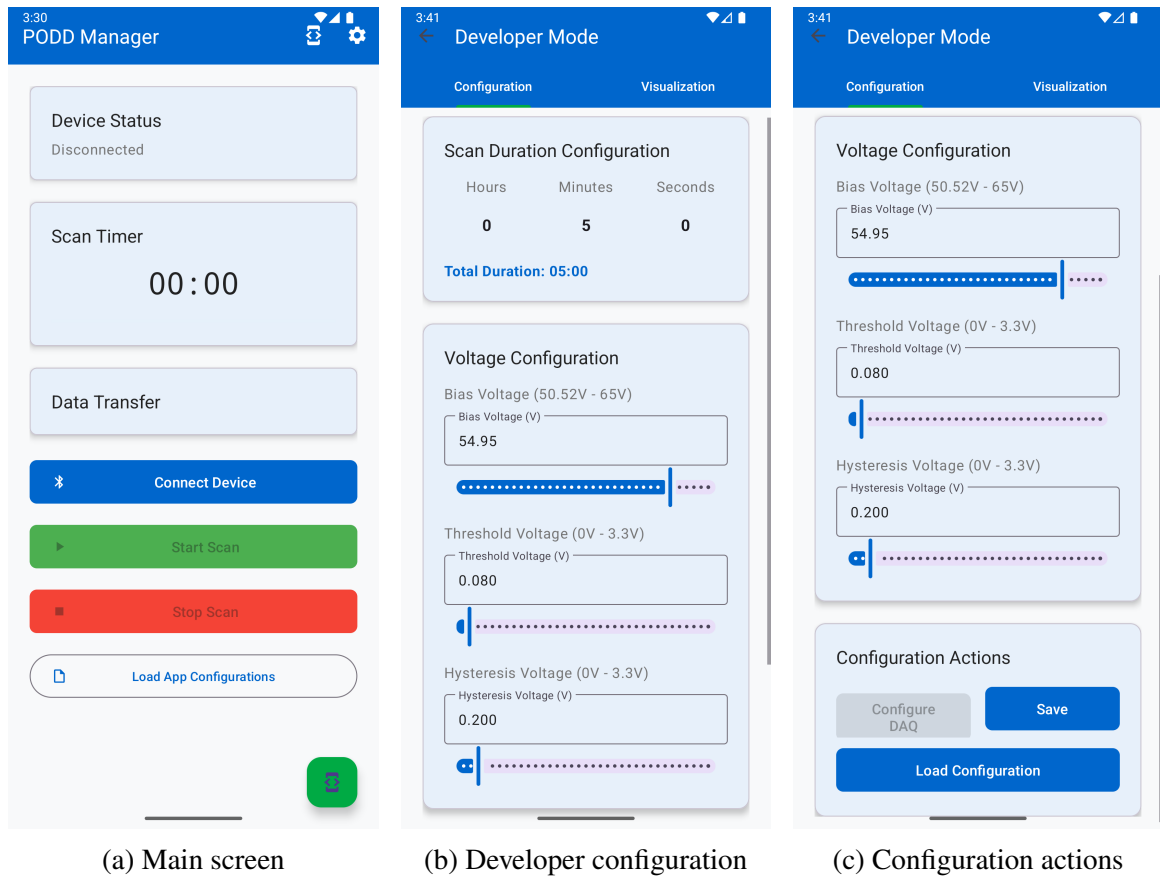


Figure 2.15: PODD mobile application screens: (a) Main screen showing device status, scan timer, and control buttons for clinical operation; (b) Developer mode configuration tab with scan duration and voltage parameter controls; (c) Configuration actions panel with buttons to apply settings to hardware, save configurations, or load previously saved presets.

The command protocol uses single-character commands for common operations (scan start, stop, settings query) and structured multi-character commands prefixed with CONFIG: for parameter updates. Bidirectional handshaking ensures parameter changes are verified before subsequent operations proceed.

2.6.2 Operational Modes

The application provides two operational modes tailored to different user requirements:

Clinical mode presents a simplified interface for medical professionals, or patients, offering one-touch scan initiation with pre-configured parameters, real-time countdown display, and streamlined result viewing. This mode minimizes cognitive load during patient measurements.

Developer mode exposes comprehensive system parameters for research and calibration purposes. Users can adjust bias voltage (50.52–65 V), threshold voltage (0–3.3 V), hysteresis voltage (0–2.8 V), and scan duration (1 second to 12 hours). The interface provides real-time histogram visualization using the MPAndroidChart library, with interactive zoom, pan, and channel selection capabilities. An automated peak identification algorithm locates characteristic photopeaks for energy calibration, and temperature profile visualization displays thermal data from the 16-channel thermistor array alongside histogram results.

Chapter 3

SYSTEM TESTING & RESULTS

This chapter presents the systematic testing methodology employed to validate each subsystem of the PODD before integration. A staged approach was adopted, beginning with isolated component-level tests and progressing to full system characterization with radioactive sources. This methodology enables identification and resolution of issues at each stage, reducing debugging complexity during system integration.

3.1 *Component-Level Testing*

3.1.1 *Analog Front-End PCB Functionality Tests*

Before introducing radioactive sources, the analog front-end PCB was validated using controlled electrical test signals. This approach allows precise characterization of circuit behavior without the variability inherent in radiation measurements.

3.1.1.1 *Test Setup*

Figure 3.1 shows the test configuration for analog board validation. The PCB was powered by an external bench supply providing the required 5 V and 3.3 V rails. A function generator (not shown) produced a low-frequency sinusoidal signal to emulate the analog output from a SiPM-scintillator detector. This signal was injected into one of the 16 input channels, and oscilloscope probes monitored signals at three key points in the signal chain: the channel input, the gain stage output, and the comparator output (ToT signal).




Figure 3.1: Analog front-end PCB under test. The board is powered via external bench supply connections (red, yellow wires). The EMCO SIP90 high-voltage DC-DC converter module is visible in the center. Oscilloscope probes (orange wires) monitor signals at various test points along the signal conditioning chain.

3.1.1.2 Channel Output Verification

The signal conditioning channel was verified by applying a sinusoidal input and measuring the response at each stage. Figure 3.2 presents the oscilloscope capture showing the three monitored signals.

The measured gain of 2.0 closely matches the designed gain of 2, confirming correct operation of the amplification stage. The green waveform (gain stage output) exhibits the expected doubling of amplitude relative to the yellow input signal while preserving the sinusoidal wave shape, indicating linear operation within the tested amplitude range.

The blue ToT output demonstrates proper comparator functionality. The digital pulses transition high when the amplified sinusoidal signal exceeds the programmed threshold

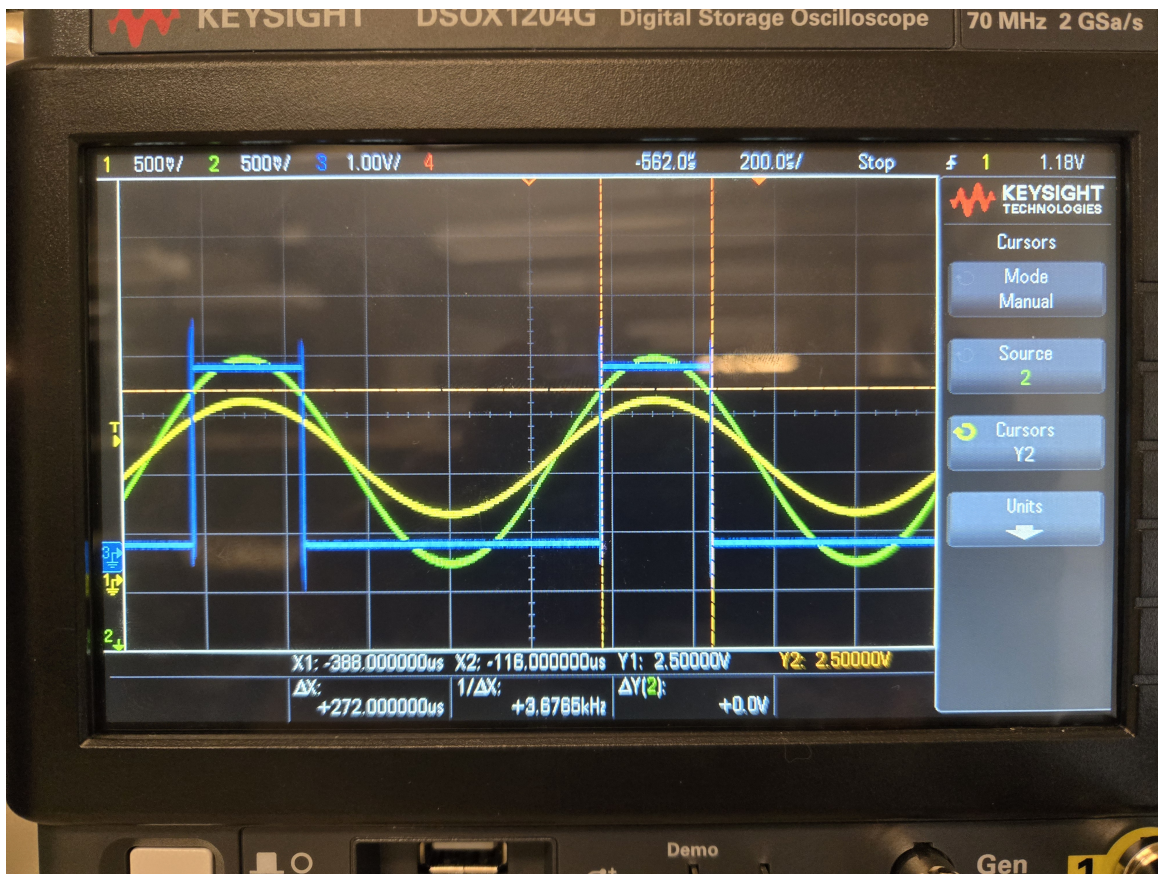


Figure 3.2: Oscilloscope capture showing signal propagation through the analog channel. Yellow (Ch1, 500 mV/div): input sinusoidal test signal with $\sim 1.0 V_{pp}$ amplitude. Green (Ch3, 1.00 V/div): gain stage output showing $\sim 2.0 V_{pp}$ amplitude, confirming the designed gain of 2. Blue (Ch2, 500 mV/div): comparator output (ToT signal) producing digital pulses when the amplified signal exceeds the threshold. Time base: $200 \mu s/div$.

voltage (approximately 2.5 V, as indicated by the oscilloscope cursor measurements Y1 and Y2). The pulse width varies with the instantaneous signal amplitude. Wider pulses correspond to higher-amplitude portions of the input waveform where the signal remains above threshold longer. This time-domain encoding forms the basis of the time-over-threshold measurement technique.

3.1.2 FPGA Logic Functionality Tests

Validating the FPGA digital logic required a two-phase approach: software simulation using ModelSim for detailed behavioral verification, followed by hardware-based FPGA-in-the-Loop (FIL) testing to confirm correct operation in the actual silicon.

3.1.2.1 ModelSim Simulations

Individual FPGA modules were verified using ModelSim simulations with targeted test vectors. This simulation-based approach enables inspection of internal signals that are inaccessible once the design is synthesized, facilitating rapid debugging of state machine logic and timing relationships.

Figure 3.3 presents a representative simulation of the BRAM read/write controller finite state machine. The testbench (TB SIGNALS section) first asserts the write enable signal (`wr_en1`) to store sequential data values (0x000–0x009) at corresponding addresses. After the write sequence completes, the read enable signal (`rd_en1`) is asserted to retrieve the stored data. The BRAM SIGNALS section shows the dual-port memory interface: Port A (`address_a`, `data_a`, `wren_a`) handles the write operations while Port B (`address_b`, `q_b`, `rden_b`) performs the read operations. The simulation confirms that data written to each address is correctly retrieved during the subsequent read phase, validating the memory controller state machine logic.

Similar simulation testbenches were developed for the UART controller module, the histogram acquisition state machine, and the clock domain crossing synchronizers. Each module was verified in isolation before integration into the complete design.

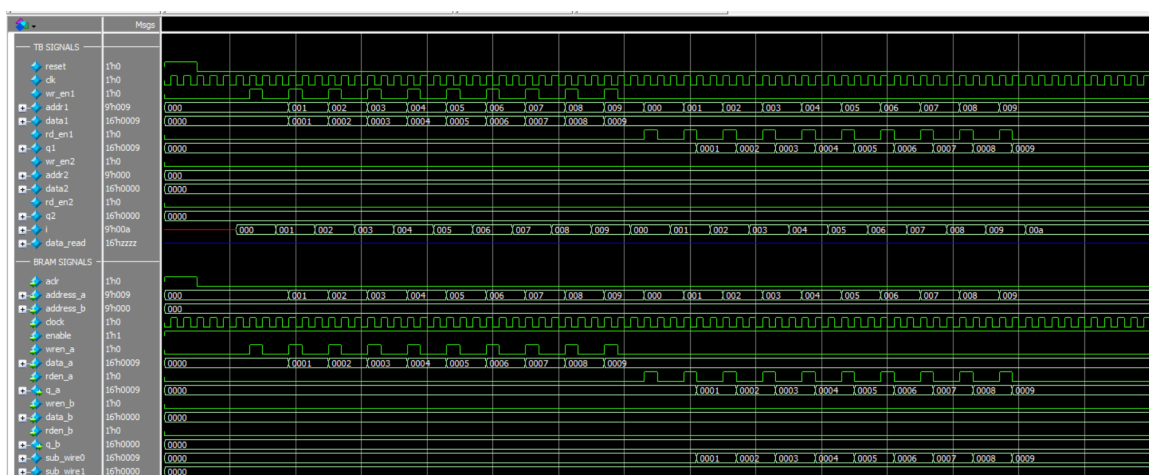


Figure 3.3: ModelSim simulation waveforms for the BRAM read/write controller. The upper section shows testbench control signals initiating sequential write operations (addresses 0x000–0x009) followed by read operations. The lower section displays the dual-port BRAM interface signals, confirming correct data storage and retrieval.

3.1.2.2 FPGA-in-the-Loop Test

While ModelSim simulations verify logical correctness, they cannot account for physical timing variations, signal integrity issues, or synthesis tool optimizations that may affect hardware behavior. FPGA-in-the-Loop (FIL) testing addresses these concerns by exercising the actual programmed FPGA with known test patterns and verifying the results against expected values.

Figure 3.4 illustrates the FIL test environment. A Python application running on a host computer communicates with the DE0-Nano FPGA board through a USB-to-UART TTL cable. The Python script implements the same UART command-response handshake protocol described in Section 2.4.7, enabling direct interaction with the FPGA’s histogram acquisition and readout logic. Programmable test pulses with known widths are generated internally by the FPGA’s test pattern generator and fed into the histogram processing pipeline.

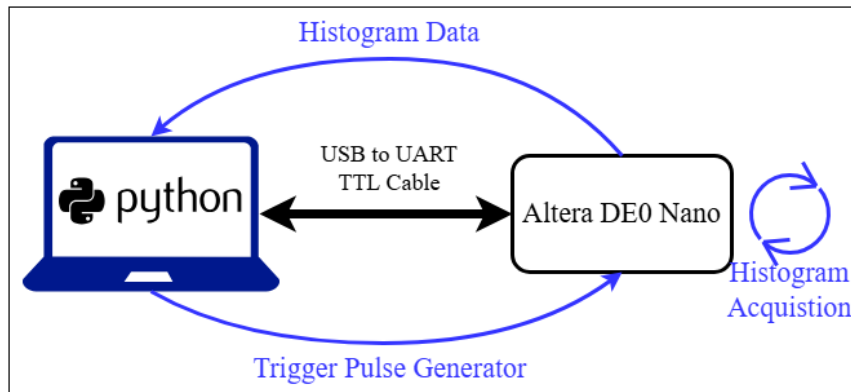


Figure 3.4: FPGA-in-the-Loop test configuration. A Python application communicates with the DE0-Nano FPGA via USB-to-UART interface, issuing commands to start/stop histogram acquisition and retrieve accumulated data. Internal test pulses exercise the histogram processing logic, and the resulting histogram is compared against expected values.

Figure 3.5 presents results from a representative FIL test using 800 pulses with varying widths distributed across the histogram bin range. The top panel shows the binary pulse waveform pattern driven into the histogram acquisition logic, a sequence of back-to-back pulses with different durations. The middle panel displays the expected histogram computed by the Python test harness based on the known pulse widths. The bottom panel shows the actual histogram data retrieved from the FPGA via UART. The exact match between expected and actual histograms confirms correct operation of the entire signal processing chain: pulse width measurement, bin address computation, BRAM increment operations, and UART data retrieval.

The UART communication log below demonstrates the command sequence executed during histogram retrieval. The test script issues `START_HISTOGRAM` to begin acquisition, `STOP_HISTOGRAM` to halt accumulation, `SET_NUM_BINS` to configure the transfer count (512 bins), `SET_BIN_ADDRESS` to specify the starting address, and finally `START_UPLOAD` to initiate data transfer. The FPGA responds to each command with an `END_COMMAND`




Figure 3.5: FPGA-in-the-Loop histogram verification results. Top: test pulse waveform with varying pulse widths. Middle: expected histogram computed from known pulse parameters. Bottom: actual histogram retrieved from FPGA.

acknowledgment, confirming successful protocol handshaking.

Listing 3.1: Python console output showing UART command sequence for histogram retrieval.

```

Available serial ports:
0: COM8    1: COM11

Select the serial port index: 0

Connected to COM8.

Sending command: START_HISTOGRAM

Sending command: END_COMMAND

Received END_COMMAND.

```

```
Sending command: STOP_HISTOGRAM
Sending command: END_COMMAND
Received END_COMMAND.

Sending command: SET_NUM_BINS
Sending Byte: 0
Sending Byte: 2
Sending command: END_COMMAND
Received END_COMMAND.

Sending command: SET_BIN_ADDRESS
Sending Byte: 0
Sending Byte: 0
Received END_COMMAND.

Calculated expected transfers: 1025

Sending command: START_UPLOAD
Sending command: END_COMMAND

Time taken to receive 1025 responses: 0.090 seconds
Histogram data saved to histogram_data.csv
```

The measured transfer time of 90 ms to retrieve one channel's histogram data (512 bins \times 16 bits = 1 KB) at 115200 baud represents an effective throughput of approximately 11.4 KB/s. This performance is consistent with the theoretical maximum of 11.52 KB/s

(115200 baud / 10 bits per UART frame), indicating minimal protocol overhead and efficient FPGA response timing. For the complete 16-channel system, full histogram retrieval requires approximately 1.44 seconds, well within acceptable limits for clinical workflow where scan durations typically range from minutes to hours.

3.1.3 Full System Test

Following successful validation of individual subsystems, the complete PODD was tested with radioactive sources to verify end-to-end system functionality. These tests confirm that the integrated system (i.e., detectors, analog front-end, FPGA processing, microcontroller coordination, and mobile application) correctly acquires and displays energy histograms as illustrated in the system overview (Figure 2.1).

3.1.3.1 Radioactive Isotopes as Radiation Source

Two radioactive isotopes with well-characterized gamma emission spectra were used to validate the PODD's spectroscopic performance: ^{177}Lu and ^{22}Na .

^{177}Lu Measurement. Lutetium-177 is the primary therapeutic isotope for which the PODD is designed. It emits gamma photons at 113 keV and 208 keV, producing two distinct photopeaks in the energy histogram. Figure 3.6(a) shows the ^{177}Lu histogram captured by the mobile application with the SiPM bias voltage set to 59 V. The peak analysis algorithm automatically identified two peaks at histogram bin indices 287 and 383, corresponding to the 113 keV and 208 keV emissions respectively. The clear separation between peaks and the low background between them indicate good energy resolution.

^{22}Na Measurement. Sodium-22 serves as a calibration source due to its well-defined 511 keV annihilation peak (from positron emission) and 1275 keV gamma emission. Figure 3.6(b) presents the ^{22}Na histogram acquired with a bias voltage of 54.95 V. The prominent

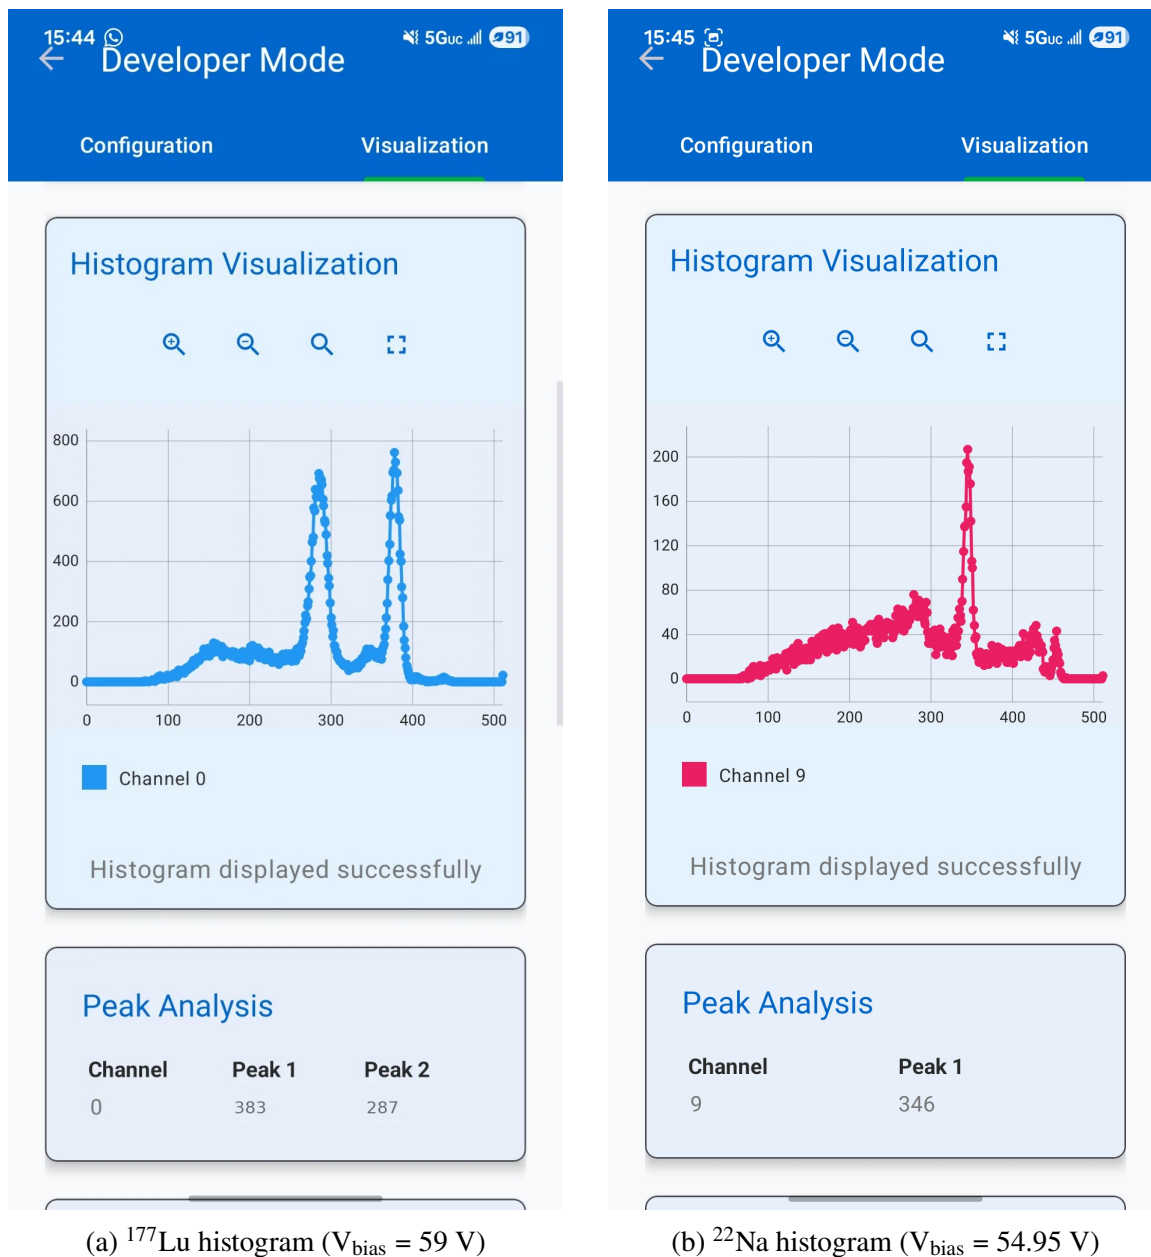


Figure 3.6: Mobile application histogram displays for radioactive source measurements. (a) ^{177}Lu showing characteristic photopeaks at 113 keV (bin 287) and 208 keV (bin 383). (b) ^{22}Na showing the 511 keV annihilation peak (bin 346). The Peak Analysis section displays automatically identified peak locations.

peak at bin index 346 corresponds to the 511 keV annihilation photons. The lower bias voltage compared to the ^{177}Lu measurement shifts the gain, positioning the higher-energy peak within the histogram range.

The successful detection of characteristic photopeaks from both isotopes validates the complete signal chain: gamma photon interaction in the GAGG:Ce scintillator, light detection by the SiPM, analog signal conditioning, time-over-threshold conversion, FPGA histogram accumulation, and wireless data transfer to the mobile application.

3.1.3.2 Tests in Temperature Controlled Environment

SiPM gain varies significantly with temperature due to the temperature dependence of the breakdown voltage (approximately $54 \text{ mV}/^\circ\text{C}$ for the Hamamatsu S13360 series). This gain variation causes photopeak positions to drift in the energy histogram, potentially affecting measurement accuracy in clinical environments where temperature may not be precisely controlled.

To characterize this effect, the PODD detector module was placed inside a temperature-controlled enclosure while the electronics remained at ambient temperature. Figure 3.7 shows the test configuration, with the detector connected to the PODD via ribbon cables routed into the temperature-controlled chamber.

Energy histograms were acquired using a ^{177}Lu source at four temperatures spanning the range expected in clinical environments: 19.70°C , 25.00°C , 30.20°C , and 35.30°C . Figure 3.8 presents the overlaid histograms from these measurements.

The results clearly demonstrate the temperature dependence of system gain. As temperature increases from 19.70°C to 35.30°C , both photopeaks shift toward lower histogram bin indices. The 113 keV peak migrates from bin 330 at the lowest temperature to bin 301 at the highest temperature, while the 208 keV peak shifts from bin 419 to bin 392 over the




Figure 3.7: Test setup for temperature-dependent measurements. The PODD electronics board operates at ambient temperature while the detector module is placed inside a temperature-controlled enclosure. Ribbon cables connect the detector to the analog front-end PCB.

same 15.6°C range. Linear regression analysis of peak centroid positions yields temperature coefficients of $-0.7\%/^{\circ}\text{C}$ for the 113 keV photopeak and $-0.5\%/^{\circ}\text{C}$ for the 208 keV photopeak.

The observed photopeak shifts can be correlated with the expected SiPM gain temperature dependence to validate the physical consistency of the measurements. For a silicon photomultiplier, the gain of each microcell is given by [1]

$$G = \frac{C_d \cdot V_{ov}}{q} \quad (3.1)$$

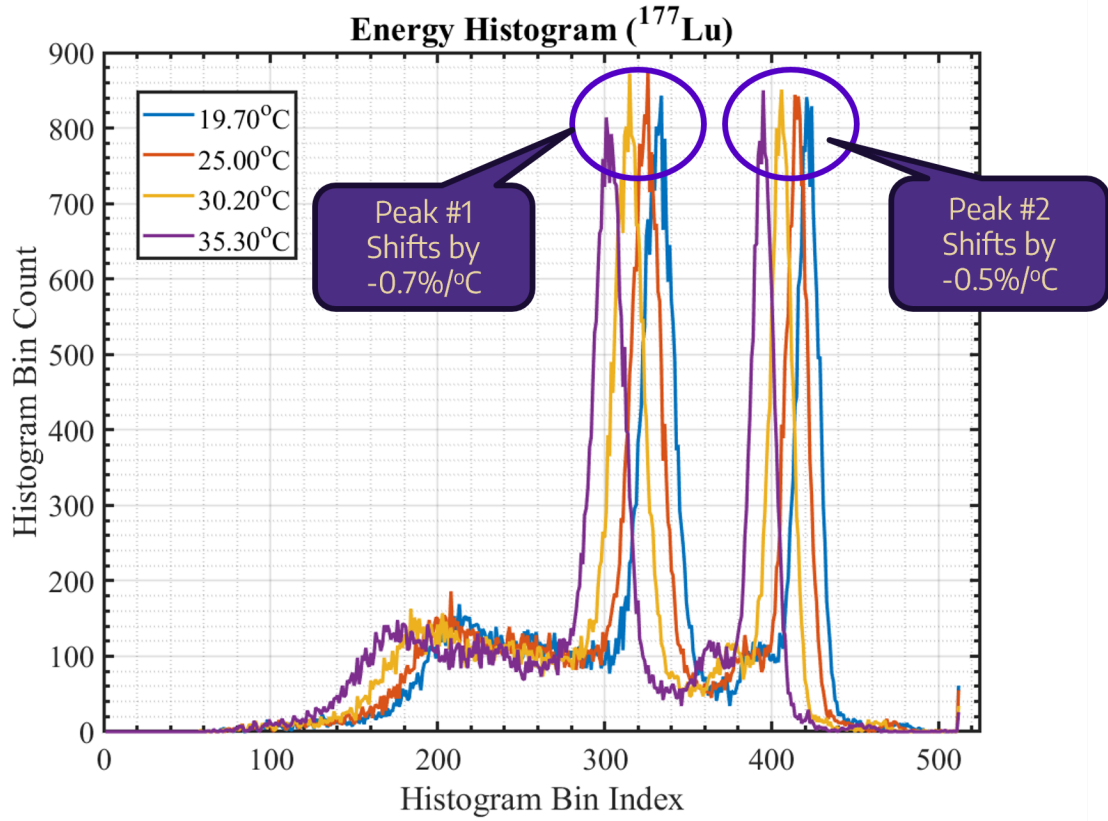


Figure 3.8: Energy histograms of ^{177}Lu acquired at four temperatures. As temperature increases, the SiPM gain decreases, causing photopeaks to shift toward lower bin indices. Linear regression analysis of peak centroid positions yields temperature coefficients of $-0.7\%/^{\circ}\text{C}$ (i.e., -2.3 bins/ $^{\circ}\text{C}$) for the 113 keV peak and $-0.5\%/^{\circ}\text{C}$ (i.e., -2.1 bins/ $^{\circ}\text{C}$) for the 208 keV peak.

where C_d is the microcell capacitance, q is the electron charge, and $V_{\text{ov}} = V_{\text{bias}} - V_{\text{BR}}$ is the overvoltage. Since the gain is directly proportional to the overvoltage, the fractional change in gain with temperature at a fixed bias voltage is

$$\frac{1}{G} \frac{dG}{dT} = \frac{1}{V_{\text{ov}}} \frac{dV_{\text{ov}}}{dT} = -\frac{1}{V_{\text{ov}}} \frac{dV_{\text{BR}}}{dT} \quad (3.2)$$

where the negative sign arises because V_{bias} is held constant while V_{BR} increases with

temperature, thereby reducing the overvoltage.

According to the manufacturer datasheet [9], the S13360-6050VE has a typical breakdown voltage of $V_{BR} = 53$ V (at 25°C) and a breakdown voltage temperature coefficient of $dV_{BR}/dT = +54$ mV/°C. For the bias voltage used in this experiment ($V_{bias} = 59$ V), the operating overvoltage is

$$V_{ov} = V_{bias} - V_{BR} = 59 - 53 = 6 \text{ V} \quad (3.3)$$

Substituting into Equation 3.2, the expected fractional gain change is

$$\frac{1}{G} \frac{dG}{dT} = -\frac{54 \text{ mV}/^{\circ}\text{C}}{6000 \text{ mV}} = -0.90\%/^{\circ}\text{C} \quad (3.4)$$

This value is consistent with the 1–3% per °C range reported by Otte [22] for SiPM gain temperature dependence. The fractional sensitivity decreases at higher overvoltages; at the recommended overvoltage of 3 V, the same calculation yields $-1.8\%/^{\circ}\text{C}$, which falls near the center of the reported range.

The experimentally observed photopeak shifts of $-0.7\%/^{\circ}\text{C}$ (113 keV) and $-0.5\%/^{\circ}\text{C}$ (208 keV) are both smaller in magnitude than the calculated gain coefficient of $-0.90\%/^{\circ}\text{C}$. This attenuation is attributable to the nonlinear relationship between SiPM gain and the measured time-over-threshold pulse width. Because the ToT method measures the duration a scintillation pulse remains above a fixed voltage threshold, the mapping from pulse amplitude to pulse width is monotonic but sub-linear and hence a given fractional change in pulse amplitude produces a smaller fractional change in pulse width. This compression effect is more pronounced for higher-energy pulses, which is consistent with the observation that the 208 keV peak exhibits a smaller fractional shift ($-0.5\%/^{\circ}\text{C}$) compared to the 113 keV peak ($-0.7\%/^{\circ}\text{C}$).

The consistency between the calculated SiPM gain temperature coefficient ($-0.90\%/^{\circ}\text{C}$) and the observed photopeak shifts (in both direction and order of magnitude) validates the temperature-dependent behavior of the detection system. The PODD's integrated thermistors (described in Section 2.3.2) provide the per-channel temperature data necessary for offline gain correction, enabling consistent energy measurements across varying environmental conditions.

Chapter 4

CONCLUSIONS

This thesis presented the design and validation of the PODD, a compact, wireless 16-channel gamma spectroscopy system for real-time monitoring of ^{177}Lu radio-pharmaceutical therapy. The work addressed the need for accessible dosimetry tools that can enable personalized treatment optimization at the point of care.

The PODD integrates SiPM-based scintillation detectors, analog signal conditioning with time-over-threshold encoding, FPGA-based histogram computation, and BLE wireless connectivity into a three-board sandwich architecture controlled through an Android mobile application. Full system testing with ^{177}Lu and ^{22}Na radioactive sources produced energy histograms with clearly resolved photopeaks, validating the complete signal chain from gamma detection through wireless data visualization.

Temperature-controlled measurements characterized the system's gain dependence, demonstrating the importance of the integrated thermistor array for accurate energy calibration across clinical operating conditions. The PODD's compact form factor, wireless operation, and intuitive interface address key barriers to widespread dosimetry adoption, representing a step toward accessible radio-pharmaceutical therapy monitoring.

To support reproducibility and further development, the complete project codebase, including FPGA Verilog source files, PSoC microcontroller firmware, and Android mobile application, is available on GitHub at <https://github.com/uw-acme/PRRT-POD1>. The PCB layout files for all boards are hosted on Altium 365 at <https://ethan-ingalls.365.altium.com/designs>.

Chapter 5

FUTURE WORK

The following enhancements are identified as future work for this project, with the aim of improving the PODD's clinical utility.

The first enhancement involves implementing real-time temperature compensation. Currently, thermistor data is recorded for offline gain correction. Future firmware development would implement real-time compensation by dynamically adjusting SiPM bias voltage or applying bin-shifting corrections based on continuous temperature readings, ensuring stable photopeak positions without post-processing.

The second enhancement involves extending the detector arrangement to a two-dimensional imaging array. The current 16-channel linear arrangement supports single-point measurements. Extending to a two-dimensional detector grid would enable spatial dose mapping, visualizing activity distribution across the measurement field for applications such as tumor dose heterogeneity assessment.

BIBLIOGRAPHY

- [1] Fabio Acerbi and Stefan Gundacker. Understanding and simulating SiPMs. *Nuclear Instruments and Methods in Physics Research Section A*, 926:16–35, 2019.
- [2] Bluetooth SIG. *Bluetooth Core Specification Version 5.0*. Bluetooth Special Interest Group, 2016. Available at: <https://www.bluetooth.com/specifications/specs/core-specification-5-0/>.
- [3] Lisa Bodei, Ken Herrmann, Heiko Schöder, Andrew M. Scott, and Jason S. Lewis. Radiotheranostics in oncology: current challenges and emerging opportunities. *Nature Reviews Clinical Oncology*, 19(8):534–550, 2022.
- [4] Werner Boeglin. Scintillation detector – modern lab experiments documentation. https://wanda.fiu.edu/boeglinw/courses/Modern_lab_manual3/scintillator.html, 2025. Florida International University, Department of Physics.
- [5] Miodrag Bolić and Vojislav Drndarević. Digital gamma-ray spectroscopy based on FPGA technology. *Nuclear Instruments and Methods in Physics Research Section A*, 482(1–2):761–766, 2002.
- [6] Clifford E. Cummings. Clock domain crossing (CDC) design & verification techniques using SystemVerilog. In *SNUG Boston*. Synopsys Users Group, 2008.
- [7] Jonathan Gear, Lois Boersma, Michael Lassmann, and Glenn D. Flux. Dosimetry in radiopharmaceutical therapy. *Journal of Nuclear Medicine*, 63(10):1467–1474, 2022.
- [8] Andrew Greensted. Comparator hysteresis network – the lab book pages. <http://www.labbookpages.co.uk/electronics/resNetworks/comparator.html>, 2014.
- [9] Hamamatsu Photonics K.K. *MPPC (Multi-Pixel Photon Counter) S13360 Series Datasheet*, 2022. Cat. No. KAPD1052E05.
- [10] Ulrike Hennrich and Klaus Kopka. Lutathera®: The first FDA- and EMA-approved radiopharmaceutical for peptide receptor radionuclide therapy. *Pharmaceuticals*, 12(3):114, 2019.

- [11] Ken Herrmann, Markus Schwaiger, Jason S. Lewis, Samuel B. Solomon, Barbara J. McNeil, Michael Baumann, Sanjiv S. Gambhir, Hedvig Hricak, and Ralph Weissleder. Radiotheranostics: a roadmap for future development. *The Lancet Oncology*, 21(3):e146–e156, 2020.
- [12] Infineon Technologies AG. *PSoC 6 MCU: CY8C62x8, CY8C62xA Datasheet*, 2021. Document No. 002-24899.
- [13] Intel Corporation. *Cyclone IV Device Handbook*, 2020. Document ID: CV-5V1.
- [14] Jan Jakubek. Precise energy calibration of pixel detector working in time-over-threshold mode. *Nuclear Instruments and Methods in Physics Research Section A*, 633:S262–S266, 2011.
- [15] Kei Kamada, Takayuki Yanagida, Toshihisa Endo, Kazuhisa Tsutumi, Yuki Usuki, Martin Nikl, Yasuhiro Fujimoto, Akihiro Fukabori, and Akira Yoshikawa. Composition engineering in cerium-doped $(\text{Lu,Gd})_3(\text{Ga,Al})_5\text{O}_{12}$ single-crystal scintillators. *Crystal Growth & Design*, 12(3):1418–1422, 2012.
- [16] Glenn F. Knoll. *Radiation Detection and Measurement*. John Wiley & Sons, Hoboken, NJ, 4th edition, 2010.
- [17] Michael Ljungberg, Anna Celler, Mark W. Konijnenberg, Keith F. Eckerman, Yuni K. Dewaraja, and Katarina Sjögren-Gleisner. MIRD pamphlet no. 26: Joint EANM/MIRD guidelines for quantitative ^{177}Lu SPECT applied for dosimetry of radiopharmaceutical therapy. *Journal of Nuclear Medicine*, 57(1):151–162, 2016.
- [18] Robert Miyaoka, Larry Pierce, Robert Harrison, Don DeWitt, Lindsey Hauck, and Hubert Vesselle. The personalized remote radiation tracking (PRRT) vest: experimental results. *Journal of Nuclear Medicine*, 62(supplement 1):1139, 2021.
- [19] Robert Miyaoka, Larry Pierce, Robert Harrison, and Hubert Vesselle. Wearable technology to enable personalization of Lu^{177} -DOTATATE therapy for neuroendocrine tumor patients. *Journal of Nuclear Medicine*, 60(supplement 1):313, 2019.
- [20] Kazuma Ogawa. Recent development of probes for radiotheranostics. *Analytical Sciences*, 40(1):1–2, 2024.
- [21] Toshiki Orita, Kenji Shimazoe, and Hiroyuki Takahashi. The current mode time-over-threshold ASIC for a MPPC module in a TOF-PET system. *Nuclear Instruments and Methods in Physics Research Section A*, 912:303–308, 2018.

- [22] A. Nepomuk Otte, Distefano Garcia, Thanh Nguyen, and Dhruv Purushotham. Characterization of three high efficiency and blue sensitive silicon photomultipliers. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 846:106–125, 2017.
- [23] Steffie M. B. Peters, Sebastiaan L. Meyer Viol, Nick R. van der Werf, Niels de Jong, Floris H. P. van Velden, Natasha Freedman, et al. A multicentre and multi-national evaluation of the accuracy of quantitative Lu-177 SPECT/CT imaging performed within the MRTDosimetry project. *EJNMMI Physics*, 8:55, 2021.
- [24] Slawomir Piatek. What is an SiPM and how does it work? Hamamatsu Photonics Technical Note, 2016. Available at: <https://hub.hamamatsu.com/us/en/technical-notes/mppc-sipms/what-is-an-SiPM-and-how-does-it-work.html>.
- [25] Mattias Sandström, Ulrike Garske-Román, Dan Granberg, Silvia Johansson, Christer Widström, Barbro Eriksson, Anders Sundin, Hans Lundqvist, and Mark Lubberink. Individualized dosimetry of kidney and bone marrow in patients undergoing ^{177}Lu -DOTA-octreotate treatment. *Journal of Nuclear Medicine*, 54(1):33–41, 2013.
- [26] George Sgouros, Lisa Bodei, Michael R. McDevitt, and Jessie R. Nedrow. Radiopharmaceutical therapy in cancer: clinical advances and challenges. *Nature Reviews Drug Discovery*, 19(9):589–608, 2020.
- [27] Kenji Shimazoe, Hiroyuki Takahashi, Bowen Shi, Toshiki Orita, Tetsuro Furumiya, Jun Ooi, Kenichi Kato, Keisuke Kamada, and Akira Yoshikawa. Dynamic time over threshold method. *IEEE Transactions on Nuclear Science*, 59(6):3213–3217, 2012.
- [28] Katarina Sjögren Gleisner, Nicolas Chouin, Pablo M. Gabina, Francesco Cicone, Silvano Gnesin, Caroline Stokke, et al. EANM dosimetry committee recommendations for dosimetry of ^{177}Lu -labelled somatostatin-receptor- and PSMA-targeting ligands. *European Journal of Nuclear Medicine and Molecular Imaging*, 49(6):1778–1809, 2022.
- [29] Jonathan Strosberg, George El-Haddad, Edward Wolin, Andrew Hendifar, James Yao, Beth Chasen, Erik Mittra, Pamela L. Kunz, Matthew H. Kulke, Heather Jacene, et al. Phase 3 trial of ^{177}Lu -Dotatate for midgut neuroendocrine tumors. *New England Journal of Medicine*, 376(2):125–135, 2017.
- [30] Jonathan R. Strosberg, Martyn E. Caplin, Pamela L. Kunz, Philippe B. Ruszniewski, Lisa Bodei, Andrew Hendifar, Erik Mittra, Edward M. Wolin, James C. Yao, Marianne E. Pavel, et al. ^{177}Lu -Dotatate plus long-acting octreotide versus high-dose

long-acting octreotide in patients with midgut neuroendocrine tumours (NETTER-1): final overall survival and long-term safety results from an open-label, randomised, controlled, phase 3 trial. *The Lancet Oncology*, 22(12):1752–1763, 2021.

[31] Terasic Technologies. *DEO-Nano User Manual*, 2012. Version 1.9.

Appendix A

ANALOG CIRCUIT DESIGN DETAILS

This appendix provides detailed circuit schematics for the analog front-end board, covering the microcontroller interface, power supply generation, and thermistor signal conditioning subsystems.

A.1 *Microcontroller Connector Circuitry*

Figure A.1 presents the interface circuitry between the microcontroller board and the analog front-end PCB. Two MIKROE-1324 headers (H1 and H2) provide connectivity for control voltages, sensor readback signals, and communication lines.

The interface carries several categories of signals:

Programmable Control Voltages. The high-voltage programming signal (HV_PGM) controls the SiPM bias voltage via the SIP90 DC-DC converter. This signal passes through an RC low-pass filter ($R9 = 1\text{ k}\Omega$, $C2 = 10\text{ }\mu\text{F}$) to attenuate PWM ripple, with a Schottky diode (D1, B5819W) providing overvoltage protection. Similarly, the hysteresis control voltage (V_HYS) and threshold reference voltage (V_REF) are filtered through RC networks to produce stable analog levels from PWM outputs.

Voltage Sense Lines. The SiPM bias voltage (V_BIAS) is scaled through a resistor divider ($R14 = 300\text{ k}\Omega$, $R15 = 10\text{ k}\Omega$) to produce V_BIAS_SENSE, enabling the microcontroller's ADC to monitor the high-voltage rail. The system supply voltage (V_SYS) is similarly monitored through V_SYS_SENSE.

Thermistor Multiplexer Control. Digital control lines (MUX_S0, MUX_S1,

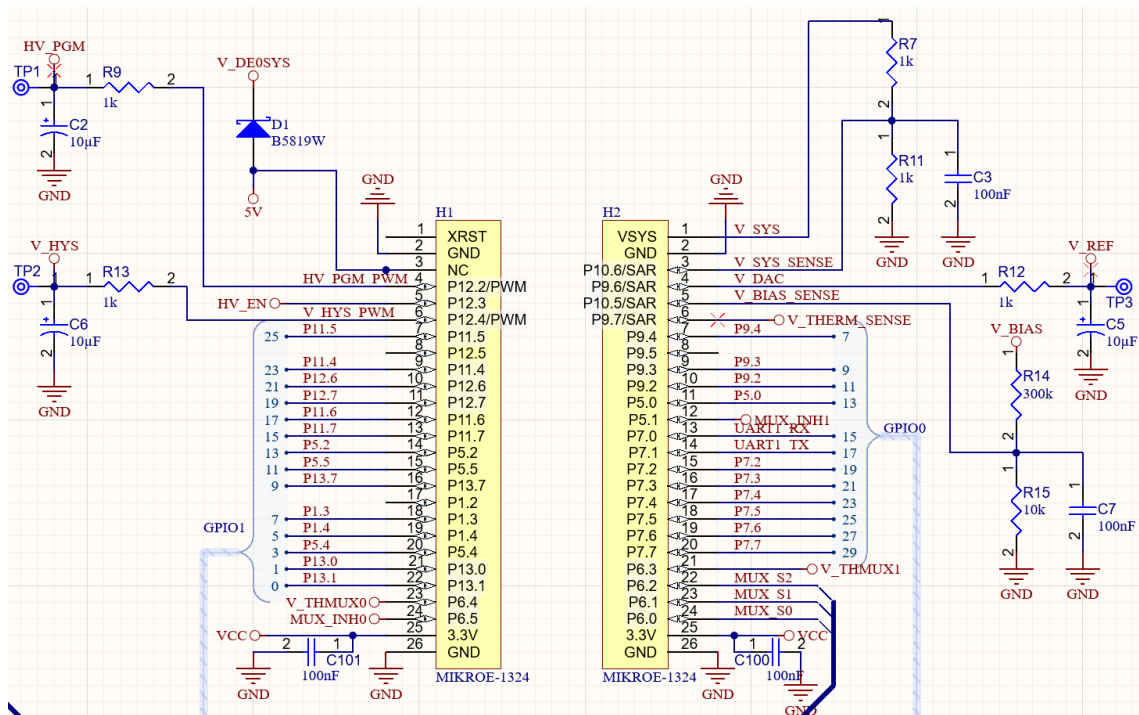


Figure A.1: Microcontroller interface circuitry showing the two board-to-board headers (H1, H2) and associated signal conditioning. Control voltages (HV_PGM, V_HYS, V_REF) pass through RC low-pass filters for noise reduction. Voltage sense lines (V_BIAS_SENSE, V_SYS_SENSE) use resistor dividers to scale high voltages for ADC measurement. Analog multiplexer control signals (MUX_S0–S2, MUX_INH0/1) and UART lines provide digital connectivity.

MUX_S2) select among the eight thermistor channels on each multiplexer, while inhibit signals (MUX_INH0, MUX_INH1) enable or disable the multiplexer outputs. The selected thermistor voltages (V_THMUX0, V_THMUX1) are routed to the microcontroller's ADC inputs.

Communication. UART transmit and receive lines (UART1_TX, UART1_RX) provide the serial communication path to the FPGA board.

A.2 Power Supply Circuitry

Figure A.2 shows the power generation section that produces the SiPM bias voltage and negative supply rail for the operational amplifiers.

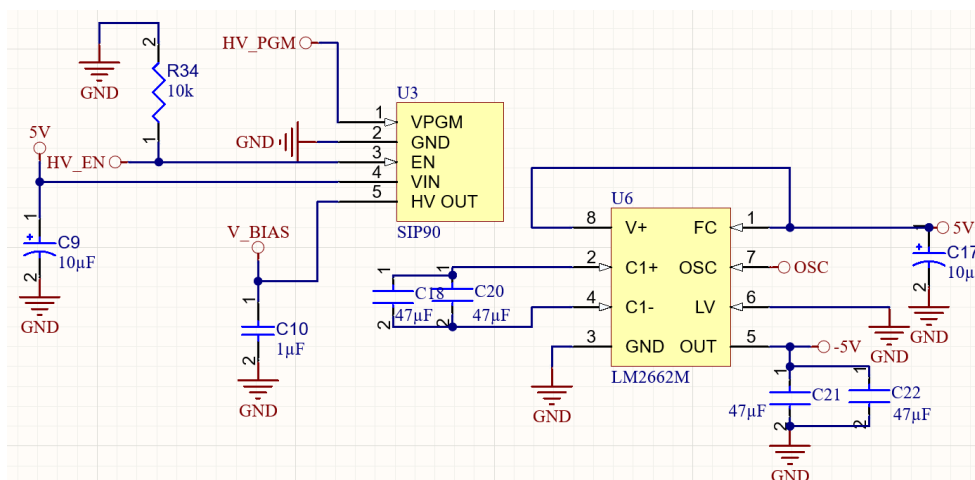


Figure A.2: Power supply circuitry featuring the EMCO SIP90 programmable high-voltage DC-DC converter (U3) and the LM2662M switched-capacitor voltage inverter (U6). The SIP90 generates the SiPM bias voltage (V_{BIAS} , 25–90 V) from the 5 V input, controlled by the VPGM input. The LM2662M produces the -5 V rail required by the operational amplifiers.

High-Voltage Bias Generation. The EMCO SIP90 (U3) is a compact DC-DC converter that generates the programmable SiPM bias voltage. The module accepts a 5 V input (V_{IN}) and produces an output (HV OUT, labeled V_{BIAS}) that varies inversely with the programming input voltage (VPGM). With $V_{PGM} = 0$ V, the output reaches approximately 90 V; with $V_{PGM} = 5$ V, the output drops to approximately 25 V. The enable input (EN) allows software control of the high-voltage supply via the HV_EN signal. Input and output capacitors (C9, C10, C18, C20) provide filtering and energy storage. A 10 k Ω pull-down resistor (R34) ensures the VPGM input does not float when the microcontroller output is in a high-impedance state.

Design Limitation. The current board configuration does not allow the full output voltage range of the SIP90 to be achieved. The programming voltage (HV_PGM) is generated by a PWM output from the PSoC 6 microcontroller, which operates at 3.3 V logic levels. Since the SIP90 requires a 0–5 V input at VPGM for full output range, the 3.3 V maximum from the microcontroller limits the achievable programming range to approximately 0–3.3 V. Given the inverse voltage relationship, this restricts the minimum achievable bias voltage to approximately 47 V (corresponding to VPGM = 3.3 V), rather than the 25 V minimum specified by the SIP90. The upper output range (toward 90 V at VPGM = 0 V) remains fully accessible. For the target SiPM operating range of 50–65 V, this limitation does not affect normal operation, but prevents bias voltages below approximately 47 V.

Negative Supply Generation. The LM2662M (U6) is a switched-capacitor voltage inverter that generates the –5 V rail required by the LT1818 operational amplifiers in the signal conditioning channels. External capacitors (C18, C20, C21, C22, each 47 μ F) serve as the flying capacitor and output filter. This inductor-free topology minimizes electromagnetic interference in the sensitive analog section.

A.3 *Thermistor and Detector Connector Circuitry*

Figure A.3 presents the interface between the analog board and the external detector array, including the thermistor signal conditioning and multiplexing circuitry.

Detector Connector Interface. The 50-pin ribbon cable connector (J1) carries signals to and from eight detector modules. Each detector channel includes: the SiPM analog output (ADC_INx) routed to the corresponding signal conditioning channel, the high-voltage bias (V_BIAS) distributed to all SiPMs, the thermistor sense line (THERMx) for temperature monitoring, and ground references.

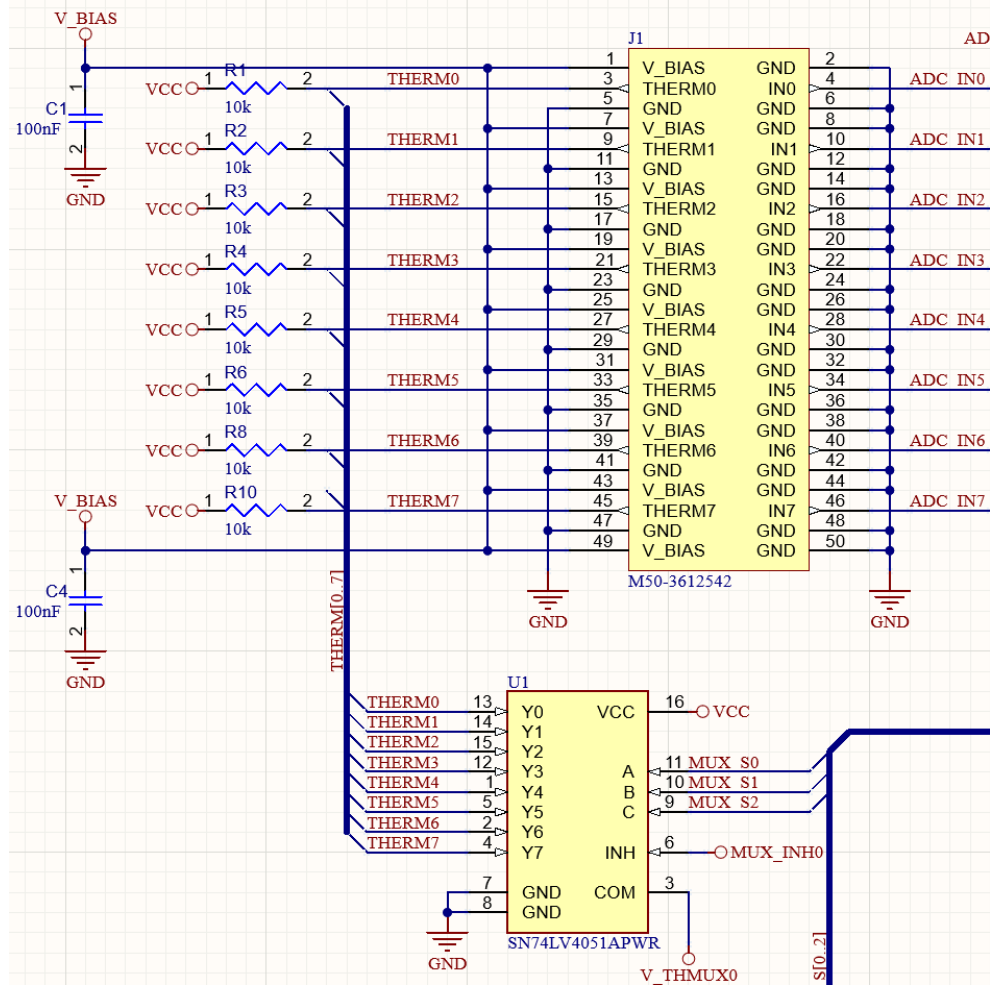


Figure A.3: Thermistor signal conditioning and detector interface circuitry. The 50-pin connector (J1, M50-3612542) interfaces with the detector ribbon cable, carrying SiPM signal inputs (ADC_IN0–7), bias voltage distribution (V_BIAS), and thermistor sense lines (THERM0–7). The resistor divider network (R1–R10, each 10 k Ω) forms the upper half of a voltage divider with the NTC thermistors located on the detector probe PCBs. An 8:1 analog multiplexer (U1, SN74LV4051APWR) routes the selected thermistor voltage to the microcontroller ADC.

Thermistor Voltage Divider Network. Each thermistor channel employs a voltage divider topology for temperature measurement. The upper resistors (R1–R10, each 10 k Ω) are located on the analog board and connect to VCC. The lower element of each divider is the NTC thermistor (RT1 in the detector probe schematic, Figure 2.5), which is physically located on the detector probe PCB rather than the analog board. This arrangement allows the temperature sensor to be co-located with the SiPM for accurate thermal monitoring. The voltage at the divider midpoint (THERM_x) varies with thermistor resistance, which in turn depends on temperature.

Analog Multiplexing. To minimize the number of ADC channels required on the microcontroller, an 8:1 analog multiplexer (U1, SN74LV4051APWR) routes the eight thermistor voltages to a single ADC input (V_THMUX0). The microcontroller selects the active channel via three address lines (MUX_S0, MUX_S1, MUX_S2) and enables the multiplexer output via the inhibit line (MUX_INH0). The firmware sequentially scans all eight channels to acquire the complete temperature profile. A second multiplexer (not shown in this view) handles thermistor channels 8–15 from the second detector connector, with its output routed to V_THMUX1.

Appendix B

UART HANDSHAKE PROTOCOL SPECIFICS

This appendix presents the command and status codes used in the FPGA-microcontroller UART communication protocol. The protocol employs a command-response architecture where each command from the microcontroller is terminated by an `END_COMMAND` byte, and the FPGA responds with appropriate acknowledgments or status codes.

B.1 Command Codes

Table B.1 lists the supported UART commands transmitted from the microcontroller to the FPGA. Each command is encoded as an 8-bit value and must be followed by an `END_COMMAND` delimiter (0xFF) to complete the transaction.

B.2 Status Codes

Table B.2 presents the status codes transmitted from the FPGA to the microcontroller in response to status queries or to indicate operational state changes.

B.3 Protocol Operation

Commands requiring parameter values use multi-byte sequences. For example, the `SET_BIN_ADDRESS` command requires a 9-bit address (0–511), transmitted as two bytes (LSB first, then MSB) followed by `END_COMMAND`. The complete transaction format is:

<COMMAND_BYTE> | <PARAM_LSB> | <PARAM_MSB> | <END_COMMAND>

Table B.1: UART command codes for FPGA-microcontroller communication.

Command	Code (Binary)	Purpose
END_COMMAND	1111_1111	Terminates the command transfer to the FPGA.
FPGA_VERSION	0000_0001	Requests the FPGA hardware version.
START_HISTOGRAM	0000_0010	Starts the histogram capture process on the FPGA.
STOP_HISTOGRAM	0000_0011	Stops the histogram capture process on the FPGA.
CLEAR_RESULTS	0000_0100	Clears the bin data for all channels on the FPGA.
START_UPLOAD	0000_0101	Initiates histogram data transfer from FPGA to MCU for the requested settings.
SET_BIN_ADDRESS	0000_0110	Sets the starting bin address for data upload to MCU.
SET_NUM_BINS	0000_0111	Sets the number of bins for data upload to MCU.
SET_CHANNEL	0000_1000	Sets the channel number for data upload to MCU.
IS_HIST_RUNNING	0000_1001	Queries the FPGA status; returns FPGA_RUNNING or FPGA_IDLE.

Table B.2: UART status codes returned by the FPGA.

Status	Code (Binary)	Purpose
FPGA_RUNNING	1111_0000	Indicates that FPGA histogram capture is in progress.
FPGA_IDLE	1111_0001	Indicates that FPGA histogram capture is halted/idle.

The FPGA acknowledges each valid command by responding with END_COMMAND (0xFF). During histogram upload, each 16-bit bin count is transmitted as two bytes, with the upper nibble of the second byte containing the channel identifier for data integrity verification.

Appendix C

MOBILE APPLICATION SPECIFICS

This appendix provides detailed documentation of additional mobile application screens for data visualization and system settings, along with the peak detection algorithm implementation. Figure C.1 presents the visualization and settings interfaces.

C.1 Histogram Visualization

The histogram visualization screen (Figure C.1a) enables real-time display of energy spectra from selected detector channels. Users can select which of the 16 channels to display, configure peak annotation to automatically identify and label characteristic photopeaks, and interact with the chart using pinch-to-zoom and pan gestures. The “Load CSV” option allows importing previously saved histogram data for comparison or offline analysis.

C.2 Temperature Profile Visualization

The temperature visualization screen (Figure C.1b) displays thermal data from the 16-channel thermistor array. Users can select specific thermistor channels for display, choose between Celsius and Fahrenheit units, and view running averages. The temperature profile chart shows temporal evolution of detector temperatures during acquisition, supporting identification of thermal drift that may affect SiPM gain calibration.

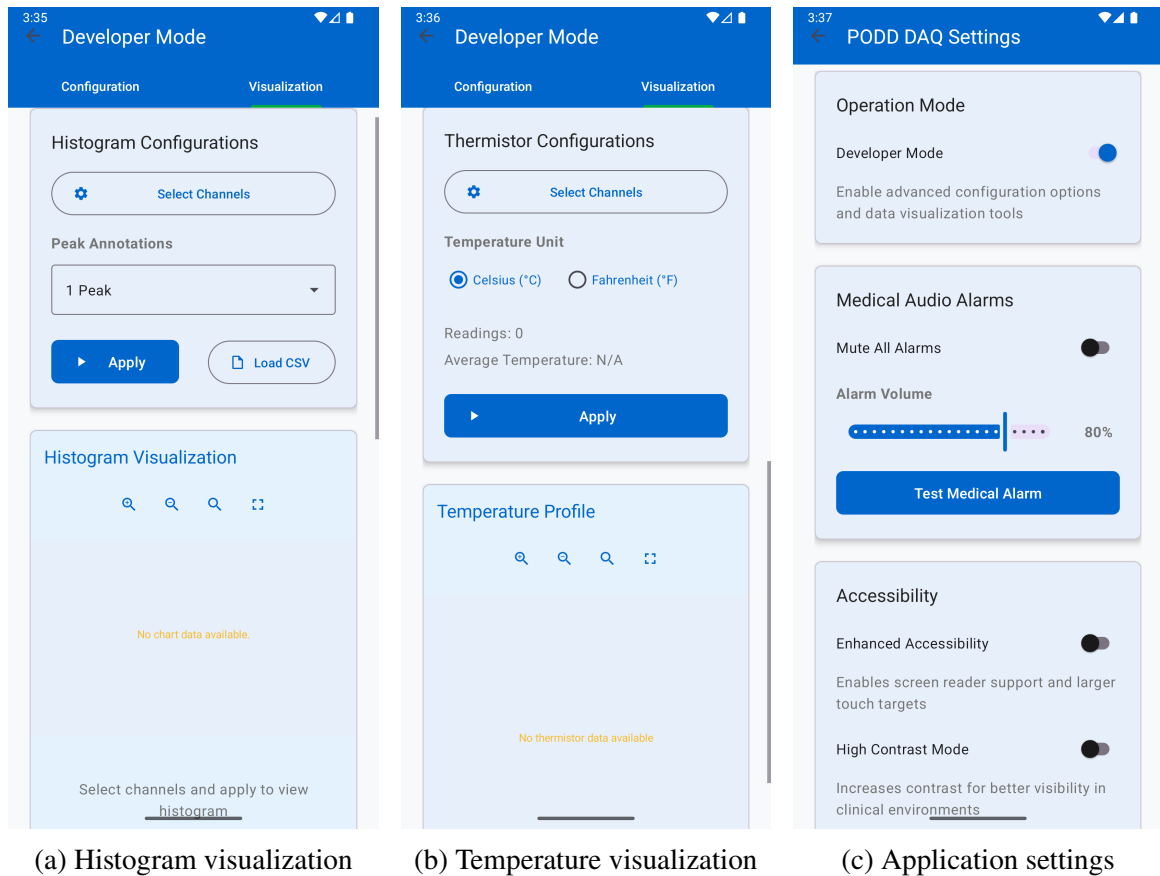


Figure C.1: Additional mobile application screens: (a) Histogram visualization tab with channel selection, peak annotation options, and interactive chart display; (b) Temperature profile visualization with thermistor channel selection and unit conversion; (c) Application settings for operation mode selection, medical audio alarms, and accessibility options.

C.3 Application Settings

The settings screen (Figure C.1c) provides system-wide configuration options. The Operation Mode toggle switches between Clinical mode (simplified interface) and Developer mode (full parameter access). Medical Audio Alarms settings control audible notifications with adjustable volume and a test function. Accessibility options include enhanced touch targets for users with motor impairments and high contrast mode for improved visibility in

clinical lighting environments.

C.4 Peak Detection Algorithm

The mobile application implements an automated peak detection algorithm to identify characteristic photopeaks in energy histograms. This capability enables automatic annotation of the 113 keV and 208 keV peaks from ^{177}Lu spectra, assisting users in verifying correct system operation and energy calibration.

C.4.1 Algorithm Overview

The algorithm is implemented in Kotlin within the `DeveloperViewModel` class. Listing C.1 presents the box-car filter implementation that applies a 5-bin moving average to reduce noise.

Listing C.1: Box-car filter for noise reduction.

```
private fun applyBoxCarFilter(data: ShortArray): IntArray {
    val filterWidthBins = 5
    val halfWidth = filterWidthBins / 2
    val smoothed = IntArray(data.size)

    for (i in data.indices) {
        var sum = 0
        var count = 0
        val start = maxOf(0, i - halfWidth)
        val end = minOf(data.size - 1, i + halfWidth)

        for (j in start..end) {
            sum += data[j].toInt()
        }
    }
}
```

```

        count++
    }
    smoothed[i] = if (count > 0) sum / count else 0
}
return smoothed
}

```

Listing C.2 shows the local maxima detection logic, which identifies bins exceeding both neighbors and applies a minimum height threshold of 2% of the spectrum maximum (or 10 counts, whichever is greater) to reject noise spikes.

Listing C.2: Local maxima detection with threshold filtering.

```

private fun findAllLocalMaxima(data: IntArray): List<Int> {
    val localMaxima = mutableListOf<Int>()
    val maxValue = data.maxOrNull() ?: 0
    val minPeakHeight = maxOf(10, (maxValue * 0.02).toInt())

    for (i in 1 until data.size - 1) {
        if (data[i] > data[i - 1] && data[i] > data[i + 1]) {
            if (data[i] >= minPeakHeight) {
                localMaxima.add(i)
            }
        }
    }

    return localMaxima
}

```

Listing C.3 shows the intensity sorting step, which converts local maxima indices to

PeakInfo objects and sorts them by count value in descending order. This ensures that the global maximum is processed first during peak separation. Note that the count values are taken from the original (unsmoothed) histogram data to preserve accurate intensity measurements, since the smoothing operation is used solely for robust peak location identification, not for quantification.

Listing C.3: Converting local maxima to PeakInfo and sorting by intensity.

```
// Convert to PeakInfo and sort by intensity (descending)
val candidatePeaks = localMaxima.map { binIndex ->
    PeakInfo(
        binIndex = binIndex,
        binValue = binIndex,
        count = data[binIndex].toInt(),
        percentage = (data[binIndex] / total) * 100f,
        peakNumber = 0,
        isLocalMaximum = true
    )
}.sortedByDescending { it.count }
```

Listing C.4 presents the peak separation algorithm, which ensures that identified peaks are at least 15 bins apart. Since candidates are sorted by intensity (descending), the global maximum is accepted first. Subsequent candidates are only accepted if they are at least minSeparation bins away from all previously accepted peaks.

Listing C.4: Peak separation constraint enforcement.

```
private fun applyPeakSeparation(
    candidatePeaks: List<PeakInfo>,
    minSeparation: Int
```

```
) : List<PeakInfo> {  
    val acceptedPeaks = mutableListOf<PeakInfo>()  
  
    for (candidate in candidatePeaks) {  
        var tooClose = false  
        for (accepted in acceptedPeaks) {  
            val distance = kotlin.math.abs(  
                candidate.binIndex - accepted.binIndex  
            )  
            if (distance < minSeparation) {  
                tooClose = true  
                break  
            }  
        }  
        if (!tooClose) {  
            acceptedPeaks.add(candidate)  
        }  
    }  
    return acceptedPeaks  
}
```

C.4.2 Peak Information Data Structure

Each detected peak is represented by a `PeakInfo` data class containing the bin index, count value, percentage of total counts, and peak ranking number. This information is displayed in the Peak Analysis section of the histogram visualization screen, enabling users to quickly identify and verify characteristic energy peaks.

Listing C.5: PeakInfo data class for peak representation.

```
data class PeakInfo(  
    val binIndex: Int,  
    val binValue: Int,  
    val count: Int,  
    val percentage: Float,  
    val peakNumber: Int = 0,  
    val isLocalMaximum: Boolean = false  
)
```

The 15-bin minimum separation distance was chosen empirically to prevent detection of multiple peaks within the same photopeak region while allowing resolution of the ^{177}Lu 113 keV and 208 keV peaks, which are typically separated by approximately 100 bins in the PODB histogram.

Appendix D

PODD SCAN STARTUP AND OPERATIONAL SEQUENCE

This appendix documents the complete startup and operational procedure for performing a radiation scan with the PODD system. It describes every stage from hardware power-on through data file generation, including the specific indicators at each stage that allow the patient to confirm correct system operation without technical assistance.

D.1 Stage 1: Hardware Power-On

The PODD hardware assembly consists of the three-board sandwich stack (Figure 2.2): the Altera DE0-Nano FPGA board (bottom), the analog front-end PCB (middle), and the Mikro-e Clicker 2 microcontroller board (top). The entire stack is powered through the Clicker 2 board's battery input, with its onboard voltage regulators distributing power to the FPGA and analog boards through the board-to-board headers.

Prerequisites. The patient must ensure the PODD unit and the Android phone are in the same room before powering on. The PSoC 6 BLE 5.0 radio supports a maximum transmit power of +4 dBm with a receiver sensitivity of -92 dBm on the LE 2M PHY [12], yielding a link budget of 96 dB that supports reliable indoor communication well beyond typical room dimensions [2]. For maximum reliability during the histogram data transfer phase, keeping the phone within a few meters of the PODD unit is recommended. The detector probe must be positioned on the patient according to the clinical setup instructions before powering on.

Indicators of successful power-on:

- The **power LED on the microcontroller board** illuminates steady, indicating the PSoC 6 is receiving power.
- The **power LED on the DE0-Nano FPGA board** illuminates steady, indicating the FPGA is powered and the configuration bitstream has been loaded from onboard EPCS flash memory.
- The **Bluetooth LED on the microcontroller board begins blinking** at a regular interval. This blink is driven by the PWM peripheral in the BLE advertising state handler (`PWM_Start()` called on `CY_BLE_EVT_GAPP_ADVERTISEMENT_START_STOP`), confirming that the BLE stack is active and the device is advertising under the name `PODD`.

If the Bluetooth LED does not begin blinking within a few seconds of power-on, the microcontroller has not completed its initialization. Power-cycle the device and retry.

D.2 Stage 2: Bluetooth Connection via the Mobile Application

With the PODD hardware advertising, the patient opens the PODD Manager application on their Android phone. The app validates its audio system on launch by verifying that the Android `ToneGenerator` is initialized and that the device volume is at least 20% of maximum, displaying a warning if either check fails.

Prerequisites. Bluetooth must be enabled on the phone. If it is not, the app will prompt the patient to enable it before proceeding. The app also requires Bluetooth scan and connect permissions; if not previously granted, a permission dialog will be presented.

Action. The patient taps the **Connect Device** button. The button label changes to “Searching for PODD. . .” and the app begins a BLE scan filtered by the device name

PODD. If the device is not found within 60 seconds, a timeout error is displayed and the scan terminates.

Connection sequence (internal). Upon locating the advertising PODD device, the app automatically stops the BLE scan and initiates a GATT connection. Once the physical link is established, the app requests a Maximum Transmission Unit (MTU) of 512 bytes to support efficient histogram data transfer in 256-byte notification packets. Following successful MTU negotiation, the app performs GATT service discovery on the PODD service (UUID: 0000B873), enumerates the Timer, Histogram, and Command characteristics, and sequentially enables notifications on the Timer and Histogram characteristics by writing to their respective Client Characteristic Configuration Descriptor (CCCD) descriptors. Finally, the app sends a settings query command ('q') to retrieve and display the current operating parameters (bias voltage, threshold, hysteresis, scan timer) without requiring manual entry.

Indicators of successful connection:

- The **device status text** in the app changes from “Not connected” to “Connected to PODD” and its color changes to green.
- The **Connect Device** button label changes to “Connected” and becomes disabled, preventing duplicate connection attempts.
- The app plays a **medium-priority connection confirmation tone** through the phone speaker, consisting of three short pulses (200 ms on, 100 ms off each), repeated twice with a 1-second pause between repetitions. The tones are generated via Android’s ToneGenerator API.
- A persistent **foreground service notification** appears in the phone’s notification shade (labeled “Connected to PODD device”), confirming the BLE connection is maintained

even if the app is running in the background.

- On the microcontroller, the **Bluetooth LED stops blinking and turns steady on**, driven by `PWM_Disable()` followed by `Cy_GPIO_Write(LED_ConnectStatus_0_PORT, ..., 0)` on the `CY_BLE_EVT_GAP_ENHANCE_CONN_COMPLETE` event.
- The **Start Scan** button in the app becomes enabled, indicating the system is ready to accept a scan command.

If the connection is not established within 60 seconds after tapping Connect Device, the app displays a “Device Not Found” error dialog and plays a medium-priority connection warning tone. The patient should verify that the PODD Bluetooth LED is still blinking and retry.

D.3 Stage 3: Scan Initiation

The patient taps the **Start Scan** button. Before initiating acquisition, the app automatically transmits a configuration command to the microcontroller over the Command BLE characteristic. This command encodes the scan duration (in seconds), SiPM bias voltage, comparator threshold, and hysteresis voltage in the format:

$$\text{CONFIG:TIMER}=\mathit{t},\text{BIAS}=\mathit{v},\text{THRESH}=\mathit{th},\text{HYST}=\mathit{hy}$$

The microcontroller parses and applies each parameter: the timer value is set via `setTimerAbsolute()`, the bias voltage via `setBiasVoltageAbsolute()` which programs the EMCO SIP90 DC-DC converter, the threshold voltage via `setThresholdPercentAbsolute()` which programs the 12-bit VDAC, and the hysteresis voltage via `setHysteresisPercentAbsolute()` which configures the filtered PWM output. The microcontroller responds with `CONFIG_APPLIED` on success or

CONFIG_ERROR: . . . identifying the failed parameter. Following successful configuration, the app sends the single-byte start command ('g') over the Command characteristic, which is dispatched to the microcontroller's FreeRTOS command queue, triggering the FPGA to begin histogram acquisition via the UART interface and starting the countdown timer.

Indicators that the scan is in progress:

- The app plays a **high-priority scan-started alert** consisting of 10 rapid pulses (150 ms on, 50 ms off) repeated three times with 1-second pauses between repetitions, using a high-frequency system beep to draw the patient's attention. After a 1-second delay, a **periodic heartbeat sound** begins (a two-pulse lub-dub pattern repeating every 2 seconds) providing continuous auditory confirmation that acquisition is ongoing.
- The **countdown timer** in the app updates each second to display the remaining scan time in MM:SS format, fed by GATT timer notifications arriving from the microcontroller's FreeRTOS timer task.
- A **progress bar** tracks the proportion of elapsed scan time.
- The status text changes to "Scan in progress. . .".
- The **Start Scan** button is disabled and the **Stop Scan** button becomes enabled.
- A **real-time temperature chart** appears on the home screen once the first thermistor data packet arrives. The chart plots temperature readings from all 16 thermistors as individual color-coded traces (representing temperature in °C vs. elapsed scan time in seconds), updating continuously as the microcontroller's thermistor task transmits new 64-byte BLE notification packets at its polling interval.

The patient should not move the detector during acquisition. If the Bluetooth connection is lost during a scan, the app plays a **connection-lost alarm** consisting of a three-tone descending sequence (200 ms per tone with 100 ms gaps), stops the heartbeat monitor and acquisition immediately, and displays a critical disconnection dialog titled “CRITICAL MEDICAL ERROR” requiring acknowledgment before the system returns to the ready state.

D.4 Stage 4: Scan Completion and Data Transfer

When the countdown timer reaches zero, the microcontroller’s timer task signals the FPGA to stop acquisition via the UART interface and initiates the histogram data upload sequence. The FPGA transmits the histogram data over UART to the microcontroller, which then transfers it to the phone over BLE as a sequence of 256-byte GATT notification packets (16 channels \times 512 bins \times 2 bytes = 16,384 bytes total, sent in 64 packets). A 4-byte completion signal is sent upon transfer completion.

Indicators of scan completion and transfer:

- The **heartbeat sound stops** immediately when the scan completes.
- The app plays a **scan-completion sound** consisting of a three-tone sequence (250 ms per tone with 75 ms gaps) using distinct system tones that step from a mid-frequency tone to a higher-mid tone and back, signaling successful acquisition.
- The **countdown timer displays 00:00**.
- The status text changes to “Scan completed. Transferring data. . .”.
- A **data transfer progress bar** becomes visible and advances from 0% to 100% as the histogram bytes are received.

- Upon receiving all histogram data and the 4-byte completion signal, the app plays a **low-priority data transfer confirmation tone** consisting of two slow pulses (300 ms on, 200 ms off), then parses the 16 channel histograms and saves them to a **CSV file** named `PODD_YYYYMMDD_HHMMSS.csv` in the phone's Documents directory. If thermistor data was recorded during the scan, a separate file `PODD_YYYYMMDD_HHMMSS_thermistor.csv` is also saved.
- A **toast notification** appears at the bottom of the screen confirming the save (e.g., "Scan saved at `PODD_20260227_143022.csv`").
- The temperature chart, progress bars, and timer reset to their idle state. The **Start Scan** button re-enables, indicating the system is ready for the next scan.

If the data transfer does not complete within 15 seconds of scan completion, the app automatically resets to the ready state. The patient should attempt the scan again or contact their clinician.

D.5 Summary of Patient-Observable Indicators

Table D.1 summarizes all patient-observable indicators at each stage of the startup and scan sequence.

Table D.1: Patient-observable indicators at each stage of the PODD scan procedure.

Stage	Indicator	Expected Observation
Power-On	MCU power LED	Steady on
	FPGA power LED	Steady on
	BLE LED (MCU)	Blinking at regular interval
BLE Connec- tion	App status text	“Connected to PODD” (green)
	App button	Connect button changes to “Connected”
	App audio	Medium-priority tone (3 short pulses, repeated twice)
	Phone notification shade	Persistent “Connected to PODD device” notification
Scan In Progress	BLE LED (MCU)	Blinking stops; LED turns steady on
	App timer	Counts down from set duration in MM:SS
	App audio	10-pulse alert \times 3 on start; then a heartbeat every 2 s
	App status text	“Scan in progress. . .”
Scan Complete / Data Transfer	Temperature chart	Live multi-channel temperature traces visible
	App timer	Displays 00:00
	App audio	Heartbeat stops; 3-tone completion sequence; then 2-pulse transfer tone
	App status text	“Scan completed. Transferring data. . .”
	Transfer progress bar	Advances 0% \rightarrow 100%
	Toast notification	“Scan saved at PODD_YYYYMMDD_HHMMSS.csv”